

Effective Game-Theoretic Motion Planning via Nested Search

Avishav Engle*, Andrey Zhitnikov, Oren Salzman,
Omer Ben-Porat, and Kiril Solovey

Technion–Israel Institute of Technology, Haifa, Israel
{avishav, andreyz}@campus.technion.ac.il, osalzman@cs.technion.ac.il,
{omerbp, kirilsol}@technion.ac.il

Abstract. To facilitate effective and safe deployment, individual robots must reason about interactions with other agents, which often occur without explicit communication. Recent work has identified game theory, particularly the concept of *Nash Equilibrium* (NE), as a key enabler for behavior-aware motion planning. Yet, existing work falls short of fully unleashing the power of game-theoretic reasoning. Specifically, optimization-based methods require simplified robot dynamics and may get trapped in local minima due to convexification. Other works that rely on the explicit enumeration of all possible trajectories suffer from poor scalability. To bridge this gap, we introduce *Game-Theoretic Nested Search* (GTNS), a scalable, and provably correct approach for computing NEs in general dynamical systems. GTNS efficiently searches the action space of all agents involved, while discarding trajectories that violate the NE constraint through an inner search over a lower-dimensional space. Our algorithm enables explicit selection among equilibria by utilizing a user-specified global objective, thereby capturing a rich set of realistic interactions. We demonstrate the approach across a variety of autonomous driving and racing scenarios, achieving solutions in mere seconds on commodity hardware.

1 Introduction

Modern robotic applications—from autonomous driving to assistive robotics—require behavior-aware motion planning that accounts for interactions between the ego robot and surrounding agents. These interactions often occur without explicit communication. Yet, to accomplish its mission effectively, the robot must reason about the behavior of other agents and the ways they respond to its actions. To address this challenge, this work seeks to develop game-theoretic motion planners that are both theoretically grounded and computationally efficient.

The majority of approaches for behavior-aware motion planning consider a two-step process where the ego robot computes a best-response trajectory after predicting the actions of surrounding agents [16, 23, 31, 32, 35]. Such approaches overlook that agents can update their actions based on other agents’ behavior, leading to overly conservative solutions.

* Corresponding author.

More recent methods aim to model interactions between agents through a game-theoretic perspective, thereby assuming that agents are rational and optimize a private objective function (e.g., travel time, safety, lead over opponent), while accounting for the fact that other agents are involved in the same process. For instance, some papers model interactions via a Stackelberg equilibrium and develop approaches based on dynamic programming to approximate the equilibrium [9]. However, Stackelberg games give an unfair advantage to one of the players [22], and hence lead to less realistic behavior. In the context of autonomous driving, such approaches can create overly aggressive maneuvers for one of the vehicles and be overly conservative for the rest.

More advanced methods model interactions via a Nash equilibrium (NE) [28], wherein no agent is better off by unilaterally deviating from the current solution. Although the game-theory community has extensively studied NE [29, 36], most computationally-efficient results consider overly simplistic models that fail to capture the intricacies of real robotic systems, including continuous action and state spaces with non-convex constraints, and long-horizon strategies. Some papers have attempted to approximate an NE using relaxations and utilizing differential dynamic programming [10], albeit with limited applicability to robotics due to the inability to capture collision-avoidance constraints. More specialized approaches that incorporate hard collision constraints and robot dynamics approximate NEs through optimization-based approaches [21, 22, 40, 43], albeit with limited theoretical justification. For example, [22] utilizes a first-order approximation and, as such, cannot guarantee finding a global NE except in some special convex settings, which rarely occur in robotics. Specifically, the presence of obstacles, nonlinear dynamics, and multiple agents gives rise to multiple homotopy classes of NE solutions, which cannot be captured in a convex approximation. An iterative approach [43], which relies on [40], falls short of the sufficient conditions for NE convergence, attaining only the necessary conditions. This is due to a lack of convexity in both cases, which follows from nonlinear dynamics and a non-convex state space.

A different work [25] considers a mathematically principled approach for efficiently obtaining a NE from a set of precomputed robot paths, under special conditions, such as right-of-way, which benefits one of the agents. Unfortunately, in a general setting, the approach boils down to a brute-force computation of the NE via the full payoff matrix, which is exponential in the number of robots. Moreover, the approach requires the explicit enumeration of all robot trajectories used, which can be costly if a rich set of strategies needs to be considered. A recent work [41] invokes the bi-matrix game reasoning of [25] in a receding horizon fashion to choose the immediate next action in a Stackelberg game, where the robots are confined to motions on a precomputed quasi-kinodynamic *graph* where edges represent Euler curves. Unfortunately, no equilibrium guarantees are given, and the use of Euler curves is highly problem- and model-specific.

An approach by LaValle and Hutchinson [20] also leverages graph search, explicitly constructing a frontier of all non-Pareto-dominated NE candidate trajectories for each node. While this yields equilibrium guarantees, it incurs a

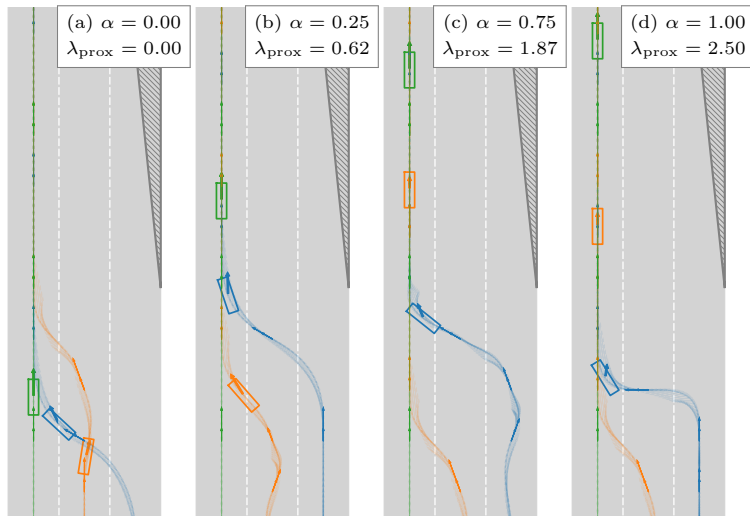


Fig. 1: Multi-lane merge: Robot 1 (blue) merges across Robot 2’s (orange) lane into Robot 3’s (green) lane. Existing algorithms typically restricted to a single homotopy class, whereas our method—via tuning of the global objective (here, decreasing Robot 1’s priority (larger α in $J=(1-\alpha)J^1+\alpha(J^2+J^3)$) and increasing proximity penalty weight λ_{prox})—yields distinct NE behaviors: (a) aggressive zip-merge; (b) zip-merge; (c) yield-then-merge; (d) over-cautious.

substantial computational burden due to the need to verify and maintain all such equilibria throughout the search.

Learning-based approaches have been applied in behavior-aware settings in the context of head-to-head racing in computer games [46], regulating mixed traffic flow [45], and autonomous driving [44], to name a few examples. However, as is typically the case with learning methods, they offer no explainability or generalizability and heavily rely on past experience for making correct decisions. Thus, their reliability remains unclear, especially in long tail events encompassing complex interactions between multiple agents [44].

Contribution. We introduce an efficient and provably correct method, termed Game-Theoretic Nested Search (GTNS), for computing NEs in motion planning. Our method guarantees convergence to an NE without restricting the problem domain (e.g., convexity) and applies to general objective functions and dynamical constraints of individual robots. Furthermore, the approach enables the explicit tuning of the NE solution retrieved via a user-specified global objective function, ensuring that the returned NE minimizes this function among all available NEs (see, e.g., Fig. 1).

At the core of our methodology is a nested-search approach consisting of an outer and inner search method. The outer search implicitly explores a tree of potential multi-robot trajectories, while the inner search is invoked to verify that a given solution satisfies the NE requirement by considering improving unilateral deviations of individual robots. For computational efficiency, our nested-search

approach implicitly explores the joint state space of the multi-robot system via graphs that efficiently encode the trajectory space of individual robots while adhering to the robots’ full dynamical constraints.

This search-based approach allows for an implicit encoding of trajectory-based constraints and incremental reuse of information, which is otherwise difficult to realize efficiently within an optimization-based approach [3] or in methods that require explicit enumeration of all possible trajectories [25, 41]. Our algorithm has conceptual similarities to the LaValle-Hutchinson approach [20], but crucially does not explicitly maintain all candidate NE trajectories, and employs a powerful pruning mechanism to discard non-NE trajectories (and their subtrees), rendering the search over complex and high-dimensional spaces tractable.

We demonstrate the approach on a variety of realistic autonomous driving and racing scenarios (including intersections, lane merges, following distance, opposing-lane passes, and racetrack overtakes) where we achieve solutions in mere seconds on commodity hardware. From the theory side, our algorithmic framework leads to a rigorous but natural proof for the attainment of an NE.

Finally, we mention that our framework directly supports scalable generation of training data for generative AI and foundation models in autonomy [12, 16, 23, 32]. It enables rich simulation of multi-agent dynamics across domains (traffic, racing, aerial swarms, aircraft sequencing, warehouses) and can efficiently synthesize rare, safety-critical cases for training and evaluation.

Organization. Sec. 2 provides preliminaries and our problem statement. Sec. 3 considers a graph-based formulation of the above problem. Sec. 4 describes our GTNS algorithm, whose theoretical guarantees are discussed in Sec. 5. Sec. 6 details experimental results. Conclusion and future work appears in Sec. 7.

2 Preliminaries

We begin by introducing the notation for kinodynamic multi-robot motion planning, then formalize the underlying game-theoretic planning problem. Consider $m \geq 2$ robots indexed by $i \in [m] := \{1, \dots, m\}$. Each robot i has a state space $\mathcal{X}^i \subseteq \mathbb{R}^{d_i}$ and a control space $\mathcal{U}^i \subseteq \mathbb{R}^{D_i}$. Its motion is governed by the dynamics

$$\dot{x}^i(t) = f^i(x^i(t), u^i(t)), \quad x^i(t) \in \mathcal{X}^i, \quad u^i(t) \in \mathcal{U}^i. \quad (1)$$

Additionally, the robot must remain within the free space $\mathcal{X}_f^i \subseteq \mathcal{X}^i$ to avoid collisions with static obstacles (collisions with other robots discussed below).

For a given planning horizon $T > 0$, let \mathcal{U}_T^i be the set of admissible control functions $u^i : [0, T] \mapsto \mathcal{U}^i$ for robot i . Given an initial state $x_0^i \in \mathcal{X}^i$ and a control function $u^i \in \mathcal{U}_T^i$, the induced trajectory is the mapping $\pi_{x_0^i, u^i}^i : [0, T] \mapsto \mathcal{X}^i$, where $\pi_{x_0^i, u^i}^i(\cdot)$ is a solution to Eq. (1) with $\pi_{x_0^i, u^i}^i(0) = x_0^i$. The set of feasible trajectories for robot i is $\Pi_T^i := \{ \pi_{x_0^i, u^i}^i \mid x_0^i \in \mathcal{X}^i, u^i \in \mathcal{U}_T^i \}$.

Next, we generalize the above definitions to encompass the full multi-robot system. The joint state space of the m robots is denoted by $\mathcal{X} := \times_{i=1}^m \mathcal{X}^i$, i.e., a point in \mathcal{X} describes the simultaneous state of all the robots in the system.

Similarly, the joint control space is $\mathcal{U} := \times_{i=1}^m \mathcal{U}^i$ and the joint control space within specific time horizon T is $\mathcal{U}_T := \times_{i=1}^m \mathcal{U}_T^i$. Given initial joint state $x_0 = (x_0^1, \dots, x_0^m) \in \mathcal{X}$ and joint control $u = (u^1, \dots, u^m) \in \mathcal{U}_T$, the induced joint trajectory $\pi_{x_0, u} = (\pi_{x_0, u}^1, \dots, \pi_{x_0, u}^m)$ specifies the trajectories of each of the m robots, and can be interpreted as the mapping $\pi_{x_0, u} : [0, T] \mapsto \mathcal{X}$. The set of all joint trajectories is $\Pi_T := \{ \pi_{x_0, u} \mid x_0 \in \mathcal{X}, u \in \mathcal{U}_T \}$.

Each robot i evaluates its performance via a cost functional $J^i : \Pi_T^i \times \Pi_T^{-i} \mapsto \mathbb{R}_{\geq 0}$, where $\Pi_T^{-i} := \times_{j \neq i} \Pi_T^j$ is the set of trajectories of the remaining robots with indices $[m] \setminus \{i\}$. For trajectories $\pi^i \in \Pi_T^i$ and $\pi^{-i} \in \Pi_T^{-i}$, we define

$$J^i(\pi^i, \pi^{-i}) := \int_0^T c^i(\pi^i(t), \pi^{-i}(t)) dt, \quad (2)$$

with stage cost c^i depending on states and controls.

2.1 Problem Definition

We seek to compute joint trajectories that satisfy the following requirement.

Definition 1 (Nash Equilibrium). *A joint trajectory $\pi = (\pi^1, \dots, \pi^m) \in \Pi_T$ is a (pure) Nash equilibrium (NE) if, for every robot $i \in [m]$ and for all alternative feasible trajectories $\tilde{\pi}^i \in \Pi_T^i$ (from the same start state), no robot can reduce its cost by deviating while the other robots fix their trajectories, i.e., $J^i(\pi^i, \pi^{-i}) \leq J^i(\tilde{\pi}^i, \pi^{-i})$.*

Moreover, we wish to compute an NE that minimizes a given global cost.

Problem 1 (Optimal NE Planning). *Fix a planning horizon $T > 0$ and consider a global cost functional*

$$J(\pi) := \int_0^T c(\pi^1(t), \dots, \pi^m(t)) dt, \quad \pi \in \Pi_T, \quad (3)$$

with stage cost c depending on joint states and controls. We wish to find a joint trajectory $\pi^ \in \Pi_T$ (and its corresponding control laws) such that: (1) π^* is a Nash equilibrium (Def. 1) that minimizes the global cost $J(\pi^*)$ among all NEs; (2) the final state of π^* is in the joint goal region, i.e., $\pi^*(T) \in \mathcal{X}_{\text{goal}} := \times_{i=1}^m \mathcal{X}_{\text{goal}}^i$; (3) $\pi(t) \in \mathcal{X}_f$ for all $t \in [0, T]$.*

Discussion. We consider an *open-loop* setting: an NE solution is computed once for a horizon and then executed, rather than replanned at every time step. This setting is common in robotics (see, e.g., [47]), motivated by the high computational demands of computing NEs, which currently preclude real-time feedback control. The resulting NE solutions promote stable execution, since local controllers cannot improve their cost by deviating from the prescribed trajectory.

We seek *global* NEs, where robots may choose any action in their control spaces. This contrasts with *local* NEs, which are implicitly considered in,




Fig. 2: Kinodynamic grid graph. At each (x, y) lattice coordinate (blue dots) there are various θ , v and δ values. Graph edges are gray; a randomly selected subset is emphasized and colored by the source state.

e.g., [10, 22, 30, 40, 43], where the solution is restricted to a subset of the joint state space (e.g. due to linearization and convexification). Local NEs may not correspond to any global NE, leading to unstable or otherwise poorly chosen execution. See illustration in App. B, Fig. 6. Another key difference from prior work is that we require the chosen NE to optimize a global objective. This promotes predictability and explainability, as (under mild conditions) a unique global optimum exists, ensuring all robots converge to the same NE.

A remaining challenge is incentivizing robots to follow the chosen solution. An NE optimizing a specific cost function J may favor certain robots, motivating others to defect. In cooperative domains (e.g., driving), J can encode social welfare or fairness and may even be set by an external arbitrator (e.g., to resolve right-of-way). In competitive settings, the more general solution concept of correlated equilibrium (CE) [2] could be applied: a mediator samples from a distribution and privately recommends actions, which can enable coordination on outcomes that are more efficient or fair. We leave this for future work.

3 Graph-Based setting

To tackle Problem 1, we draw inspiration from previous work on multi-robot motion planning, wherein the curse of dimensionality, inherent in exploring the joint state space, is alleviated by first capturing the state space of the individual robots via discrete graphs, and then planning their collective motion on an implicitly-defined tensor graph [20, 38, 39]. In our setting, however, we utilize specialized constructions of those graphs to accommodate differential constraints (1) and, most importantly, game-theoretic aspects. In preparation for our algorithm, which is described in the next section, we describe here the graph-based setting, as well as a graph-based counterpart of Problem 1.

For each robot $i \in [m]$, we construct a *directed* graph $G^i := (V^i, E^i)$, that embodies its dynamically feasible motions. The vertex set $V^i \subseteq \mathcal{X}^i$ consists of sampled states, including the start state x_0^i and states in the goal region $\mathcal{X}_{\text{goal}}^i$. Each edge $e^i = (x^i, y^i) \in E^i$ corresponds to a dynamically feasible local trajectory $\pi_{e^i}^i$ from the state x^i to the state y^i obtained under some control law $u_{e^i}^i \in \mathcal{U}_{\Delta t}^i$, executed over a fixed duration $\Delta t > 0$ from the state $x^i \in \mathcal{X}^i$.

Every edge, therefore, represents a time-parameterized trajectory of duration Δt , making G^i a kinodynamic variant of the probabilistic roadmap (PRM) [17], as used in, e.g., [33, 34]. We note that each edge in the graph is a curve (kinodynamic-adhering trajectory) induced by the corresponding control function and not a

straight line. We consider an explicit representation of each G^i ; more efficient semi-implicit constructions [15, 37] are left for future work.

From the individual graphs $\{G^i\}_{i \in [m]}$, we form the *tensor-product directed graph* $G = (V, E)$ with $V = \times_{i=1}^m V^i$, and $E = \times_{i=1}^m E^i$. A joint vertex $x = (x^1, \dots, x^m) \in V$ encodes a joint state of the m robots, while a joint edge $e = (e^1, \dots, e^m) \in E$ represents a local joint trajectory $\pi_e = (\pi_{e^1}^1, \dots, \pi_{e^m}^m)$ synchronized motion where each robot i follows the trajectory $\pi_{e^i}^i$ associated with the edge $e^i \in E^i$ during the same duration Δt . Since $|V|$ and $|E|$ grow exponentially in m , we never construct G explicitly; instead, we explore it implicitly by combining neighbors from the explicit graphs $\{G^i\}$, where an A*-search directs which portion of G to explore next (see below).

Next, we adapt Problem 1 to the graph-based setting. Fix a number of steps $n \in \mathbb{N}_+$, such that $n\Delta t$ is equal to the planning horizon $T > 0$. Let a trajectory π_n^i of length n consist of a sequence of n edges on the graph G_i . That is, given the graph path e_1^i, \dots, e_n^i , the trajectory π_n^i is the concatenation of the trajectories $\pi_{e_1^i}^i, \dots, \pi_{e_n^i}^i$. We denote by Π_n^i the set of n -step trajectories of robot i . We similarly define the joint trajectory set $\Pi_n := \times_{i=1}^m \Pi_n^i$. Notice that each joint trajectory $\pi_n \in \Pi_n$ corresponds to a path on the graph G .

Definition 2 (Graph Nash Equilibrium). *For a horizon of n steps, a joint discrete trajectory $\pi_n \in \Pi_n$ is a graph Nash equilibrium (gNE) if, for every robot $i \in [m]$ and every alternative trajectory $\tilde{\pi}_n^i \in \Pi_n^i$ (from the same start state), no robot can reduce its cost by unilaterally deviating to another feasible n -step trajectory, i.e., $J^i(\pi_n^i, \pi_n^{-i}) \leq J^i(\tilde{\pi}_n^i, \pi_n^{-i})$.*

Problem 2 (Optimal gNE Planning). *Fix $n > 0$. Given a global cost $J : \Pi_n \mapsto \mathbb{R}_{\geq 0}$, the objective is to find $\pi_n^* \in \Pi_n$ such that: (1) π_n^* is a gNE (Def. 2) that minimizes the global cost $J(\pi_n^*)$ among all gNEs; (2) π_n^* reaches the joint goal region $\mathcal{X}_{\text{goal}}$ in n steps; (3) π_n^* is collision-free, i.e., all intermediate joint states, and the continuous motions along edges, respect \mathcal{X}_f .*

Discussion. We discuss the interrelation between Problem 1 and Problem 2. Previous work has shown that for the geometric setting (i.e., without kinodynamic constraints), the graphs G^1, \dots, G^m can be constructed such that any feasible multi-robot trajectory in the continuous state space of the robot group can be approximated arbitrarily well in the graph G [5, 38]. The extension to the kinodynamic setting can be done via recent analysis techniques for asymptotically optimal sampling-based planning [18]. We leave this technical proof for future work. We note that the points highlighted in the discussion in Sec. 2, are equally applicable to the graph setting.

4 Algorithm for graph NE

We describe our algorithmic approach, Game-Theoretic Nested Search (GTNS), to tackle Problem 2. The algorithm employs a nested approach, consisting of an

outer search (Alg. 1) and an inner search (Alg. 2). Both levels of search are performed by A*-based search algorithms [13], albeit on different graphs and for different purposes. We describe the high-level structure of the algorithm (see additional implementation optimizations in App. A).

Outer search. The GTNS algorithm (Alg. 1) explores the implicit tensor-product graph G induced by the per-robot graphs $\{G^i\}_{i \in [m]}$ each produced by `generateRobotRoadmap`. These individual graphs can be precomputed once (offline) and reused across scenarios for the same robot platform. By construction, motion along edges of G is dynamically feasible for all robots; we additionally enforce collision avoidance and the gNE condition.

The A* search maintains a priority queue `OPEN` keyed by $J(\pi_x) + H(x)$ (where $H(\cdot)$ is a cost-to-go heuristic) and a predecessor map P for path reconstruction. At each iteration, the algorithm: (i) extracts the minimum-key node x (line 4); (ii) recovers the current best path π_x from the root using P (line 5); and (iii) enumerates neighbors X_{near} of x in G (line 6). Here, a joint node $y = (y^1, \dots, y^m) \in V$ belongs to X_{near} iff for every i there is an edge $(x^i, y^i) \in E^i$.

For each potential neighbor $x_{\text{new}} \in X_{\text{near}}$, the algorithm (a) verifies that the local edge trajectory $\pi_{x \rightarrow x_{\text{new}}}$ is collision-free (line 8); (b) forms the candidate path $\pi_{x_{\text{new}}} := \text{concatenate}(\pi_x, \pi_{x \rightarrow x_{\text{new}}})$ (line 9); and (c) checks improvement w.r.t. any incumbent path to x_{new} (line 11), where the incumbent is recovered via P (line 10) and updated (line 13). It then calls the predicate `ISNASHEQUILIBRIUM` on $\pi_{x_{\text{new}}}$ (line 12); only if the candidate path satisfies the gNE condition and offers an improved cost over the incumbent path (similar to standard A* dominance checks), is x_{new} inserted into `OPEN` with key $J(\pi_{x_{\text{new}}}) + H(x_{\text{new}})$. If a goal node is popped and validated, the corresponding trajectory is returned.

Inner search. `ISNASHEQUILIBRIUM` is a best-response oracle: it solves a single-robot motion planning problem on G^i to see whether any agent has a strictly cheaper unilateral deviation of the same horizon that is dynamically feasible and collision-free w.r.t. others' fixed paths. For each candidate joint trajectory $\pi_{x_{\text{cand}}} = (\pi_{x_{\text{cand}}}^i, \pi_{x_{\text{cand}}}^{-i})$ encountered by GTNS, `ISNASHEQUILIBRIUM` (Alg. 2) runs, for each $i \in [m]$, an A* search on G^i subject to three simple constraints: (1) The search goal is set (line 3) to the *specific* terminal configuration x_{cand}^i appearing at the end of $\pi_{x_{\text{cand}}}^i$ (no goal *region*); (2) Only paths of *exactly* n edges (garnered by the method `numEdges`) are considered competitive (the same length as $\pi_{x_{\text{cand}}}$). Nodes are only considered until the partial path length d reaches n ; (3) Edges for robot i are collision-checked (`collisionFreeEdge`, line 13) against the *fixed, time-parameterized* trajectories of the other robots $\pi_{x_{\text{cand}}}^{-i}$.

Operationally, for each i the routine initializes `OPEN` and a predecessor map P^i , then performs the usual A* node selection, neighbor expansion, concatenation, and incumbent comparison on the single-robot graph G^i (with key $J^i(\cdot) + H^i(\cdot)$), exactly mirroring the outer loop but *only* for agent i and under the constraints above. If it ever reaches x_{cand}^i in exactly n steps with a trajectory $\tilde{\pi}_{x_{\text{cand}}}^i$

Algorithm 1: GTNS

```

1 OPEN  $\leftarrow \{x_0\}$ ;  $P(x) \leftarrow \emptyset, \forall x \in V$ 
2  $\{G^i \leftarrow \text{generateRobotRoadmap}(f^i, J^i, x_0^i, \mathcal{X}_{\text{goal}}^i)\}_{i \in [m]}$ 
3 while not OPEN.empty() do
4      $x \leftarrow \text{OPEN.extractMin}()$ 
5      $\pi_x \leftarrow \text{trajFromRoot}(P, x)$ 
6      $X_{\text{near}} \leftarrow G.\text{neighbors}(x)$ 
7     for  $x_{\text{new}} \in X_{\text{near}}$  do
8         if collisionFreeEdge( $\pi_{x \rightarrow x_{\text{new}}}$ ) then
9              $\pi_{x_{\text{new}}} \leftarrow \text{concatenate}(\pi_x, \pi_{x \rightarrow x_{\text{new}}})$ 
10             $\pi_{x_{\text{new}}}^{\text{old}} \leftarrow \text{trajFromRoot}(P, x_{\text{new}})$   $\triangleright$  incumbent traj. to  $x_{\text{new}}$ 
11            if notYetVisited( $x_{\text{new}}$ ) or  $J(\pi_{x_{\text{new}}}) < J(\pi_{x_{\text{new}}}^{\text{old}})$  then
12                if ISNASHEEQUILIBRIUM( $\pi_{x_{\text{new}}}$ ) then
13                     $P(x_{\text{new}}) \leftarrow x$   $\triangleright$  store parent
14                    OPEN.insert( $x_{\text{new}}, J(\pi_{x_{\text{new}}}) + H(x_{\text{new}})$ )
15                    if  $x_{\text{new}} \in \mathcal{X}_{\text{goal}}$  then
16                        return  $\pi_{x_{\text{new}}}$ 
17 return  $\emptyset$ 

```

yielding $J^i(\tilde{\pi}_{x_{\text{cand}}}^i) < J^i(\pi_{x_{\text{cand}}}^i)$, the routine immediately returns FALSE (the candidate joint trajectory does not meet the gNE conditions). If no robot can improve, the routine returns TRUE.

5 Theoretical guarantees

We prove that Alg. 1 provides an optimal solution to Problem 2. In preparation, we obtain the following lemma that states that the NE property is monotone, i.e., every subtrajectory of a NE trajectory is also NE. Given a trajectory π_n of n edges, the notation $\pi_{n'}$ for some $0 < n' < n$ denotes the subtrajectory of π_n consisting of the first n' edges. Namely, $\pi_{n'}$ has a duration of $n' \cdot \Delta t$ and $\pi_{n'}(\tau) = \pi_n(\tau)$ for any $0 \leq \tau \leq \Delta t \cdot n'$.

Lemma 1 (Monotonicity of Nash Equilibrium). *If a trajectory π_n is a gNE (Def. 2) then any subtrajectory $\pi_{n'}$, where $0 \leq n' < n$, is a gNE too.*

Proof. First, recall that we consider additive cost functions $J, \{J^i\}$. Suppose that π_n is a gNE, namely for any robot $i \in [m]$, and any alternative trajectory $\tilde{\pi}_n^i \in \Pi_n^i$, with the same start and final \mathcal{X}^i state, i.e., $\pi_n^i(0) = \tilde{\pi}_n^i(0)$, and $\pi_n^i(n \cdot \Delta t) = \tilde{\pi}_n^i(n \cdot \Delta t)$, respectively, it holds that $J^i(\pi_n^i, \pi_n^{-i}) \leq J^i(\tilde{\pi}_n^i, \pi_n^{-i})$.

Assume, by contradiction, that the subtrajectory $\pi_{n'}$, for some $0 < n' < n$, is not a gNE: for some robot $i \in [m]$, there exists $\tilde{\pi}_{n'}^i \in \Pi_{n'}^i$ satisfying the collision constraints, $\pi_{n'}^i(0) = \tilde{\pi}_{n'}^i(0)$, and $\pi_{n'}^i(n' \cdot \Delta t) = \tilde{\pi}_{n'}^i(n' \cdot \Delta t)$ and inequality $J^i(\pi_{n'}^i, \pi_{n'}^{-i}) > J^i(\tilde{\pi}_{n'}^i, \pi_{n'}^{-i})$. Consider the remaining portion of π_n from the $n'+1$ edge, denoted by $\pi_{n'+1:n}$. I.e., $\pi_{n'+1:n}$ has a duration of $(n-n') \cdot \Delta t$

Algorithm 2: ISNASH EQUILIBRIUM ($\pi_{x_{\text{cand}}}$)

```

1  $n \leftarrow \text{numEdges}(\pi_{x_{\text{cand}}})$ 
2 for  $i \in [m]$  do
3    $x_{\text{cand}}^i \leftarrow \text{last vertex in } \pi_{x_{\text{cand}}^i}^i$   $\triangleright \pi_{x_{\text{cand}}} = (\pi_{x_{\text{cand}}^i}^i, \pi_{x_{\text{cand}}^{-i}}^{-i})$ 
4   OPEN  $\leftarrow \{x_0^i\}$ ;  $P^i(x^i) \leftarrow \emptyset$ ,  $\forall x^i \in V^i$ 
5   while not OPEN.empty() do
6      $x^i \leftarrow \text{OPEN.extractMin}()$ 
7      $\tilde{\pi}_{x^i}^i \leftarrow \text{trajFromRoot}(P^i, x^i)$ 
8      $d \leftarrow \text{numEdges}(\tilde{\pi}_{x^i}^i)$ 
9     if  $d \geq n$  then
10      continue
11     $X_{\text{near}}^i \leftarrow G^i.\text{neighbors}(x^i)$ 
12    for  $x_{\text{new}}^i \in X_{\text{near}}^i$  do
13      if collisionFreeEdge( $\tilde{\pi}_{x^i \rightarrow x_{\text{new}}^i}^i, \pi_{x_{\text{cand}}^{-i}}^{-i}$ ) then
14         $\tilde{\pi}_{x_{\text{new}}^i}^i \leftarrow \text{concatenate}(\tilde{\pi}_{x^i}^i, \tilde{\pi}_{x^i \rightarrow x_{\text{new}}^i}^i)$ 
15         $\tilde{\pi}_{x_{\text{new}}^i}^{\text{old}} \leftarrow \text{trajFromRoot}(P^i, x_{\text{new}}^i)$ 
16        if not yet visited( $x_{\text{new}}^i$ ) or  $J^i(\tilde{\pi}_{x_{\text{new}}^i}^i) < J^i(\tilde{\pi}_{x_{\text{new}}^i}^{\text{old}})$  then
17           $P^i(x_{\text{new}}^i) \leftarrow x^i$ 
18          OPEN.insert( $x_{\text{new}}^i, J^i(\tilde{\pi}_{x_{\text{new}}^i}^i) + H^i(x_{\text{new}}^i)$ )
19          if  $x_{\text{new}}^i = x_{\text{cand}}^i$  and  $d + 1 = n$  then
20            if  $J^i(\tilde{\pi}_{x_{\text{cand}}^i}^i) < J^i(\pi_{x_{\text{cand}}^i}^i)$  then
21              return FALSE
22 return TRUE

```

and $\pi_{n'+1:n}(\tau) = \pi_n(\tau + n' \cdot \Delta t)$. Next, denote by $\tilde{\pi}_n^i$ the concatenation of $\tilde{\pi}_{n'}^i$ and $\pi_{n'+1:n}$. This new trajectory is valid as $\tilde{\pi}_{n'}^i$ ends where $\pi_{n'+1:n}$ begins, and both those subtrajectories are collision free. By exploiting the additivity of the cost function J^i , we upper bound the J^i cost resulting from $\tilde{\pi}_n^i$:

$$\begin{aligned}
J^i(\tilde{\pi}_n^i, \pi_n^{-i}) &= J^i(\tilde{\pi}_{n'}^i, \pi_{n'}^{-i}) + J^i(\pi_{n'+1:n}^i, \pi_{n'+1:n}^{-i}) \\
&< J^i(\pi_{n'}^i, \pi_{n'}^{-i}) + J^i(\pi_{n'+1:n}^i, \pi_{n'+1:n}^{-i}) = J^i(\pi_n^i, \pi_n^{-i}).
\end{aligned}$$

This contradicts π_n being a gNE. Thus, the gNE property is monotone. \square

Combining the above lemma, with the optimality guarantee of A*, which is the backbone of both Alg. 1 and 2, we establish the correctness of GTNS.

Theorem 1. *GTNS finds the optimal solution to Problem 2.*

Proof. The A* algorithm returns an optimal solution when running on a search problem P using an admissible heuristic [13]. Recall that a search problem is a tuple $P = \langle S, \text{Succ}, s_{\text{begin}}, S_{\text{goal}} \rangle$, where S is a set of (abstract) states, $\text{Succ} : S \rightarrow 2^S$ is a successor function mapping each state to a set of neighboring states, $s_{\text{begin}} \in S$ is the initial state and $S_{\text{goal}} \subset S$ a set of goal states.

We start by formalizing our problem as a search problem and then continue to show that Alg. 1, which can be seen as an adaptation of A*, inherits its optimality. Specifically, consider the search problem defined over the set of states $S = V$ (i.e., each abstract state represents a joint vertex, which in itself represents a joint robot state). The successor function is defined such that for vertices $v, v' \in V$, we have that $v' \in \text{Succ}(v)$ if and only if $(v, v') \in G$. Finally, the start and goal states are x_0 and states satisfying Problem 2.

Alg. 1 can be seen as an adaptation of A* where nodes are only added to the OPEN list if they satisfy equilibrium conditions (line 12) and using an admissible heuristic H . Pruning search nodes in this search problem does not hinder optimality of the search due to the Nash monotonicity property (Lemma 1). That is, search nodes that are pruned cannot contribute to a longer trajectory that is a NE. Thus, optimality follows from the optimality of A*. As noted, Alg. 2 is too an adaptation of A*, and thus when used with an admissible heuristic, returns an optimal, collision-free path of upto length n . \square

Time complexity. The branching factor of node expansion in GTNS grows exponentially with the number of robots. Even so, our approach efficiently copes with realistic problem settings in a few seconds (see Sec. 6), due to the aggressive pruning of ISNASHEQUILIBRIUM, which substantially reduces the search space, as compared to the non-game-theoretic and centralized multi-robot motion planning setting [42]. We discuss additional speed-up techniques in Sec. 7.

6 Experiments and Results

We provide an experimental evaluation of our game-theoretic approach in driving and racing scenarios. We begin with implementation details and experiment setup, and then proceed to a qualitative evaluation showcasing the behaviors of the approach. Next, we evaluate the computational aspects of our approach and conclude with a comparison with LaValle-Hutchinson [20]. An additional experiment, which studies the effect of NE mismatch, is found in App. B.1.

6.1 Implementation details and problem setting

The single-robot graphs G_i are built offline in Python using CasADi Opti [1] for BVP solving (see more information below). GTNS is implemented in C++ using the LEMON [8] graph library. Simulations were run on a 64-bit Win11 laptop with an Intel Core i9-14900HX CPU (2.20 GHz) and 32 GB RAM.

For the robot model, we consider the second-order bicycle model [19]. The robot has the states, controls, and dynamics, of the form

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ \theta \\ v \\ \delta \end{bmatrix}, \quad \mathbf{u} = \begin{bmatrix} a \\ \omega \end{bmatrix}, \quad \dot{\mathbf{x}} = \begin{bmatrix} v \cos \theta \\ v \sin \theta \\ \frac{v}{L} \tan \delta \\ a \\ \omega \end{bmatrix},$$

respectively. Here, x , y and θ are the 2-D positions and heading of the car, v , and δ are its speed and steering angle, and a and ω control their rate of change. L is the distance between front and rear axles (see visualization in App. B, Fig. 8).

Kinodynamic graphs. For the single-robot graphs G^i , we consider two constructions. The *grid graph* (Fig. 2), which is task-agnostic, samples a spatial lattice in (x, y) (with $\Delta x = \Delta y = 1$ m) and expands said lattice with headings, speeds, and steering angles on discrete sets (θ, v, δ) . From each node, we attempt directed connections to *nearby* nodes within a k -hop neighborhood of the lattice. The *track graph* leverages the environment’s topology and samples along a track centerline with lateral offsets (App. B, Fig. 7). Unless otherwise stated, robots share identical dynamics and graph G^i .

Graph size is determined by the sampling resolutions $(\Delta x, \Delta y, \Delta \theta, \Delta v, \Delta \delta)$ and the connection parameter k . We choose these to meet a computational budget while preserving sufficient fidelity for the maneuvers of interest. Concretely, we set spatial and angular resolutions and k to reach a desired ballpark of nodes/edges, then prune infeasible nodes/edges via kinodynamic feasibility and obstacle inflation. The resulting graph sizes are reported in Table 1.

Edge generation is performed in the following manner. For each candidate source→target pair of nodes, we solve a fixed-duration two-point boundary-value problem (BVP) over $\Delta t = 1$ s with hard input limits $|a| \leq 5$ m/s², $|\omega| \leq 2$ rad/s. If the BVP is feasible, we add the corresponding directed edge and store its continuous trajectory. Edges are discarded if their continuous trajectory intersects obstacles inflated by r_{coll} . Unless otherwise specified, each edge has unit cost.

Search heuristics. As the inner search ISNASH EQUILIBRIUM may be called on any $\pi_{x_{\text{cand}}}$, we require an admissible heuristic to every node x , and not just to $\mathcal{X}_{\text{goal}}$, thus for each robot i we precompute all-pairs shortest-path (APSP) distances on G^i and use $H(x) = \sum_i H^i(x^i)$ (or the weighted equivalent when $J(\cdot)$ is weighted, see below) where H^i is the single-robot shortest distance from x^i to $\mathcal{X}_{\text{goal}}^i$. Because single-robot distances satisfy the triangle inequality and joint costs add across robots, H is consistent (and hence admissible) on the tensor-product graph G . Moreover, during inner A* searches for robot i on G^i with dynamic obstacles induced by a fixed plan π^{-i} , H^i remains a valid lower bound on i ’s cost-to-go: the added constraints only restrict feasible motion and cannot reduce the optimal single-robot cost.

6.2 Qualitative analysis

We present a set of multi-agent scenarios mimicking interactions in autonomous driving and racing and discuss the results obtained by our algorithm. These experiments showcase the generality of our approach, and demonstrate the diversity of equilibria available, which are hard to capture with previous optimization-based methods.

Across tasks we vary two interpretable knobs that shape the resulting NE: The global objective aggregates the individual costs $J(\pi_n) = \sum_{i=1}^m \alpha_i J^i(\pi_n^i, \pi_n^{-i})$ to encode priorities via the weights $\alpha_i \in [0, 1]$. Additionally, each J^i includes a

state-based penalty that encourages separation from other robots:


$$c_{\text{prox}}^i(\mathbf{x}^i, \mathbf{x}^{-i}) = \frac{\lambda_{\text{prox}}}{\max(\min_{j \in -i} \|\text{proj}(\mathbf{x}^i) - \text{proj}(\mathbf{x}^j)\|_2, \varepsilon)},$$

where $\lambda_{\text{prox}} \geq 0$, $\text{proj} : \mathcal{X} \rightarrow \mathbb{R}^2$ returns the planar position (x, y) of a state and $\varepsilon = 10^{-3}$ prevents singularities (yielding a bounded, well-behaved, Lipschitz stage cost). Larger λ_{prox} promotes greater standoff. In our experiments, each J^i is the sum over n edges of π_n^i of unit edge cost plus the proximity term for each of the $n + 1$ nodes along the way:

$$J^i(\pi_n^i, \pi_n^{-i}) = n + \sum_{\mathbf{x}^i, \mathbf{x}^{-i} \in (\pi_n^i, \pi_n^{-i})} (c_{\text{prox}}^i(\mathbf{x}^i, \mathbf{x}^{-i})).$$

Next, we discuss the scenarios and the solution obtained by GTNS. An additional experiment that showcases the effect of λ_{prox} on following distance in a semi-collaborative setting, can be found in App. B, Fig. 5.

Four-way intersection. We consider a single-lane four-way intersection with perpendicular approaches by two robots, objective $J = \alpha J^1 + (1 - \alpha) J^2$, $\alpha \in [0, 1]$, $\lambda_{\text{prox}} = 0$. For $\alpha \geq 0.5$ the NE selected by our planner has Robot 1 (blue) crossing from left to right while Robot 2 (orange) yields and waits for the intersection to clear before proceeding from bottom to top. For $\alpha < 0.5$, the NE is symmetric: Robot 2 crosses first and Robot 1 yields. This result illustrates the ability of our NE planner to select a safe NE in an explainable manner (via tuning of the objective function J), without requiring clipping of the dynamics or costs, or initial guesses (as commonly considered in previous work).



Racetrack overtake. We consider a racetrack segment to showcase priority-driven overtaking decisions between two robots. This setting probes how a scalar priority α biases raceline choice and pass timing within realistic behavioral dynamics trajectories. We set $J = \alpha J^1 + (1 - \alpha) J^2$, $\alpha \in [0, 1]$, $\lambda_{\text{prox}} = 0$. Fig. 3 shows that the tilt of robot priority determines which robot targets the inside line at the final turn, and ultimately wins the race. Note that even at high α , Robot 1 is blocked at the early corner—illustrating that even in the best NE for Robot 1, overtaking early is not feasible due to kinodynamic/safety constraints. Thus, via a tuned NE our algorithm can robustly express the lower bound on the cost of a given robot’s solution (within the space of all NEs).

Multi-lane highway merge. We consider a three-lane highway merge with three robots (Fig. 1) to explore homotopy classes of merge options as α and λ_{prox} trade priority against risk. We set $J = (1 - \alpha) J^1 + \alpha (J^2 + J^3)$, $\alpha \in [0, 1]$, and $\lambda_{\text{prox}} \in [0, 2.5]$. As priority for Robot 1 decreases and λ_{prox} increases, the NE transitions through distinct maneuvers. This shows the multiplicity of homotopy classes that often (if not always) exist in complex interactions, that are often impossible to capture with local NE solvers.




Fig. 3: Racetrack overtake. Prioritized robot takes inside lane at final turn and wins the race. Robot 1 (blue) always blocked at first turn by Robot 2 (orange) despite priority.

Opposing-lane overtake. We consider a two-lane bidirectional road with oncoming traffic and three robots to simulate risk-aware passing with heterogeneous capabilities (G^1 has $v \in [0, 10]$, $k = 5$, G^2, G^3 have $v \in [0, 4]$, $k = 3$). We also test how α and λ_{prox} modulate dangerous overtake decisions and select among oncoming-lane vs. in-lane homotopy classes. We set $J = (1 - \alpha)J^1 + \alpha(J^2 + J^3)$, $\alpha \in [0, 1]$, and $\lambda_{\text{prox}} \in [0, 1]$. In Fig. 4 we see that with high priority and low λ_{prox} , Robot 1 overtakes at the first viable opportunity. Increasing λ_{prox} and lowering priority, however, delays the pass until oncoming traffic clears, and beyond a certain threshold, the NE forgoes the overtake entirely. This shows behavioral changes across the different homotopy classes.

6.3 Running time evaluation

Table 1 reports wall-clock solve times (not including preprocess time for graph generation) for GTNS across several maps, goal distances and number of robots. Goal distance D is a proxy for planning horizon T as a time cap T restricts search to the subgraph reachable within $\lceil T/\Delta t \rceil$ edges from the start; by instead varying the goal distance D , we probe increasingly far targets without imposing this cap.

We note that increasing D does not substantially increase running time—a considerable strength of our approach. Unlike optimization-based approaches—




Fig. 4: Opposing-lane overtake. Robot 1 (blue) tries to overtake Robot 2 (orange) via Robot 3’s (green) opposing lane. Different NE overtake strategies chosen as priority and proximity penalty are modulated.

Table 1: Running time evaluation for GTNS. Each row corresponds to a map, number of robots (m). The rightmost block lists time in seconds with the associated goal distance D in parentheses. The *Generic Free* map is an obstacle-free workspace used to stress test expansion (maximal branching). *Generic Obstacles* uses the same spatial dimensions but has random obstacles, pruning many nodes/edges (lower branching).

Map	$ V $	$ E $	m	Solve Times (sec) by Distance D						
Generic Free	7,560	51,733	2	3.31(7)	0.046(9)	2.11(11)	1.57(13)	0.063(15)	4.73(17)	
			3	396.1(7)	3.7(9)	59.1(11)	16.2(13)	9(15)	178.5(17)	
Generic Obstacles	6,300	25,818	2	0.051(7)	0.039(9)	1.28(11)	0.666(13)	0.797(15)	0.136(17)	
			3	15.1(7)	34(9)	3.79(11)	12.2(13)	4.58(15)	12(17)	
Two-Lane Road	5,138	39,618	2	0.034(7)	1.02(9)	0.984(11)	0.501(13)	4.59(15)	2.77(17)	
			3	10.7(7)						
Track	693	4,359	2	0.097(7)	0.24(9)	0.435(11)	0.865(13)	0.247(15)	1.55(17)	
			3	60.3(7)	2.22(9)	32.1(11)	55.2(13)	525.9(15)		
Crossroads	1,428	550	2	0.005(2)	0.006(4)	0.004(6)	0.007(7)	0.007(8)		
			3	0.046(2)	0.05(4)	0.045(6)	0.037(7)	0.041(8)		
			4	0.162(2)	0.111(4)	0.154(6)	1.53(7)	0.283(8)		
			5	0.724(2)	1(4)	12.2(6)	4.93(7)	12.1(8)		

where running time typically scales with horizon due to the growing number of decision variables—our graph-based formulation keeps per-step local problems fixed-size and searches over a discrete action set, so increasing D does not balloon the optimizer. By contrast, running time generally increases with the number of robots. The joint neighbor set grows combinatorially, interaction geometry becomes richer (more constraints and blocking situations), and the equilibrium check ISNASH-EQUILIBRIUM triggers more single-robot A* calls (a principal computational bottleneck) as m grows. An exception is the *Crossroads* map: its branching factor is small and interactions are simple, so adding robots does not markedly increase search effort.

Across all settings, *instance difficulty* (map layout and interaction geometry) dominates running time. Tight merges, opposing-lane passes, and narrow corridors can make some two-robot cases slower—or even infeasible at a given discrete horizon—than multi-robot cases on simpler maps with few viable options. This can also explain why *Generic Obstacles* is faster on average than *Generic Free*: moderate obstacle density reduces branching without making the graph too sparse. See App. B, Fig. 9 for visualization.

6.4 Comparison to LaValle-Hutchinson method

As noted in Sec. 1, the LaValle-Hutchinson method [20] also leverages graph search techniques and yields equilibrium guarantees. Here, we study the computational burden of generating, pairwise verifying, and explicitly maintaining all non-dominated candidate trajectories (up to one representative per equivalence class) within that approach. Unsurprisingly, in maps with low branching, and

short horizons (e.g., Crossroads, Track), and two robots, the two algorithms are almost comparable: our method achieves a modest speedup of roughly $1.5\text{-}2\times$. For dense graphs with high branching factors and long horizons (e.g., Two-Lane-Road, Generic Free), however, the speedup reaches $2\text{-}3$ orders of magnitude.

The main reason is that the LaValle-Hutchinson algorithm must carry *every* non-dominated partial trajectory (equivalency class) at each joint state. As the branching factor and horizon grow, the number of distinct cost vectors at a single node can grow combinatorially, so each additional layer multiplies both the number of wavefront paths and the cost of dominance checks. In contrast, our planner explores the composite state space with a best-first strategy guided by single-robot distance heuristics and aggressive dominance, so it (a) explores more promising areas first (only) and (b) effectively prunes entire sub-trees. Thus, its complexity scales primarily with the number of joint states actually visited, rather than with the total number of distinct joint trajectories, which makes high-branching and large-horizon regimes much easier to handle in practice.

7 Conclusion and future work

We presented Game-Theoretic Nested Search (GTNS), a scalable framework for computing global Nash equilibria for game-theoretic motion planning with general dynamics. GTNS enjoys strong theoretical guarantees, while also allowing explicit selection among multiple equilibria via a user-specified global objective. Across driving and racing scenarios, GTNS finds collision-free, dynamically feasible interactions within seconds.

In the future, we plan to address a key limitation of the approach, which is the reliance on high-fidelity kinodynamic roadmaps, whose construction requires many calls to a BVP solver; although performed offline, and can be reused in different scenarios, the compute cost and memory footprint grow with state discretization and connection density, and are a key practical bottleneck (e.g., map construction for the overtake scenario (Fig. 4) required several hours to compute). One possible direction could be substituting such graphs with kinodynamic trees that are computed on-the-fly by sampling-based planners [11, 14, 24]. This modification should be done with care as to ensure that the game-theoretic guarantees hold. We also plan to explore improving the algorithm’s scalability with respect to the number of robots. Although the tensor graph is explored implicitly, joint branching still grows with the number of robots; moreover, each equilibrium check triggers single-robot searches that can dominate running time in tightly coupled scenes, as well as parallelizing the single-robot searches.

In the longer run, we plan to extend the approach to the closed-loop setting and study principled tie-breaking and fairness- and risk-aware objectives, as well as coordination mechanisms when several NEs exist.

Acknowledgments. This work was supported by the Israel Ministry of Science and Technology grant no. 2034500 and the Technion Autonomous Systems Program grant no. 2072638. The AI system ChatGPT was used for light editing and grammar enhancement, as well as a preliminary literature review.

References

1. Andersson, J.A.E., Gillis, J., Horn, G., Rawlings, J.B., Diehl, M.: CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation* **11**(1), 1–36 (2019). <https://doi.org/10.1007/s12532-018-0139-4>
2. Aumann, R.J.: Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics* **1**(1), 67–96 (1974)
3. Burger, C., Fischer, J., Bieder, F., Taş, O.S., Stiller, C.: Interaction-aware game-theoretic motion planning for automated vehicles using bi-level optimization. In: *International Conference on Intelligent Transportation Systems*. pp. 3978–3985 (2022)
4. Cohen, B., Phillips, M., Likhachev, M.: Planning single-arm manipulations with n-arm robots. In: *the International Symposium on Combinatorial Search*. vol. 6, pp. 226–227 (2015)
5. Dayan, D., Solovey, K., Pavone, M., Halperin, D.: Near-optimal multi-robot motion planning with finite sampling. *IEEE Transactions on Robotics* (2023)
6. Dellin, C., Srinivasa, S.: A unifying formalism for shortest path problems with expensive edge evaluations via lazy best-first search over paths with edge selectors. In: *International Conference on Automated Planning and Scheduling*. vol. 26, pp. 459–467 (2016)
7. Department for Transport: Road safety factors: initial analysis. Official statistics, UK Department for Transport, London, UK (May 2024), <https://www.gov.uk/government/statistics/road-safety-factors-initial-analysis/road-safety-factors-initial-analysis>, published 30 May 2024
8. Dezső, B., Jüttner, A., Kovács, P.: LEMON—an open source c++ graph template library. *Electronic Notes in Theoretical Computer Science* **264**(5), 23–45 (2011)
9. Fisac, J.F., Bronstein, E., Stefansson, E., Sadigh, D., Sastry, S.S., Dragan, A.D.: Hierarchical game-theoretic planning for autonomous vehicles. In: *International conference on robotics and automation*. pp. 9590–9596. IEEE (2019)
10. Fridovich-Keil, D., Ratner, E., Peters, L., Dragan, A.D., Tomlin, C.J.: Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In: *International Conference on Robotics and Automation*. pp. 1475–1481 (2020)
11. Fu, M., Solovey, K., Salzman, O., Alterovitz, R.: Toward certifiable optimal motion planning for medical steerable needles. *Int. J. Robotics Res.* **42**(10), 798–826 (2023)
12. Gao, Y., Piccinini, M., Zhang, Y., Wang, D., Moller, K., Brusnicki, R., Zarrouki, B., Gambi, A., Totz, J.F., Storms, K., et al.: Foundation models in autonomous driving: A survey on scenario generation and scenario analysis. *arXiv preprint arXiv:2506.11526* (2025)
13. Hart, P.E., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* **4**(2), 100–107 (1968)
14. Hauser, K., Zhou, Y.: Asymptotically optimal planning by feasible kinodynamic planning in a state-cost space. *IEEE Trans. Robotics* **32**(6), 1431–1443 (2016)
15. Hönig, W., de Haro, J.O., Toussaint, M.: db-A*: Discontinuity-bounded search for kinodynamic mobile robot motion planning. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 13540–13547. IEEE (2022)
16. Jiang, C., Cornman, A., Park, C., Sapp, B., Zhou, Y., Anguelov, D., et al.: Motion-diffuser: Controllable multi-agent motion prediction using diffusion. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9644–9653 (2023)

17. Kavraki, L.E., Svestka, P., Latombe, J.C., Overmars, M.H.: Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* **12**(4), 566–580 (1996)
18. Kleinbort, M., Granados, E., Solovey, K., Bonalli, R., Bekris, K.E., Halperin, D.: Refined analysis of asymptotically-optimal kinodynamic planning in the state-cost space. In: *IEEE International Conference on Robotics and Automation*. pp. 6344–6350 (2020)
19. LaValle, S.M.: *Planning algorithms*. Cambridge university press (2006)
20. LaValle, S.M., Hutchinson, S.A.: Optimal motion planning for multiple robots having independent goals. *IEEE Transactions on Robotics and Automation* **14**(6), 912–925 (2002)
21. Le Cleac’h, S., Schwager, M., Manchester, Z.: Lucidgames: Online unscented inverse dynamic games for adaptive trajectory prediction and planning. *IEEE Robotics and Automation Letters* **6**(3), 5485–5492 (2021)
22. Le Cleac’h, S., Schwager, M., Manchester, Z.: Algames: a fast augmented lagrangian solver for constrained dynamic games. *Autonomous Robots* **46**(1), 201–215 (2022)
23. Lee, N., Choi, W., Vernaza, P., Choy, C.B., Torr, P.H., Chandraker, M.: Desire: Distant future prediction in dynamic scenes with interacting agents. In: *IEEE conference on computer vision and pattern recognition*. pp. 336–345 (2017)
24. Li, Y., Littlefield, Z., Bekris, K.E.: Asymptotically optimal sampling-based kinodynamic planning. *International Journal of Robotics Research* **35**(5), 528–564 (2016)
25. Liniger, A., Lygeros, J.: A noncooperative game approach to autonomous racing. *IEEE Transactions on Control Systems Technology* **28**(3), 884–897 (2019)
26. Mandalika, A., Choudhury, S., Salzman, O., Srinivasa, S.: Generalized lazy search for robot motion planning: Interleaving search and edge evaluation via event-based toggles. In: *International Conference on Automated Planning and Scheduling*. vol. 29, pp. 745–753 (2019)
27. Najm, W.G., Ranganathan, R., Srinivasan, G., Smith, J.D., Toma, S., Swanson, E., Burgett, A.: Description of light-vehicle pre-crash scenarios for safety applications based on vehicle-to-vehicle communications. Technical Report DOT HS 811 731, U.S. Department of Transportation, National Highway Traffic Safety Administration, Washington, DC (2013), <https://www.nhtsa.gov/sites/nhtsa.gov/files/811731.pdf>
28. Nash Jr, J.F.: Equilibrium points in n-person games. *Proceedings of the national academy of sciences* **36**(1), 48–49 (1950)
29. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.): *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK (2007)
30. Rowold, M., Langmann, A., Lohmann, B., Betz, J.: Open-loop and feedback nash trajectories for competitive racing with iLQGames. In: *IEEE International Conference on Intelligent Transportation Systems*. pp. 1827–1834 (2024)
31. Rudenko, A., Palmieri, L., Herman, M., Kitani, K.M., Gavrila, D.M., Arras, K.O.: Human motion trajectory prediction: A survey. *International Journal of Robotics Research* **39**(8), 895–935 (2020)
32. Salzmann, T., Ivanovic, B., Chakravarty, P., Pavone, M.: Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In: *European Conference on Computer Vision*. pp. 683–700. Springer (2020)
33. Schmerling, E., Janson, L., Pavone, M.: Optimal sampling-based motion planning under differential constraints: the drift case with linear affine dynamics. In: *IEEE Conference on Decision and Control*. pp. 2574–2581 (2015)

34. Schmerling, E., Janson, L., Pavone, M.: Optimal sampling-based motion planning under differential constraints: The driftless case. In: IEEE International Conference on Robotics and Automation. pp. 2368–2375 (2015)
35. Schmerling, E., Leung, K., Vollprecht, W., Pavone, M.: Multimodal probabilistic model-based planning for human-robot interaction. In: IEEE International Conference on Robotics and Automation. pp. 3399–3406 (2018)
36. Shoham, Y., Leyton-Brown, K.: Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. Cambridge University Press, Cambridge, UK (2009)
37. Shome, R., Kavradi, L.E.: Asymptotically optimal kinodynamic planning using bundles of edges. In: IEEE International Conference on Robotics and Automation. pp. 9988–9994 (2021)
38. Shome, R., Solovey, K., Dobson, A., Halperin, D., Bekris, K.E.: dRRT^{*}: Scalable and informed asymptotically-optimal multi-robot motion planning. *Auton. Robots* **44**(3-4), 443–467 (2020)
39. Solovey, K., Salzman, O., Halperin, D.: Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. *International Journal of Robotics Research* **35**(5), 501–513 (2016)
40. Spica, R., Cristofalo, E., Wang, Z., Montijano, E., Schwager, M.: A real-time game theoretic planner for autonomous two-player drone racing. *IEEE Transactions on Robotics* **36**(5), 1389–1403 (2020)
41. Tikna, A., Roveri, M., Fontanelli, D., Palopoli, L.: When graphs meet game theory: a scalable approach for robotic car racing. In: IEEE Computers, Software, and Applications Conference. pp. 1–8. IEEE (2023)
42. Wagner, G., Choset, H.: Subdimensional expansion for multirobot path planning. *Artif. Intell.* **219**, 1–24 (2015)
43. Wang, M., Wang, Z., Talbot, J., Gerdes, J.C., Schwager, M.: Game-theoretic planning for self-driving cars in multivehicle competitive scenarios. *IEEE Transactions on Robotics* **37**(4), 1313–1325 (2021)
44. Wang, Y., Xing, S., Can, C., Li, R., Hua, H., Tian, K., Mo, Z., Gao, X., Wu, K., Zhou, S., You, H., Peng, J., Zhang, J., Wang, Z., Song, R., Yan, M., Zimmer, W., Zhou, X., Li, P., Lu, Z., Chen, C.J., Huang, Y., Rossi, R.A., Sun, L., Yu, H., Fan, Z., Yang, F.H., Kang, Y., Greer, R., Liu, C., Lee, E.H., Di, X., Ye, X., Ren, L., Knoll, A., Li, X., Ji, S., Tomizuka, M., Pavone, M., Yang, T., Du, J., Yang, M.H., Wei, H., Wang, Z., Zhou, Y., Li, J., Tu, Z.: Generative ai for autonomous driving: Frontiers and opportunities (2025)
45. Wu, C., Kreidieh, A.R., Parvate, K., Vinitsky, E., Bayen, A.M.: Flow: A modular learning framework for mixed autonomy traffic. *IEEE Transactions on Robotics* (2021)
46. Wurman, P.R., Barrett, S., Kawamoto, K., MacGlashan, J., Subramanian, K., Walsh, T.J., Capobianco, R., Devlic, A., Eckert, F., Fuchs, F., et al.: Outracing champion gran turismo drivers with deep reinforcement learning. *Nature* **602**(7896), 223–228 (2022)
47. Zhu, E.L., Borrelli, F.: A sequential quadratic programming approach to the solution of open-loop generalized nash equilibria. In: IEEE International Conference on Robotics and Automation. pp. 3211–3217. IEEE (2023)