Cooperation-Aware Neuromorphic Decision Dynamics for Path-Preserving Navigation in Heterogeneous Crowds

Himani Sinhmar[1*], Antoun Salib[1], and Naomi Ehrich Leonard[1]

***Abstract*— Robots sharing a crossing must avoid collisions without becoming trapped in mutual yielding or repeatedly reversing who passes first. Unknown neighboring control policies make this coordination difficult: a robot must determine when it can rely on others to help resolve a conflict. We present Cooperation-Aware Neuromorphic Decision Dynamics (NDD), a decentralized architecture that infers cooperation from observed motion and coordinates passage by adjusting speed along prescribed paths. We tune the response dynamics to exploit small motion changes before the crossing geometry changes appreciably. Each speed adjustment both responds to a neighbor and changes the evidence for assessing that neighbor's contribution to avoidance. Agents thus establish a traversal order through physical interaction, without exchanging decisions or assigning global priorities. Across 250 simulation rollouts, NDD agents reach their goals without safety violations or path deviation. TurtleBot experiments demonstrate selective yielding and passage around constant-speed neighbors, rather than requiring every controlled robot to wait behind them. The analysis gives conditions for safety and sustained progress in reciprocal NDD groups; extending these guarantees to cooperation-dependent switching and external controllers remains open. The results support an interaction-based view of cooperation and show how traversal timing can accommodate heterogeneous behavior while preserving planned routes.**

I. Introduction

Robots navigating shared spaces encounter people and autonomous systems with different goals and control policies. At a crossing, several agents may approach the same region simultaneously. Each speed adjustment changes the clearance and arrival times available to neighboring agents. Navigation therefore requires both collision avoidance and coordination of access to contested space. Mutual yielding can lead to deadlocks, and repeated changes in passing order can lead to livelocks [1]. A persistent traversal order provides a basis for resolving these interactions. Human-facing navigation also requires attention to how people infer robot intent from motion and respond to communication cues [2], [3].

Negotiating a traversal order requires interpreting a neighbor's speed response in relation to the pair's collision geometry. Slowing down, maintaining speed, or accelerating can each improve or reduce clearance. Cooperation can therefore be represented through the effect of observed motion on collision risk [4], [5]. The robot's speed modulation changes the collision geometry and the evidence available for cooperation inference. Cooperation states and speed commands therefore evolve together in a continuous feedback loop throughout the encounter.

Crowding forces each agent to satisfy several collision constraints through one motion response. A maneuver that clears one neighbor can obstruct another. Even under a shared reciprocal avoidance protocol, ORCA can lack feasible velocities in crowds [6]. Within this reciprocal framework, VR-ORCA allows unequal responsibility, requiring each agent to fulfill its assigned share [7]. NH-ORCA extends reciprocal avoidance to nonholonomic motion without fully guaranteeing goal convergence [8]. Non-cooperative motion can violate the assumptions supporting these safety guarantees. MPC-CBF enforces predictive safety constraints without generally guaranteeing continued feasibility [9]. To address progress explicitly, liveness-constrained control prevents deadlock in social mini-games within specified local geometries [10]. For connected teams, hierarchical control resolves deadlocks through leader assignments over a communication graph [11]. Reinforcement-learning methods train navigation policies in simulation to balance goal-directed progress against collision risk. CADRL approximates multi-agent decisions using a two-agent value function [12]. SARL captures additional crowd interactions through attentive aggregation [13]. Learning from demonstrations also supports prediction of cooperative joint trajectories [14]. Opinion-driven passing uses decision dynamics to resolve pairwise decision deadlock around a selected human [15]. In mixed crowds, AVOCADO adapts to uncertain cooperation, reporting collisions among several competing planners [5]. These advances do not establish persistent collective traversal order in congested heterogeneous groups. Safe, deadlock-free traversal requires compatible responses across simultaneous conflicts.

To address this collective coordination problem, we introduce Cooperation-Aware Neuromorphic Decision Dynamics (NDD), a decentralized architecture for making safe traversal a tunable emergent property of local physical interaction. A slow planner provides each agent's path and nominal motion. The fast embodied layer augments the physical state with pairwise cooperation states c_i^j, a graded proceed-or-yield state z_i, adaptive attention states that regulate sensitivity [16], [17]. These internal states co-evolve continuously with physical motion. Cooperation-dependent conflict evidence drives z_i to modulate speed along the prescribed path. The resulting motion updates the observations driving c_i^j, closing the inference–decision–control loop through physical interaction.

This recurrent feedback exploits transient coordination cues while negotiation is still unfolding. Fast inference,

[1]Himani Sinhmar, Antoun Salib, and Naomi Ehrich Leonard are with the Department of Mechanical and Aerospace Engineering, Princeton University, NJ, USA.

[*]Corresponding author is himani.sinhmar@princeton.edu.

decision, and tracking dynamics integrate small changes in speed and projected arrival time before the crossing geometry changes appreciably. Each agent acts on local egocentric motion observations, without exchanging decisions, assigning global priorities, synchronizing updates, or knowing neighboring control policies. Conflicts activate local responses; their appearance and clearance reconfigure the interacting groups, while disjoint groups negotiate concurrently. The architecture therefore supports decentralized coordination through motion, with sensing and response requirements determined by local congestion and independent of total population size.

The central design mechanism is a controlled loss of stability of collective indecision. Conflict pressure and attention tune a group-level bifurcation: weak encounters leave nominal motion stable, while sufficiently strong conflict evidence amplifies small motion asymmetries into differentiated proceed-and-yield responses. These responses create temporal separation and an acyclic traversal order acorss a locally engaged group. As conflicts clear, the same feedback releases the agents toward nominal motion. Cooperation inference makes this response sensitive to whether each neighbor contributes to collision avoidance, allowing controlled agents to negotiate around non-cooperative agents.

We characterize how local feedback design governs order formation, persistence, and release. For admissible all-NDD groups, the analysis establishes continuous safety, finite-time traversal order formation, and deadlock- and livelock-free progress under bounded errors and changing interaction topologies. Section III specifies the sensing, feedback, tracking, and actuation conditions, together with the additional requirements introduced by cooperation inference and external controllers. Randomized and dense-crossing simulations evaluate NDD's scalability and path preservation. Hardware experiments with TurtleBots demonstrate cooperation-aware temporal coordination around connon-cooperative uncontrolled agents.

II. Cooperation-Aware Decision and Control

A. Local interaction and time-scale design

Each agent must follow its prescribed path toward its goal while maintaining a required center-to-center clearance $d_{\mathrm{safe}} > 0$ from its neighbors. A conflict arises when the agents' current motion predicts a violation of this clearance. At a crossing, agents can resolve the conflict by passing through the shared space at different times. A traversal order specifies which agent passes first in each conflicting pair. We call the process of establishing this order through mutually observed speed adjustments *negotiation*: one agent's motion changes the passage opportunities available to others, whose responses in turn change its own.

These responses convey two kinds of information: whether a neighbor contributes to increasing clearance, which we interpret as cooperation, and whether its motion favors passing ahead or yielding. An agent can use this information before the neighbor completes its maneuver. For example, a small reduction in speed delays the arrival predicted from current motion and may allow another agent to pass first; that agent's ensuing acceleration then reinforces the separation in arrival times. We design cooperation inference and speed decisions to act on these evolving cues, rather than wait for a completed maneuver or an explicit agreement.

This feedback must operate while the observed crossing remains relevant. Let T^{neg} denote the interval from conflict detection to establishment of a traversal order, and let T_{trav} denote the time over which positions and local path directions change appreciably. Let T_{fast} budget the response needed to measure motion, infer cooperation, form a speed decision, and track it. We tune the closed-loop response to satisfy

$$T_{\mathrm{fast}} \lesssim T^{\mathrm{neg}} \ll T_{\mathrm{trav}}, \tag{1}$$

by choosing internal response rates and sensing range jointly with the available speed and braking authority. Earlier detection provides time both to establish an order and to execute the required motion.

Agent i has position $\boldsymbol{q}_i$, velocity $\boldsymbol{v}_i = \dot{\boldsymbol{q}}_i$, heading ψ_i, and speed $v_i = \|\boldsymbol{v}_i\| \in [0, v_{\max}]$, where $v_{\max}$ is the speed limit. A slower planner supplies nominal commands $U_i^{\mathrm{nom}} = (\psi_i^{\mathrm{nom}}, v_i^{\mathrm{nom}})$, with heading tangent to the prescribed path (Fig. 1(A)). Perception constructs the local observation $\mathcal{X}_i(\boldsymbol{q}, \boldsymbol{v})$ from ego motion, sensed relative states, and recent motion history. Here $\boldsymbol{q}$ and $\boldsymbol{v}$ stack the physical states; each agent uses only its local measurements. NDD adjusts speed while low-level control tracks the nominal heading.

To assess neighbor j's contribution to collision avoidance, agent i uses a forward kinematic projection of both agents' current motion. Define the relative position $\boldsymbol{q}_{ij} = \boldsymbol{q}_j - \boldsymbol{q}_i$, the relative velocity $\boldsymbol{v}_{ij} = \boldsymbol{v}_j - \boldsymbol{v}_i$, and the angle $\theta_{ij} \in [0, \pi]$ between them. A pair is closing when $\boldsymbol{q}_{ij}^{\mathsf{T}} \boldsymbol{v}_{ij} < 0$. For a closing pair with $\|\boldsymbol{q}_{ij}\| > d_{\mathrm{safe}}$, this projection gives the closest-approach distance $d_{ij}^{\mathrm{ca}} = \|\boldsymbol{q}_{ij}\| \sin\theta_{ij}$. The *collision cone* consists of relative-velocity directions that would bring the agents within d_{safe}. We define the dimensionless signed directional margin m_{ij} to its boundary as

$$m_{ij} = \sqrt{1 - \left(d_{\mathrm{safe}}/\|\boldsymbol{q}_{ij}\|\right)^2} + \cos\theta_{ij}. \tag{2}$$

For a closing pair, the margin is positive when projected separation exceeds d_{safe}, zero at tangency, and negative when the projection predicts a clearance violation. Agent i continuously evaluates the margin on the perception time scale.

B. Motion-based cooperation evidence

The margin lets agent i distinguish a helpful speed response from one that worsens the encounter. Over a measurement window $T_{\mathrm{fd}} > 0$, it records $\Delta v_j = v_j(t) - v_j(t - T_{\mathrm{fd}})$ and evaluates the sensitivity $g_{ij} = \partial m_{ij}/\partial v_j$. This partial derivative holds positions, ego velocity, and neighbor heading fixed to isolate the neighbor's speed contribution. We choose the window short enough that these quantities vary negligibly, but long enough to resolve speed changes above measurement noise. If g_{ij} retains its sign over the window,

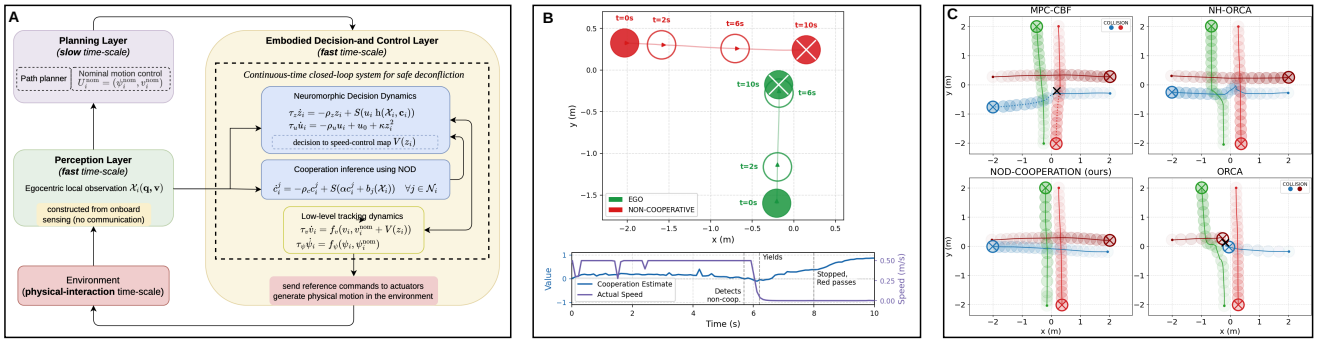


Fig. 1. **Cooperation-aware decision and control: architecture and hardware encounters. (A)** Slow planning supplies nominal motion; fast cooperation inference and NDD regulate speed. Realized motion closes the feedback loop through local observations. **(B)** The green robot brakes after its cooperation estimate becomes negative near $t = 6$ s, then revises the estimate as yielding changes clearance to the red constant-speed robot. **(C)** Each trial contains two controlled and two red constant-speed robots. Faded disks show successive positions; black crosses mark collisions. Cooperation-aware control preserves straight paths. The displayed MPC-CBF and ORCA trials contain collisions; NH-ORCA clears with heading changes.

the sign of $g_{ij}\Delta v_j$ identifies whether the response improves the margin under this approximation. Cooperation therefore depends on collision geometry: acceleration can help, and deceleration can hinder.

When a neighbor maintains speed, this response-based test provides little information. We therefore combine a speed-response cue A_{ij} with a current-clearance cue B_{ij}, using a weight $x_{ij} \in [0,1]$ that grows with the magnitude of the measured speed change. A near-stop cue D_{ij} adds positive evidence for yielding below a speed threshold $\varepsilon_v > 0$. With positive gains $\kappa_s, \kappa_g, \kappa_m, \kappa_x, \kappa_0$ controlling the cue transitions, agent i constructs its evidence $b_j(\mathcal{X}_i)$ about neighbor j as

$$\begin{aligned} A_{ij} &= \tanh(\kappa_s \Delta v_j)\tanh(\kappa_g g_{ij}), \\ B_{ij} &= \tanh(\kappa_m m_{ij}), \qquad x_{ij} = |\tanh(\kappa_x \Delta v_j)|, \\ D_{ij} &= \tfrac{1}{2}\big[1 + \tanh(\kappa_0(\varepsilon_v - v_j))\big], \\ b_j(\mathcal{X}_i) &= x_{ij}A_{ij} + (1 - x_{ij})B_{ij} + D_{ij}. \end{aligned} \tag{3}$$

The blend distinguishes corrective action from an already safe trajectory. A neighbor can have $B_{ij} < 0$ yet $A_{ij} > 0$: its present motion still predicts a clearance violation, but its speed adjustment is reducing that risk. A sufficiently large response shifts the evidence toward the corrective action. Conversely, when Δv_j is small, B_{ij} distinguishes benign constant-speed motion from motion that sustains a conflict. Thus, the robot need not wait for a safe projected trajectory before recognizing a contribution to avoidance.

The near-stop cue encodes a yielding convention rather than a clearance guarantee. It can make b_j positive despite adverse motion evidence when $D_{ij} > |x_{ij}A_{ij} + (1-x_{ij})B_{ij}|$. A stopped neighbor that occupies the prescribed path still requires separate handling.

Agent i accumulates this evidence in a scalar cooperation state c_i^j, rather than changing its assessment with every instantaneous cue. A positive state labels the neighbor cooperative; a nonpositive state withholds reliance on its cooperation. The magnitude expresses persistence of the assessment, not a calibrated probability. Because both agents' motion changes m_{ij} and g_{ij}, this label describes the current interaction rather than a fixed property of the neighbor's controller.

C. Cooperation-dependent decisions and coupled dynamics

To choose who passes first, agent i converts the clearance requirement into occupancy times. The agents' current positions and path headings define projected rays with crossing angle β_{ij}. On each ray, the segment within d_{safe} of the other ray forms a *contested segment*, extending $d_{\text{safe}}/|\sin\beta_{ij}|$ on either side of the intersection. We restrict this construction to crossings bounded away from parallel. For positive speeds, along-path distances to the segment boundaries divided by speed give projected entry and exit times τ_i^{enter} and τ_i^{exit}, with the neighbor index suppressed. We set entry time to zero for an agent already inside. Overlapping occupancy intervals identify a conflict; an exit before the other's entry defines a separated passage. Stopped agents require explicit occupancy handling rather than division by zero, and stopping inside a contested segment does not clear it.

We encode each pair's passing preference in a gate G_{ij}: positive values favor advancing, and negative values favor yielding. For a cooperative neighbor, agent i relies on reciprocal adjustment, in which both agents help separate their passage times. The smooth gate $G_{ij}^{\text{rec}} = \tanh(\kappa_\tau(\tau_j^{\text{enter}} - \tau_i^{\text{enter}}))$, with gain $\kappa_\tau > 0$, favors the earlier projected arrival and vanishes at a tie. Without positive cooperation evidence, agent i assumes the separation burden: it advances only if its projected exit precedes the neighbor's entry by a temporal margin $\varepsilon_\tau > 0$. These two rules give

$$G_{ij}(c_i^j) = \begin{cases} G_{ij}^{\text{rec}}, & c_i^j > 0, \\ +1, & c_i^j \le 0,\ \tau_i^{\text{exit}} + \varepsilon_\tau < \tau_j^{\text{enter}}, \\ -1, & \text{otherwise.} \end{cases} \tag{4}$$

This gate allows passage ahead of a non-cooperative neighbor without assuming that the neighbor will slow down. The test remains conditional on the validity of the motion prediction and tracking bounds; yielding also requires sufficient stopping distance.

Agent i must reconcile its pairwise preferences through one speed command. Let $\mathcal{N}_i$ denote its sensed interaction set and $\boldsymbol{c}_i = (c_i^j)_{j\in\mathcal{N}_i}$. For each neighbor, a nonnegative pressure P_{ij} measures conflict severity, and a nonnegative weight W_{ij} allocates influence among simultaneous conflicts. We normalize $\sum_j W_{ij} = 1$ for a nonempty set; a softmax over pressure provides one choice. Their weighted sum with G_{ij} defines the aggregate input $h(\mathcal{X}_i, \boldsymbol{c}_i)$, with $h = 0$ for an empty set. Separating severity, influence, and direction lets us tune urgency without prescribing who must pass first.

The scalar decision state z_i integrates h; positive z_i commands a speed above nominal and negative z_i commands a speed below nominal. Two adaptive attention states regulate the nonlinear response: α_i^j reinforces the current cooperation assessment, and u_i regulates sensitivity to the aggregate passing preference. For cooperation, we use $\tanh$ and the logistic function $\mathrm{sigm}(a) = (1 + e^{-a})^{-1}$. For decisions, we use a bounded odd sigmoid S. A bounded, continuous, increasing map V, with $V(0) = 0$, converts z_i into a speed offset that respects the speed limits. The low-level laws f_v and f_ψ track commanded speed v_i^{cmd} and nominal heading, respectively.

Let $\rho_c, \rho_\alpha, \rho_z, \rho_u > 0$ be decay coefficients and $\tau_z, \tau_u, \tau_v, \tau_\psi > 0$ be response-scale parameters. The positive gains κ_α and κ govern clearance-dependent cooperation attention and decision-dependent attention reinforcement, respectively; $u_0 > 0$ supplies basal decision attention. With $\alpha_i^j(0) \geq 0$, the complete inference–decision–control loop is

$$\dot{c}_i^j = -\rho_c c_i^j + \tanh\Big(\alpha_i^j c_i^j + b_j(\mathcal{X}_i)\Big), \tag{5a}$$

$$\dot{\alpha}_i^j = -\rho_\alpha \alpha_i^j + \mathrm{sigm}\big(\kappa_\alpha(d_{\mathrm{safe}} - d_{ij}^{\mathrm{ca}})\big), \tag{5b}$$

$$h(\mathcal{X}_i, \boldsymbol{c}_i) = \sum_{j\in\mathcal{N}_i} P_{ij} W_{ij} G_{ij}(c_i^j), \tag{5c}$$

$$\tau_z \dot{z}_i = -\rho_z z_i + S(u_i h(\mathcal{X}_i, \boldsymbol{c}_i)), \tag{5d}$$

$$\tau_u \dot{u}_i = -\rho_u u_i + u_0 + \kappa z_i^2, \tag{5e}$$

$$v_i^{\mathrm{cmd}} = v_i^{\mathrm{nom}} + V(z_i), \tag{5f}$$

$$\tau_v \dot{v}_i = f_v(v_i, v_i^{\mathrm{cmd}}), \tag{5g}$$

$$\tau_\psi \dot{\psi}_i = f_\psi(\psi_i, \psi_i^{\mathrm{nom}}). \tag{5h}$$

The cooperation dynamics balance decay toward neutrality against reinforcement of the current assessment. Worsening projected clearance raises the input to α_i^j, strengthening that reinforcement. To interpret its effect, consider a labeling interval short enough that α_i^j and b_j vary negligibly. This requires tuning their variation relative to the convergence of c_i^j, not merely assuming slow spatial motion. At zero evidence, increasing α_i^j through ρ_c destabilizes the neutral equilibrium and creates two attracting equilibria of opposite signs—a pitchfork bifurcation. Nonzero evidence biases the response toward one sign. Reliable labeling requires persistent valid evidence and an initial state in the corresponding attraction basin, the set of states that converge to that equilibrium. Once established, a label can survive brief evidence reversals; this hysteresis provides persistence without making the label permanent.

Decision attention has a complementary role. The resting value u_0/ρ_u sets sensitivity before a decision forms, while κz_i^2 reinforces an emerging proceed-or-yield response. As conflicts clear and h vanishes, z_i decays toward zero and the command returns to nominal speed. We tune the decay coefficients, attention gains, and response scales together with sensing and actuation to meet (1), including the time needed to establish a cooperation label.

The loop therefore uses motion both to assess cooperation and to establish a passing order. A change in c_i^j can switch the gate, but the decision and tracked motion states remain continuous. The resulting speed change revises occupancy times, projected clearance, and the next cooperation evidence. Agents influence one another through this physical feedback without exchanging internal states or waiting for final decisions.

III. Safe Deconfliction and Scope of Guarantees

Avoiding a collision does not by itself ensure that any agent can advance. We therefore examine how local speed responses establish a compatible passing order across simultaneous conflicts. Define a physical conflict graph with node set $\mathcal{V} = \mathcal{V}_c \cup \mathcal{V}_e$, where $\mathcal{V}_c$ contains NDD-controlled agents and $\mathcal{V}_e$ contains agents with external controllers. An undirected edge joins agents whose contested segments require temporal separation. A connected component collects agents linked directly or through a chain of conflicts; we call it an engaged group. The traversal-order graph directs each edge $i \to j$ when i passes before j with the required clearance margin. If this graph has no directed cycle, at least one agent has no predecessor. Progress then requires that this leading agent can move and clear a conflict.

First consider a group in which every agent uses NDD and every gate uses the reciprocal rule. At a group tie, all pairwise arrival preferences vanish. The equilibrium $\boldsymbol{z} = \boldsymbol{0}$, with nominal speeds and attention $u_i = u_0/\rho_u$, represents *collective indecision*: no agent has a passing preference. Under (1), we hold positions, path headings, and the interaction graph locally fixed while retaining feedback through changing speeds. Let $\mathcal{X}_{\boldsymbol{z}}$ denote observations at speeds $v_i^{\mathrm{nom}} + V(z_i)$, and let $\boldsymbol{h}$ stack the agents' aggregate inputs. For locally differentiable evidence and speed maps, $\boldsymbol{A} = D_{\boldsymbol{z}}\boldsymbol{h}(\mathcal{X}_{\boldsymbol{z}})|_{\boldsymbol{z}=\boldsymbol{0}}$ measures how a small decision-induced speed change alters the group's evidence. Linearizing the reduced equilibrium equation at indecision gives the Jacobian $\boldsymbol{J}$ and, for $S'(0) > 0$, the critical eigenvalue λ_*:

$$\boldsymbol{J} = -\rho_z \boldsymbol{I} + S'(0)\frac{u_0}{\rho_u}\boldsymbol{A}, \qquad \lambda_* = \frac{\rho_z \rho_u}{S'(0) u_0}. \tag{6}$$

Suppose the dominant eigenvalue of $\boldsymbol{A}$ is real and simple. Below λ_*, damping suppresses small decision differences; above it, motion-induced feedback amplifies them. We tune conflict pressure and basal attention so that weak encounters retain nominal motion, while sufficiently strong conflicts destabilize indecision. Small asymmetries in arrival times, geometry, or inherited decisions select a differentiated response. The order-formation result requires this response to

approach an attracting branch, a family of stable equilibria with different decisions across conflicting agents, within T^{neg}.

For positive nominal speeds, the admissible speed maps produce distinct fractional speeds v_i/v_i^{nom} across conflict edges on this branch. Under the temporal-clearance conditions, these differences establish the passing order: fractional speed decreases along each directed edge, so a directed cycle cannot occur. A leading agent with a positive lower speed bound clears its finite contested segment in finite time. Removing that conflict preserves the remaining orders, and successive clearances resolve a fixed finite conflict set. When new conflicts appear or groups merge, agents renegotiate on the updated graph. Finite conflict activations and event times that do not accumulate in finite time allow the episode to terminate. Persistent orders prevent repeated reversals on unresolved conflicts; leading movers prevent mutual waiting. After the final conflict clears, $z_i \to 0$, $u_i \to u_0/\rho_u$, and speed returns to nominal.

Safety requires more than this ordering argument. Valid nonparallel projections, bounded sensing and tracking errors, bounded local congestion, and attainable speed changes must support the required temporal margins. Sensing must provide enough time and distance for each potential follower to yield or stop before entering its contested segment while the order forms. Together with sufficiently fast attraction to the differentiated response, these conditions give continuous clearance throughout negotiation and graph changes, not only after agents establish an order. Cooperation inference adds the measurement window T_{fd} and the time needed to establish a reliable label to this response budget. We must therefore tune sensing and response rates jointly with braking or passage requirements, while ensuring that the evidence remains valid during labeling.

The reciprocal guarantee does not automatically extend to cooperation-dependent gate switching. A label change alters the feedback that establishes and preserves the order; the full loop therefore requires a separate verification of clearance during labeling and compatibility of the resulting passing decisions. External controllers add two further limits. Controlled robots cannot enforce safety between two external agents, and a speed-only yielding policy cannot ensure progress if an external agent stops on the prescribed path. External motion must remain within the prediction bounds and eventually clear the contested region. Proving safety and progress under these conditions for specified external policies remains open. The experiments assess the resulting behavior without treating successful encounters as a general guarantee.

IV. Simulation and Hardware Evaluation

We use simulation to test whether speed-based coordination sustains progress across many simultaneous conflicts, and hardware to examine how cooperation inference changes the response to independently controlled neighbors. In simulation, homogeneous Dubins agents encounter the same 50 randomized start–goal scenarios for each group size $N \in \{4, 8, 12, 24, 32\}$. Across 250 rollouts, every NDD agent reaches its goal without a safety violation and preserves its prescribed path. NDD also exceeds ORCA, LivePoint, and decentralized CBF-QP in lane-normalized safe throughput, which measures violation-free goal arrivals per unit time and per lane. Thus, path preservation does not come at the expense of lower safe throughput in these tests. In a dense ten-agent near-antipodal exchange, NDD again resolves overlapping conflicts through speed differentiation. These simulations evaluate the decision–control mechanism, rather than isolate the contribution of cooperation inference.

The hardware experiments use TurtleBot 3 Waffle robots at shared crossings. Vicon motion capture provides every robot with all neighbors' measured positions and velocities, so these tests assess motion-based feedback under external state measurement rather than onboard perception. The controller implements the inference and speed-regulation structure of (5), with pressure weighting and a discrete reciprocal arrival-order gate that resolves near-simultaneous arrivals through tie breaking, instead of the smooth reciprocal example in (4). The comparison controllers—ORCA, NH-ORCA, and MPC-CBF—can change both speed and heading.

The two-agent encounter in Fig. 1(B) shows the coupling between inference and action. The controlled robot initially moves near nominal speed while its cooperation estimate remains near zero. Around $t = 6$ s, the estimate becomes negative and the robot brakes to a stop. The estimate subsequently rises as yielding changes projected clearance and lets the constant-speed neighbor pass. The neighbor never changes its policy. The changing estimate therefore reflects the evolving interaction, including the controlled robot's own response, rather than a change in the neighbor's willingness to cooperate.

The four-robot trials in Fig. 1(C) test this response across overlapping conflicts with two controlled and two constant-speed robots. In the cooperation-aware trial, the blue robot passes ahead of one constant-speed robot, while the green robot yields to the blue robot and the other constant-speed robot. Both controlled robots preserve their straight paths. The response is selective: the controller does not force every robot to wait behind every non-cooperative neighbor, but permits passage when the projected timing allows it. The displayed MPC-CBF and ORCA trials contain collisions with constant-speed robots; NH-ORCA clears the encounter by changing headings. These runs do not establish hardware success rates or identify which controller component causes the differences; repeated trials and an ablation of cooperation inference are needed for those conclusions.

V. Discussion

Cooperation-aware NDD treats a neighbor's motion as both evidence for inference and part of the feedback that shapes the next action. Its purpose is not to assign permanent behavioral types, but to decide when an agent can rely on reciprocal adjustment and when it must assume the separation burden. In reciprocal groups, tunable feedback amplifies small motion differences into a persistent passing order and releases agents toward nominal motion after conflicts clear.

Cooperation inference adapts the pairwise response to heterogeneous behavior, while the unresolved switching conditions delimit the corresponding guarantees. The broader design principle is to shape collective coordination through the strength and timing of local feedback rather than prescribe a global sequence of actions.

Separating planning from fast interaction assigns path selection and traversal timing to different layers. A slower planner can handle static obstacles, direct head-on encounters, and blocked paths that speed regulation alone cannot resolve. For road vehicles, aerial robots, or manipulators following prescribed motions, analogous timing control could coordinate access to shared lanes, corridors, or workspaces. Such extensions require evidence maps that represent the relevant conflicts and control maps that respect platform-specific actuation. They also require new checks of sensing range, response time, and order persistence; the navigation guarantees do not transfer solely because the architecture has the same form.

Additional observations could let the evidence map account for intent, mobility constraints, social conventions, or explicit communication. Learning could adapt the associated parameters while enforcing verified safety and response-time conditions. In human–robot encounters, evaluation must test how people interpret the resulting motion, how variable responses and onboard sensing errors affect the evidence, and whether the speed adjustments preserve comfort. Acceleration, jerk, stopping duration, and passing-order reversals would quantify motion fluidity beyond collision avoidance and goal completion. These measurements would test whether timely, persistent speed decisions make coordination more predictable as well as physically feasible.

REFERENCES

[1] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 797–803.

[2] A. D. Dragan, K. C. T. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction*, 2013, pp. 301–308.

[3] Y. Che, A. M. Okamura, and D. Sadigh, "Efficient and trustworthy social navigation via explicit and implicit robot–human communication," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 692–707, 2020.

[4] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human-robot cooperation," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.

[5] D. Martinez-Baselga, E. Sebastián, E. Montijano, L. Riazuelo, C. Sagüés, and L. Montano, "Avocado: Adaptive optimal collision avoidance driven by opinion," *IEEE Transactions on Robotics*, vol. 41, pp. 2495–2511, 2025.

[6] J. van den Berg, S. J. Guy, M. C. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research*, ser. Springer Tracts in Advanced Robotics. Springer, 2011, vol. 70, pp. 3–19.

[7] K. Guo, D. Wang, T. Fan, and J. Pan, "Vr-orca: Variable responsibility optimal reciprocal collision avoidance," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4520–4527, 2021.

[8] J. Alonso-Mora, A. Breitenmoser, M. Rufli, P. Beardsley, and R. Siegwart, "Optimal reciprocal collision avoidance for multiple non-holonomic robots," in *Distributed Autonomous Robotic Systems*, ser. Springer Tracts in Advanced Robotics. Springer, 2013, vol. 83, pp. 203–216.

[9] J. Zeng, B. Zhang, and K. Sreenath, "Safety-critical model predictive control with discrete-time control barrier function," in *2021 American Control Conference (ACC)*, 2021, pp. 3882–3889.

[10] R. Chandra, V. Zinage, E. Bakolas, P. Stone, and J. Biswas, "Deadlock-free, safe, and decentralized multi-robot navigation in social mini-games via discrete-time control barrier functions," *Autonomous Robots*, vol. 49, no. 2, p. 12, 2025.

[11] K. Garg, S. Hamilton, and C. Fan, "Deadlock resolution of connected multi-agent systems using hierarchical control," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 2024, pp. 1275–1282.

[12] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 285–292.

[13] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6015–6022.

[14] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.

[15] C. Cathcart, M. Santos, S. Park, and N. E. Leonard, "Proactive opinion-driven robot navigation around human movers," 2023. [Online]. Available: https://arxiv.org/abs/2210.01642

[16] A. Bizyaeva, A. Franci, and N. E. Leonard, "Nonlinear opinion dynamics with tunable sensitivity," *IEEE Transactions on Automatic Control*, vol. 68, no. 3, pp. 1415–1430, 2023.

[17] N. E. Leonard, A. Bizyaeva, and A. Franci, "Fast and flexible multiagent decision-making," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, pp. 19–45, 2024. [Online]. Available: https://www.annualreviews.org/content/journals/10.1146/annurev-control-090523-100059