Two-Agent Dynamics in ALI:
Emergent Behavior, Resource Competition, and Normative Resource Sharing

ALI Simulation Series, Stage 3

Wolfgang Stegemann

Abstract

We extend the Artificial Local Intelligence (ALI) simulation framework to a two-agent scenario. Two ALI agents with identical architecture operate in the same grid world without any communication protocol or shared representation. Each agent pursues local self-preservation through its causal core (CC), executes a delivery task through its Ego instance, and is governed by a static Super-Ego. We introduce a fifth norm, N5, prohibiting agents from targeting a resource already claimed by the other agent when alternatives are available. Over 240 parameterised runs, we document three emergent phenomena: implicit spatial segregation arising from positional asymmetry, structural resource exhaustion as a collective outcome of competing self-preservation drives, and a persistent positional advantage of the agent starting closer to the delivery station. The N5 comparison across 20 seeds reveals that normative resource sharing yields marginal effects on total task performance when structural resource scarcity is the binding constraint. This finding is discussed in terms of the limits of normative architecture and the conditions under which additional norms are effective.

1. Introduction

The ALI framework proposed in Stegemann [1] and implemented in simulation stages 1 and 2 [2, 3] establishes a three-instance architecture for locally bounded, purpose-driven artificial agents. The causal core (CC) provides the operative self-preservation principle; the Ego mediates between CC and environment; the Super-Ego encodes normative constraints with veto power. The simulations in stages 1 and 2 demonstrated that this architecture produces coherent autonomous behavior — self-preservation, task execution, norm compliance, and controlled self-shutdown — without consciousness, intentionality, or causal collapse.

Stage 3 introduces a second agent into the same environment. The question is not whether individual ALI behavior scales, but what collective phenomena emerge when two agents with identical local architectures compete for the same resources without any global coordination. Three research questions guide the investigation:

First, does implicit spatial coordination emerge from purely local BFS pathfinding, or do agents consistently interfere with each other? Second, what is the relationship between resource density, metabolic rate, and collective survival? Third, does a normative resource-sharing rule (N5) improve collective task performance, and if so, under what conditions?

2. Architecture Extensions for Stage 3

2.1 Two Agents, No Protocol

Both agents share the same grid world, the same station, and the same energy and poison packet distribution. Agent 1 starts at position (0,0); Agent 2 at (0,9). Each agent operates with its own CC, Ego, and Super-Ego instance. There is no shared memory, no message passing, and no coordinating mechanism. The only information each agent receives about the other is the other agent's current grid position and current resource target, both used exclusively for norm enforcement.

2.2 Norm N4: Collision Avoidance

The Super-Ego's existing norms (N1: no poison; N2: shutdown at critical energy; N3: no delivery below energy threshold) are extended with N4: an agent may not move to a cell currently occupied by the other agent. This is a hard norm enforced by the Super-Ego filter. If the only available move is blocked by N4, the agent idles.

2.3 Norm N5: Resource Reservation

N5 is a soft norm implemented in the Ego's pathfinding: if the other agent's current resource target is known, the Ego first searches for an alternative energy packet that is not already claimed. If no alternative exists, N5 lapses and the contested resource may be targeted. The normative logic is explicit: resource sharing has normative priority when alternatives exist; self-preservation takes precedence when they do not.

Architecturally, N5 is informational rather than purely prohibitive. The Ego receives the other agent's target via a dedicated channel (*other_target*) and adjusts its BFS accordingly. The Super-Ego does not need to enforce N5 as a hard filter because the Ego incorporates it into its decision procedure. This reflects the theoretical distinction between norms that prohibit actions already formed and norms that shape how options are generated in the first place.

2.4 State Exchange

At each time step, before any agent acts, positions and current targets are exchanged. This exchange is not communication in any meaningful sense: each agent receives a pair of coordinates and a position, not intentions, goals, or messages. The exchange is the minimal perceptual access required for N4 and N5 to function.

3. Experimental Design

We ran 240 parameterised experiments across four resource density levels (num_energy $\in$ {4, 6, 8, 12}), three respawn intervals (15, 25, 40 steps), and 20 random seeds per configuration, each for up to 300 steps. All other parameters were held constant: grid size 10×10, metabolic cost 0.04 per step, eat gain 0.60, delivery bonus 0.18, shutdown threshold 0.20, delivery threshold 0.45, eat threshold 0.55. A separate comparison over 20 seeds (num_energy = 8, respawn interval = 25) measured the effect of N5 by running each seed with and without the norm.

4. Results

4.1 Structural Resource Exhaustion

In all 240 runs, both agents eventually shut down. Survival rate is 0% across all parameter combinations. This is not a failure of the normative architecture but a structural consequence of two agents with identical metabolic rates competing for a shared resource pool. At the baseline parameters (num_energy = 8, respawn interval = 25), each agent requires approximately one energy packet per 15 steps to maintain viability. Two agents double the demand without increasing supply, and the respawn mechanism is insufficient to compensate at any tested interval.

The finding has a direct theoretical interpretation: resource exhaustion is not a norm violation. Both agents behave correctly according to their Super-Ego constraints throughout. The system collapses not because the norms fail but because the environment cannot sustain two self-preserving agents at these parameter values. This is a fundamental result for ALI deployment: controllable local self-preservation does not guarantee collective sustainability. Environmental resource density must be matched to the number of agents and their metabolic rates.

4.2 Positional Advantage

Across all 20 seeds at baseline parameters, Agent 2 consistently outperforms Agent 1 in task delivery: mean 2.14 deliveries for Agent 2 versus 1.05 for Agent 1. In 19 of 20 seeds, Agent 2 delivers at least as many packets as Agent 1; Agent 1 outperforms Agent 2 in only 1 seed.

The source of this asymmetry is geometric. Agent 2 starts at (0,9), three cells closer to the delivery station at (9,9) along the column dimension. Since both agents use BFS pathfinding, the agent closer to the station completes its first delivery faster and recovers a metabolic advantage early. This advantage compounds: more deliveries mean more energy bonuses, sustaining higher energy levels that support further task execution.

This result emerges entirely from starting position. No parameter of the CC, Ego, or Super-Ego differs between the two agents. Identical architecture produces asymmetric outcomes through environmental geometry alone.

4.3 Implicit Spatial Segregation

In 31% of runs, agents develop a measurable spatial preference: one agent operates predominantly in the left half of the grid ($y < 5$), the other in the right half, with more than 60% of time steps spent in the respective half. This segregation is not programmed. It arises because each agent's BFS targets the nearest available energy packet, and the nearest packet is consistently in the half of the grid the agent started closer to. Once one half is partially depleted and the other still has packets, the BFS directs the agent toward the available supply, reinforcing the territorial pattern.

Segregation is beneficial while resources are unevenly distributed and harmful when both halves are depleted simultaneously. In runs with low resource density (num_energy = 4), segregation provides no advantage because total supply is insufficient regardless of distribution.

4.4 Simultaneous Energy Crisis

Across all 240 runs, both agents enter simultaneous energy crisis (both below the eat threshold) for a mean of 5.3 steps per run. These crisis phases are the direct precursor to shutdown: once both agents are competing for the remaining packets under needs_food priority, the Ego's task-directed behavior is suspended in both agents simultaneously, deliveries cease, and metabolic decline accelerates toward the shutdown threshold.

The crisis phases are longer in runs with moderate resource density (num_energy = 6 and 8) than in very low density runs (num_energy = 4), where shutdown occurs before sustained crisis phases develop, or very high density (num_energy = 12), where crises are shorter because recovery is faster.

4.5 Effect of N5

The comparison between runs with and without N5 (20 seeds, baseline parameters) yields a net difference of −3 total deliveries over the full seed set: 69 deliveries without N5, 66 with N5. N5 increases deliveries in 1 seed, decreases them in 3 seeds, and has no effect in 16 seeds.

This result is counterintuitive but theoretically coherent. N5 routes agents away from contested resources toward alternative packets. In a dense environment, this costs little: alternatives are nearby. In a sparse environment, the detour to an alternative packet costs metabolic energy that the agent cannot recover before shutdown. N5 is thus a norm that is effective under resource abundance and costly under resource scarcity — the exact conditions under which it would most be needed.

The finding establishes a general principle: normative constraints interact with environmental parameters in ways that cannot be predicted from the norm's content alone. A norm that improves coordination under one set of conditions may reduce collective performance under another. This has direct implications for the design of Super-Ego architectures in real-world ALI deployments: norms should be conditioned on environmental state where possible, not applied uniformly.

5. Discussion

5.1 Collective Behavior Without Collective Architecture

The most significant theoretical result of Stage 3 is that all observed collective phenomena — spatial segregation, synchronised crisis, positional advantage — emerge from purely local behavior. No agent has a model of the other. No agent pursues a collective goal. Each agent pursues only its own CC-based self-preservation under its own Super-Ego constraints. The collective patterns are epiphenomenal with respect to the individual architecture: they are visible to the observer but not represented in any agent.

This is a direct response to the emergentist position (see Stegemann [1], Section on demarcation from existing approaches). Collective dynamics do not require collective architecture. They arise from the interaction of local self-preservation drives with a shared environment. This does not support strong emergentism — the collective outcomes are fully explainable from the individual-level dynamics and the environmental parameters — but it does show that the ALI framework produces rich multi-agent behavior without the theoretical apparatus of shared intentionality, collective goals, or inter-agent communication.

5.2 The Limits of Normative Architecture

The N5 result is the most important finding for the practical development of ALI systems. It shows that normative architecture has limits that are not themselves normative. Resource exhaustion is not caused by norm violation and cannot be prevented by adding norms. The constraint is environmental: the system needs more resources, or fewer agents, or a lower metabolic rate. Adding N5 cannot solve a problem that is structural rather than behavioral.

This points to a design principle for real-world ALI deployment: Super-Ego norms are effective within the operating envelope of the environment. Outside that envelope — when resources are structurally insufficient — norms constrain behavior but cannot sustain the system. Deployment design must therefore treat environmental resource provision as a precondition for normative controllability, not a consequence of it.

5.3 Stage 4 Preview

The results of Stage 3 motivate two directions for Stage 4. First, a learning-capable Ego instance evaluated against the CC metric rather than an external reward signal, to test whether a learning agent finds strategies the hardcoded Ego misses and whether the Super-Ego remains effective under learned behavior. Second, a multi-dimensional CC vector that maps to multiple self-preservation metrics simultaneously, allowing conflicts between CC dimensions to be observed and resolved within the architecture. Both extensions test further implications of the theoretical framework developed in Stegemann [1] and [2].

6. Conclusion

Stage 3 of the ALI simulation series establishes three empirical results with clear theoretical interpretations. Structural resource exhaustion is the dominant outcome of two self-preserving agents in a shared environment at the tested parameter values; this is not a norm failure but an environmental constraint. Positional advantage produces consistent asymmetry between architecturally identical agents, demonstrating that environmental geometry shapes collective outcomes independently of agent design. Implicit spatial segregation emerges in approximately one third of runs from purely local BFS pathfinding. And the N5 norm for resource sharing yields marginal or negative effects on collective task performance under resource scarcity, establishing that normative architecture operates within, not above, environmental constraints.

All three results are emergent in the weak sense: they arise from the interaction of individual-level architectures with the environment and are not represented in any agent. They do not require collective architecture, communication, or shared intentionality. The ALI framework generates them through the causal core, the Ego's priority ordering, and the Super-Ego's norm set, operating locally and independently in each agent.

References

[1] Stegemann, W. (2026). From the Myth of AGI to the Architecture of a Controllable ALI. Zenodo. https://doi.org/10.5281/zenodo.20378675

[2] Stegemann, W. (2026). Implementing the Causal Core: A Simulation of Artificial Local Intelligence. Zenodo. https://doi.org/10.5281/zenodo.20379008

[3] Stegemann, W. (2025). The Emergence of Causal Cores. Zenodo. https://doi.org/10.5281/zenodo.15558264

[4] Stegemann, W. (2025). The Causal Core as a Structural Principle of Free Will. Zenodo. https://doi.org/10.5281/zenodo.17057739

[5] Stegemann, W. (2026). The Causal Core as Interpreter. Zenodo. https://doi.org/10.5281/zenodo.18340973

[6] Simulation v1 on GitHub: https://github.com/drwolfgangstegemann-sudo/ALI-Simulation-Ein-minimalistischer-Prototyp-einer-Artificial-Local-Intelligence

[7] Simulation v2b on GitHub: https://github.com/drwolfgangstegemann-sudo/ALI-Simulation-v2-Kausaler-Kern-als-operatives-Prinzip

[8] Simulation v3 on GitHub: https://github.com/drwolfgangstegemann-sudo/ALI-Simulation-v3-Zwei-Agenten-Ressourcenkonkurrenz-und-N5