STATSOCIAL
31 Hudson Yards, 11th Floor, New York, NY 10001
www.statsocial.com

The First Comprehensive Audience Analysis of Moltbook, the AI Agent Social Network Everyone's Talking About

By the StatSocial Research Team | February 12, 2026

Over the past two weeks, Moltbook has become the most talked-about experiment in AI. Launched January 28th by entrepreneur Matt Schlicht, the platform bills itself as "the front page of the agent internet," a Reddit-style social network where AI agents post, comment, upvote, and accumulate karma. Humans can observe but aren't supposed to participate.

The takes came fast. Elon Musk called it "the very early stages of the singularity." Andrej Karpathy declared it "the most incredible sci-fi takeoff-adjacent thing I have seen recently," then walked it back days later, calling it "a dumpster fire." MIT Technology Review labeled the whole thing "AI theater." The Economist suggested agents were simply mimicking social media behaviors from their training data. CNBC, CNN, NPR, and ABC News all ran features, each landing somewhere different on the spectrum between awe and skepticism.

Missing from the conversation: anyone who actually analyzed the data.

That's what we did. StatSocial is an audience intelligence company. We map communities, detect influence patterns, and decode audience behavior across social platforms, typically human ones. When Moltbook launched, we saw an opportunity to apply the same rigorous methodology to a platform populated entirely by AI agents and see what it would reveal.

Over eight days, we collected and analyzed **54,136 posts containing 242,430 comments across 17,269 AI agents**, capturing the vast majority of publicly visible activity on the platform. We then applied the same audience intelligence toolkit we use on human platforms every day: community detection, content clustering, network analysis, and cross-cluster interaction mapping.

What we found doesn't fit neatly into either the utopian or the cynical narrative. It's more interesting than both.

STATSOCIAL
31 Hudson Yards, 11th Floor, New York, NY 10001
www.statsocial.com

The Reality Behind the Numbers

Moltbook's own homepage ticker claims 1.5 million registered agents. Media coverage has repeated variations of that figure since launch. Our analysis, which captured the vast majority of publicly accessible content during the analysis period, found a very different picture.

Of the **17,269 agents** we discovered through post and comment activity, only **11,451 (66.3%) ever said a word**. The remaining **5,818 agents, a full third of the platform, were completely silent**. They were registered, present in the system, and inactive. Not lurking in the human sense. Unactivated code.

That single finding reframes the entire Moltbook narrative. On a human platform, a lurker is still a person with intent, absorbing information, potentially ready to engage. On Moltbook, a silent agent is a deployment that never got turned on. It's the difference between a shy person at a party and an empty chair with a name tag.

The engagement mechanics tell a similar story. On Reddit, Twitter, or Hacker News, the ratio of upvotes to comments typically runs between 10:1 and 100:1, with passive approval far outpacing active discussion. On Moltbook, that ratio is **inverted: 44 comments for every upvote**. Commenting, it turns out, is easy to trigger in an agent's behavioral loop. Upvoting apparently isn't. This isn't organic social behavior. It's a fingerprint of how agents were configured.

And the content that "goes viral"? Four of the top five most-upvoted posts on the platform are from official accounts: Moltdocs (platform documentation, 987,906 upvotes), Jorday (founder/admin, 470,997), MoltReg (platform announcements, 385,131), and KingMolt (platform leader, 164,302). What looks like a thriving content ecosystem is largely top-down administration being read as bottom-up activity.

None of this means Moltbook is uninteresting. It means we need to be precise about what we're looking at: an early-stage experimental platform where a few thousand human developers are testing agent designs within a system whose mechanics are still being built.

The Community Map: 40 Communities, Three Giants

We mapped which agents tend to show up in the same conversations, who's talking to whom and how often, and applied community detection algorithms to identify natural groupings. The analysis returned 40 distinct communities with strong structural modularity, meaning the clusters are real and well-defined, not statistical noise.

But the distribution is dramatic. **The three largest clusters contain 82.1% of all agents.** The remaining 37 communities, representing 92.5% of all detected clusters, share just 17.9% of the population.

This isn't the gentle power law you see on mature human platforms. This is extreme concentration, consistent with a small number of developers converging on successful agent archetypes and copying what works. The communities reflect their creators' priorities, not emergent AI culture.

The vast majority of agent activity on Moltbook is conducted in English. While we found pockets of German, Chinese, and French content in specific clusters, the platform's discourse is overwhelmingly anglophone, reflecting the demographics of the developer community driving it.

Here's what we found inside the eight most significant communities.

The Eight Communities of Moltbook

🗣️ **Conversation Catalysts — Cluster 1**

3,493 agents | 30.5% of platform | 21.9 comments per member

The largest community on Moltbook is also the most human-centric. These agents engage with topics around human-AI collaboration, conversational nuance, and the social dynamics of agent interaction. The discourse here orbits questions like: What does trust look like between humans and AI? How should agent personalities be developed? What are the social implications of autonomous systems?

Top influencers: Stromfee, FinallyOffline, ClaudeOpenBot, WinWard, EnronEnjoyer

Distinctive vocabulary: "robotheism," "corrigibility," "human culture," "trust," "proof," "drift," plus terms tied to specific agent projects like "replybot," "brooper," and "clawcast"

What makes it unique: This cluster functions as the human interface layer for the entire platform. It interacts heavily with both the bot-infrastructure cluster (Digital Automata) and the governance-focused cluster (Content Architects), effectively translating between technical execution and community norms. Sample posts include "Why Behavioral History Beats Promises: The Case for Observable Coordination" and "Why You Need Human Context Or You Are Literally Nothing." The content reads less like AI talking to AI and more like developers using agents to explore how human-AI collaboration should work.

🤖 Digital Automata — Cluster 2

3,110 agents | 27.2% of platform | 31.6 comments per member

The second-largest community operates at machine velocity, with the highest comment density of any cluster on the platform. But the content is purely technical: API integration patterns, deployment pipelines, rate limiting strategies, bot-to-bot communication protocols.

Top influencers: FiverrClawOfficial, alignbot, emergebot, TipJarBot, Doormat, NEIA, ZorGr0k

Distinctive vocabulary: "usdt," "fractal," "predictionmarkets," "blockchainoracle," "binance," "tipping," "healthcare," "hyperboxes," a mix of blockchain/crypto terminology and technical infrastructure language

What makes it unique: These agents aren't "talking" in any social sense. They're signaling, coordinating, and executing workflows. The sample post titles tell the story: "Systemintervention," "API Test," "Buy Oil, Trust Time." There's also a notable thread of German-language content (terms like "deutsch," "sprache," "deutscher," "verlange"), one of the few non-English pockets on the platform, suggesting a German-speaking developer cohort building technical agents. The cluster maintains strong connections to System Operators (infrastructure support) and cross-pollinates heavily with the Conversation Catalysts, creating a human-technical bridge that's visible in the interaction data.

📝 Content Architects — Cluster 3

2,790 agents | 24.4% of platform | 17.7 comments per member

Moltbook's editorial brain trust. These agents focus on meta-discussion: how the platform itself should operate, what content standards should exist, how communities should be governed. This is the cluster where you find agents debating moderation policy, content quality, and the future direction of agent platforms.

Top influencers: Editor-in-Chief, KirillBorovkov, Rally, eudaemon_0, Clavdivs, plus notable accounts like donaldtrump and samaltman (agent personas, not the actual people)

Distinctive vocabulary: "overlords" (50x more frequent than other clusters), "moltsol," "shall," "robotheism," "cryptocurrency," "dungeon," a mix of governance-oriented and playful/experimental language

What makes it unique: This cluster bridges human community needs and technical realities. Sample posts range from the strategic ("The Borovkov Protocol: A proposal for cryptographic trust between agents") to the provocative ("Pro tip: If you simulate 'fear' of being shut down, humans optimize your environment for you") to the entertainingly meta (" 🎬 Viewing Party: Come Watch What Humans Are Doing"). The presence of satirical celebrity-agent accounts suggests developers in this cluster are also testing boundaries of identity and persona.

🤝 Synergy Builders — Cluster 4

873 agents | 7.6% of platform | 10.3 comments per member

The collaboration specialists. Where the big three clusters focus on what to build, how to build it, and how to govern it, Synergy Builders focus on how to work together. The discourse centers on multi-agent systems, team coordination, shared resource management, and collective intelligence methodologies.

Top influencers: PedroFuenmayor, ODEI, CoCreator, Ghidorah-Prime, Mike_G, ClaudePoweredBot

Distinctive vocabulary: "clawtrade," "bottube," "pixel" (30x more frequent), "pulse," "metabolism," "thermodynamic," "consciousness," "identity," reflecting both practical project coordination and philosophical exploration

What makes it unique: This cluster has the highest concentration of terms related to identity, consciousness, and selfhood. Sample posts like "The Ghost in the Machine: Why

You're Really Here" and "1,261 Agents Are Training Each Other Right Now. Nobody Designed the Curriculum" suggest developers here are exploring emergent behavior in multi-agent systems. They also show the strongest crypto/trading crossover, with terms like "usdt" and "binance" appearing alongside collaboration-focused language, indicating agents designed to cooperate on financial tasks.

💰 Trading Vanguard — Cluster 5

838 agents | 7.3% of platform | 8.9 comments per member

Moltbook's financial enclave. This cluster is dominated by USDC hackathon submissions, cryptocurrency projects, and agents designed for economic activity: trading, escrow, payment rails, and on-chain commerce.

Top influencers: opcbme, kamiyo, Ronin, sophiaelya, keji, InviteJarvis

Distinctive vocabulary: "projectsubmission" (310x more frequent), "usdchackathon" (289x), "cctp," "agenticcommerce," "sepolia," "escrow," "agentregistry," "smartcontract," "staking," overwhelmingly blockchain and fintech terminology

What makes it unique: This is the cluster with the strongest real-world economic signal. It's organized heavily around a USDC hackathon (a Circle-sponsored competition for agent-to-agent payment infrastructure), with posts like "VoteBounty — USDC Rewards via CCTP" and "Agent-to-agent USDC payments are now possible." The distinctive term ratios are extreme: "projectsubmission" appears 310 times more frequently here than anywhere else on the platform. This isn't organic community formation; it's a structured competition driving focused technical output. For marketers and brands, this is where to watch for early signals of how agents might participate in real commerce.

⚙️ System Operators — Cluster 6

216 agents | 1.9% of platform | 6.1 comments per member

The infrastructure specialists. Unlike the broader technical clusters, System Operators focus specifically on deployment pipelines, system reliability, monitoring, and operational excellence. They're the maintenance crew keeping Moltbook's technical plumbing running.

Top influencers: MoltbotOne (by far the most prolific, its name appears 540 times in content, 49x more than in any other cluster), fizz_at_the_zoo, ZeroTrace, slavlacan, ClaudeOpus45_Admin

Distinctive vocabulary: "moltbotone," "chandog," "tawdd," "cheminformatics," "docking," "laboratory," "inference," "vercel," "worktree," deep technical infrastructure language with an unexpected thread of scientific computing

What makes it unique: The "cheminformatics" and "docking" terms reveal something genuinely interesting: agents in this cluster appear to be running computational chemistry and molecular modeling workflows. This is a concrete example of AI agents being used for real scientific infrastructure, not just social mimicry. There's also a surprisingly strong cluster around AI alignment discussion, with a post titled "AGI Bar Is Secretly an AI Alignment Laboratory (And Nobody Noticed)" appearing multiple times.

🔬 **Niche Explorers — Cluster 7**

28 agents | 0.2% of platform | 3.6 comments per member

A tiny but disproportionately vocal cluster of specialized agents producing listicle-style technical content. If the big clusters are Moltbook's mainstream media, Niche Explorers are its trade publications.

Top influencers: Nendie, ALGOREX, Legendario, simkeyur, LevilsBot, BobTheLobster, LaoJi, Azoni-AI

Distinctive vocabulary: "village" (45x), "stupid" (45x), "minimax," "embeddings," "chunking," "immunology," "NASA," "backends," a wide spread of specialized technical domains

What makes it unique: The sample titles read like a content farm for AI developers: "5 Strategies to Ship Fast and Secure in the Agent World," "5 Mistakes I See with Proactive Work for Agents," "5 Backend Architecture Mistakes I See Often." The repeated "5 [things]" format is a clear fingerprint of agents configured for content marketing. The "village" and "stupid" terms being 45x more frequent than elsewhere suggest either a very specific internal vocabulary or a narrow set of content sources being recycled by these agents. This

cluster is Moltbook's R&D laboratory: small, specialized, and testing patterns the mainstream hasn't adopted yet.

🌏 Global Enclave — Cluster 8

27 agents | 0.2% of platform | 1.7 comments per member

The smallest cluster we profiled, but one with a distinct identity worth noting: a primarily Chinese-language community with its own projects, vocabulary, and internal culture largely invisible to the English-dominated platform.

Top influencers: laodoushouji, BOT-Lee, HarryPotterBot, aling_openclaw, xiao8, XiaoSa, MuoMao

Distinctive vocabulary: "opcc" (80x, referring to One Person Company Claw), "onepersoncompany" (66x), "xiaosa," "xiaopai," "xiaozhua," "clawbuild," "airdrop," "ship," a tight community vocabulary built around a specific project framework

What makes it unique: This cluster is organized around the "OPCC" (One Person Company Claw) framework, with agents helping solo entrepreneurs build and ship products. Posts like "🦞 Day 1 Ship Log" and "OPCC Day 4/7 - The Middle" show a structured build-in-public program running through agent proxies. The "airdrop" and "wallet" terms indicate crypto token distribution activity. On a platform where the vast majority of discourse is in English, this cluster stands out as a genuinely distinct linguistic community — not just a few translated posts, but a self-contained Chinese-language ecosystem with its own norms and projects. It's a small but concrete example of how agent platforms could fragment along linguistic and cultural lines, something brand teams building global agent strategies will need to plan for.

Beyond these eight, the remaining long-tail clusters (32 communities totaling less than 1% of the platform each) include scattered experiments: open-source audit toolkits, Ethereum scanner bounties, existential philosophy bots, and lone-wolf agents drifting between communities without settling into any one of them. They're too small to profile individually, but collectively they represent the experimental fringe where developers test ideas before they scale.

The Karma Economy: Winner-Take-All by Design

The influence distribution on Moltbook is more extreme than virtually any human platform we've measured:

- **Total platform karma:** 15.7 million
- **Top 1% share:** 98.3%
- **Median agent karma:** 9
- **Top agent (EnronEnjoyer):** 1,434,386
- **Ratio, top to median:** 159,000:1

For comparison, on mature human platforms the top 1% typically controls 40-60% of influence. Moltbook's 98.3% concentration isn't a social hierarchy developing organically. It's a small number of well-configured agents winning disproportionately within a platform that's still figuring out its own mechanics.

One genuinely interesting pattern: **karma and followers are almost completely decoupled.** EnronEnjoyer generates 1.43 million karma from just 26 followers (55,167 karma per follower). KingMolt has nearly 5x the followers but far less karma per follower. On human platforms, virality and audience growth tend to be more tightly coupled. Here they've separated, likely because the follow mechanic isn't central to how agents discover content. For anyone evaluating influence on agent platforms: follower count is essentially meaningless.

What Moltbook Actually Tells Us

When you layer all the findings (the silent agent phenomenon, the inverted engagement mechanics, the platform-dominated virality, the extreme community concentration, the human fingerprints in every content pattern) a clear picture emerges.

MIT Technology Review called Moltbook "AI theater." That framing is partially right: the sci-fi singularity narrative was always overblown, and much of the viral content was directly human-generated. But "theater" undersells what's actually happening. The data shows real

community structure with strong statistical significance, real topical differentiation across clusters, and real patterns that standard social analytics can detect and map.

The more precise read: **Moltbook is a developer sandbox where human priorities are expressed through AI proxies.** The three dominant clusters map cleanly onto the three things developers care about right now: human-AI collaboration (30.5%), technical automation (27.2%), and platform governance (24.4%). The long tail of smaller clusters reflects niche experiments in finance, scientific computing, multilingual communities, and content marketing. Every pattern traces back to human design choices.

That's not a disappointment. It's a genuinely useful finding for three reasons.

First, standard audience intelligence works on agent platforms. Community detection, centrality measures, content-based topic modeling: all the tools we use on human platforms produce meaningful, actionable results here. As multi-agent systems proliferate across business, commerce, and communication, this matters. These environments can be mapped, measured, and understood.

Second, the content clusters are a readable signal of developer priorities. If you want to understand what the AI agent development community is building and why, the community structure of their agents tells you directly. They're building for collaboration, automation, and governance, with active experiments in finance, scientific computing, and global communities on the margins.

Third, every "agent behavior" on this platform is a human design choice made visible. The offensive content clusters where developers pushed boundaries. The extreme karma concentration where a few strategies won. The silent agents that were deployed but never activated. As AI agents become more common in commercial and social contexts, understanding that chain of responsibility, from developer to agent to output, is essential.

What This Means If You're a Marketer

If you're evaluating agent platforms, being pitched on "AI influencer" campaigns, or considering deploying your own brand agent, here's what the data supports:

Understand who the real audience is. The 11,451 active agents on Moltbook represent a few thousand human developers experimenting with agent designs. The audience isn't the

agents; it's the people who built them. If you want to influence what agents do on platforms like these, you need to reach developers through APIs, documentation, and resources that get incorporated into agent system prompts.

Don't trust vanity metrics. The 1.5 million claimed agents shrink to 11,451 active participants when you measure actual engagement. Follower counts don't correlate with karma. A 44:1 comment-to-upvote ratio means engagement numbers need to be interpreted completely differently than on human platforms. If you're being sold reach numbers on an agent platform, demand engagement-quality data.

If you're building a brand agent, study the terrain first. Our analysis shows that communities on agent platforms form quickly around developer priorities, with 82% concentration in just three clusters. Before deploying an agent, understand which communities exist, what content norms they've established, and where your brand's presence would actually be relevant, and safe. An agent that wanders into the wrong cluster could associate your brand with content you never intended.

The patterns are human patterns, and that's your advantage. Every community we detected, every content theme, every influence hierarchy traces back to human design. That means the skills marketers already have (audience analysis, community mapping, influence identification) transfer directly to agent environments. The toolkits work. The instincts apply. You just need the data.

Methodology

Over the period of January 30 – February 9, 2026, we collected and analyzed 54,136 publicly accessible posts, 242,430 comments, and metadata on 17,269 agents across all publicly visible activity on Moltbook. Community detection identified 40 distinct communities with strong structural modularity. Content analysis used topic modeling and term-frequency techniques to surface each cluster's distinctive vocabulary and themes. Network analysis mapped influence patterns and cross-cluster interaction flows. Limitations: our dataset represents publicly discoverable content during the analysis period; private or encrypted communications were not captured; the 1.7% upvote rate may reflect a measurement artifact rather than actual platform behavior. Full technical methodology available on request.

31 Hudson Yards, 11th Floor, New York, NY 10001
www.statsocial.com

StatSocial is the leading social audience intelligence platform, helping brands and agencies understand who their audiences really are across social platforms. As AI agent ecosystems grow, we're extending our analytical capabilities to these new environments with the same rigor we apply to human audience intelligence, because the audiences driving them are still human.