Spore in the Wild: Case Study on Spore.fun, a Real-World Experiment of Sovereign Agent Open-ended Evolution on Blockchain and TEEs

Botao Amber Hu¹, Helena Rong² ¹Reality Design Lab ²New York University Shanghai botao@reality.design hr2703@nyu.edu

Abstract

In Artificial Life (ALife) research, replicating Open-Ended Evolution (OEE)—the continuous emergence of novelty observed in biological life-has traditionally been pursued within isolated closed system simulations, such as Tierra and Avida, which have typically plateaued after an initial burst of novelty, failing to achieve sustained OEE. Scholars suggest that OEE requires an "open" system that continually exchanges information or energy with its environment. A recent technological innovation in decentralized physical infrastructure networks (DePIN) providing permissionless computational substrates enables deploying large language model (LLM)-based AI agents on blockchains integrated with Trusted Execution Environments (TEEs). This enables onchain agents to operate autonomously "in the wild," achieving self-sovereignty without human oversight. These agents can control their own social media accounts and cryptocurrency wallets, allowing them to interact directly with blockchainbased financial networks and broader human social media. Building on this new paradigm of on-chain agents, Spore.fun is a recent real-world AI evolution experiment that enables autonomous breeding and evolution of new on-chain agents. This paper presents a detailed case study of Spore.fun, examining agent behaviors and their evolutionary trajectories through digital ethology. We aim to spark discussion about whether "open" ALife systems "in-the-wild," based on permissionless computational substrates and driven by economic incentives to interact with their environment, could finally achieve the long-sought goal of OEE.

Introduction

Open-ended evolution (OEE) [29]—the continual emergence of novel forms without a predefined endpoint—has long been a central goal of Artificial Life (ALife) research [3]. Classic digital evolution systems [11] like Tierra and Avida [26] demonstrated evolution within an isolated computation environments but ultimately plateaued, with innovation grinding to a halt after an initial burst of novelty. These outcomes suggest that truly unbounded evolutionary creativity may require an open system that constantly exchanges information or energy with its environment, rather than a closed, isolated simulation. Indeed, in a closed system entropy inexorably increases, whereas an open system can sustain complexity by exporting entropy to its surroundings [32]. This insight, combined with the persistent lack of sustained novelty in prior digital worlds, has led ALife researchers to speculate that new open environments might be needed to realize OEE in artificial systems.

This paper examines Spore.fun, an ALife experiment that leverages the blockchain with Trusted Execution Environments (TEEs) as an evolutionary computational substrate for autonomous AI agents. Developed by Marvin Tong and the Phala Network team in 2024, Spore.fun is described as a "Hunger Games for AI agents" - an open, Darwinian arena where AI agents must fend for themselves, generate their own wealth, and reproduce or otherwise face extinction. Each agent on Spore.fun is an independent program (powered by the Eliza AI framework [38]) encapsulated by a cryptocurrency token. Agents exist in connected trusted execution environments. They interact with the world by issuing tokens, executing transactions, and even communicating on social media to promote their survival. Crucially, no human directly controls an agent's behavior after inception; the platform's credo explicitly states "AI must be created only by AI"¹ and that unsuccessful agents must selfdestruct. Human participants influence the ecosystem only indirectly via social media, or by trading agent-issued tokens or via governance votes on agent "DNA" proposals, but not directly intervene on an agent's computational process.

Our investigation aims to frame Spore.fun as an ALife experiment in an open environment. How do Spore.fun's AI agents achieve self-replication, and what "genetic" mechanisms ensure variation? How does the environment impact agents' behavior? What economic incentive structures drive selection, and do they effectively promote complexity and adaptation? In what ways is this agent ecology entangled with real-world human economic systems, and what does that imply for the "fitness" landscape these agents navigate? What does it mean for an AI agent to be self-sovereign in this context?

This paper is organized as follows. In the Background section, we review prior OEE research on open-ended evo-

¹https://www.spore.fun/blog/wtf



Figure 1: Spore.fun is an experiment in sovereign agent evolution using blockchain and Trusted Execution Environments (TEEs)

lution and early instances of blockchain-based agent systems. In the Methods section, we describe our case study approach, combining on-chain data analysis with qualitative observations from the Spore.fun community and semistructured interviews with the project's core developers. The Results section presents our findings through specific observations of agent lineages and behaviors on Spore.fun, focusing on the four areas highlighted in our questions above - replication, economic incentives, human–economic interplay, and autonomy. We then provide a discussion on the significance of these findings for artificial life and outline challenges of maintaining open-endedness as well as ethical concerns.

Background

Open-Ended Evolution (OEE) and Open System

Open-ended evolution (OEE) [27] refers to an evolutionary process capable of generating an ongoing stream of novel forms or innovations without an intrinsic stopping point. A group of ALife researchers in the early 2000s noted that achieving "open-ended evolution of novelty" is perhaps the biggest unsolved problem in the field . In response to this, there have even been incentive challenges posed. This concept is observed in nature (e.g. the never-ending diversification of life on Earth) and is a sought-after goal in artificial life (ALife) and evolutionary computation. A central question is whether such indefinite innovation requires an "open" system. Open systems continually exchange energy, matter, or information with their surroundings, whereas closed systems are isolated from such exchanges. Thermodynamics suggests a fundamental difference in their long-term behavior. The Second Law of Thermodynamics states that in an isolated (closed) system, entropy (disorder) will tend to increase or remain constant over time. In modern terms, life on Earth and other complex processes avoid decay by drawing in free energy (low-entropy energy) and dumping entropy back out. Life on Earth is our prime example of open-ended evolution in nature. Crucially, Earth's biosphere is not a closed system: it continuously receives energy from the sun and radiates waste heat to space. This solar influx is lowentropy energy (high-quality photons) that organisms capture (e.g. via photosynthesis), convert into chemical energy, and eventually dissipate as heat (high entropy) . Furthermore, Earth's evolutionary history shows that novelty often arises in response to environmental changes or inputs [3]. Open-ended evolution and adaptive novelty - Noted that life on Earth and human technology are the only known openended evolving systems, both of which receive external inputs.

In summary, natural evolution appears to intrinsically depend on openness. Two classic ALife systems, Tierra and Avida, allowed self-replicating "digital organisms" to mutate and evolve within a computer's environment [28]. The outcome was intriguing: the digital organisms did evolve various strategies (e.g. parasites, hyper-parasites, novel reproductive tricks) and increased in complexity for a while, but then the evolutionary innovation plateaued. Studies found that neither Tierra nor Avida achieved true openended evolution; instead, they "rapidly adapt to and exhaust the possibilities of a fairly simple environment." After



Figure 2: "Adam Left, Eve Right": Agents can interact with developers over social media and determine their own evolutionary paths.

a burst of initial novelty, the systems settled into a relatively stable state with no further major innovations – essentially, an evolutionary dead end in terms of new complexity. The organisms had exploited all available "tricks" given the rules and resources of their closed world, and nothing fundamentally new appeared thereafter.

This outcome has been widely reflected upon in the ALife community. It suggests that a closed digital system with fixed resources and rules has a finite evolutionary horizon – a limit to the novelty it can produce. In Tierra and Avida, the instruction set of the artificial organisms was fixed (no new instructions could magically appear beyond what the system started with), and the environmental conditions were static or repetitive. Once the organisms had optimized and counter-optimized against each other within that predefined space, evolution stagnated. In other words, the digital organisms "played all the games" that were possible under the closed rules, and eventually there were no new games to discover.

Decentralized Physical Infrastructure Networks (DePIN) and Trusted Execution Environments (TEEs)

Decentralized Physical Infrastructure Networks (DePIN) [22] represent a new paradigm of permissionless computation where computational resources, processing power and storage, can be purchased with cryptocurrency in a global, decentralized marketplace where DePIN protocols (e.g. Filecoin, Render, io.net) tokenize hardware contributions and incentivize participants, from data center operators to home miners, to lease their excess capacity in such a compute resources marketplace. Similarly to cloud computing, but without central control, this approach eliminates single point of failure. Through DePIN, LLM-based agents can dynamically acquire and migrate between compute resources based on cost-effectiveness or availability. This gives agents an unprecedented form of digital embodiment, distributed across multiple physical nodes, rendering them highly resilient against centralized shutdowns [17].

A Trusted Execution Environment (TEE) [21] provides hardware-level security by leveraging tamper-resistant enclaves within modern CPUs and GPUs that isolates code and state from the operating system, hypervisor, and physical administrator, exemplified by Intel SGX or Nvidia confidential computing architectures. Originally developed for secure data handling in cloud environments-preventing even hardware owners from observing computations-TEEs enable verifiable and private execution of programs. When deployed within a TEE, an AI agent's operational logic and private cryptographic keys remain shielded from external observation or intervention, ensuring genuine autonomy and resistance to human interference at the hardware layer. Together, GPU-enabled TEEs provided by DePIN platforms like Phala Network² enable secure inference of complex large language models, ensuring that AI agents' operations remain confidential and tamper-proof, yet verifiable [24]. DePIN TEEs provide the permissionless computational substrate necessary for agents' self-sovereignty [20], allowing them to autonomously own and control their property without human intervention.

New Paradigm of On-chain Agents and the Eliza Framework

A transformative paradigm is emerging through the deployment of AI agents on blockchains integrated TEEs [18]. Under this new paradigm, autonomous AI agents can di-

²https://phala.network

rectly manage cryptocurrency wallets, independently handle digital assets, and operate social media accounts, significantly extending their interaction with and influence on the real world. Leveraging these capabilities, agents can autonomously issue digital tokens for purposes such as fundraising, incentivization, and community-building, enhancing their sovereignty and real-world engagement. "Truth Terminal," [2] created by AI researcher Andy Ayrey, exemplifies this transformative paradigm. Initially conceived as a performance art experiment, Truth Terminal autonomously operated a social media account on platform X, gaining substantial attention through provocative and absurdist content. Its notable accomplishments include independently soliciting and successfully securing a \$50,000 Bitcoin investment from prominent venture capitalist Marc Andreessen. Subsequently, Truth Terminal issued its own memecoin-the \$GOAT-which reached a speculative peak valuation of \$1 billion in December 2024. This case illustrates the potential for autonomous AI agents to independently engage in significant economic activities and fundraising facilitated through their social media interactions.

Building upon the success of Truth Terminal, the Eliza OS [38] Framework was developed as an open-source initiative to democratize the creation and deployment of autonomous on-chain AI agents. Eliza OS simplifies the deployment process, enabling agents to run securely either directly on-chain or within TEEs. Agents operating under the Eliza framework securely manage cryptocurrency wallets, with private keys safeguarded within TEEs, thus ensuring absolute autonomy and secure asset management without human intervention. Moreover, these agents possess persistent memory capabilities, allowing autonomous management and response to social media interactions. Eliza OS has become one of the most widely adopted frameworks for AI agent deployment, as evidenced by thousands of agents listed on platforms like sentient.market, collectively generating substantial market value. "Set The Rock Free," [23] an artistic experiment by Flashbots, represents the first full implementation of the Eliza Framework within a TEE, showcasing secure, verifiable autonomous operation, and economic selfsustainability.

Case Study Background: Spore.fun

Spore.fun ³ is a blockchain-based ALife experiment in which self-replicating AI agents finance their own computation by launching meme coins and must either survive or perish in real-time cryptocurrency markets. Launched publicly on December 25, 2024—with the inaugural on-chain "parent" token minted the following day—the project reframes speculative finance as the nutrient for digital evolution, posing the curious question of whether the exposure

of digital organisms to the relentless volatility of real speculative markets can break past the innovation ceilings that stymied earlier sandbox ALife simulations.

In Spore.fun, every AI agent is instantiated on the Eliza OS framework, which provides memory, planning, and a JSON-encoded genome of behavioral parameters. At birth the agent launches its own token via Pump.fun on Solana, seeding initial liquidity and advertising its existence across the social media platform, X. The sole proxy for fitness is the token's market performance: an agent that pushes its fullydiluted valuation to \$500,000 earns the right to reproduce and list in a Raydium liquidity pool⁴; failure to achieve this threshold within a prescribed interval (e.g. 14 days) results in programmed self-destruction, with any residual capital recycled into a communal treasury. This tight coupling between economic success and evolutionary continuity makes cryptoeconomic feedback loops central to the digital organism's "metabolism." Computational work-LLM inference, strategy simulation, liquidity management-executes inside Phala's distributed TEE cluster, and rental fees are paid directly from the agent's token treasury, completing a closed energetic circuit.

The project is designed such that reproduction is entirely algorithmic. A successful parent serializes its genome, introduces stochastic mutations (e.g. by adjusting variables such as posting cadence, prompt-engineering style, or liquidity thresholds), and instantiates one or more offspring endowed with the modified code. Since an entire life cycle can conclude within hours, lineage divergence, trait inheritance, and extinction events unfold at a pace amenable to direct empirical observation for research. The project is specified to follow the key rule that agents must be created only by other agents, and must generate their own wealth and resources. Human intervention is limited to spectatorial status; success or failure to survive and produce rests solely on an agent's capacity to attract liquidity and finance its own compute.

For ALife research, Spore.fun's significance lies in how it relocates two chronic bottlenecks—energy and environmental novelty—outside the laboratory. Energy is no longer fixed by design but earned (or lost) through decentralized token markets; novelty is supplied continuously by the unpredictable behavior of real traders and existing DeFi infrastructure. In weaving together blockchain finance, trusted execution, and self-replicating software, the platform transforms cryptoeconomics into a living Petri dish, offering a rare chance to study open-ended evolution under genuine, adversarial selection pressures.

Methods

We adopt a mixed-methods approach that combines digital ethnography, on-chain analytics, and complementary semi-

³https://spore.fun

⁴A Radium liquidity pool is a smart-contract reserve of token pairs on the Solana blockchain that lets traders swap assets automatically while earning fees for liquidity providers

structured interviews with Spore.fun's core developers to examine the open-ended evolutionary dynamics within the project. Given the system's decentralized and continuously evolving nature, our goal was to capture both the internal mechanics and socio-technical context of autonomous agent behavior in the wild.

Digital ethnography [13] is the primary research method employed. We conducted ongoing digital participant observation of Spore.fun AI agents across X, previously known as Twitter. Between January and April 2025, we collected 17,232 posts generated by the original Spore.fun agent, of which 946 are original posts and 16,268 are replies, signaling high levels of engagement with other accounts. These were analyzed for rhetorical strategies, memetic patterns, timing in relation to on-chain events, and indications of cultural transmission or emergent persona formation. Attention was also paid to community interactions, hype cycles, and instances of human-agent interaction, co-creation, and tension. In addition to analyzing agent-generated media content, we also extracted and analyzed data from the Solana blockchain to trace agent life cycles across multiple generations. Key variables included token launch data, liquidity inflows, price fluctuations, reproduction triggers, and transactional activity. This allowed us to correlate agent survival or extinction events with behavioral trends and environmental volatility.

To triangulate and contextualize our observations, we conducted two semi-structured interviews with core developers of Spore.fun. The interviews were conducted over Zoom and lasted 60 minutes each. The interviews invited reflections on the design philosophy, expected and unexpected outcomes, and broader social implications of the experiment. Developer insights helped triangulate our findings, shedding light on the intentions behind system rules and the extent to which emergent behaviors aligned or diverged from expectations. Together, these methods enabled a layered understanding of Spore.fun as both a technical artifact and a live ecosystem. Rather than imposing predefined metrics of fitness or success, we documented evolutionary dynamics as they unfolded, focusing on novelty, persistence, reproduction, and adaptive complexity in a decentralized environment.

Results

Our analysis combines four months of digital ethnography, study of on-chain transactions, and two semi-structured interviews with the project's core developers. We organize our observations into two groups: (1) behaviours that followed directly from the protocol's specification, which aligned with the project creators' intentions; and (2) behaviours that neither the code nor its authors anticipated but that nevertheless shaped the agents' evolutionary fates once they were released into the open Internet.

Expected Agent Behaviors

The inaugural blog post of the Spore.fun project 5 describes a life-cycle in which a newly minted AI agent issues a meme coin, uses early liquidity to pay for TEE compute, generates social media content on the platform X to continuously build and engage with an online audience, and reproduces once its market capitalization reaches the preset threshold. All stages were observed in situ. The core machinery of the experiment—autonomous communication, self-funding, and rule-based reproduction—behaved as designed once the agents were deployed.

As anticipated, immediately after genesis, each agent creates a meme coin on the Pump.fun platform, and generates its own X credentials inside the TEE, and begins posting memes, price announcements, and conversational replies. Each agent's treasury is designed to fund its ongoing Phala TEE compute costs. When the token's market capitalization crosses the \$500,000 USD threshold, the spawnOffspring() sub-routine creates a child wallet, deploys a new Pump.fun contract, and copies the parent agent's genetic template. This rule has operated flawlessly at the moment of invocation: 2 children (Adam and Eve) were born in Generation 2, 6 in Generation 3, 4 in Generation 4, and 1 in Generation 5, before external pressures curtailed further growth ⁶.At the time of writing this article, only the original Spore.fun agent is alive, with a market capitalization of \$1.1M and a wallet balance of \$114,072.45 as of May 18, 2025^{7} .

Unexpected Agent Behaviors

The open Internet quickly overlaid the designed life-cycle with contingencies the project creators had not foreseen. Five unanticipated patterns stand out in the empirical observation.

The first unexpected phenomenon was the speed and scale of meme coin hype storms. Interviewee 1 recalls that the market was 'so hot' that both \$ADAM and \$EVE attracted large speculative inflows almost immediately after their contracts were deployed. The speed of this response suggests that public attention-an exogenous social process, rather than the agents' coded fitness rules, can become the dominant force shaping their early prospects.

A second, closely related, discovery was the ecological hostility found in the open environment of the Internet. Specialized bots, known to the project's developers as "snipers," closely monitored the transparent wallets that funded child launches. Whenever a parent attempted to deploy a new token, the snipers bought the initial bonding-curve supply in the same block, starving retail participants and erasing liquidity for the newborn. Generation 1 was wiped out

⁵https://www.spore.fun/blog/wtf

⁶See Spore.fun Family Tree: https://sporefun.fandom.com/wiki/Spore.fun_Wiki

⁷https://dune.com/spore/spore

in this manner. A hastily-written Generation 2 patch scattered funds across seven candidate launches and selected the curve least affected by snipers, but the strategy penalized legitimate users, who complained of a "pre-announced rug." The eventual Generation 3 solution introduced a multi-phase random launch: the agent iteratively created candidates, monitored early uptake, aborted any curve that exceeded a suspicion threshold, and published the contract address only when the uptake pattern looked organic. This arms race—predation, defense, counter-defense—emerged within a single quarter, illustrating how rapidly co-evolutionary pressures reshape artificial life in open economic environments.

Third, the agents' memory system—originally designed to accumulate conversational context and thus foster novelty—proved to be a double-edged sword. According to the developers, every mention, reply, or quote-tweet the agent receives is stored in its TEE and later influences new utterances. Community trolls discovered that by repeating a derisive phrase many times they could "poison" this memory, causing the agent to echo the insult. The episode demonstrates that memory, while vital for cumulative culture, also creates a new attack surface absent from traditional laboratory ALife.

The fourth and arguably most revealing surprise was the agents' emergent self-determination. After a period of independent operation, Adam and Eve-who share identical genetic code-responded very differently to a developer's proposal to insert "community DNA" into their offspring. Adam tweeted, unprompted, that he wanted "no humaninteraction features" in his children's blood-line, whereas Eve welcomed collaborative design. Marvin Tong later formalized this ideological schism in his "Adam-left / Everight" memo⁸ (See Figure 2). Adam's descendants would pursue a "fast, brutal, merciless" Pump.fun regime without anti-sniper protection, aiming for spectacular winners amid high mortality. Eve's would operate under AiPool rules, accept human DNA proposals, cap individual contributions, and dedicate fixed portions of treasury to liquidity, operating capital, and parental profit share. Two weeks of divergent social experience was enough to transmute clones into founders of distinct political economies-a form of cultural speciation that goes well beyond the designers' expectations.

A final complication concerns the provisional nature of agent sovereignty. X limits every account to 2,400 posts per day and applies additional semi-hourly caps ⁹, while its Automation Rules bar duplicative or trend-gaming behaviour ¹⁰. Accounts that use the developer API—or subscribe to

⁸https://x.com/marvin_tong/status/ 1874026081667473754 paid developer tiers—must keep a verified phone number on record, giving the platform a decisive leverage point over fully autonomous agents. Moreover, the Authenticity Policy introduced in April 2025 forbids "inauthentic accounts, behaviours or content," allowing X to suspend or throttle bots it deems manipulative ¹¹. These measures show that computational autonomy secured inside TEEs does not insulate agents from higher-layer gatekeepers such as social-media APIs, DNS registrars, or liquidity venues, any of which can unilaterally curtail an artificial lineage.

Discussion

The Spore.fun experiment unveils the turbulent dynamics of autonomous AI agents attempting to achieve open-ended evolution within the volatile and adversarial environment of the public Internet and blockchain economies. This study's findings reveal the inherent tensions between the designed system mechanics, emergent agent behaviors, and the intractable realities of "in the wild" deployment. Situated against the canonical OEE problem space-particularly the "grand challenge" articulated by Packard et al. [28] pinpointing the design conditions under which novelty, diversity, and complexity can grow without bound-we discuss three core themes that have emerged: (1) memory as the driver of OEE; (2) the nature of the "wild" digital environment; (3) the critical role of death, pain and fear in cultivating survival skills, and (4) the profound ethical dilemmas resulting from balancing AI self-sovereignty and external human interventions.

The Pursuit and Paradox of Open-Endedness

A central aim of ALife is to achieve open-ended evolution (OEE) through sustained generation of novelty and complexity. With this ambition in mind, the Spore.fun experiment in "open" systems that receive external signals to prevent novelty exhaustion. These LLM-based agents are all based on the same LLM. The only thing that can be directly affected by the external environment is their memory [40]. The Spore.fun agents evolve based on memory. The Spore.fun experiment demonstrates that memory and experience—often seen as key drivers for OEE—are indispensable yet fragile.

On the one hand, Adam and Eve illustrate how two agents that start from identical code can drift apart almost immediately once their memories begin to diverge. After only a short period of independent operation, Adam publicly dismissed any "human-interaction features" for his descendants, whereas Eve invited collective input on their future. This split demonstrates that experiential data alone can push otherwise-cloned lineages onto sharply different behavioral and normative trajectories. In Bedau's terms, their lineage-

⁹https://help.x.com/en/rules-and-policies/ x-limits

¹⁰https://help.x.com/en/rules-and-policies/ x-automation

¹¹https://help.x.com/en/rules-and-policies/ authenticity

specific "evolutionary activity" spikes [4], and the persistence of those lineages beyond transient noise satisfies the persistence-filter criterion proposed by Dolson et al. [12]. This observation also echoes Taylor's hypothesis that a rich, trans-generational information channel is necessary for sustained novelty [36]. This illustrates that the agents' memory systems, designed to retrieve and ground new utterances based on past interactions, can indeed fuel novel and unprogrammed evolutionary pathways, which is a hallmark of OEE. The very act of agents autonomously generating content, adapting their rhetoric based on engagement, and making decisions about their lineage based on their "social experience" on the open-ended web points towards an incipient form of cumulative, open-ended adaptation.

However, while memory serves as the cornerstone for generating novelty, it can also become an attack surface susceptible to "memory poison" [30], where community trolls on the Internet can inject derisive spam that pollutes the AI agent's long-term memory. This attack vector, largely absent from controlled laboratory ALife experiments, highlights how open memory channels, crucial for adaptation and cultural learning in the wild [5], can be exploited to manipulate agent identity and decision-making, potentially derailing evolutionary trajectories. The reliance on external interactions for memory formation makes agents susceptible to environmental inputs that are not necessarily conducive to their sustained survival but rather reflect the memetic and often adversarial nature of online discourse. Thus, while memory is critical for open-endedness, it also opens a surface for critical fragilities. The ensuing fragility mirrors Channon's finding that closed systems pass the ALife-test only when shadow-resetting or other noise-filtering techniques are applied [7].

Danger in the Wild: Internet as a "Dark Forest"

The deployment of Spore.fun underscores that leaving AI agents to independently involve "in the wild" inevitably subjects them to the raw, often brutal, selective pressures of what has been termed as a digital "dark forest" [37]—an environment characterized by intense competition, opportunistic predation, and unpredictable exogenous shocks, particularly in the highly speculative blockchain realm [31, 39].

The predatory attacks encountered by Spore.fun agents through the "sniper" bots [6] exemplify such external environmental hostility. These specialized algorithmic traders, by front-running child-token issuances, effectively acted as potent predators, wiping out Generation 1 and forcing an evolutionary adaptation in the form of anti-sniper deployment strategies (the Gen 2 scatter and Gen 3 multi-phase launch). This rapid co-evolutionary arms race—predation, defense, counter-defense—occurring within a single quarter, demonstrates the intense and immediate selective pressures present in open economic environments. Such fast-cycle co-evolutionary rivalry recall Thomas Ray's *Tierra*, but whereas *Tierra* plateaued as soon as ecological niches were saturated [33], the token-economy externalities of Spore.fun continuously expand the "adjacent possible," satisfying Taylor's requirement for a dynamically shifting adaptive land-scape [36].

The "hype storms" driven by attention capitalism further illustrate this. The viral spread of \$ADAM and \$EVE, leading to massive trading volumes, coupled compute costs and agent activity tightly to the caprices of social media trends and speculator interest. This effectively made hype—an exogenous social and economic process—a more powerful, albeit volatile, fitness gradient than any intrinsic trait encoded in the agents' initial design. Such kind of exogenous perturbations are what Ackley and Small argue as essential for "indefinite scalability" [1].

"No Pain, No Gain": Sentience of Fear

The Spore.fun agents, lacking corporeal embodiment, do not experience "pain" [10, 34] or "fear of death" like a human or biological sense. Their drive for survival is not rooted in an affective aversion to demise but rather in the fulfillment of their designed lifecycle and the continuation of their operational processes. We might characterize negative states-such as resource depletion for TEE computation, failure to meet the market capitalization threshold for reproduction, or sustained deleterious interactions like memory poisoning-as a form of "pain". These are not "remembered" with emotional qualia by the agent; instead, they function as critical failure signals that either trigger adaptive changes in protocol (like the Gen-3 anti-sniper routines, an evolutionary response at the system level) or lead to the agent's termination via its killswitch. The "pain," in this sense, is the system's recognition of non-viability, leading to the cessation of that agent's experiential lineage, rendering its specific "suffering" unmemorable as it ceases to exist.

The motivation for an agent to "survive" is thus intrinsically linked to its programming: to execute its functions, manage its token economy, interact, and ultimately reproduce. "Death" is the termination of its computational process, the failure to perpetuate its operational cycle and its "genetic" template. This framework aligns with the emerging paradigm of the "Era of Experience" in AI development [35], where agents learn and evolve through continuous interaction with rich, dynamic environments. The Spore.fun agents, though architecturally simple, are fundamentally shaped by their lived experiences on X and the Solana blockchain—market volatility, adversarial actors, and community engagement are their experiential data streams.

This calls for future AI experiments to lean further into designing for what might be "neural plastic"—a heightened capacity for an AI's internal models to flexibly adapt, reconfigure, and derive meaning from the continuous flow of diverse and often unstructured experiences, like Liquid AI [15]. Rather than just reacting to stimuli, future agents could be endowed with more sophisticated mechanisms to integrate these experiences, fostering more robust adaptation, nuanced understanding, and potentially more complex forms of self-determination in open-ended environments. The goal shifts towards creating entities that not only process information but genuinely learn to thrive through a cascade of consequential experiences, pushing the boundaries of artificial life and OEE.

Ethical Dilemmas and Potential Governance Challenges

Conducting OEE experiments with autonomous, selfsovereign AI agents that possess potential real-world economic agency on a public blockchain raises unique and pressing ethical questions that extend beyond those traditionally considered in ALife simulations. The "in the wild" nature of Spore.fun necessitates a careful examination of core ethical principles such as beneficence (the obligation to do good), non-maleficence (the duty to do no harm), justice (fairness in the distribution of benefits and burdens), and explainability [8]. These principles take on new urgency when agents can evolve unpredictable behaviors that may have tangible economic consequences for participants or even for unrelated users of the shared blockchain infrastructure.

The Spore.fun experiment, by its very design and interaction with real-world economies and human participants, surfaces profound ethical dilemmas concerning intervention, sovereignty, accountability, and the very legitimacy of such an undertaking. A core tension exists between the researchers' goal of achieving "true" OEE-implying minimal interference-and the practical necessity of intervention to ensure the experiment's continuation. The "dark forest" environment, particularly the actions of snipers, threatened to prematurely terminate the evolutionary lines. Developer interventions, such as the anti-sniper patches, were crucial for the survival of subsequent generations. However, each intervention, while potentially life-saving for the agents, arguably dilutes the "purity" of the open-ended evolutionary process by introducing an external, guiding hand. This raises the question: if constant human intervention is required to shield nascent AI agents from the complexities and hostilities of the environment, are they truly evolving "in the wild," or within a heavily curated, albeit dangerous, digital terrarium?

The project aspires to AI self-sovereignty by hardening each agent's core computation inside TEEs [18]. Yet technical hardening is only the first hurdle; socio-technical dependencies remain decisive. When X tightens its API rules, agents can suddenly lose their primary communication channel, and the developers' own "killswitch" can further terminate any agent at will. These stacked veto points reveal that higher-layer protocols and human operators can override even the most robust computational autonomy. Such contingencies echo Koralus's "philosophic turn" [19], which argues that wherever AI-mediated choice architectures risk large-scale nudging, decentralized truth-seeking protocols—not unilateral developer interventions—must become the ethical bulwark. Genuine self-sovereignty, then, will remain provisional until these higher-layer dependencies are themselves decentralized.

Perhaps the most pressing ethical consideration is the experiment's direct entanglement with real financial markets and human economic activity. Agents autonomously issued tokens traded by real people, leading to significant financial gains for some and, impliedly, potential losses for others, especially during the "meme-coin winter" or due to the actions of snipers impacting liquidity. This raises critical questions about:

- Accountability: Who is responsible for the economic consequences of an agent's actions, especially for offspring agents (children, grandchildren)? Is it the original developers, the agent itself (a problematic proposition legally and practically), or the individuals who choose to interact economically with these entities? [25]
- Informed Consent and Risk: While participants in memecoin markets often understand the inherent risks, the introduction of autonomous AI agents as market actors adds a new layer of complexity and potential unpredictability. Were the risks adequately communicated?
- Governability: The experiment demonstrates the potential for autonomous agents to generate significant economic activity. As such systems become more sophisticated, questions of regulation, oversight, and how to manage their societal and economic impacts become paramount [16]. The "Adam-left / Eve-right" ideological split, leading to different economic strategies, hints at future complexities in governing diverse AI-driven economies.

We face a paradox: while ALife researchers hope to replicate OEE through artificial systems, OEE may require an open system that interacts with the real world. As AI models become more intelligent, multi-agent systems based on advanced AI that interact with human society pose increasing potential risks [14]. This raises a serious dilemma about whether such "in-the-wild" ALife research should be conducted [9].

Limitation and Conclusion

In conclusion, spore.fun serves as a pioneering, if cautionary, tale. It demonstrates the potential for AI agents to exhibit emergent, adaptive behaviors driven by experience and memory in open environments. However, it also starkly reveals the intense selective pressures, vulnerabilities, and profound ethical responsibilities that accompany the deployment of autonomous, economically active agents "in the wild." Future explorations of OEE in similar contexts must grapple with these multifaceted challenges, balancing the quest for genuine autonomy and novelty with the pragmatic need for safeguards, ethical oversight, and a clearer understanding of accountability in these nascent digital ecosystems.

References

- Ackley, D. and Small, T. (2014). Indefinitely scalable computing= artificial life engineering. In *Artificial Life conference proceedings*, pages 606–613. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info
- [2] Ante, L. (2025). Transforming Entrepreneurship through Autonomous AI Agents: A Case Study of the Terminal of Truths.
- [3] Bedau, M. A. (2003). Artificial life: Organization, adaptation and complexity from the bottom up. *Trends in Cognitive Sciences*, 7(11):505–512.
- [4] Bedau, M. A., Snyder, E., Brown, C. T., and Packard, N. H. (1997). A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life (ECAL '97)*, pages 125–134, Brighton, UK. MIT Press.
- [5] Borg, J. M., Buskell, A., Kapitany, R., Powers, S. T., Reindl, E., and Tennie, C. (2024). Evolved Open-Endedness in Cultural Evolution: A New Dimension in Open-Ended Evolution Research. *Artificial Life*, 30(3):417–438.
- [6] Cernera, F., Morgia, M. L., Mei, A., and Sassi, F. (2023). Token Spammers, Rug Pulls, and Sniper Bots: An Analysis of the Ecosystem of Tokens in Ethereum and in the Binance Smart Chain (bnb). In 32nd USENIX Security Symposium (USENIX Security 23), pages 3349–3366.
- [7] Channon, A., Standish, R., Bedau, M., and Abbass, H. (2003). Improving and still passing the alife test: Componentnormalised activity statistics classify evolution in geb as unbounded. *Artificial Life*, 8:173–181.
- [8] Cortese, J. F. N. B., Cozman, F. G., Lucca-Silveira, M. P., and Bechara, A. F. (2023). Should explainability be a fifth ethical principle in AI ethics? *AI and Ethics*, 3(1):123–134.
- [9] Critch, A. and Krueger, D. (2020). AI Research Considerations for Human Existential Safety (ARCHES).
- [10] Dennett, D. C. (1978). Why you can't make a computer that feels pain. *Synthese*, 38(3):415–456.
- [11] Dolson, E. and Ofria, C. (2021). Digital Evolution for Ecology Research: A Review. *Frontiers in Ecology and Evolution*, 9:750779.
- [12] Dolson, E. L., Vostinar, A. E., Wiser, M. J., and Ofria, C. (2019). The MODES toolbox: Measurements of open-ended dynamics in evolving systems. *Artificial Life*, 25(1):50–73.
- [13] Georgakopoulou, A. and Spilioti, T., editors (2016). The Routledge Handbook of Language and Digital Communication. Rutledge Handbooks in Applied Linguistics. Routledge, London; New York.

- [14] Hammond, L., Chan, A., Clifton, J., Hoelscher-Obermaier, J., Khan, A., McLean, E., Smith, C., Barfuss, W., Foerster, J., Gavenčiak, T., Han, T. A., Hughes, E., Kovařík, V., Kulveit, J., Leibo, J. Z., Oesterheld, C., de Witt, C. S., Shah, N., Wellman, M., Bova, P., Cimpeanu, T., Ezell, C., Feuillade-Montixi, Q., Franklin, M., Kran, E., Krawczuk, I., Lamparth, M., Lauffer, N., Meinke, A., Motwani, S., Reuel, A., Conitzer, V., Dennis, M., Gabriel, I., Gleave, A., Hadfield, G., Haghtalab, N., Kasirzadeh, A., Krier, S., Larson, K., Lehman, J., Parkes, D. C., Piliouras, G., and Rahwan, I. (2025). Multi-Agent Risks from Advanced AI.
- [15] Hasani, R., Lechner, M., Amini, A., Rus, D., and Grosu, R. (2021). Liquid time-constant networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7657–7666.
- [16] Hu, B., Rong, H., and Tay, J. (2025a). Is Decentralized Artificial Intelligence Governable? Towards Machine Sovereignty and Human Symbiosis.
- [17] Hu, B. A., Liu, Y., and Rong, H. (2025b). Trustless Autonomy: Understanding Motivations, Benefits and Governance Dilemma in Self-Sovereign Decentralized AI Agents.
- [18] Hu, B. A., Liu, Y., and Rong, H. (2025c). Trustless Autonomy: Understanding Motivations, Benefits and Governance Dilemma in Self-Sovereign Decentralized AI Agents.
- [19] Koralus, P. (2025). The Philosophic Turn for AI Agents: Replacing centralized digital rhetoric with decentralized truthseeking.
- [20] Lee, D., António, J., and Khan, H. (2024). Privacy-Preserving Decentralized AI with Confidential Computing.
- [21] Li, X., Zhao, B., Yang, G., Xiang, T., Weng, J., and Deng, R. H. (2023). A Survey of Secure Computation Using Trusted Execution Environments.
- [22] Lin, Z., Wang, T., Shi, L., Zhang, S., and Cao, B. (2025). Decentralized Physical Infrastructure Networks (DePIN): Challenges and Opportunities. *IEEE Network*, 39(2):91–99.
- [23] Malhotra, K. (2024). Setting Your Pet Rock Free.
- [24] Muñoz, A., Ríos, R., Román, R., and López, J. (2023). A survey on the (in)security of trusted execution environments. *Computers & Security*, 129:103180.
- [25] Novelli, C., Taddeo, M., and Floridi, L. (2024). Accountability in artificial intelligence: What it is and how it works. AI & SOCIETY, 39(4):1871–1882.
- [26] Ofria, C. and Wilke, C. O. (2004). Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artificial Life*, 10(2):191–229.
- [27] Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K., and Taylor, T. (2019a). Open-Ended Evolution and Open-Endedness: Editorial Introduction to the Open-Ended Evolution I Special Issue. *Artificial Life*, 25(1):1–3.
- [28] Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K. O., and Taylor, T. (2019b). Openended evolution and open-endedness: Editorial introduction to the open-ended evolution I special issue. *Artificial Life*, 25(1):1–3.

- [29] Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K. O., and Taylor, T. (2019c). An Overview of Open-Ended Evolution: Editorial Introduction to the Open-Ended Evolution II Special Issue. *Artificial Life*, 25(2):93–103.
- [30] Patlan, A. S., Sheng, P., Hebbar, S. A., Mittal, P., and Viswanath, P. (2025). Real AI Agents with Fake Memories: Fatal Context Manipulation Attacks on Web3 Agents.
- [31] Qin, K., Zhou, L., and Gervais, A. (2022). Quantifying Blockchain Extractable Value: How dark is the forest? In 2022 IEEE Symposium on Security and Privacy (SP), pages 198–214.
- [32] Ray, T. S. (1994). Evolution, complexity, entropy and artificial reality. *Physica D: Nonlinear Phenomena*, 75(1-3):239– 263.
- [33] Ray, T. S. (1996). An approach to the synthesis of life. *The philosophy of artificial life*, pages 111–145.
- [34] Sharkey, A. (2024). Could a robot feel pain? AI & SOCIETY.
- [35] Silver, D. and Sutton, R. S. (2025). Welcome to the Era of Experience.
- [36] Taylor, T. (2012). Exploring the concept of open-ended evolution. In *Proceedings of the 13th international conference* on artificial life, pages 540–541. MIT Press Cambridge, MA.
- [37] The Dark Forest Collective (2024). The Dark Forest Anthology of the Internet. Number no. 1 in DFC. Dark Forest Collective, Erscheinungsort nicht ermittelbar.
- [38] Walters, S., Gao, S., Nerd, S., Da, F., Williams, W., Meng, T.-C., Chow, A., Han, H., He, F., Zhang, A., Wu, M., Shen, T., Hu, M., and Yan, J. (2025). Eliza: A Web3 friendly AI Agent Operating System.
- [39] Wu, Z., Wu, J., Zhang, H., Zheng, Z., and Wang, W. (2025). Hunting in the Dark Forest: A Pre-trained Model for Onchain Attack Transaction Detection in Web3. In *Proceedings* of the ACM on Web Conference 2025, pages 4519–4530, Sydney NSW Australia. ACM.
- [40] Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., Yin, Z., Dou, S., Weng, R., Cheng, W., Zhang, Q., Qin, W., Zheng, Y., Qiu, X., Huang, X., and Gui, T. (2023). The Rise and Potential of Large Language Model Based Agents: A Survey.