

A Real-Time Generalized Nash Equilibrium Framework for Interaction-Aware Autonomous Driving in Mixed Traffic

Nouhed Naidja

Institut VEDECOM, Versailles, France
nihed.aidja@vedecom.fr

Mohamed-Cherif Rahal

Institut VEDECOM, Versailles, France
mohamed.rahah@vedecom.fr

Steve Pechberti

Institut VEDECOM, Versailles, France
steve.pechberti@vedecom.fr

Stéphane Font

Laboratoire des Signaux et Systèmes (L2S)
stephane.font@centralesupelec.fr

Guillaume Sandou

Laboratoire des Signaux et Systèmes (L2S)
guillaume.sandou@centralesupelec.fr

Marc Revilloud

Dotflow
marc.revilloud@dotflow.fr

Abstract—Safe and efficient navigation in mixed-traffic environments remains a critical challenge for Autonomous Vehicles (AVs), primarily due to the complex interdependence between the AV’s decisions and the unpredictable reactions of human drivers. This paper introduces a comprehensive decision-making framework that formulates the driving interaction as a Generalized Nash Equilibrium Problem (GNEP). Unlike decoupled optimization approaches, this framework explicitly models shared safety and geometric constraints, ensuring that the feasibility of the AV’s strategy is dynamically linked to the opponent’s actions. To solve this non-convex problem in real-time, we propose a dedicated solver based on Particle Swarm Optimization (PSO). The complete architecture was validated on a test track using a real autonomous Renault Zoé interacting with a human driver. Experimental results demonstrate the system’s ability to handle critical scenarios by generating comfortable, human-like trajectories. Benchmarks confirm the solver’s operational feasibility, achieving convergence in under 50 ms.

I. INTRODUCTION

As the transition towards autonomous mobility progresses, the interaction between autonomous vehicles (AVs) and human drivers will create unique traffic dynamics. In this context, the paper in [1] lists a number of human drivers’ expectations regarding the behavior of autonomous vehicles. They expect autonomous vehicles to: (i) be aware of their intentions, (ii) act consistently with respect to their expectations, and also (iii) exhibit a behavior recognizable and understandable by human drivers. Failure to meet these expectations could lead to mistrust, reticence, and even a rejection of cohabitation with autonomous vehicles. In such a context, it is no longer sufficient for autonomous vehicles to merely adhere to rules or optimize predefined objectives independently. They need to be able to understand, predict, and adapt to the behavior of others. These observations highlight the need for a decision-making framework that will enable autonomous vehicles to choose actions that are not only safe and efficient but also interaction-aware. Game-theoretic-based approaches have particularly drawn researchers’ interest, aiming to formalize and analyze the interdependencies that may arise in transportation systems [2]. In line with this research direction, this paper proposes a decision-making framework, grounded in game theory, that shifts the paradigm from egocentric optimization to multi-agent interactive planning. By explicitly modeling the

coupled nature of traffic, we aim to bridge the gap between algorithmic safety and human-like negotiation.

II. BACKGROUND

Driving is inherently game-theoretic: each driver’s decisions depend on predictions of others’ reactions, creating a complex interdependent decision-making environment [2]. Early formulations model such interactions as two-player zero-sum games [3], [4], combining proactive and reactive strategies for robust decision-making. However, traffic scenarios such as merging or intersection crossing do not fit the zero-sum framework, as mutual benefit is possible. Non-zero-sum dynamic games [5] better capture these collaborative aspects by allowing simultaneous gains for all players. Within game theory, the Nash equilibrium defines a strategy profile where no player can unilaterally improve their payoff given fixed strategies of others [6]. Yet, this classical formulation assumes independent strategy sets, which is rarely satisfied in traffic due to shared safety constraints and road geometry [7]. To address these dependencies, the Nash equilibrium is extended to Generalized Nash Equilibrium Problems (GNEPs), where both objective functions and feasible strategy sets depend on the opponents’ strategies [8].

The remainder of this paper is organized as follows: Section III details the proposed decision-making model, Section IV develops the PSO-based solver, and Section V assesses its performance through simulation experiments.

III. METHODS

A. Multi-Agent Interaction as a Game

We frame the autonomous vehicle’s decision problem as a two-player, non-zero-sum game. The Ego vehicle, denoted v , chooses a strategy $s_v \in \mathbb{R}^n$ within a feasible strategy set $\mathcal{S}_v$. The opponent’s strategy is denoted $s_o \in \mathbb{R}^m$. Where n and m are the decision variables, characterizing the vehicle’s motion (control point of trajectory, velocity profile parameters). Considering this mutual interdependence, we define a joint strategy space $\mathcal{S}_{\text{joint}} = \mathcal{S}_v \times \mathcal{S}_o \subseteq \mathbb{R}^{n+m}$ that encompasses all possible pairs of feasible strategies (s_v, s_o) , and provides the domain in player’s interactions can be jointly analyzed. Within this strategic space, each player $i \in \{v, o\}$ incurs a

cost $Q_i(s_v, s_o) : \mathbb{R}^{n+m} \rightarrow \mathbb{R}$ that it aims to minimize over its feasible strategy set, given the strategy of the other player.

Within this strategic space, each player $i \in \{v, o\}$ incurs a cost $Q_i(s_v, s_o)$ that it seeks to minimize. In our formulation, the structure of Q_i is derived directly from our previous work [9], where we employ a multi-criteria approach to evaluate the quality of a generated trajectory.

Specifically, each cost function combines three components:

- **Safety:** geometric interaction metrics between adaptive elliptical safety zones, capturing minimum-distance and overlap severity.
- **Comfort:** longitudinal and lateral jerk indicators evaluated along the trajectory.
- **Efficiency:** global travel time to complete the maneuver.

These components are aggregated via a weighted sum, complemented by jerk constraint penalty terms. The functions Q_v and Q_o share this structure, each evaluated from the respective agent's perspective given the opponent's predicted strategy.

Then, we assume that each player is motivated to choose the best response available to optimize its own utility [10]. Let the joint best response operator BR define a mapping which assigns to each pair of strategies the corresponding pair of optimal responses.

$$BR : S_v \times S_o \rightarrow S_v \times S_o \quad (1)$$

$$(s_v, s_o) \mapsto (BR_v(s_o), BR_o(s_v))$$

Considering the best response concept, BR_v and BR_o correspond to the individual optimization problems solved by each player. Specifically, given the opponent's choice, the best responses are:

$$BR_v(s_o) = \arg \min_{s_v \in S_v} Q_v(s_v, s_o) \quad (2)$$

$$BR_o(s_v) = \arg \min_{s_o \in S_o} Q_o(s_v, s_o). \quad (3)$$

Here BR_v defines a mapping from S_o to S_v , assigning to each strategy of player o the optimal response of player v . Likewise, BR_o defines a mapping from S_v to S_o . The joint best response operator in Equation (1) aggregates the two individual best responses Equations (2) and (3). It assigns to every pair (s_v, s_o) the corresponding pair of mutually optimal responses.

B. Shared Constraints and Feasible Strategy Sets

Unlike standard optimization problems, where decision variables are independent, in our problem, the strategies of the ego and opponent vehicles are coupled through safety, kinematic, and road geometry constraints. This means that the feasibility of a strategy of an agent depends directly on the choice of the other one. This coupling is expressed through a set of shared inequality constraints of the form:

$$h(s_v, s_o) \leq 0, \quad (4)$$

where $h : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^k$ denotes a vector-valued function capturing all constraint components (collision avoidance, road boundaries, dynamic limits, etc.). Such constraints directly

involve both players' strategy variables, and therefore make each player's feasible strategy set explicitly dependent on the opponent's decision. Formally, the feasible sets become:

$$S_v(s_o) = \{s_v \in S_v \mid h(s_v, s_o) \leq 0\}, \quad (5)$$

$$S_o(s_v) = \{s_o \in S_o \mid h(s_v, s_o) \leq 0\}. \quad (6)$$

The sets $S_v(s_o)$ and $S_o(s_v)$ arise from the shared constraints (4) and illustrate the essence of a Generalized Nash Equilibrium Problem: each player optimizes over a feasible region that itself depends on the opponent's strategy.

Considering that each player's strategy set S_i may depend on the strategies of other players s_{-i} , a strategy profile $s^* = (s_1^*, s_2^*, \dots, s_n^*)$ is a generalized Nash Equilibrium if, for every player $i \in \{v, o\}$:

$$\forall s_i \in S_i(s_{-i}^*), \quad s_i^* \in S_i(s_{-i}^*), \quad Q_i(s_i^*, s_{-i}^*) \geq Q_i(s_i, s_{-i}^*) \quad (7)$$

A generalized Nash equilibrium is then characterized by a pair $(s_v^*, s_o^*) \in S_v \times S_o$ of this mapping satisfying:

$$(s_v^*, s_o^*) = BR(s_v^*, s_o^*) \quad (8)$$

This definition ensures that each player's optimal strategy is not only the best response to the strategies of others but also respects the shared constraints that link the players' strategy.

C. Characterizing the Generalized Nash Equilibrium

1) Deviation Analysis and Equilibrium Consistency:

Now, we want to evaluate whether an agent could unilaterally improve its outcome by deviating. To capture this dynamic, we assess how an arbitrary strategy pair (s_v, s_o) relates to the equilibrium condition (s_v^*, s_o^*) . If (s_v, s_o) coincide with (s_v^*, s_o^*) , then agents' unilateral deviation is not beneficial. To capture this dynamic, we define a deviation measure $\Delta(s_v, s_o)$, one for each player, to measure the largest possible payoff improvement that a player could obtain by switching from a strategy s_i to an alternative one $\bar{s}_v$, while the other player keeps its strategy fixed. The deviation measures are defined as follows:

$$\Delta_v(s_v, s_o) = \sup_{\bar{s}_v \in S_v} (Q_v(s_v, s_o) - Q_v(\bar{s}_v, s_o)) \quad (9)$$

$$\Delta_o(s_v, s_o) = \sup_{\bar{s}_o \in S_o} (Q_o(s_v, s_o) - Q_o(s_v, \bar{s}_o)) \quad (10)$$

Equation (9) represents the largest possible payoff improvement for player **v** if it switches from its current strategy s_v to another strategy $\bar{s}_v$, while the player **o** keeps its strategy fixed at s_o . Similarly, the measure (10) represents the payoff improvement for player **o** when switching from s_o to $\bar{s}_o$, while keeping s_v unchanged. From the player **v**'s perspective, if the difference $(Q_v(s_v, s_o) - Q_v(\bar{s}_v, s_o))$ is positive, it indicates that it could benefit (minimize its cost) from deviating to an alternative strategy $\bar{s}_v$. Conversely, a negative difference indicates that the current strategy s_v is already preferable to $\bar{s}_v$, and thus there is no incentive to deviate. The same applies to the player **o** in Equation (10).

2) **Deviation Measures under Shared Feasibility:** In the generalized Nash setting, the feasible strategy sets depend on the opponent's decision through the shared constraints introduced earlier. To remain consistent with this structure, we restrict unilateral deviations to strategies that satisfy the coupled feasibility conditions. Accordingly, the deviation measures are extended as follows:

$$\Delta_v(s_v, s_o) = \begin{cases} \sup_{\bar{s}_v \in S_v(s_o)} (Q_v(s_v, s_o) - Q_v(\bar{s}_v, s_o)), & \text{if } (s_v, s_o) \in \mathcal{S}_{\text{joint}}, \\ +\infty, & \text{otherwise.} \end{cases} \quad (11)$$

With an analogous definition for $\Delta_o(s_v, s_o)$. This compact formulation enforces that only deviations respecting the shared feasibility constraints are admissible, while infeasible configurations are penalized.

When considered separately, the deviation measures $\Delta_v(s_v, s_o)$ and $\Delta_o(s_v, s_o)$ only capture each player's unilateral incentive to deviate, and provide information on a strategy profile concerning a single agent at a time. Thus, they don't provide information regarding the joint dynamics under which the two players simultaneously converge towards equilibrium. However, a Nash equilibrium is defined precisely as a configuration in which neither player has an incentive to deviate simultaneously. Consequently, combining both perspectives into a unified metric is necessary. In light of these considerations, the next section introduces an approach that simultaneously searches for a pair of strategies (s_v^*, s_o^*) satisfying both players' optimality conditions and the joint feasibility constraints.

D. Characterizing Nash Equilibria via a Joint Criterion

We propose formulating the autonomous vehicle's decision-making task as a joint optimization problem. To this end, we formulate the generalized Nash equilibrium such that the ego player, the vehicle v , conceptually solves the optimization problems not only for itself but also for its opponent, the vehicle o . In this framework, we define a joint criterion $J(s_v, s_o)$, expressed as the aggregation of two deviation measures Δ_v and Δ_o , as follows:

$$J(s_v, s_o) = \Delta_v(s_v, s_o) + \Delta_o(s_v, s_o) \quad (12)$$

In the definition (12), each deviation measure $\Delta_v(s_v, s_o)$ and $\Delta_o(s_v, s_o)$ is defined using a supremum over all possible strategies, including the current strategy s_v (respectively s_o). This guarantees that the difference $Q_v(s_v, s_o) - Q_v(s_v, s_o) = 0$ is taken into account, thus ensuring that the supremum is non-negative. Thus, we have:

$$\Delta_v(s_v, s_o) \geq 0 \quad \text{and} \quad \Delta_o(s_v, s_o) \geq 0 \quad (13)$$

$$J(s_v, s_o) = \Delta_v(s_v, s_o) + \Delta_o(s_v, s_o) \geq 0 \quad (14)$$

To make this property explicit, we introduce the maximum operator $\max(\cdot, 0)$ as follows:

$$J(s_v, s_o) = \max(\Delta_v(s_v, s_o), 0) + \max(\Delta_o(s_v, s_o), 0) \quad (15)$$

The criterion $J(s_v, s_o)$ Eq. (15) represents the sum of the maximum unilateral improvements available to both players; it measures the improvement of deviating from the actual strategy profile. Consequently, when $J(s_v, s_o)$ is close to zero, the opportunity for improvement is minimal for both players. In this sense, the value of $J(s_v, s_o)$ provides an intuitive notion of the distance between the current strategy profile and a Nash equilibrium.

Consequently, the search for an equilibrium can be framed as an optimization problem over the joint strategy space as follows:

$$\min_{\substack{\bar{s}_v \in S_v, \bar{s}_v \neq s_v \\ \bar{s}_o \in S_o, \bar{s}_o \neq s_o}} J(s_v, s_o) \quad (16)$$

The optimization problem Eq. (16) may be computationally intensive, since the criterion involves an evaluation over two nested supremum operations, leading to a minimax formulation, which can be very costly. To accommodate the criterion for numerical evaluation, we introduce a sequence of simplification and approximation.

We first exclude the trivial comparison $\bar{s}_v = s_v$ and $\bar{s}_o = s_o$, therefore reducing redundancy in the deviation measures. We redefine the criterion as:

$$J(s_v, s_o) = \underbrace{\max \left(\sup_{\bar{s}_v \in S_v, \bar{s}_v \neq s_v} \{Q_v(s_v, s_o) - Q_v(\bar{s}_v, s_o)\}, 0 \right)}_{\Delta_v(s_v, s_o)} + \underbrace{\max \left(\sup_{\bar{s}_o \in S_o, \bar{s}_o \neq s_o} \{Q_o(s_v, s_o) - Q_o(s_v, \bar{s}_o)\}, 0 \right)}_{\Delta_o(s_v, s_o)} \quad (17)$$

Since the supremum Eq. (17) is evaluated over continuous strategy spaces, it may be computationally demanding. This is particularly problematic in optimization contexts, where the evaluation of J may need to be repeated during the search for an optimal solution. To tackle this problem, we suggest approximating the criteria by considering a finite set of sampled strategies. Formally, we replace the supremum operators with a maximization over sampled subsets $\hat{S}_v \subset S_v$ and $\hat{S}_o \subset S_o$. This yields the following discrete deviation measures:

$$\hat{\Delta}_v(s_v, s_o) = \max \left(\max_{\bar{s}_v \in \hat{S}_v, \bar{s}_v \neq s_v} (Q_v(s_v, s_o) - Q_v(\bar{s}_v, s_o)), 0 \right) \quad (18)$$

$$\hat{\Delta}_o(s_v, s_o) = \max \left(\max_{\bar{s}_o \in \hat{S}_o, \bar{s}_o \neq s_o} (Q_o(s_v, s_o) - Q_o(s_v, \bar{s}_o)), 0 \right) \quad (19)$$

From these, the approximated joint criterion is given by:

$$\hat{J}(s_v, s_o) = \hat{\Delta}_v(s_v, s_o) + \hat{\Delta}_o(s_v, s_o) \quad (20)$$

Since the maximization is performed over subsets of the original domains, the discrete values form the lower bounds of their exact counterparts: $\hat{\Delta}_v(s_v, s_o) \in [0, \Delta_v(s_v, s_o)]$, $\hat{\Delta}_o(s_v, s_o) \in [0, \Delta_o(s_v, s_o)]$. As a consequence, the global criterion also satisfies $\hat{J}(s_v, s_o) \in$

$[0, J(s_v, s_o)]$. This method yields a stochastic estimator of the exact criterion $J(s_v, s_o)$, it reduces the computational cost of the evaluation process while preserving information and informative content. The characteristics of the designed estimator $\hat{J}(s_v, s_o)$ suggest using evolutionary algorithms, which are well-suited for optimization problems with nondifferentiable and highly irregular cost functions.

IV. SOLVING THE GAME VIA META-HEURISTICS

When modeling complex game structures with nonlinear or non-convex functions, traditional solvers may not guarantee global convergence. Recent works [11], [12] address this by employing an augmented Lagrangian formulation to solve constrained multi-player general-sum dynamic games, seeking points where the gradients of all augmented Lagrangian functions vanish. While offering near real-time performance and stable convergence, these solvers only guarantee local convergence satisfying first-order optimality conditions.

In light of these limitations, metaheuristic approaches emerge as a promising alternative. As emphasized in [13], combining game theory with metaheuristics yields strong results for complex equilibrium problems, offering a balanced compromise between optimality and individual objectives [14]. Among these, Particle Swarm Optimization (PSO) has proven particularly effective for multiplayer games with interconnected objectives. A concrete example is provided in [15], where game-theoretic PSO improves autonomous navigation in robotic search and rescue scenarios.

A. Problem Reformulation in the PSO Context

In this section, we propose an adaptation of the Particle Swarm Optimization (PSO) algorithm to search for a Nash equilibrium. The goal of the PSO-based solver is to minimize $\hat{J}(s_v, s_o)$ in Eq. (20). Let a swarm of P particles explore a search space of dimension D . Each particle $p \in \{1, \dots, P\}$ represents a candidate strategy profile of the form (s_v, s_o) within the search space. Since we are seeking to compute a Nash equilibrium, each particle can be considered as a potential equilibrium candidate. We denote its associated strategic profile by:

$$\mathbf{X}_p = (s_v, s_o) \in S_v \times S_o. \quad (21)$$

We decompose each strategy into two components, namely, finite-dimensional vectors $\mathcal{P}_p^{v,k}$ and $\mathcal{P}_p^{o,k}$ that parametrize the geometric paths of the autonomous vehicle (agent v) and the human driver o , respectively, and the corresponding velocity vectors $\mathcal{V}_p^{v,k}$ and $\mathcal{V}_p^{o,k}$. The explicit value of particle p at generation k is then defined as $\mathbf{X}_p^k = (\mathcal{P}_p^{v,k}, \mathcal{V}_p^{v,k}, \mathcal{P}_p^{o,k}, \mathcal{V}_p^{o,k})$. The objective of the PSO algorithm is to iteratively update these vectors according to PSO dynamics, in order to minimize the joint deviation measure $\hat{J}$. For conciseness, we introduce the shorthand notation

$$z_p^{v,k} = (\mathcal{P}_p^{v,k}, \mathcal{V}_p^{v,k}), \quad z_p^{o,k} = (\mathcal{P}_p^{o,k}, \mathcal{V}_p^{o,k}),$$

so that $\mathbf{X}_p^k = (z_p^{v,k}, z_p^{o,k})$ denotes the full strategy profile of particle p at iteration k . The criterion to minimize is then:

$$\hat{J}(\mathbf{X}_p^k) = \max \left\{ \max_{\substack{\bar{z}^{v,k} \in \hat{S}_v \\ \bar{z}^{v,k} \neq z_p^{v,k}}} [Q_v(z_p^{v,k}, z_p^{o,k}) - Q_v(\bar{z}^{v,k}, z_p^{o,k})], 0 \right\} \\ + \max \left\{ \max_{\substack{\bar{z}^{o,k} \in \hat{S}_o \\ \bar{z}^{o,k} \neq z_p^{o,k}}} [Q_o(z_p^{v,k}, z_p^{o,k}) - Q_o(z_p^{v,k}, \bar{z}^{o,k})], 0 \right\} \quad (22)$$

With $\hat{S}_v$ and $\hat{S}_o$ are finite-dimensional subsets obtained via random sampling from S_v and S_o , respectively.

$\hat{J}(\mathbf{X}_p^k)$ is an approximation of the continuous criterion $J(\mathbf{X}_p^k)$, which would result from computing the true supremum over S_v and S_o , rather than through sampled maxima. Moreover, we note the important property that $\hat{J}(\mathbf{X}_p^k) \in [0, J(\mathbf{X}_p^k)]$, which implies that the expectation $\mathbb{E}[\hat{J}(\mathbf{X}_p^k)] \in [0, J(\mathbf{X}_p^k)]$.

Definition of Personal Best and Global Best:

We now outline the personal and global best solutions in the game-theoretic context:

- **Personal Best** ($p_{\text{best},p}^k$): the best-known strategy profile found by particle p , i.e., the one that has yielded the smallest deviation from a Nash equilibrium:

$$p_{\text{best},p}^k = (\mathcal{P}_p^{v,j}, \mathcal{V}_p^{v,j}, \mathcal{P}_p^{o,j}, \mathcal{V}_p^{o,j}) \quad (23)$$

Where $j = \arg \min_{1 \leq i \leq k} \hat{J}(\mathcal{P}_p^{v,i}, \mathcal{V}_p^{v,i}, \mathcal{P}_p^{o,i}, \mathcal{V}_p^{o,i})$

- **Global Best** (g_{best}^k): the best-known strategy profile across the entire swarm up to iteration k :

$$g_{\text{best}}^k = (\mathcal{P}_p^{v,j}, \mathcal{V}_p^{v,j}, \mathcal{P}_p^{o,j}, \mathcal{V}_p^{o,j}), \quad (24)$$

Where $(p, j) = \arg \min_{\substack{1 \leq i \leq k \\ 1 \leq q \leq P}} \hat{J}(\mathcal{P}_q^{v,i}, \mathcal{V}_q^{v,i}, \mathcal{P}_q^{o,i}, \mathcal{V}_q^{o,i})$

Using these precise definitions of $p_{\text{best},p}^k$ and g_{best}^k , each particle's velocity and position from iteration k to $k+1$ is governed by the following:

$$\mathbf{V}_p^{k+1} = \omega \mathbf{V}_p^k + c_1 r_1 \Delta \mathbf{p}_p^k + c_2 r_2 \Delta \mathbf{g}_p^k, \\ \mathbf{X}_p^{k+1} = \mathbf{X}_p^k + \mathbf{V}_p^{k+1} \quad (25)$$

where

$$\Delta \mathbf{p}_p^k = \mathbf{p}_{\text{best},p}^k - \mathbf{X}_p^k, \quad \Delta \mathbf{g}_p^k = \mathbf{g}_{\text{best}}^k - \mathbf{X}_p^k.$$

Here: ω is the inertia weight that governs the velocity's influence. c_1, c_2 are cognitive and social learning coefficients [16]. $r_1, r_2 \sim \mathcal{U}(0, 1)$ are stochastic, uniformly distributed, D -dimensional random vectors.

V. EXPERIMENTAL EVALUATION AND RESULTS

A. Experimental Protocol and Scenario

The experimental validation was conducted on a closed test track reproducing a real-world unsignalized four-way intersection (lane width: 3.15 m), located in Versailles-Satory, France. The setup involves two vehicles:

- **Ego Vehicle (Autonomous, in red):** An experimental Renault Zoé provided by the VEDECOM Institute, operating in autonomous mode, equipped with LiDAR, Radar, and RTK-GNSS/INS localization. See Figure (1).

- **Opponent Vehicle (Human-Driven, in blue):** A standard vehicle driven by a human, introducing realistic behavioral variability and non-deterministic dynamics into the decision-making loop.



Fig. 1: Experimental Automated Vehicle of VEDECOM

B. Scenario Configuration

Both vehicles approach the intersection from opposing directions (Ego from the Northeast, Opp from the Southwest) to perform simultaneous left-turn maneuvers, creating a critical conflict zone that requires strategic coordination for collision-free passage.

C. Strategic Adaptation to Human Behavior

The framework was tested in scenarios requiring distinct strategic responses to the opponent's trajectory. Ego adapts its decisions based on Opp's predicted trajectory and movement through the conflict zone.

1) **Simultaneous Crossing with Low-Speed Cooperative Adjustment:** This scenario is illustrated in Figures 2. It considers a simultaneous crossing maneuver where both agents negotiate access to a shared conflict zone. At $t \approx 2.0$ s, Ego anticipates that Opp will clear the conflict area before its arrival, which is confirmed as Opp exits the shared zone by $t \approx 5.9$ s. Ego then performs a controlled left-turn, crosses while Opp is already downstream ($t \approx 6.8$ s), and completes the maneuver by $t \approx 9.0$ s. The longitudinal profiles, illustrated in Figures 2 (e–f), highlight the cooperative nature of this interaction: Opp follows a typical deceleration–acceleration pattern, while Ego maintains a steady low speed of ≈ 2 m/s after an initial slowdown, avoiding unnecessary stops while preserving safety. The slight lateral deformation in Ego's trajectory reflects the safety-oriented cost terms, increasing lateral clearance without compromising comfort. Overall, this scenario illustrates anticipation-driven, non-aggressive coordination consistent with our game-theoretic framework.

2) **Early Crossing Uninterrupted Left Turn:** In this scenario, illustrated in Figures 3, Ego executes an early crossing while the human-driven Opp vehicle approaches from the Southwest. Initially, Opp is far from the intersection and does not create an immediate conflict. Ego begins decelerating to 3.5 m/s, adapting its speed to the intersection geometry. Ego predicts its entry into the conflict zone around $t \approx 6.8$ s, at




Fig. 2: Simultaneous Crossing with Low-Speed Maneuver

which point Opp remains sufficiently upstream. Figures 3 (c–d) show that Ego exits the intersection before Opp reaches the critical area.




Fig. 3: Early Crossing Uninterrupted Left Turn by Ego

D. Convergence Behavior and Runtime Assessment of the PSO-Based Solver

We assess the robustness, convergence behavior, and real-time feasibility of the proposed PSO-based solver across the two interaction scenarios studied above. Figure (4) illustrates convergence toward the Nash equilibrium, tracking both the global best particle cost g_{best} and the swarm-wide average objective $\bar{J}_k = \frac{1}{P} \sum_{j=1}^P f(X_j^k)$.




Fig. 4: Convergence Profile – Use Case 2: Fast convergence with preserved exploration

The solver exhibits fast initial convergence within the first 10 iterations, after which g_{best} stabilizes near zero while the average objective $\bar{J}^k$ continues to decrease smoothly, reflecting preserved swarm exploration. Both metrics converge to a consistent Nash equilibrium, confirming solver robustness.

To further assess real-time feasibility, Table (I) reports the mean convergence time over 100 independent runs across configurations ranging from 50 to 160 strategies per agent, along with the 95% confidence interval.

TABLE I: Computation Time and Equilibrium Accuracy of the PSO-Based Solver

Scenario	Strategy Grid Size	Evaluations Used	Nash Cost $\bar{J}(s_v, s_o)$	Time (s) $\pm$ Average (95% CI)
1 (100 Runs)	50 $\times$ 50	863/2500	0.0000	0.0195 $\pm$ 0.0012
2 (100 Runs)	100 $\times$ 100	4603/10000	0.0000	0.0203 $\pm$ 0.0020
3 (100 Runs)	100 $\times$ 160	5432/16000	0.0000	0.0312 $\pm$ 0.0024
4 (100 Runs)	160 $\times$ 160	17200/25921	0.0000	0.0485 $\pm$ 0.0024

Note: The PSO solver consistently converges to a valid Nash equilibrium, with near real-time feasibility for strategy sets up to 160 $\times$ 160.

As expected, larger discretization grids yield longer convergence times. Notably, even for the largest tested configuration (160 $\times$ 160), convergence is achieved in under 50 ms, confirming real-time suitability. Unlike fixed-grid approaches, the solver relies on random sampling, enabling stochastic estimation of the exact evaluation criteria within the game-theoretic framework.

VI. CONCLUSION

This paper presented the experimental validation of a game-theoretic decision-making framework for autonomous vehicles in mixed traffic. Modeling interactions as a Generalized Nash Equilibrium Problem (GNEP) and solving it via a tailored Particle Swarm Optimization (PSO) algorithm, the approach generates interaction-aware strategies beyond traditional rule-based methods. Real-world experiments with an autonomous

Renault Zoé at an unsignalized intersection confirmed the system’s ability to negotiate simultaneous crossings and early merging scenarios under coupled safety constraints. The results validate both the GNEP formulation for capturing complex traffic dependencies and the solver’s real-time efficiency (< 50 ms convergence).

Extension to Multi-Agent Scenarios : The current framework is formulated as a two-player game, covering the most critical interaction dyad in intersection crossing. Extending it to $N > 2$ players raises two challenges: the PSO search space grows combinatorially to $\sum_{i=1}^N d_i$ dimensions, and the shared constraint set $h(s_1, \dots, s_N) \leq 0$ must encode pairwise collision avoidance between all agents, potentially rendering the joint feasible set non-convex. Both aspects will be investigated in future work through simulation in multi-vehicle interaction. scenarios.

REFERENCES

- [1] A. Mahajan, T. Kumano, and Y. Yasui, “Intention estimation and controllable behaviour models for traffic merges,” *Journal of Control, Measurement, and System Integration*, 2021.
- [2] Z. Qin, A. Ji, Z. Sun, G. Wu, P. Hao, and X. Liao, “Game theoretic application to intersection management: A literature review,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [3] Z. Zhang and J. F. Fisac, “Safe occlusion-aware autonomous driving via game-theoretic active perception,” in *17th Robotics: Science and Systems, RSS 2021*. MIT Press Journals, 2021.
- [4] M. Bacharach, “Zero-sum games,” *Game theory*, pp. 253–257, 1989.
- [5] A. Liniger and J. Lygeros, “A noncooperative game approach to autonomous racing,” *IEEE Transactions on Control Systems Technology*, vol. 28, no. 3, pp. 884–897, 2020.
- [6] J. F. Nash Jr, “Equilibrium points in n-person games,” *Proceedings of the national academy of sciences*, vol. 36, no. 1, pp. 48–49, 1950.
- [7] J. Lei and U. V. Shanbhag, “Stochastic nash equilibrium problems: Models, analysis, and algorithms,” *IEEE Control Systems Magazine*, vol. 42, no. 4, pp. 103–124, 2022.
- [8] F. Facchinei and C. Kanzow, “Generalized nash equilibrium problems,” *Annals of Operations Research*, vol. 175, no. 1, 2010.
- [9] N. Naidja, S. Font, M. Revilloud, and G. Sandou, “Toward a holistic multi-criteria trajectory evaluation framework for autonomous driving in mixed traffic environment,” in *2025 8th International Conference on Intelligent Robotics and Control Engineering (IRCE)*. IEEE, 2025, pp. 1–10.
- [10] G. Askari, M. E. Gordji, and C. Park, “The behavioral model and game theory,” *Palgrave Communications*, vol. 5, no. 1, 2019.
- [11] S. Le Cleac’h, M. Schwager, and Z. Manchester, “Lucidgames: Online unscented inverse dynamic games for adaptive trajectory prediction and planning,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5485–5492, 2021.
- [12] —, “Algames: a fast augmented lagrangian solver for constrained dynamic games,” *Autonomous Robots*, vol. 46, no. 1, pp. 201–215, 2022.
- [13] D. Greiner, J. Periaux, J. M. Emperador, B. Galván, and G. Winter, “Game theory based evolutionary algorithms: a review with nash applications in structural engineering optimization problems,” *Archives of Computational Methods in Engineering*, vol. 24, pp. 703–750, 2017.
- [14] M. N. Vrahatis, P. Kontogiorgos, and G. P. Papavasilopoulos, “Particle swarm optimization for computing nash and stackelberg equilibria in energy markets,” in *SN Operations Research Forum*, vol. 1, no. 3. Springer, 2020, p. 20.
- [15] K. A.-R. Youssefi, M. Rouhani, H. R. Mashhadi, and W. Elmenreich, “A swarm intelligence-based robotic search algorithm integrated with game theory,” *Applied Soft Computing*, vol. 122, p. 108873, 2022.
- [16] J. Gou, Y.-X. Lei, W.-P. Guo, C. Wang, Y.-Q. Cai, and W. Luo, “A novel improved particle swarm optimization algorithm based on individual difference evolution,” *Applied Soft Computing*, vol. 57, pp. 468–481, 2017.