

Recursive Coordination and Functional Stability in Multi-Agent Systems

Testing the Three-Layer Model under Decentralized Interaction

Abstract

Decentralized multi-agent systems can preserve locally admissible behavior while failing to maintain collective progress. Reciprocal collision avoidance provides a useful case: individual agents can continue to satisfy immediate safety constraints even when their repeated interaction produces oscillation, deadlock, or loss of goal-directed progress.

This study uses documented multi-robot collision-avoidance and deadlock cases to stress-test the Three-Layer Model (TLM), which distinguishes three layers according to their functional roles in information-processing ecosystems: the Syntactic Layer as representation and structural availability, the Semantic Layer as contextual interpretation, and the Operative Layer as decision, execution, and consequential action. The analysis asks whether these distinctions remain coherent when the consequences of one agent's operation become part of the represented conditions processed by another, and whether Cross-Layer Stabilization preserves a diagnostic distinction after possible localization to an isolated layer has been tested. A conventional state/input/output description is retained throughout as the reduction hypothesis.

The analysis produces a mixed result. The distinction between represented state and context-conditioned significance remains reproducible, although conventional state- and constraint-based descriptions can capture the same relational dependence. Recursive role transitions remain compatible with the functional organization of the TLM but provide no novel account of distributed interaction. Cross-Layer Stabilization receives conditional support only as a diagnostic classification: locally admissible functional processing can remain coherent with an immediate constraint while repeated interaction fails to preserve recursive functional compatibility with continued collective operation. The study does not support Cross-Layer Stabilization as an independent control mechanism, dynamical principle, or substitute for established multi-agent system models. The results therefore narrow rather than expand the claim space of the TLM.

1. Introduction

Decentralized coordination creates a distinction between what is locally admissible and what remains collectively viable. An autonomous agent may act on available information, satisfy an immediate operating constraint, and produce an action that is locally acceptable without ensuring that repeated interaction among multiple such agents will preserve continued progress toward a shared or concurrent system objective. In the present study, **local validity** is therefore used in a restricted sense: compatibility with the relevant local constraint. **Collective viability** refers to continued functioning relative to the system objective being considered. Neither term implies global correctness or universal system quality.

Multi-robot navigation provides a concrete setting in which this distinction becomes observable. Decentralized reciprocal collision-avoidance methods allow independently operating robots to derive collision-avoiding actions from information about their own state and the states or movements of nearby robots without requiring a single centralized controller (van den Berg et al., 2011; Raibail et al., 2022). Reciprocal navigation methods have also explicitly addressed oscillatory interaction produced when agents repeatedly adjust their trajectories in response to one another (Snape et al., 2011).

Collision avoidance, however, is not equivalent to continued progress. Contemporary multi-robot navigation research distinguishes objectives including safety and liveness rather than treating successful avoidance and successful navigation as the same requirement (Chandra et al., 2026). Reactive decentralized controllers can additionally reach deadlock equilibria in which robots come to a standstill before reaching their goals (Grover et al., 2021). A system can therefore continue to satisfy an immediate safety requirement while failing to preserve progress toward its navigation objective.

The relevant interaction is recursive. An agent acts on its current situation; that action changes the relational conditions encountered by another agent; the second agent responds; and the consequences of that response can subsequently alter the conditions encountered by the first. Such recursive interaction is not itself theoretically novel. Established models of concurrent and distributed systems already describe components whose outputs or actions participate in the subsequent behavior of other components (Lynch & Tuttle, 2026). The fact that an action or output associated with one component can influence another component's subsequent processing therefore cannot itself constitute evidence for a new model of distributed interaction.

The question addressed here is narrower. The Three-Layer Model (TLM) distinguishes three layers according to their functional roles in information-processing ecosystems: the **Syntactic Layer** concerns representation and structural availability, the **Semantic Layer** concerns contextual interpretation, and the **Operative Layer** concerns decision, execution, and consequential action. These layers are functional classifications rather than mathematical objects or required technical components. Cross-Layer Stabilization is treated separately as a cross-cutting stabilization condition and not as a fourth layer (Rights of Robots, 2026a).

This functional organization creates two immediate points of exposure in decentralized multi-agent systems. First, the distinction between representation and contextual interpretation must remain analytically meaningful even where no technical component explicitly performs a separate semantic operation. Second, functional classification must remain coherent when the consequence of an operation in one processing context later forms part of the represented conditions available in another. This does not introduce a new rule of agent-relative layer assignment. It tests whether the existing distinction between functional role, mathematical object, and technical component remains coherent under recursive interaction.

The first research question is therefore:

Can the Three-Layer Model distinguish local representation, interpretation, and operation from a recursive stabilization failure that emerges only through interaction across agents and time?

A second question concerns Cross-Layer Stabilization:

Does such a case provide evidence that Cross-Layer Stabilization identifies a functional condition not reducible to an isolated failure of the Syntactic, Semantic, or Operative Layer?

These questions are evaluated against a deliberately strong reduction hypothesis:

A conventional state/input/output description is sufficient; TLM terminology adds no analytically relevant distinction.

Established control and distributed-system models retain explanatory priority for the technical behavior examined in the study. The purpose of the TLM analysis is not to replace those descriptions, predict deadlock, or introduce a new collision-avoidance mechanism. The test is whether the TLM preserves a diagnostically relevant functional distinction after the conventional account has already identified the states, interactions, actions, constraints, and failure conditions involved.

The study therefore uses a structured analytical case study. A documented two-robot deadlock provides the principal failure case (Grover et al., 2021). The interaction is reconstructed first in established technical terms and only then examined through the functional distinctions among representation, context-conditioned significance, and operation. Three stress tests are applied. The first examines whether representation and contextual interpretation remain reproducibly distinguishable without assuming a dedicated semantic component. The second tests whether the functional classification remains coherent when the consequences of one agent's operation subsequently enter another agent's represented conditions. The third asks whether locally valid functional processing can persist while repeated interaction fails to preserve recursive functional compatibility.

The results are mixed. The Semantic Layer remains functionally distinguishable in the selected case, but its additional diagnostic value remains exposed to reduction by conventional state- and constraint-based descriptions. Recursive role transition is compatible with the functional organization of the TLM but provides no novelty over established distributed-system models. Cross-Layer Stabilization receives conditional support only as a diagnostic classification distinguishing local functional validity from recursive functional compatibility; it is not supported as an independent control mechanism or dynamical principle.

The contribution of the study is therefore deliberately bounded. It does not propose a new collision-avoidance mechanism, control architecture, swarm model, or formal theory of multi-agent coordination. It examines whether the Three-Layer Model retains diagnostic distinctions under recursive decentralized interaction and whether Cross-Layer Stabilization identifies a failure condition not attributable to an isolated layer. Reduction, partial support, and theoretical constraint are treated as legitimate outcomes rather than as failures of the analysis.

2. Related Work

2.1 Decentralized Reciprocal Collision Avoidance

Decentralized reciprocal collision avoidance provides the technical reference setting for the present study. Reciprocal collision-avoidance methods allow multiple autonomous robots or agents to derive collision-avoiding motions from information about their own movement and the movement of nearby agents without requiring a single centralized controller (van den Berg et al., 2011). Such approaches form part of a broader body of decentralized multi-robot collision-avoidance research in which agents make navigation decisions from locally available information and interaction constraints (Raibail et al., 2022).

Reciprocal approaches are particularly relevant because each agent acts within a situation partly shaped by the actions of others. In reciprocal velocity-based methods, the relative positions and velocities of interacting agents contribute to the calculation of admissible motion, while reciprocal adjustment by other agents is incorporated into collision avoidance (van den Berg et al., 2011). A movement by one robot therefore alters the relational conditions encountered by another robot, whose subsequent response can alter those conditions again.

The Hybrid Reciprocal Velocity Obstacle was developed to address oscillatory behavior associated with reciprocal navigation decisions and to support collision-free, oscillation-free navigation among independently acting robots (Snape et al., 2011). Oscillatory reciprocal interaction is therefore

already an established technical problem and cannot itself constitute evidence for a new TLM mechanism.

The present study does not evaluate the comparative performance of collision-avoidance algorithms. Reciprocal collision avoidance is used because it provides an established technical setting in which local action, relational dependence, and repeated decentralized interaction can be examined without introducing a new control architecture.

2.2 Oscillation, Deadlock, and Loss of Progress

Collision avoidance and continued navigation are related but distinct requirements. Contemporary multi-robot navigation research distinguishes objectives including safety and liveness, allowing collision avoidance and continued progress to be evaluated separately (Chandra et al., 2026).

Deadlock provides a particularly useful case for examining this distinction. Grover et al. analyze deadlock in multi-robot systems using reactive safety controllers and characterize deadlock as an equilibrium of the system dynamics in which robots come to a standstill before reaching their goals. Their analysis further characterizes the deadlock condition through force equilibrium and locates the deadlock set on the boundary of the safe set (Grover et al., 2021).

The important distinction for the present study is therefore not simply between stable and unstable behavior. A system can reach an equilibrium while failing to maintain progress toward its navigation objective. Conversely, continued satisfaction of a safety constraint does not establish successful completion of the broader navigation task.

In this study, the terms **local functional validity** and **collective viability** are introduced as analytical terms for the TLM test. They are not terminology adopted from the collision-avoidance literature.

2.3 Formal Models of Distributed Interaction

Established models of concurrent and distributed systems already represent interactions in which actions or outputs associated with one component affect the subsequent behavior of another. Input/Output Automata provide a formal framework for modeling concurrent and distributed discrete-event systems and their composition through input, output, and internal actions (Lynch & Tuttle, 2026).

This establishes an important reduction boundary. The observation that an action or output associated with one component can contribute to another component's subsequent processing does not constitute a novel account of recursive interaction. The role of the formal comparison is therefore not to reproduce the selected case mathematically, but to prevent an established interaction structure from being reclassified as novel merely because different terminology is applied to it.

The relevant question is instead whether the TLM preserves a diagnostically relevant functional distinction that would otherwise remain undifferentiated in the description of the case. Established formal models retain explanatory priority for the underlying system dynamics.

2.4 Three-Layer Model as Functional Classification

The Three-Layer Model distinguishes three layers according to their functional roles in information-processing ecosystems: the Syntactic Layer concerns representation and structural availability, the Semantic Layer concerns contextual interpretation, and the Operative Layer concerns decision, execution, and consequential action. These layers are functional classifications rather than mathematical objects or required technical components. Cross-Layer Stabilization is treated separately as a cross-cutting stabilization condition and not as a fourth layer (Rights of Robots, 2026a).

Prior validation of the model limits the contribution that can legitimately be claimed. The TLM is not established as a universal architecture, unique formal verification method, or replacement for established technical models. Its bounded contribution lies in integrative functional classification across information-processing contexts (Rights of Robots, 2026b).

The multi-agent case therefore creates two separate tests. The first asks whether the three functional roles remain reproducibly distinguishable under recursive decentralized interaction. The second asks whether Cross-Layer Stabilization retains a diagnostic distinction after failures attributable to representation, contextual interpretation, operation, or established coupled-system behavior have already been considered.

Neither test assumes a positive result.

3. Method

This study uses a structured analytical case study rather than a new experiment, simulation, controller, or mathematical model. Documented behavior from decentralized multi-robot collision avoidance and deadlock research provides the technical reference setting (van den Berg et al., 2011; Grover et al., 2021).

The analysis is deliberately asymmetric. Established technical descriptions are applied first. The TLM is introduced only after the relevant system states, interactions, constraints, actions, and outcomes have been characterized through the existing technical account. This prevents TLM terminology from determining the description of the case it is intended to test.

Four analytical steps are used:

1. conventional reconstruction of the documented behavior;
2. functional reconstruction using the three TLM layers;
3. reduction against established state, interaction, constraint, and control descriptions;
4. evaluation against predefined result conditions.

The method is conceptual rather than mathematical. Formal and technical models serve as analytical controls; they are not replaced by a new formal representation of the case.

3.1 Conventional Reconstruction

The selected interaction is first reconstructed without TLM terminology.

The reconstruction identifies the locally available system conditions, the relevant collision-avoidance and navigation objectives, the actions generated by the participating controllers, the resulting changes in relational conditions, and the collective trajectory that develops through repeated interaction.

For the principal failure case, the endpoint is a documented deadlock equilibrium in which robots come to a standstill before reaching their goals (Grover et al., 2021).

This conventional account retains explanatory priority. If it sufficiently explains how the system reaches the observed condition, the TLM cannot claim an alternative technical explanation merely by assigning functional labels to the same process.

3.2 Functional Reconstruction

Only after the conventional reconstruction is established is the same case examined through the TLM.

For each relevant processing context, three questions are asked:

Representation: What condition is available for processing?

Contextual interpretation: What functional significance does that represented condition acquire under the current relational conditions, objectives, and constraints?

Operation: What decision or action follows, and how does it alter subsequent conditions?

The classification follows functional role rather than permanent object identity. Consequently, an operation performed by one agent may change the shared relational state, and that resulting condition may subsequently become part of the represented situation available to another agent.

This is not an additional rule of agent-relative layer assignment. It is a stress test of the existing functional classification under recursive interaction.

3.3 Reduction Test

Functional distinguishability alone is insufficient to establish additional diagnostic value.

For each proposed TLM distinction, the analysis asks:

Does the TLM distinguish something diagnostically relevant that would otherwise remain undifferentiated in the description of the case?

This is the canonical reduction criterion used throughout the study.

Established technical descriptions provide the reduction control. Distributed composition and output/input relationships are already represented by established formal models (Lynch & Tuttle, 2026), while reciprocal interaction and deadlock already possess technical explanations within multi-robot navigation and control research (van den Berg et al., 2011; Grover et al., 2021).

A conventional technical account may therefore remain sufficient for explaining the system dynamics while a TLM distinction remains analytically useful for a different classificatory purpose. Technical explanatory sufficiency does not by itself eliminate a functional distinction.

The stronger reduction occurs when the proposed TLM distinction merely renames a relation that is already equally differentiated for the analytical purpose under consideration. In that case, no additional diagnostic support is obtained.

3.4 Analytical Trajectories

Three trajectories are used to separate successful interaction, localizable functional divergence, and collective failure under recursive coupling.

Trajectory I — Successful Recursive Coordination

Locally generated actions repeatedly alter the conditions encountered by other agents, but the interaction remains compatible with continued collective progress.

This trajectory establishes that recursive interaction and functional role transition do not themselves constitute a failure condition.

Trajectory II — Local Functional Divergence

A relevant difference can be localized to representation, contextual interpretation, or operation.

This trajectory tests whether the three functional roles remain distinguishable when a failure or divergence can be associated with an individual processing function.

Trajectory III — Recursive Collective Failure

Representation, contextual interpretation, and operation can remain locally coherent with the relevant immediate constraint while repeated interaction fails to preserve continued collective progress.

This is the principal stress case for Cross-Layer Stabilization.

The term *failure* is used deliberately rather than *instability*. The principal deadlock case is characterized in the technical literature as an equilibrium of the system dynamics (Grover et al., 2021). The relevant problem is therefore loss of collective viability with respect to the continuing navigation objective, not necessarily dynamical instability.

3.5 Predefined Evaluation Conditions

The three stress tests are evaluated independently.

For the **Semantic Layer**, support requires a reproducible distinction between represented state and context-conditioned significance. If the distinction merely renames conventional state or constraint processing without preserving an additional diagnostic distinction, no additional support is obtained.

For **recursive role transition**, the question is whether the existing functional classification remains coherent when an operation changes conditions that are subsequently represented in another processing context. Successful classification supports internal applicability only; it does not establish novelty.

For **Cross-Layer Stabilization**, local validity is evaluated relative to the relevant immediate constraint. Conditional support is possible only if locally coherent representation, contextual interpretation, and operation can coexist with a failure of recursive functional compatibility that is not adequately localized to an isolated layer.

If the relevant failure can instead be sufficiently localized to representation, contextual interpretation, an unsuitable operative policy, or another established functional cause, no additional Cross-Layer Stabilization classification is required.

3.6 Anti-Confirmation Boundary

The analysis does not infer support for a TLM distinction from the mere presence of recursive interaction, context dependence, oscillation, deadlock, or successful recovery.

These phenomena or interaction structures already have independent technical treatment in the relevant literature (Snape et al., 2011; Grover et al., 2021; Lynch & Tuttle, 2026).

Accordingly:

- recursive interaction is not evidence of TLM novelty;
- context dependence alone is not evidence of an irreducible Semantic Layer;
- deadlock is not automatically a Cross-Layer Stabilization failure;
- recovery through additional coordination is not evidence that Cross-Layer Stabilization is a technical mechanism.

Any implications for other theoretical or contextual frameworks are considered only after the three primary stress tests have been evaluated.

4. Case Construction

The case is constructed from documented decentralized multi-robot collision avoidance, deadlock, and coordination-recovery research. Its purpose is not to evaluate a particular navigation algorithm or reconstruct controller dynamics mathematically. It provides a technically established interaction sequence against which the functional distinctions of the TLM can be tested.

4.1 Baseline Coordination

The baseline condition is successful decentralized coordination.

In reciprocal collision avoidance, robots operate from locally available information concerning their own movement and the relevant states or movements of nearby robots. Collision-avoiding motion can be determined without centralized coordination among all participating robots (van den Berg et al., 2011).

A locally generated movement changes the relational conditions encountered by nearby robots. Those robots can subsequently alter their own trajectories while preserving collision avoidance and continued navigation. Recursive interaction therefore does not inherently produce failure.

This baseline is analytically necessary because it separates recursion from failure. Recursive interaction alone is not treated as evidence of a Cross-Layer Stabilization problem.

4.2 Reciprocal Interaction

The relevant technical information in reciprocal collision avoidance is relational. Velocity-obstacle-based approaches evaluate potential collision conditions through relations among the positions and velocities of interacting agents and their prospective movements. Reciprocal approaches additionally account for expected adjustment by other agents (van den Berg et al., 2011). HRVO specifically addresses oscillatory behavior associated with reciprocal navigation decisions (Snape et al., 2011).

The position or velocity of another robot therefore does not establish its action relevance in isolation. Its relevance depends on its relation to the receiving robot's own state, prospective movement, navigation objective, and collision-avoidance conditions.

A represented external condition can consequently remain the same or sufficiently comparable while its relevance for subsequent action changes with the surrounding relational configuration.

This final distinction is an analytical observation used for the subsequent TLM test. It is not attributed to the collision-avoidance literature.

4.3 Two-Robot Deadlock

The principal failure case is the deadlock condition analyzed by Grover et al. (2021). Their work characterizes deadlock as an equilibrium of the system dynamics in which robots come to a standstill before reaching their goals and further analyzes the condition through force equilibrium and its relation to the safe set.

As interacting robots approach a collision-sensitive configuration, local safety constraints influence the actions produced by their controllers. Each response changes the relational conditions encountered by the other robot. Repeated reciprocal adjustment can lead the coupled system to the documented deadlock equilibrium.

At that point, two conditions coexist:

collision avoidance remains preserved;

while

goal-directed progress has ceased.

The case must therefore not be characterized simply as system instability. The technically established condition is an equilibrium associated with loss of progress.

This distinction provides the central analytical contrast of the study.

4.4 Local Validity and Collective Viability

The underlying technical case separates continued safety from loss of goal-directed progress (Grover et al., 2021). This is consistent with the broader distinction between safety and liveness objectives in contemporary multi-robot navigation research (Chandra et al., 2026).

The present study uses this technical separation to distinguish two levels of functional assessment.

At the local level, an action can remain admissible with respect to the immediate collision-avoidance constraint. In this restricted analytical sense, the action is **locally valid**.

At the collective level, repeated interaction can nevertheless prevent continued progress toward the navigation goals. In this objective-relative analytical sense, the resulting trajectory is **functionally non-viable**.

The distinction is therefore:

Local validity with respect to an immediate constraint does not guarantee collective viability with respect to the continuing system objective.

These are analytical terms introduced by the present study rather than terminology attributed to the underlying robotics literature.

Local validity does not mean optimality, global correctness, or sufficiency for the complete navigation task. Collective non-viability likewise does not imply that every local state, interpretation, or action is erroneous.

4.5 Coordination Recovery

Additional coordination provides a useful comparison. Contemporary multi-robot navigation research includes hybrid approaches in which locally affected agents can enter an additional coordination or planning regime when ordinary decentralized navigation encounters conflict or deadlock and subsequently return to decentralized operation (Chandra et al., 2026).

The recovery case is used only for this contrast. It does not demonstrate that the TLM supplies a coordination mechanism, and it does not establish Cross-Layer Stabilization as the process through which recovery occurs.

The relevant observation is narrower:

Changing the organization of interaction can change the collective trajectory even when the underlying navigation problem and participating agents remain substantially the same.

4.6 Case Boundary

Before TLM terminology is applied, the case establishes five source-grounded observations: agents can act from locally available conditions without centralized coordination; their actions alter relational conditions encountered by other agents; repeated interaction can preserve collision

avoidance while losing goal-directed progress; the principal deadlock case is an equilibrium rather than a condition that requires an assumption of dynamical instability; and additional coordination can alter the collective trajectory (van den Berg et al., 2011; Grover et al., 2021; Chandra et al., 2026).

None of these observations is attributed to the Three-Layer Model. They constitute the technical case against which the model is tested.

5. Analysis I — Semantic Layer Legitimacy

The first stress test concerns whether representation and contextual interpretation remain reproducibly distinguishable without assuming a dedicated semantic component.

The question is not whether semantic terminology can be applied to collision avoidance. It is whether the distinction identifies a functional relation that remains analytically meaningful after the conventional technical description is acknowledged.

5.1 Represented State and Context-Conditioned Significance

Reciprocal collision avoidance operates on relational information concerning interacting agents, including their positions, velocities, and prospective movements (van den Berg et al., 2011). The technical relevance of such information depends on its relation to the receiving robot's own state and prospective motion rather than on the external state in isolation.

This permits two different functional questions to be asked:

Representation: What relational condition is available for processing?

Contextual interpretation: What functional significance does that condition acquire under the current relations, objectives, and constraints?

The distinction is not between raw data and human-like understanding. It is between the availability of a represented condition and the functional consequence that condition acquires within a particular relational context.

5.2 Reduction Against the Conventional Description

Conventional collision-avoidance models already capture relational dependencies among agent states, prospective trajectories, and collision constraints (van den Berg et al., 2011). A separate Semantic Layer is therefore not technically required in order to calculate action-relevant relations.

The reduction question is narrower:

Does the TLM preserve a diagnostically relevant distinction between what is represented and what that representation functionally implies under the current relational conditions?

If the distinction merely renames an already equally differentiated relation in the conventional account, no additional diagnostic support follows.

5.3 Same or Comparable Representation, Different Functional Consequence

The strongest case for functional distinguishability occurs when the same or sufficiently comparable represented external condition acquires different action relevance under changed relational conditions.

Another agent's position or velocity may remain unchanged or closely comparable while changes in the receiving agent's own position, prospective movement, objective, or surrounding configuration alter whether that external state is collision-relevant.

The analytical distinction is therefore reproducible:

represented condition and **context-conditioned significance** need not vary together.

This establishes functional distinguishability. It does not establish irreducibility.

5.4 No Semantic Component Is Required

A reactive controller can calculate action-relevant relations directly from state and constraint information without implementing a technical component identified as semantic.

The absence of such a component therefore does not by itself invalidate a functional distinction between representation and context-conditioned significance. Conversely, the ability to describe that distinction does not establish that the controller contains an independent semantic process.

The test concerns functional classification, not component decomposition.

5.5 Anthropomorphic Boundary

The term *interpretation* is used functionally.

The analysis does not claim that a robot understands another robot, experiences risk, assigns human-like meaning, or performs subjective interpretation.

The appropriate description is:

A represented state becomes collision-relevant under the current relational configuration.

The Semantic Layer therefore concerns context-conditioned functional significance rather than anthropomorphic cognition.

5.6 Result

The selected case supports a reproducible distinction between represented relational state and the context-conditioned significance of that state for action.

This provides **partial support for the functional distinguishability of the Semantic Layer.**

The result remains exposed to reduction. Conventional state-, relation-, constraint-, and control-based descriptions can capture the same dependency, and the case does not establish that the Semantic Layer is an irreducible or technically necessary category.

6. Analysis II — Recursive Role Transition

The second stress test concerns the coherence of the functional classification under recursive interaction.

The underlying interaction pattern is already established in distributed-system models. Input/Output Automata, for example, explicitly provide for composition among components through input and output actions (Lynch & Tuttle, 2026). The test therefore makes no novelty claim about output-to-input relations or recursive coupling.

6.1 From Operation to Subsequent Representation

An operation performed by Agent A changes some aspect of the shared relational situation.

That changed condition may subsequently become available to Agent B through sensing or another established information process. For Agent B, it can then form part of the represented conditions relevant to subsequent processing.

The relation is therefore:

operation by A → changed relational condition → represented condition for B.

The operation itself does not literally become a representation. Its consequence changes the situation that is subsequently represented within another processing context.

6.2 Functional Role Without Permanent Object Identity

This transition does not require a concrete event or state change to possess one permanent layer classification across all processing contexts.

Within Agent A's processing chain, an action functions operatively. The resulting change may later participate in the represented conditions available within Agent B's processing chain.

This follows from classification by functional role rather than permanent object identity or technical component. It does not introduce an additional principle of agent-relative layer assignment.

The question is only whether the existing functional classification remains coherent when processing chains interact recursively.

6.3 Reduction Against Established Distributed Interaction

The interaction structure itself provides no novelty. Established distributed-system models already accommodate cases in which an output associated with one component participates as an input action within a composed system (Lynch & Tuttle, 2026).

The reduction test therefore asks only whether applying the TLM functional distinctions to such an interaction produces a contradiction or collapse.

It does not.

An operative role in one processing context and a representational role associated with the resulting condition in another can coexist without requiring either role to become an intrinsic property of the underlying event.

6.4 What the Test Establishes

The result is a consistency result.

The TLM remains functionally applicable when recursively interacting agents alter one another's subsequent processing conditions. Its classifications do not require every event or state change to occupy a single global functional position.

This supports internal coherence under the selected recursive interaction condition.

It does not establish superiority over distributed-system models, a new theory of recursive interaction, or additional diagnostic value merely from the existence of the role transition itself.

6.5 Recursive Interaction Does Not Imply Failure

The same interaction structure occurs during successful reciprocal coordination. Agents can repeatedly alter one another's subsequent navigation conditions while maintaining collision avoidance and continued progress (van den Berg et al., 2011).

Recursive role transition therefore does not itself indicate coordination failure, loss of collective viability, or a Cross-Layer Stabilization problem.

Any Cross-Layer Stabilization result must depend on what happens to functional compatibility across repeated interaction, not on recursion alone.

6.6 From Local Functional Chains to Collective Trajectories

The first two analyses concern functions within individual processing contexts.

Recursive interaction adds a different level of assessment because locally coherent processing chains can influence one another repeatedly. The consequence of one local operation changes the conditions of another local process, whose operation can then alter subsequent conditions again.

The validity of each local classification therefore does not determine the collective trajectory.

This does not imply that the TLM explains the resulting dynamics. It establishes only the analytical setting required for asking whether locally valid functional processes remain compatible across repeated interaction.

6.7 Result

The functional classification remains coherent under recursive role transition.

An operation performed in one processing context can change the relational situation that is subsequently represented in another without requiring a permanent global layer identity for the underlying event or state change.

This is **support for internal applicability, not novelty**. Established distributed-system models already describe comparable interaction structures (Lynch & Tuttle, 2026).

7. Analysis III — Cross-Layer Stabilization

The third stress test concerns Cross-Layer Stabilization.

Deadlock does not itself support a Cross-Layer Stabilization classification. The failure must first be accepted under its established technical explanation and then tested for localization within the individual TLM layers.

7.1 Conventional Diagnosis

The selected case can be described conventionally: the coupled decentralized controllers reach a deadlock equilibrium in which the robots come to a standstill before reaching their goals (Grover et al., 2021).

This technical account remains sufficient for characterizing the underlying system behavior.

Cross-Layer Stabilization does not explain why the controllers reach the equilibrium, predict when it will occur, calculate its dynamics, or provide the mechanism through which it should be resolved.

The TLM analysis therefore begins only after the conventional diagnosis has been accepted.

7.2 Testing for an Isolated Layer Failure

The first question is whether the relevant failure can be localized to one of the three TLM layers.

Syntactic Layer.

The documented deadlock does not require the assumption that locally relevant relational state information is unavailable or incorrectly represented.

Semantic Layer.

The case does not require a misclassification of the immediate collision relation. The represented conditions can acquire action relevance coherently with the applicable local safety constraint.

Operative Layer.

The controllers generate the actions through which the deadlock develops and are therefore causally involved in the resulting equilibrium. Grover et al. analyze the deadlock within the controller dynamics themselves (Grover et al., 2021). The operative reduction must therefore remain fully available.

Causal involvement, however, is not identical to local functional invalidity. An action can remain admissible with respect to the immediate collision-avoidance constraint while contributing through repeated interaction to loss of progress toward the navigation objective.

If an unsuitable local action policy sufficiently localizes the relevant failure for the analytical purpose under consideration, no additional Cross-Layer Stabilization classification is required.

7.3 Local Validity Is Not Global Correctness

The phrase *locally valid* remains intentionally restricted.

It means that the relevant local processing is admissible with respect to the immediate constraint being examined. In the selected case, that constraint is principally collision avoidance.

It does not mean that the action is optimal, globally correct, or sufficient for completing the navigation objective.

The distinction therefore does not transfer responsibility for a defective local decision to the collective system. It permits only the narrower possibility that locally admissible actions can interact in a way that prevents continued progress toward another objective.

7.4 Recursive Functional Compatibility

Each agent can operate through a locally coherent sequence of represented conditions, context-conditioned action relevance, and admissible operation. The resulting operation changes the relational situation encountered by another agent. That agent can likewise process the changed conditions coherently and act within its local constraint.

Repeated coupling can nevertheless lead to the deadlock equilibrium documented in the technical case (Grover et al., 2021).

The relevant distinction is therefore between:

validity within a local functional process

and

recursive functional compatibility among locally valid processes.

In this study, **recursive functional compatibility** refers to the preservation of compatibility among locally valid functional processes across repeated interaction with respect to the continuing collective objective under consideration.

This definition is diagnostic. It does not introduce a new dynamical variable, coordination mechanism, or technical system property.

7.5 Strong Reduction Test

The strongest reduction remains the conventional system-level account.

The deadlock can already be technically described and analyzed as an equilibrium produced by coupled multi-robot controller dynamics (Grover et al., 2021). The TLM therefore receives no support merely by redescribing that outcome as a lack of stabilization.

A separate diagnostic distinction remains only if the analysis needs to differentiate:

an isolated failure within representation, contextual interpretation, or operation

from

a loss of recursive functional compatibility among processes that remain locally valid with respect to the relevant immediate constraint.

The selected case supports this distinction at the classificatory level. It does not show that the distinction is required for technical explanation.

Cross-Layer Stabilization is therefore supported only conditionally: as a diagnostic classification of compatibility across recursively coupled functional processes, not as an alternative explanation of the underlying system dynamics.

7.6 Equilibrium and Functional Viability

Deadlock is particularly discriminating because it prevents a simple identification of stabilization with successful collective operation.

The technical case reaches an equilibrium while progress toward the agents' goals has ceased (Grover et al., 2021).

The relevant distinction is therefore:

equilibrium does not by itself establish functional viability.

Functional viability remains relative to the objective being assessed. A deadlocked configuration can remain satisfactory with respect to immediate collision avoidance while failing with respect to continued goal-directed progress.

This prevents Cross-Layer Stabilization from being reduced to a claim that the system should merely become dynamically stable.

7.7 Recovery Does Not Prove the Classification

Additional coordination can change the collective trajectory and restore progress in multi-robot navigation settings (Chandra et al., 2026).

That observation does not prove Cross-Layer Stabilization. The recovery mechanism remains technically describable through the coordination method that produces it.

The comparison demonstrates only that a change in interaction organization can alter whether locally generated behavior remains compatible with continued collective progress.

It provides a contrast for the diagnostic analysis, not a mechanism supplied by the TLM.

7.8 Result

The selected deadlock case supports a distinction between **local functional validity** and **recursive functional compatibility**.

No necessary failure must be assigned to representation or contextual interpretation, and individual operations can remain admissible with respect to the immediate safety constraint while repeated reciprocal interaction prevents continued progress toward the navigation goals.

Cross-Layer Stabilization receives **conditional support as a diagnostic classification of this distinction**.

It does not receive support as an independent control mechanism, dynamical law, predictive principle, or alternative technical explanation of deadlock. Established control and multi-agent models remain sufficient for explaining the underlying system behavior.

8. Results

The three stress tests produce a mixed result. The selected case supports functional distinctions within the Three-Layer Model, but the strength and analytical consequence of that support differ across the tests.

8.1 Semantic Layer

The distinction between represented relational state and context-conditioned significance remains reproducible in the selected collision-avoidance case.

Result: The Semantic Layer remains functionally distinguishable without requiring a dedicated semantic technical component.

Boundary: Conventional state-, relation-, constraint-, and control-based descriptions can capture the same dependency. The case therefore provides partial support for functional distinguishability but does not establish the Semantic Layer as irreducible or technically necessary.

8.2 Recursive Role Transition

The functional classification remains coherent when an operation by one agent changes the relational situation that is subsequently represented within another agent's processing context.

Result: Recursive interaction does not cause the distinction among representation, contextual interpretation, and operation to collapse.

Boundary: Comparable transitions among outputs, inputs, states, and subsequent actions are already represented in established distributed-system models (Lynch & Tuttle, 2026). The result supports internal applicability of the TLM under recursive interaction, not novelty.

8.3 Cross-Layer Stabilization

The deadlock case separates local functional validity from recursive functional compatibility.

Locally relevant information can remain available, context-conditioned action relevance can remain coherent with the immediate relational situation, and individual operations can remain admissible

with respect to the collision-avoidance constraint while repeated interaction prevents continued progress toward the navigation goals.

Result: The case supports a diagnostic distinction between validity within locally evaluated functional processes and compatibility among those processes across repeated interaction.

Boundary: Cross-Layer Stabilization does not explain the deadlock mechanism, predict its occurrence, provide a recovery method, or replace established control and multi-agent system descriptions. Its support is conditional and classificatory.

8.4 Combined Result

Stress Test	Result	Boundary
Semantic Layer	Partial support for functional distinguishability between represented state and context-conditioned significance.	Conventional descriptions can capture the same dependency; irreducibility and technical necessity are not established.
Recursive Role Transition	Functional classification remains coherent under recursively interacting processing contexts.	The interaction structure is established in distributed-system models; no novelty is supported.
Cross-Layer Stabilization	Conditional support for distinguishing local functional validity from recursive functional compatibility.	Diagnostic classification only; no independent mechanism, dynamical law, predictive principle, or alternative technical explanation is established.

The combined result is therefore bounded:

Locally valid functional processing can be preserved while repeated interaction fails to preserve recursive functional compatibility with the continuing collective objective.

The TLM can distinguish these levels of functional assessment in the selected case. Established technical models remain sufficient for explaining the underlying system dynamics.

9. Discussion

The results constrain the interpretation of the Three-Layer Model rather than extending it. They also provide a basis for reassessing three unpublished exploratory investigations developed within the broader research program.

These investigations are not part of the evidence used to obtain the results in Section 8. They are not established components of the TLM, and the present study does not validate them by association. The comparison proceeds only from the completed multi-agent analysis to the earlier exploratory concepts.

9.1 Oscillatory Stabilization Hypothesis

An unpublished exploratory hypothesis considered whether recursive semantic systems might pass through oscillatory states during processes of stabilization (Rights of Robots, 2026c).

The present case imposes an important restriction on that idea. Reciprocal multi-agent navigation already provides documented cases in which oscillatory interaction is a problem to be suppressed rather than evidence of successful progression toward a viable state (Snape et al., 2011).

Oscillation therefore cannot by itself be interpreted as evidence of progression toward functional stabilization.

Deadlock establishes a second boundary. The selected technical case reaches an equilibrium before the navigation goals are achieved (Grover et al., 2021). Convergence to or persistence in an equilibrium therefore does not by itself establish functional viability.

The present findings consequently require three properties to remain distinct:

oscillatory behavior;

dynamical stabilization or convergence;

functional viability relative to an objective.

No necessary progression from the first to the second or from the second to the third is supported by the selected case.

The unpublished hypothesis is therefore constrained rather than validated. Oscillation may occur within a trajectory that eventually becomes functionally viable, but the present analysis provides no basis for treating oscillation itself as evidence that such an outcome is occurring.

This also separates the hypothesis from Cross-Layer Stabilization. Oscillatory behavior describes a possible pattern in system-state development. Cross-Layer Stabilization, under the limited result obtained here, concerns diagnostic assessment of functional compatibility across recursively coupled processes. The present study provides no basis for equating them.

9.2 Adaptive Boundaries and Selective Permeability

A separate unpublished theoretical investigation explored whether some recursive systems might preserve coherence through boundaries that selectively regulate relational exchange (Rights of Robots, 2026d).

The present multi-agent case provides no direct evidence for such a mechanism.

Decentralized robots can operate from local rather than globally centralized information (van den Berg et al., 2011; Raibail et al., 2022). That condition alone does not establish a selectively regulating boundary:

partial observability $\neq$ selective permeability.

Likewise, additional coordination among locally affected agents in hybrid navigation approaches can change the organization of interaction (Chandra et al., 2026). Selective participation in such coordination does not demonstrate that a boundary selectively regulates relational exchange:

selective participation $\neq$ selective permeability.

Evidence for the unpublished boundary concept would require an identifiable boundary that selectively and state-dependently regulates relational exchange, together with an identifiable functional consequence of that regulation. The selected case establishes none of these conditions.

The resulting constraint is therefore:

Adaptive coordination should not be classified as an adaptive boundary mechanism unless selective, state-dependent regulation of relational exchange can be demonstrated.

This negative result is independent of the conditional Cross-Layer Stabilization result. Recursive functional incompatibility can be identified without assuming any boundary mechanism.

9.3 Dynamic and Multimodal Context

A third unpublished investigation considered contextual information in dynamic and multimodal systems, including possible forms of contextual representation (Rights of Robots, 2026e). The present case addresses only a limited part of that problem.

The preceding analysis supports a distinction between represented relational state and context-conditioned significance. It also shows that relevant relational conditions can change through prior operations: one agent changes the situation encountered by another, whose response can subsequently alter the conditions encountered by the first.

The case therefore supports a limited observation:

The functional significance of represented state can depend on relational conditions produced through prior interaction.

This is compatible with the study of dynamic context but does not validate a particular context architecture. Conventional controllers can process such relational dependencies without a separate contextual infrastructure.

The case also provides no substantive test of multimodality. It does not investigate integration among visual, linguistic, auditory, symbolic, environmental, or other heterogeneous forms of information as a distinct problem.

Accordingly, the present results support neither a technical requirement for separate context infrastructure nor the multimodal dimension of the unpublished investigation.

9.4 Status of the Unpublished Concepts

The three unpublished investigations receive different outcomes from the present analysis.

The Oscillatory Stabilization Hypothesis is **constrained**: oscillation is not evidence of progression toward functional stabilization, and equilibrium does not establish functional viability.

The adaptive-boundary investigation receives **no direct support**: local information and temporary coordination do not establish selective permeability.

The dynamic-context investigation receives **limited conceptual support** for the dependence of functional significance on changing relational conditions, while multimodality remains untested.

None of these outcomes modifies the Three-Layer Model. None introduces an additional layer or establishes a new mechanism of multi-agent coordination. None is required for the results reported in Section 8.

The comparison therefore functions as a theoretical boundary test rather than an extension of the present model. Any subsequent development of these unpublished concepts would require separate investigation.

10. Implications for Human Context Representation

The findings of the study identify contextual requirements that follow from the distinction between locally valid functional processing and recursive functional compatibility. These implications concern the representation and interpretation of recursively interacting systems. They do not prescribe a particular technical implementation.

10.1 Local State and Collective Trajectory

The selected case demonstrates that locally admissible behavior does not by itself determine the collective trajectory produced through repeated interaction. In decentralized collision avoidance, agents can operate from locally available information and satisfy immediate collision-avoidance constraints without centralized coordination (van den Berg et al., 2011). The deadlock case

demonstrates that such interaction can nevertheless reach an equilibrium before the navigation goals are achieved (Grover et al., 2021).

A contextual representation intended to support interpretation of such behavior must therefore distinguish between the state of a local process and the collective trajectory in which that process participates.

The relevant contextual information is not an alternative description of the local state. It concerns the relation between that state, other interacting processes, and the continuing collective objective.

10.2 Relational Context

The analysis further shows that the functional significance of represented conditions depends on relational context.

Reciprocal collision avoidance evaluates potential conflicts through relations among the positions, velocities, and prospective movements of interacting agents (van den Berg et al., 2011). The significance of another agent's represented state for subsequent action therefore depends on its relation to the receiving agent's own state, prospective movement, objective, and applicable constraints.

For contextual representation, the relevant information is consequently not exhausted by isolated state descriptions. Relations among represented conditions, constraints, objectives, and interacting processes can themselves be necessary for interpreting the functional significance of those conditions.

Such contextual relations do not constitute an alternative system state. They describe how already represented conditions become relevant within a particular interaction context.

10.3 Context Changes Through Recursive Interaction

The relevant context is also dynamic.

An operation by one agent changes the relational condition encountered by another. The resulting operation can alter subsequent conditions again. Reciprocal navigation therefore produces interaction trajectories in which the conditions relevant to later processing are partly consequences of earlier operations (van den Berg et al., 2011).

For contextual interpretation, the relation

operation → changed relational condition → subsequent representation

can therefore be relevant in addition to the individual states involved.

This does not require the operation, resulting condition, and subsequent representation to be reproduced as independent versions of the same information. The contextual significance lies in their relation across the interaction trajectory.

10.4 Constraint-Relative Validity and Objective-Relative Viability

The analysis establishes a further contextual distinction.

Local functional validity is meaningful only relative to the constraint under which it is assessed. Collective viability is meaningful only relative to the continuing objective under consideration.

The technical case makes this distinction observable. Collision avoidance can remain preserved while progress toward navigation goals ceases, and the resulting deadlock can constitute an

equilibrium of the system dynamics (Grover et al., 2021). Contemporary multi-robot navigation research likewise distinguishes safety and liveness as different objectives rather than treating them as interchangeable measures of successful navigation (Chandra et al., 2026).

A contextual representation of validity must therefore retain the constraint relative to which validity is asserted. Likewise, interpretation of a collective state requires the objective relative to which viability is assessed.

Accordingly:

local admissibility $\neq$ dynamical equilibrium $\neq$ functional viability.

These properties can coexist in different combinations and should not be collapsed into a single characterization of system state.

10.5 Contextual Requirements Derived from the Analysis

The preceding findings identify a bounded set of contextual requirements for representing recursively interacting systems.

Interpretation may require relations between local states and collective trajectories; between represented conditions and the constraints under which they acquire functional significance; between locally admissible operations and the objectives against which their collective consequences are evaluated; and between successive operations and the relational conditions they produce.

The central implication is therefore:

Recursive multi-agent interaction can require relational and trajectory-level context to make the collective significance of locally valid functional states explicit.

This requirement follows from the analytical distinction established in the study between local functional validity and recursive functional compatibility.

The analysis determines neither a unique serialization nor a unique technical architecture for representing that context. Those are subsequent implementation questions. What the present result establishes is the contextual distinction that such an implementation would need to preserve if it is intended to represent the analytical findings of this study.

11. Limitations

The conclusions of the study are bounded by the analytical design, the selected technical case, and the explanatory priority assigned to established technical models.

11.1 Single Case Setting

The analysis focuses on decentralized multi-robot coordination, with reciprocal collision avoidance and deadlock as the principal technical setting.

The resulting distinctions should not be generalized automatically to all multi-agent systems, distributed software systems, biological collectives, economic systems, or other forms of recursive interaction.

The study tests the TLM under one demanding interaction setting; it does not establish universal applicability.

11.2 Analytical Rather Than Experimental Study

The study does not introduce a new simulation, experimental dataset, controller implementation, or quantitative evaluation.

It reconstructs behavior documented in the technical literature and uses that behavior as an analytical stress test.

Accordingly, the results concern functional classification and reduction. They do not provide new empirical measurements of collision avoidance, deadlock frequency, convergence, controller performance, or recovery effectiveness.

11.3 Dependence on Established Technical Descriptions

The technical account is deliberately given explanatory priority.

Reciprocal collision avoidance, deadlock formation, distributed interaction, and coordination recovery already possess established technical descriptions in their respective research literatures (van den Berg et al., 2011; Grover et al., 2021; Chandra et al., 2026; Lynch & Tuttle, 2026).

The TLM result must therefore remain compatible with the possibility that the technical behavior is fully explicable without TLM terminology.

The study establishes no alternative model of the underlying system dynamics.

11.4 Semantic Layer Reduction Remains Open

The distinction between represented condition and context-conditioned significance is reproducible in the selected case.

That does not establish irreducibility.

A conventional controller or state-transition description may encode the same functional dependency through state variables, relations, constraints, cost functions, or transition rules without introducing an explicit semantic category.

The study therefore supports functional distinguishability, not technical necessity.

11.5 Recursive Role Transition Provides No Novel Interaction Model

The transition

operation by A → changed relational condition → represented condition for B

is compatible with the TLM functional classification.

It is not novel as an interaction structure. Distributed-system formalisms already accommodate component outputs and actions that participate in the subsequent processing of other components (Lynch & Tuttle, 2026).

The result therefore concerns internal coherence of the functional classification only.

11.6 Cross-Layer Stabilization Remains Diagnostic Classification

The selected case conditionally supports Cross-Layer Stabilization as a distinction between local functional validity and recursive functional compatibility.

It does not establish Cross-Layer Stabilization as an independent control mechanism, dynamical law, predictive variable, coordination algorithm, formal stability criterion, or alternative explanation of multi-agent deadlock.

The underlying deadlock remains technically explicable through established multi-robot system models (Grover et al., 2021).

11.7 Local Validity Is Constraint-Relative

Local validity is not an intrinsic property of an action.

It depends on the constraint selected for evaluation. An action can be locally admissible with respect to collision avoidance while being inadequate with respect to efficiency, progress, global coordination, or another system objective.

Any extension of the analysis must therefore specify the constraint relative to which local validity is asserted.

11.8 Functional Viability Is Objective-Relative

Functional viability likewise depends on the continuing objective under consideration.

In the selected case, the relevant objective is continued progress toward navigation goals. A deadlock equilibrium can remain collision-free while being non-viable relative to that objective (Grover et al., 2021).

The concept should therefore not be generalized into an absolute property of a system state.

11.9 Representational Implications Do Not Determine Implementation

The contextual requirements derived in Section 10 do not determine a unique technical architecture.

The study identifies relations that become relevant when locally valid processing is interpreted within a recursively coupled collective trajectory. It does not establish how those relations must be serialized, exposed, stored, or integrated into a particular implementation.

The distinction is important because the representational implications follow from the analytical results, whereas a specific implementation would introduce additional architectural decisions that have not been tested by the present study.

11.10 Unpublished Exploratory Concepts Remain Unvalidated

The exploratory concepts examined in Section 9 were considered only after the primary results had been established.

The Oscillatory Stabilization Hypothesis is constrained rather than validated. The adaptive-boundary concept receives no direct support from the selected case. The dynamic-context investigation receives limited conceptual support for context-conditioned significance under changing relational conditions, while multimodal context remains untested.

None of these concepts is required to obtain or explain the primary results of the present study.

12. Conclusion

This study tested the Three-Layer Model under recursive decentralized interaction using an established multi-robot coordination case in which locally admissible behavior can coexist with collective loss of progress.

Three stress tests were applied.

First, the distinction between represented relational state and context-conditioned significance remained functionally reproducible. This provides partial support for the Semantic Layer as a functional classification. The result does not establish irreducibility or technical necessity, because conventional state-, relation-, constraint-, and control-based descriptions can represent the same dependency.

Second, the functional classification remained coherent under recursive role transition. An operation by one agent can change a relational condition that subsequently becomes represented within another processing context. This result supports internal applicability of the TLM but provides no novelty over established distributed-system models.

Third, the selected deadlock case supported a distinction between local functional validity and recursive functional compatibility. Locally relevant information can remain available, context-conditioned significance can remain coherent with the immediate relational situation, and individual operations can remain admissible with respect to an immediate safety constraint while repeated interaction prevents continued progress toward the navigation goals.

The central result is therefore:

Locally valid functional processing can be preserved while repeated interaction fails to preserve recursive functional compatibility with the continuing collective objective.

Cross-Layer Stabilization is supported in this case only as a diagnostic classification of compatibility across recursively coupled functional chains. It is not supported as an independent control mechanism, dynamical principle, predictive theory, or substitute for established multi-agent system models.

The technical behavior remains explicable through established models of decentralized collision avoidance, multi-robot deadlock, and distributed interaction. The TLM contribution is therefore classificatory rather than mechanistic.

The study also sharpens several boundaries. Equilibrium does not imply functional viability. Local validity does not imply global correctness. Recursive interaction does not itself indicate failure. Context-conditioned significance does not require anthropomorphic interpretation. A functional distinction does not become a technical component merely because it remains analytically useful.

The resulting contextual implications follow from these findings. Interpreting recursively interacting systems can require relational and trajectory-level context that makes explicit how locally valid functional states relate to constraints, objectives, other processes, and the collective trajectory. The present analysis establishes those contextual requirements without determining their subsequent technical implementation.

The study therefore provides a constrained validation result rather than an expansion of the model. Its contribution lies in testing the existing functional classification under recursive interaction and identifying both where its distinctions remain analytically useful and where established technical descriptions continue to set the explanatory boundary.

References

Chandra, R., Singh, S., Luo, W., & Sycara, K. (2026). Multi-robot navigation in social mini-games: Definitions, taxonomy, and algorithms. *Autonomous Robots*, 50, Article 20. <https://doi.org/10.1007/s10514-026-10251-w>

Grover, J. S., Liu, C., & Sycara, K. (2021). Deadlock analysis and resolution for multi-robot systems. In Proceedings of the 14th International Workshop on the Algorithmic Foundations of Robotics (WAFR) (pp. 294–312). Extended version: arXiv:1911.09146.

Lynch, N. A., & Tuttle, M. R. (2026). An introduction to Input/Output Automata. *Formal Aspects of Computing*, 38(3), Article 25, 1–22. <https://doi.org/10.1145/3788687>

Raibail, M., Abd Rahman, A. H., AL-Anizy, G. J., Nasrudin, M. F., Mohd Nadzir, M. S., Noraini, N. M. R., & Yee, T. S. (2022). Decentralized multi-robot collision avoidance: A systematic review from 2015 to 2021. *Symmetry*, 14(3), 610. <https://doi.org/10.3390/sym14030610>

Snape, J., van den Berg, J., Guy, S. J., & Manocha, D. (2011). The Hybrid Reciprocal Velocity Obstacle. *IEEE Transactions on Robotics*, 27(4), 696–706. <https://doi.org/10.1109/TRO.2011.2120810>

van den Berg, J., Guy, S. J., Lin, M., & Manocha, D. (2011). Reciprocal n-body collision avoidance. In C. Pradalier, R. Siegwart, & G. Hirzinger (Eds.), *Robotics Research* (Springer Tracts in Advanced Robotics, Vol. 70, pp. 3–19). Springer. https://doi.org/10.1007/978-3-642-19457-3_1

Rights of Robots. (2026a). From Structure to Action — A Layer Model for Information-Processing Ecosystems (Version 1.1, revised August 29, 2026). Rights of Robots.

Rights of Robots. (2026b). Validation and Scope of the Three-Layer Model (Version 1.0, August 29, 2026). Rights of Robots.

Rights of Robots. (2026c). Oscillatory Stabilization Hypothesis (Appendix A). Unpublished exploratory manuscript.

Rights of Robots. (2026d). Adaptive Membranes and Selective Permeability: Boundary Formation, Relational Filtering, and Stability Preservation in Recursive Semantic Ecosystems. Unpublished manuscript.

Rights of Robots. (2026e). Multimodal Context Infrastructure: Dynamic Media, Human Interpretation Layers, and Semantic Rooting in Distributed Context Systems. Unpublished manuscript.