

Game-Theoretic Autonomous Driving: A Graphs of Convex Sets Approach

Nikolaj Käfer^{1,2*}, Ahmed Khalil^{2*}, Edward Huynh^{2*}, Efstathios Bakolas², and
David Fridovich-Keil²

¹ ETH Zurich, Zurich, ZH 8092, Switzerland
nkaefer@ethz.ch

² The University of Texas at Austin, Austin, TX 78712, USA
{akhalil, edhuynh, dfk}@utexas.edu, bakolas@austin.utexas.edu

Abstract. Multi-vehicle autonomous driving couples strategic interaction with hybrid (discrete-continuous) maneuver planning under shared safety constraints. We introduce IBR-GCS, an Iterative Best Response (IBR) planning approach based on the Graphs of Convex Sets (GCS) framework that models highway driving as a generalized noncooperative game. IBR-GCS integrates combinatorial maneuver reasoning, trajectory planning, and game-theoretic interaction within a unified framework. The key novelty is a vehicle-specific, strategy-dependent GCS construction. Specifically, at each best-response update, each vehicle builds its own graph conditioned on the current strategies of the other vehicles, with vertices representing lane-specific, time-varying, convex, collision-free regions and edges encoding dynamically feasible transitions. This yields a shortest-path problem in GCS for each best-response step, which admits an efficient convex relaxation that can be solved using convex optimization tools without exhaustive discrete tree search. We then apply an iterative best-response scheme in which vehicles update their trajectories sequentially and provide conditions under which the resulting inexact updates converge to an approximate generalized Nash equilibrium. Simulation results across multi-lane, multi-vehicle scenarios demonstrate that IBR-GCS produces safe trajectories and strategically consistent interactive behaviors.

Keywords: Multi-Vehicle Systems · Game-Theoretic Motion Planning · Autonomous Driving

1 Introduction

The deployment of self-driving cars promises improvements in safety, efficiency, and comfort in day-to-day transportation, but it also introduces fundamental challenges in planning with interactions. In mixed traffic, coordination is decentralized, and vehicle interactions combine cooperative and competitive elements: all vehicles must comply with shared traffic rules to ensure safety, yet strategic behavior naturally arises in scenarios such as merging, lane changing, and overtaking. To model the pursuit of individual objectives under shared constraints,

*Equal contribution

noncooperative game theory has been widely used to cast autonomous driving as a complete-information game [9, 14, 24]. Compared to model predictive control approaches that ignore strategic interaction, game-theoretic formulations can better capture anticipatory behavior and improve performance in interactive scenarios [25, 28]. A particularly useful subclass is (generalized) potential games, where unilateral incentives are aligned with a single scalar potential function, yielding algorithmic and theoretical tractability [4, 10, 13, 29]. The versatility of potential games has led to their adoption in diverse multi-agent settings [2, 12, 26, 27].

A key difficulty in multi-vehicle planning is that it is inherently hybrid: discrete decisions arise from lane selection and maneuver choices, while continuous decisions are governed by vehicle dynamics. This hybrid structure is often modeled via mixed-integer programs [5, 6, 8, 20], but such formulations can tightly couple discrete and continuous variables, limiting scalability and making it difficult to exploit problem structure in large multi-vehicle settings [7, 19].

To address these challenges, we leverage Graphs of Convex Sets (GCS) [15, 17], which is a recently proposed optimization framework that has been successfully applied to trajectory optimization and motion planning problems with collision avoidance [11, 16, 22]. In GCS, discrete decisions are represented by a graph topology: vertices correspond to convex feasible regions (*e.g.*, safe sets for each vehicle), and edges encode feasible transitions between them (*e.g.*, collision-free trajectories from one feasible region to another). This representation preserves convexity within regions and along transitions, and it admits a convex relaxation that is often tight in practice [17], enabling efficient computation while capturing combinatorial structure.

In this work, we present IBR-GCS, a GCS-based method for multi-vehicle highway driving that models interaction as a noncooperative game with potentially conflicting objectives while respecting traffic rules. The proposed approach, IBR-GCS, integrates combinatorial maneuver reasoning, trajectory planning, and game-theoretic interaction within a unified framework. During each iteration of IBR-GCS, all vehicles update their trajectories sequentially. Each vehicle i constructs a GCS whose vertices represent lane-specific collision-free regions over time and whose edges encode dynamically feasible, safety-preserving transitions. Given the current trajectories of the other vehicles, the resulting best-response computation reduces to a shortest-path problem (SPP) in GCS, yielding a mixed-integer convex formulation with an efficient convex relaxation that is often tight in practice. We embed these updates in an Iterative Best-Response (IBR) scheme and provide conditions under which the resulting inexact best-response dynamics converge to an approximate generalized Nash equilibrium (GNE). We evaluate IBR-GCS in numerical simulations of multi-lane highway scenarios, demonstrating safe trajectories and strategically consistent behaviors.

The remainder of the paper is organized as follows. Section 2 reviews graphs of convex sets and generalized potential games. Section 3 formalizes the multi-vehicle driving game and states the key assumptions. Section 4 presents IBR-GCS, including the strategy-dependent GCS construction, the iterative best-response algorithm, and the accompanying error analysis. Section 5 presents simulation results. Finally, Section 6 discusses conclusions.

2 Preliminaries

This section reviews GCS and generalized potential games, which form the basis for the proposed approach.

2.1 Shortest Path Problem in Graphs of Convex Sets

The GCS framework provides a graph-based representation for optimization problems that combine continuous decision variables with a combinatorial structure. The framework was originally introduced for trajectory optimization and motion planning with obstacle avoidance. A detailed formulation can be found in [15, 17].

Definition 1 (Graph of Convex Sets). *A graph of convex sets is a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with vertex set \mathcal{V} and edge set \mathcal{E} , where each vertex $V \in \mathcal{V}$ is associated with a convex set $\mathcal{X}_V \subset \mathbb{R}^n$, and each directed edge $e = (U, W) \in \mathcal{E}$ represents an admissible one-step move from vertex U to vertex W (i.e., between the sets \mathcal{X}_U and \mathcal{X}_W).*

In the relevant literature, vertices typically correspond to convex feasible regions that are subsets of the state or trajectory space, while edges encode admissible transitions between regions, such as continuity, dynamical feasibility, or safety constraints [11, 16].

Each vertex $V \in \mathcal{V}$ is associated with a continuous decision variable $x_V \in \mathcal{X}_V$. For each edge $e = (U, W) \in \mathcal{E}$, convex constraints may couple (x_U, x_W) via $(x_U, x_W) \in \mathcal{X}_e$, where $\mathcal{X}_e \subset \mathbb{R}^n \times \mathbb{R}^n$ is a convex set. Each edge may also be assigned a convex nonnegative cost function $c_e : \mathcal{X}_e \rightarrow \mathbb{R}_{\geq 0}$. Additionally, vertices may be assigned convex nonnegative costs $c_V : \mathcal{X}_V \rightarrow \mathbb{R}_{\geq 0}$.

We recall the definition of a path in a graph of convex sets.

Definition 2 (Path). *Given a source vertex V_0 and a target vertex $V_K \neq V_0$, a path is a sequence of vertices (V_0, \dots, V_K) and directed edges such that $(V_{k-1}, V_k) \in \mathcal{E}$ for all $k = 1, \dots, K$.*

To find the shortest path through \mathcal{G} , we introduce binary decision variables $z_e \in \{0, 1\}$ for all $e \in \mathcal{E}$ to indicate whether an edge e is selected, and vertex-selection variables $y_V \in \{0, 1\}$ for all $V \in \mathcal{V}$ (equal to 1 for vertices V on the selected path and 0 otherwise). Since the integrality of y_V is implied by the flow constraint (1c), the constraints on y_V are relaxed to $y_V \in [0, 1]$ for all $V \in \mathcal{V}$. For each vertex $V \in \mathcal{V}$, let $\mathcal{E}_V^{\text{in}} := \{(U, V) \in \mathcal{E}\}$ and $\mathcal{E}_V^{\text{out}} := \{(V, W) \in \mathcal{E}\}$ denote its sets of incoming and outgoing edges, respectively. A standard mixed-integer

formulation of the GCS shortest path problem is as follows:

$$\begin{array}{l} \text{minimize} \\ \{x_V, y_V\}_{V \in \mathcal{V}}, \\ \{z_e\}_{e \in \mathcal{E}}, \end{array} \quad \sum_{e=(U,W) \in \mathcal{E}} z_e c_e(x_U, x_W) + \sum_{V \in \mathcal{V}} y_V c_V(x_V) \quad (1a)$$

$$\text{subject to} \quad \sum_{e \in \mathcal{E}_V^{\text{in}}} z_e + \delta_{V, V_0} = \sum_{e \in \mathcal{E}_V^{\text{out}}} z_e + \delta_{V, V_K}, \quad \forall V \in \mathcal{V}, \quad (1b)$$

$$y_V = \delta_{V, V_0} + \sum_{e \in \mathcal{E}_V^{\text{in}}} z_e, \quad \forall V \in \mathcal{V}, \quad (1c)$$

$$x_V \in \mathcal{X}_V, \quad \forall V \in \mathcal{V}, \quad (1d)$$

$$z_e = 1 \Rightarrow (x_U, x_W) \in \mathcal{X}_e, \quad \forall e = (U, W) \in \mathcal{E}, \quad (1e)$$

$$z_e \in \{0, 1\}, \quad \forall e \in \mathcal{E}, \quad (1f)$$

$$0 \leq y_V \leq 1, \quad \forall V \in \mathcal{V}. \quad (1g)$$

Here, $\delta_{V,W}$ denotes the Kronecker delta, where $\delta_{V,W} = 1$, if $V = W$, and $\delta_{V,W} = 0$, otherwise. In (1b), the terms δ_{V, V_0} and δ_{V, V_K} therefore create the required unit flow imbalance at the source vertex V_0 and target vertex V_K . In contrast to the classical SPP [3], the GCS SPP is NP-hard in general, *cf.* [17, Theorem 3.1].

Convex Relaxation A standard convex relaxation replaces (1f) with the box constraint $0 \leq z_e \leq 1$ and converts the implications in (1e) into a convex perspective (or conic) formulation; see [15, 17] for canonical constructions.

This relaxation can be tight in many motion-planning instances, but in general it may be loose [15, Proposition 8.1]. Existing formal tightness guarantees are limited to restrictive settings [15, Proposition 8.2]. In this paper, we empirically observe tightness in our setting (*cf.* Section 5), but we do not claim a general tightness guarantee.

2.2 Generalized Potential Games

We consider a game with N vehicles, indexed by $\mathcal{N} := \{1, \dots, N\}$. Vehicle $i \in \mathcal{N}$ chooses a strategy θ_i , which in our driving setting represents a finite-horizon motion plan (trajectory and lane-change actions) and will later be encoded as a path through a vehicle-specific graph of convex sets. Vehicle i incurs a cost $J_i(\theta_i, \theta_{-i})$, where $\theta_{-i} := (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_N)$ denotes the strategies of all other vehicles. In our setting, feasibility depends on other vehicles through coupled constraints, so we use a *generalized* strategy set of the form

$$\Theta_i(\theta_{-i}) := \{\theta_i \mid (\theta_i, \theta_{-i}) \in \Theta\}, \quad (2)$$

where Θ is the joint feasible set. Each vehicle's cost is a function $J_i : \Theta \rightarrow \mathbb{R}$, evaluated on feasible joint strategies $\theta = (\theta_i, \theta_{-i}) \in \Theta$.

Collecting the vehicles, costs, and coupled feasibility constraints yields a game in which each vehicle solves an optimization problem whose feasible set depends on the other vehicles' strategies.

Definition 3 (Game). *The game is denoted by G and is specified by the following set of coupled optimization problems:*

$$G := \begin{cases} \underset{\theta_i}{\text{minimize}} & J_i(\theta_i, \theta_{-i}) \\ \text{subject to} & \theta_i \in \Theta_i(\theta_{-i}), \end{cases} \quad \forall i \in \mathcal{N}. \quad (3)$$

A central solution concept for games of this form is the Generalized Nash Equilibrium (GNE), in which no vehicle can unilaterally reduce its cost while satisfying the feasibility constraints imposed by the other vehicles' strategies.

Definition 4 (Generalized Nash Equilibrium, GNE). *A strategy profile $\theta^* \in \Theta$ is a GNE of G if, for every vehicle $i \in \mathcal{N}$,*

$$J_i(\theta_i^*, \theta_{-i}^*) \leq J_i(\theta_i, \theta_{-i}^*), \quad \forall \theta_i \in \Theta_i(\theta_{-i}^*). \quad (4)$$

To analyze the existence of equilibria and the behavior of best-response dynamics, it is often useful to identify games whose incentives can be summarized by a single scalar function. This motivates the notion of a generalized potential game, which extends exact potential games to settings with coupled feasibility constraints [18].

Definition 5 (Generalized potential game). *The game G is a generalized potential game if there exists a function $\Phi : \Theta \rightarrow \mathbb{R}$ such that for all $i \in \mathcal{N}$ and all θ_{-i} for which $\Theta_i(\theta_{-i}) \neq \emptyset$, it holds for all $\theta_i, \theta'_i \in \Theta_i(\theta_{-i})$ that*

$$J_i(\theta'_i, \theta_{-i}) - J_i(\theta_i, \theta_{-i}) = \Phi(\theta'_i, \theta_{-i}) - \Phi(\theta_i, \theta_{-i}). \quad (5)$$

The function Φ is called a generalized potential.

We next introduce the iterative best-response (IBR) dynamics that will be used in Section 4. Given θ_{-i} , the set of best responses for each vehicle $i \in \mathcal{N}$, denoted $\text{BR}_i(\theta_{-i})$, is given by

$$\text{BR}_i(\theta_{-i}) := \underset{\theta_i \in \Theta_i(\theta_{-i})}{\text{argmin}} J_i(\theta_i, \theta_{-i}). \quad (6)$$

Starting from an initial feasible profile $\theta^0 \in \Theta$, IBR updates vehicles sequentially. We define

$$\theta^{k,i} := (\theta_1^{k+1}, \dots, \theta_i^{k+1}, \theta_{i+1}^k, \dots, \theta_N^k), \quad i = 1, \dots, N, \quad (7)$$

for the intermediate profile after vehicles $1, \dots, i$ have been updated during iteration k . With this, the update rule is

$$\theta_i^{k+1} \in \text{BR}_i(\theta_{-i}^{k,i-1}), \quad i = 1, \dots, N, \quad k \geq 0, \quad (8)$$

with $\theta^{k,0} = \theta^k$ at the beginning of iteration k and $\theta^{k,N} = \theta^{k+1}$ at the end of iteration k . In a generalized potential game, any exact best-response update is a descent step for the potential function Φ . A standard result is given below, cf. [18].

Proposition 1 (Potential descent under exact IBR). *Suppose G is a generalized potential game with generalized potential Φ (Definition 5). If, at some intermediate profile $\theta^{k,i-1}$, vehicle $i \in \mathcal{N}$ performs an exact best-response update (6), then*

$$\Phi(\theta^{k,i}) \leq \Phi(\theta^{k,i-1}).$$

Consequently, if all vehicles update with exact best responses, then the sequence $\{\Phi(\theta^k)\}_{k \geq 0}$ is monotonically non-increasing along IBR, and therefore convergent, if Φ is bounded below on Θ .

3 Problem Formulation

We consider a multi-vehicle highway driving scenario with N vehicles indexed by $\mathcal{N} := \{1, \dots, N\}$ operating on a set of straight lanes $\mathcal{L} \subset \mathbb{N}$, as illustrated in Fig. 1. Planning is performed over a finite-horizon $\mathcal{T} := \{0, \dots, T-1\}$ with a discretization step $\Delta t \in \mathbb{R}_{>0}$ and transition index set $\mathcal{T}^- := \{0, \dots, T-2\}$. Vehicle $i \in \mathcal{N}$ has a continuous longitudinal position $s_i(t)$, speed $v_i(t)$ and acceleration $a_i(t)$, as well as a discrete lane index $z_i(t) \in \mathcal{L}$ and blinker state (lane-change action) $b_i(t) \in \mathcal{B} := \{-1, 0, 1\}$, where $b_i(t) = -1$ indicates an imminent lane change to the right, $b_i(t) = 0$ indicates staying in lane, and $b_i(t) = 1$ indicates a lane change to the left.




Fig. 1: Vehicles operating on a highway with vehicle-specific variables (left) and longitudinal distance (right).

Dynamics and bounds: The motion of vehicle $i \in \mathcal{N}$ is modeled by the discrete-time dynamics:

$$\begin{aligned} s_i(t+1) &= s_i(t) + \Delta t v_i(t), \\ v_i(t+1) &= v_i(t) + \Delta t a_i(t), \end{aligned} \quad \forall t \in \mathcal{T}^-. \quad (9)$$

We enforce state bounds:

$$s_i(t) \in [\underline{s}, \bar{s}], \quad v_i(t) \in [\underline{v}_i, \bar{v}_i], \quad \forall t \in \mathcal{T}, \quad (10)$$

where \underline{s} and \bar{s} denote the endpoints of the road segment and $\underline{v}_i, \bar{v}_i$ are the minimum and maximum allowable speeds for vehicle i . Acceleration is bounded by actuation limits

$$a_i(t) \in [\underline{a}_i, \bar{a}_i], \quad \forall t \in \mathcal{T}^-, \quad (11)$$

with $\underline{a}_i, \bar{a}_i$ denoting the minimum and maximum admissible accelerations of vehicle i .

Lane evolution: Lane changes are restricted to adjacent lanes and are commanded by the blinker:

$$z_i(t+1) = z_i(t) + b_i(t), \quad b_i(t) \in \mathcal{B}, \quad z_i(t+1) \in \mathcal{L}, \quad \forall t \in \mathcal{T}^-. \quad (12)$$

Relative coordinates: For any pair of vehicles $i, j \in \mathcal{N}$, define the longitudinal and lateral offsets:

$$d_{i,j}(t) := s_j(t) - s_i(t), \quad z_{i,j}(t) := z_j(t) - z_i(t), \quad (13)$$

as shown in Fig. 1. The longitudinal offset, in view of (9), evolves as:

$$d_{i,j}(t+1) = d_{i,j}(t) + \Delta t(v_j(t) - v_i(t)), \quad \forall t \in \mathcal{T}^-. \quad (14)$$

Safety rules: To prevent collisions, we impose longitudinal safety for vehicles in the same lane and lateral safety for vehicles in adjacent lanes.

Rule 1 (Longitudinal safety in the same lane). If $z_{i,j}(t) = 0$, then vehicles must maintain a minimum longitudinal separation, *i.e.*,

$$|d_{i,j}(t)| \geq d_i^s, \quad \forall t \in \mathcal{T}, \quad (15)$$

where $d_i^s > 0$ is the required safety gap for vehicle i .

Rule 2 (Lateral safety across adjacent lanes). If vehicles are side-by-side in adjacent lanes, *i.e.*,

$$|d_{i,j}(t)| \leq d_i^s, \quad |z_{i,j}(t)| = 1, \quad \forall t \in \mathcal{T}, \quad (16)$$

then simultaneous swaps into each other's lanes are forbidden:

$$z_i(t+1) \neq z_j(t), \quad z_j(t+1) \neq z_i(t), \quad \forall t \in \mathcal{T}^-. \quad (17)$$

Preferences and objectives: Each vehicle $i \in \mathcal{N}$ is assigned a desired cruising speed $v_i^{\text{des}} \in [v_i, \bar{v}_i]$ and a desired lane $\ell_i^{\text{des}} \in \mathcal{L}$. We also fix weights $w_{i,v}, w_{i,\ell}, w_{i,a}, w_{i,b} \in \mathbb{R}_{>0}$ that trade off speed tracking, lane preference, control effort, and lane-switching, respectively. These parameters are used to define the edge and terminal costs in the GCS construction in Section 4.

3.1 Individual vehicle strategy and optimization problem

A vehicle i 's finite-horizon strategy is the collection of its state and input trajectories,

$$\theta_i := \left(\{s_i(t), v_i(t), z_i(t)\}_{t \in \mathcal{T}}, \{a_i(t), b_i(t)\}_{t \in \mathcal{T}^-} \right), \quad (18)$$

and we write $\theta := (\theta_1, \dots, \theta_N)$ for the joint strategy profile. Given initial conditions $\{s_i(0), v_i(0), z_i(0)\}_{i \in \mathcal{N}}$, each vehicle $i \in \mathcal{N}$ seeks a feasible strategy that respects the shared safety rules while optimizing its own preferences (*e.g.*, tracking a desired speed v_i^{des} and lane ℓ_i^{des} , as introduced earlier). We proceed by making the following assumption.

Assumption 1 (Separable objectives) *Vehicle objectives are separable across vehicles, i.e., $J_i(\theta_i) = J_i(\theta_i, \theta_{-i})$, and interaction between vehicles occurs only through the coupled feasibility constraints (Rules 1–2).*

Assumption 1 captures the modeling choice that vehicles do not directly penalize (or reward) other vehicles in their objectives; rather, strategic coupling arises solely because each vehicle’s feasible set depends on the others through shared safety rules. Formally, vehicle i solves:

$$\begin{aligned} & \underset{\theta_i}{\text{minimize}} && J_i(\theta_i) \\ & \text{subject to} && s_i(0), v_i(0), z_i(0) \text{ given,} \\ & && (9)-(12), \\ & && (15), (17), \quad \forall j \in \mathcal{N} \setminus \{i\}. \end{aligned} \tag{19}$$

The constraints in Equations (9)-(12) enforce discrete-time longitudinal dynamics (9), bounded roadway and speed domains (10), actuation limits (11), and lane-change feasibility (12). Vehicle interactions enter only through the coupled collision-avoidance Rules 1 and 2: (15) enforces longitudinal separation between vehicles in the same lane, while (17) prevents simultaneous lane swaps when vehicles are side-by-side in adjacent lanes.

To connect each vehicle’s optimization problem in (19) with the GCS construction in Section 4, we impose the following mild modeling assumptions.

Assumption 2 (Convexity of continuous costs) *For each vehicle $i \in \mathcal{N}$ and for any fixed lane and blinker sequences (z_i, b_i) , the objective $J_i(\theta_i)$ is proper, closed, and convex in the continuous variables $\{s_i(t), v_i(t)\}_{t \in \mathcal{T}}$ and $\{a_i(t)\}_{t \in \mathcal{T}^-}$, and takes values in $\mathbb{R}_{\geq 0}$.*

Assumption 2 ensures that once the discrete lane/blinker sequence is fixed, the remaining optimization over continuous variables is a convex program. This property is essential for modeling each maneuver option as a convex set (vertex) and each time-adjacent transition as a convex constraint (edge) in a GCS.

Example objective function. A concrete cost function consistent with the preferences and assumptions introduced earlier, and used later in Section 4, is the following cost function:

$$\begin{aligned} J_i(\theta_i) := & \sum_{t \in \mathcal{T}^-} \left[w_{i,v} (v_i(t+1) - v_i^{\text{des}})^2 + w_{i,\ell} (z_i(t+1) - \ell_i^{\text{des}})^2 \right. \\ & \left. + w_{i,a} a_i(t)^2 + w_{i,b} b_i(t)^2 \right] + w_{i,v} (v_i(T-1) - v_i^{\text{des}})^2, \end{aligned} \tag{20}$$

which penalizes deviations from the desired speed and lane, as well as acceleration and lane changes. For any fixed lane and blinker sequences (z_i, b_i) , the cost in (20) is convex in the continuous variables (s_i, v_i, a_i) .

Let Θ denote the joint feasible set of all strategy profiles θ satisfying the dynamics, bounds, lane evolution, and safety rules for every vehicle. Under Assumption 1, the induced game admits a generalized potential formulation (Definition 5) with potential:

$$\Phi(\theta) := \sum_{i \in \mathcal{N}} J_i(\theta_i). \tag{21}$$

Problem 1. Given initial conditions $\{s_i(0), v_i(0), z_i(0)\}_{i \in \mathcal{N}}$ and vehicle preferences (e.g., v_i^{des} and ℓ_i^{des}), find a feasible joint strategy profile $\theta^* \in \Theta$ such that no vehicle can unilaterally decrease its cost while respecting the coupled feasibility constraints, i.e., θ^* is a generalized Nash equilibrium of the induced game.

4 Iterative Best Response Graphs of Convex Sets

We address Problem 1 using IBR. At each iteration, vehicles update their strategies sequentially by solving a single-vehicle motion planning problem while treating the trajectories of all other vehicles as fixed.

For vehicle $i \in \mathcal{N}$, we construct a directed graph of convex sets $\mathcal{G}_i = (\mathcal{V}_i, \mathcal{E}_i)$ tailored to the current best-response subproblem. Vertices represent collision-free convex regions at each time step, and directed edges represent dynamically feasible transitions between time-adjacent vertices.




Fig. 2: A set of vehicles driving on a highway with the red vehicle's safe gaps at time t in green.




Fig. 3: Time expanded graph.

4.1 Vertex Construction

Fix an IBR iteration and treat the other vehicles' trajectories as known. For each lane $\ell \in \mathcal{L}$ and time step $t \in \mathcal{T}$, define the unsafe longitudinal interval around each vehicle $j \neq i$ that is in lane ℓ at time step $t \in \mathcal{T}$:

$$\mathcal{K}_i^j(\ell, t) := (s_j(t) - d_i^s, s_j(t) + d_i^s). \quad (22)$$

Let the union of unsafe intervals in lane ℓ at time step $t \in \mathcal{T}$ be

$$\mathcal{F}_i(\ell, t) := \bigcup_{\substack{j \in \mathcal{N} \setminus \{i\} \\ z_j(t) = \ell}} \mathcal{K}_i^j(\ell, t). \quad (23)$$

Over a bounded road segment $[\underline{s}, \bar{s}]$, the collision-free set is

$$\mathcal{S}_i(\ell, t) := [\underline{s}, \bar{s}] \setminus \mathcal{F}_i(\ell, t). \quad (24)$$

Since $\mathcal{F}_i(\ell, t)$ is a finite union of open intervals, $\mathcal{S}_i(\ell, t)$ can be expressed as a finite union of disjoint closed intervals (safe gaps)

$$\mathcal{S}_i(\ell, t) = \bigcup_{g \in \mathcal{H}_i(t, \ell)} \mathcal{I}_i(t, \ell, g), \quad (25)$$

where $\mathcal{H}_i(t, \ell)$ is the finite index set of connected components of $\mathcal{S}_i(\ell, t)$ and $\{\mathcal{I}_i(t, \ell, g)\}_{g \in \mathcal{H}_i(t, \ell)}$ denotes the corresponding family of pairwise-disjoint closed intervals. Equivalently, each $g \in \mathcal{H}_i(t, \ell)$ labels one collision-free longitudinal “gap” in lane $\ell \in \mathcal{L}$ at time step $t \in \mathcal{T}$ between consecutive unsafe regions induced by other vehicles.

Each safe interval $\mathcal{I}_i(t, \ell, g)$, as illustrated in Fig. 2, induces a vertex $V_i(t, \ell, g) \in \mathcal{V}_i$ whose continuous state variable is

$$x_{V_i(t, \ell, g)} := \begin{bmatrix} s_{V_i(t, \ell, g)} \\ v_{V_i(t, \ell, g)} \end{bmatrix} \in \mathbb{R}^2, \quad (26)$$

and whose convex constraints are

$$s_{V_i(t, \ell, g)} \in \mathcal{I}_i(t, \ell, g), \quad v_{V_i(t, \ell, g)} \in [\underline{v}_i, \bar{v}_i]. \quad (27)$$

The discrete lane choice is encoded by the vertex lane index $\ell \in \mathcal{L}$.

Initial condition. To impose initial state conditions of vehicle i , choose any $g_0 \in \mathcal{H}_i(0, z_i(0))$ such that $s_i(0) \in \mathcal{I}_i(0, z_i(0), g_0)$ (ties may be broken arbitrarily), and enforce

$$x_{V_i(0, z_i(0), g_0)} = \begin{bmatrix} s_i(0) \\ v_i(0) \end{bmatrix}. \quad (28)$$

This fixes a valid source vertex consistent with the initial state.

4.2 Edge Construction

Directed edges connect time-adjacent vertices:

$$e = (V_i(t, \ell, g), V_i(t+1, \ell', g')), \quad t \in \mathcal{T}^-, \quad (29)$$

as shown in Fig. 3. Only edges between adjacent lanes are permitted, *i.e.*, $|\ell' - \ell| \leq 1$. Each edge enforces dynamic feasibility of the form (9) with a decision variable $a_i(t) \in [\underline{a}_i, \bar{a}_i]$ satisfying (11) and feasible state variables at times t and $t+1$. The lane-change action on the edge is $b_i(t) = \ell' - \ell \in \{-1, 0, 1\}$, consistent with (12). Lateral safety (Rule 2) is enforced by excluding lane-change edges that would violate (17) given the fixed trajectories and lane-change actions of other vehicles. Each edge $e = (V_i(t, \ell, g), V_i(t+1, \ell', g'))$ is assigned a convex quadratic edge cost $c_e(\cdot)$ that depends on the decision variables on that edge, namely the successor velocity $v_{V_i(t+1, \ell', g')}$ and the control $a_i(t)$ (with $b_i(t) = \ell' - \ell$):

$$c_e(x_{V_i(t, \ell, g)}, x_{V_i(t+1, \ell', g')}, a_i(t)) = w_{i,v} (v_{V_i(t+1, \ell', g')} - v_i^{\text{des}})^2 + w_{i,\ell} (\ell' - \ell_i^{\text{des}})^2 + w_{i,a} a_i(t)^2 + w_{i,b} b_i(t)^2, \quad (30)$$

where $w_{i,v}, w_{i,\ell}, w_{i,a}, w_{i,b} \in \mathbb{R}_{>0}$ are tunable weights. A terminal penalty may be included as a vertex cost at $t = T - 1$, e.g.,

$$c_V(x_{V_i(T-1,\ell,g)}) = w_{i,v} (v_{V_i(T-1,\ell,g)} - v_i^{\text{des}})^2. \quad (31)$$

Equipped with this construction, vehicle $i \in \mathcal{N}$'s single-vehicle planning problem (19) can be cast as a GCS shortest path problem on \mathcal{G}_i .

Proposition 2 (Single-vehicle optimality under tightness). *Fix the trajectories of all vehicles $j \neq i$ and construct \mathcal{G}_i accordingly. If the convex relaxation of the induced GCS shortest path problem is tight, i.e., it admits an optimal solution with $z_e \in \{0, 1\}$, then solving the relaxed problem yields a globally optimal solution of vehicle i 's mixed-integer best-response subproblem on \mathcal{G}_i . Consequently, under tightness, the resulting update is an exact best response, so each IBR step is a potential-descent step in the sense of Proposition 1.*

Remark 1. If the relaxation is not tight, the computed update can be interpreted as an approximate best response. Section 4.4 quantifies the effect of such inexactness.

4.3 Iterative Best-Response on Graphs of Convex Sets Algorithm

We solve the multi-vehicle game in Problem 1 using an IBR scheme, summarized in Algorithm 1.

Algorithm 1: Iterative Best-Response on Graphs of Convex Sets (IBR-GCS)

Input: Initial feasible strategies $\theta^0 = (\theta_1^0, \dots, \theta_N^0) \in \Theta$, maximum sweeps K_{\max} , tolerance ϵ
Output: Final strategy profile $\theta^{\text{out}} = (\theta_1^{\text{out}}, \dots, \theta_N^{\text{out}})$

- 1 Compute $\Phi(\theta^0)$;
- 2 **for** $k = 0$ **to** $K_{\max} - 1$ **do**
- 3 Set $\theta^{k,0} \leftarrow \theta^k$; // start of sweep
- 4 **for** $i = 1$ **to** N **do**
- 5 Construct vehicle i 's GCS \mathcal{G}_i^k using the other vehicles' strategies $\theta_{-i}^{k,i-1}$;
- 6 Compute a (possibly inexact) best response by solving the single-vehicle GCS problem on \mathcal{G}_i^k to obtain θ_i^{k+1} ;
- 7 Update the intermediate profile $\theta^{k,i} \leftarrow (\theta_1^{k+1}, \dots, \theta_i^{k+1}, \theta_{i+1}^k, \dots, \theta_N^k)$;
- 8 Set $\theta^{k+1} \leftarrow \theta^{k,N}$; // end of sweep
- 9 Compute potential change $\Delta \leftarrow |\Phi(\theta^{k+1}) - \Phi(\theta^k)|$;
- 10 **if** $\Delta < \epsilon$ **then**
- 11 **break**;
- 12 **return** $\theta^{\text{out}} \leftarrow \theta^{k+1}$;

The algorithm proceeds in *sweeps* indexed by k : during sweep k , vehicles update sequentially, and each vehicle i computes a best response to the most

recent strategies of the other vehicles, denoted $\theta_{-i}^{k,i-1}$. Concretely, at its update, vehicle i constructs a GCS instance \mathcal{G}_i^k that encodes collision-free regions and feasible transitions given $\theta_{-i}^{k,i-1}$, solves the resulting single-vehicle GCS SPP, and records the resulting strategy θ_i^{k+1} (lines 5–7 of Algorithm 1). After all vehicles have updated, the sweep terminates, and the algorithm produces the next joint strategy profile θ^{k+1} (line 8).

Under exact best-response updates in a generalized potential game (see Section 2.2), Proposition 1 implies that the generalized potential is monotonically nonincreasing along the intermediate profiles $\theta^{k,i}$, and hence along the sweep iterates θ^k . In our implementation, each best response is computed by solving a convex relaxation of the underlying GCS mixed-integer problem. Consequently, the update may be inexact and should be interpreted as an approximate best response. The resulting equilibrium error and its relationship to the per-update suboptimality are quantified in Section 4.4. We terminate IBR when the decrease in the potential across a sweep is below a specified tolerance (lines 8–11).

4.4 Error Quantification

This section quantifies the extent to which inexact best-response updates degrade the limiting quality of the strategy profiles produced by the IBR procedure. Since each vehicle solves a relaxed GCS problem in practice, the resulting update may not be an exact best response to the other vehicles' fixed strategies. We therefore measure the suboptimality of each update relative to the true best-response problem and translate this into an approximate equilibrium guarantee.

We recall the definitions of the joint feasible set Θ , the individual feasible sets $\Theta_i(\theta_{-i})$, and the intermediate profiles $\theta^{k,i}$ introduced in Equations (2) and (7). We note that for the remainder of this section, we revert to the cost function notation used prior to Assumption 1, *i.e.*, $J_i(\theta_i, \theta_{-i}) = J_i(\theta_i)$.

Approximate best responses. At each update, vehicle $i \in \mathcal{N}$ aims to minimize its cost over the feasible set induced by the current strategies of the other vehicles. We measure the quality of the computed update by comparing its cost to the optimal value of this best-response problem.

Definition 6 (Approximate best response). *At sweep k and vehicle $i \in \mathcal{N}$, we say the computed update θ_i^{k+1} has best-response error at most $\epsilon_i^k \geq 0$ if*

$$J_i(\theta_i^{k+1}, \theta_{-i}^{k,i-1}) \leq \inf_{\theta_i \in \Theta_i(\theta_{-i}^{k,i-1})} J_i(\theta_i, \theta_{-i}^{k,i-1}) + \epsilon_i^k. \quad (32)$$

Approximate equilibrium. A natural benchmark for IBR is a generalized Nash equilibrium (Definition 4), in which no vehicle can reduce its cost through a feasible unilateral deviation. Since we allow inexact best responses, we use the standard relaxation of this notion, *cf.* [21, Definition 3.4.7].

Definition 7 (ϵ -GNE). *A strategy profile $\theta \in \Theta$ is an ϵ -GNE if*

$$J_i(\theta_i, \theta_{-i}) \leq \inf_{\hat{\theta}_i \in \Theta_i(\theta_{-i})} J_i(\hat{\theta}_i, \theta_{-i}) + \epsilon, \quad \forall i \in \mathcal{N}. \quad (33)$$

To connect update-level errors to equilibrium quality, we introduce each vehicle's *regret*, *i.e.*, the gap between its current cost and its best feasible unilateral deviation.

$$r_i(\theta) := J_i(\theta_i, \theta_{-i}) - \inf_{\hat{\theta}_i \in \Theta_i(\theta_{-i})} J_i(\hat{\theta}_i, \theta_{-i}) \geq 0. \quad (34)$$

Our analysis is based on two standard technical assumptions. The first prevents the potential from decreasing without bound, while the second postulates that the inexactness of each best-response computation is uniformly bounded.

Assumption 3 *The potential function Φ is bounded below on Θ , and Θ is nonempty and compact.*

Assumption 4 *The IBR iterates satisfy (32) with errors uniformly bounded by $\epsilon_i^k \leq \bar{\epsilon} < \infty$ for all $i \in \mathcal{N}$ and all $k \geq 0$.*

Assumption 3 is standard in our setting, as physical constraints and finite discrete decisions ensure the feasible strategy space is a nonempty compact set over which the nonnegative potential function is bounded below. However, Assumption 4 is stronger and generally difficult to verify, requiring a uniform bound of the suboptimality of the GCS convex relaxation that is often observed in motion-planning problems (see Remark 3), but is difficult to guarantee theoretically (*cf.*, [15, Proposition 8.1]).

The following result states that if each IBR update is an ϵ -approximate best response with a uniform error bound, then the iterates asymptotically lie in an approximate equilibrium set: no vehicle can improve by more than $\bar{\epsilon}$ in the limit.

Theorem 1 ($\bar{\epsilon}$ -equilibrium under IBR-GCS). *Suppose G is a generalized potential game with potential Φ (Definition 5). Under Assumptions 3–4,*

$$\limsup_{k \rightarrow \infty} \max_{i \in \mathcal{N}} r_i(\theta^k) \leq \bar{\epsilon}.$$

In particular, the iterates asymptotically lie in the $\bar{\epsilon}$ -GNE set (Definition 7).

Proof. Fix a sweep index k and consider the update of vehicle $i \in \mathcal{N}$ from the intermediate profile $\theta^{k,i-1}$ to $\theta^{k,i}$. Since Φ is a generalized potential (Definition 5), changes in vehicle i 's cost match changes in Φ under feasible unilateral deviations. Combining this property with the approximate best-response condition (32) yields

$$\begin{aligned} \Phi(\theta^{k,i-1}) - \Phi(\theta^{k,i}) &= J_i(\theta_i^{k,i-1}, \theta_{-i}^{k,i-1}) - J_i(\theta_i^{k,i}, \theta_{-i}^{k,i-1}) \\ &\geq r_i(\theta^{k,i-1}) - \epsilon_i^k \geq r_i(\theta^{k,i-1}) - \bar{\epsilon}. \end{aligned} \quad (35)$$

We proceed by contradiction. Suppose there exists $\eta > 0$ and an infinite subsequence of sweeps $\{k_m\}_{m \in \mathbb{N}}$ such that, for each $m \in \mathbb{N}$, there exists an index $i_m \in \mathcal{N}$ with

$$r_{i_m}(\theta^{k_m, i_m - 1}) \geq \bar{\epsilon} + \eta. \quad (36)$$

Applying (35) at the update of vehicle i_m during sweep k_m gives

$$\Phi(\theta^{k_m, i_m-1}) - \Phi(\theta^{k_m, i_m}) \geq r_{i_m}(\theta^{k_m, i_m-1}) - \bar{\epsilon} \geq \eta. \quad (37)$$

In particular, since θ^{k_m, i_m} occurs within sweep k_m , we have $\Phi(\theta^{k_m+1}) \leq \Phi(\theta^{k_m}) - \eta$. Iterating along the subsequence implies $\Phi(\theta^{k_m}) \rightarrow -\infty$ as $m \rightarrow \infty$, contradicting the boundedness of Φ below on Θ (Assumption 3). Therefore,

$$\limsup_{k \rightarrow \infty} \max_{i \in \mathcal{N}} r_i(\theta^k) \leq \bar{\epsilon}, \quad (38)$$

which proves the claim. \square

Remark 2. As noted earlier, GCS relaxations are often tight in practice (*i.e.*, $\epsilon_i^k = 0$), but they can be arbitrarily loose in special cases [15, Proposition 8.1], in which case Assumption 4 may fail. For instance, this may happen when the costs and edge weights are symmetric [15, 17]. In such cases, the result above should be interpreted as a conditional guarantee: whenever the per-update best-response error remains uniformly bounded, the IBR iterates converge to a correspondingly bounded approximate equilibrium set.

5 Simulation Results

The performance of the proposed method is illustrated in a representative driving scenario. The parameters for the driving problem are $N = 6$, $|\mathcal{L}| = 4$, $T = 30$, and discretization step of $\Delta t = 0.3$ s. The scenario is set up with two vehicles merging from a terminating lane into a three-lane highway, where four other vehicles operate. The algorithm converges in 2 iterations, with the resulting equilibrium visualized in Fig. 4. All simulations were performed using the GCSOPT library [23] on a laptop with a AMD Ryzen™ 7 5800U (8-core, 16-thread) CPU and solved using the MOSEK 11.0.29 solver [1], leading to a total wall-clock time of 297.24 s for this scenario. The resulting strategies θ^* drive each vehicle $i \in \mathcal{N}$ safely over the chosen time horizon \mathcal{T} to a target lane ℓ_i^{des} , while tracking a desired velocity v_i^{des} . In doing so, the vehicles perform multi-layered overtaking maneuvers while accelerating and ensuring both lateral and longitudinal safety.




Fig. 4: Trajectories of six vehicles on a highway with one merging lane. The initial and final positions of each vehicle are shown without transparency, while positions at every third time step are shown with transparency.

Additionally, multiple randomized simulations with $N = 4$, $|\mathcal{L}| = 3$, $T = 30$, and $\Delta t = 0.3$ s were conducted, with parameters sampled from the following uniform distributions:

$$\begin{aligned} v_i^{\text{des}} &\sim \mathcal{U}\left(\frac{80}{3.6}, \frac{160}{3.6}\right), & \ell_i^{\text{des}} &\sim \mathcal{U}(\{1, \dots, |\mathcal{L}|\}), & w_{i,v} &\sim \mathcal{U}(0.1, 1.0), \\ w_{i,\ell} &\sim \mathcal{U}(5, 25), & w_{i,b} &\sim \mathcal{U}(5, 10), & w_{i,a} &\sim \mathcal{U}(0.1, 0.5). \end{aligned}$$

Initial states and lane configurations were randomly sampled according to:

$$s_i^0 \sim \mathcal{U}(0.0, 200), \quad v_i^0 \sim \mathcal{U}\left(\frac{60}{3.6}, \frac{130}{3.6}\right), \quad \ell_i^0 \sim \mathcal{U}(\{1, 2, 3\}),$$

until a collision-free arrangement was obtained. Figure 5 illustrates the evolution of the potential function $\Phi(\theta^k)$ over the iterations k . As predicted, the potential function values are decreasing, with a strictly positive decrease. Consequently, each iteration of the algorithm moves the system closer to a ϵ -GNE, providing empirical validation of the proposed approach. Our implementation of IBR-GCS, including an example, is available at https://github.com/nkaefer/IBR_GCS.




Fig. 5: Distributions of potential function evolution over iterations for 100 randomized simulation setups with the mean in orange. As shown, the potential function is always smaller at iteration k than at $k - 1$.

Remark 3. Our simulations consistently yielded tight convex relaxations, allowing us to find locally optimal solutions for complex driving scenarios. Although we do not provide theoretical guarantees for this tightness, our empirical results suggest the presence of additional underlying structure, motivating future work to establish formal proofs.

6 Conclusion

This paper presented IBR-GCS, a graphs of convex sets (GCS) formulation for multi-vehicle highway autonomous driving with strategic interaction. For each vehicle, collision-free maneuver options and dynamically feasible transitions are encoded as a shortest-path problem (SPP) on a vehicle-specific GCS. This representation separates continuous trajectory optimization (handled within convex

vertex/edge constraints and costs) from discrete maneuver selection (encoded in the graph topology) and admits a convex relaxation that is often tight. Since each vehicle’s feasible maneuver graph depends on the current strategies of the other vehicles, a single centralized GCS is not naturally available; instead, IBR-GCS employs an iterative best-response (IBR) procedure in which vehicles repeatedly solve their individual GCS subproblems while holding the strategies of the other vehicles fixed.

From a game-theoretic perspective, the coupled feasibility constraints induce a generalized Nash game with a generalized potential structure that can be exploited to interpret IBR as an (approximate) descent method on a global potential function. Under exact best responses, the potential is monotonically non-increasing (Proposition 1). When best responses are computed inexactly (*e.g.*, due to relaxation looseness), our error analysis provides conditions under which the iterates converge to a neighborhood of the generalized Nash equilibrium set, with a bound on the maximum unilateral improvement available to any vehicle (Theorem 1).

Several practical considerations affect performance and the particular equilibrium reached. For example, the limiting equilibrium point can depend on initialization and on the vehicle update order. In the current implementation, each vehicle reconstructs its GCS at every update, even though successive graphs often change only locally as other vehicles move. An important direction for future work is to warm-start and reuse graph structure across sweeps (*e.g.*, caching safe-gap decompositions and updating only affected vertices/edges), which could substantially reduce per-iteration overhead. Another direction is to incorporate randomized or priority-based update schedules to mitigate order sensitivity, and to refine termination criteria based on per-vehicle regret (Section 4.4) rather than only on potential decrease.

Finally, while convex relaxations are frequently tight in practice, relaxation looseness can degrade best-response accuracy. Developing problem-specific certificates of tightness for driving instances remains an important avenue for future work.

Acknowledgement

This work was supported in part by the National Science Foundation under Grants 2211548 and 2336840.

References

1. ApS, M.: The MOSEK Python Fusion API manual. Version 11.0. (2025), <https://docs.mosek.com/latest/pythonfusion/index.html>
2. Bakolas, E., Lee, Y.: Decentralized Game-Theoretic Control for Dynamic Task Allocation Problems for Multi-Agent Systems. In: 2021 American Control Conference (ACC). pp. 3228–3233 (2021). <https://doi.org/10.23919/ACC50511.2021.9483030>
3. Bazaraa, M., Jarvis, J., Sherali, H.: Linear Programming and Network Flows. Wiley (2009)

4. Bhatt, M., Jia, Y., Mehr, N.: Strategic Decision-Making in Multiagent Domains: A Weighted Constrained Potential Dynamic Game Approach. *IEEE Transactions on Robotics* **41**, 2749–2764 (2025). <https://doi.org/10.1109/tro.2025.3552325>, <http://dx.doi.org/10.1109/TRO.2025.3552325>
5. Deits, R., Tedrake, R.: Efficient mixed-integer planning for UAVs in cluttered environments. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 42–49 (2015). <https://doi.org/10.1109/ICRA.2015.7138978>
6. Earl, M.G., D’Andrea, R.: Modeling and Control of a Multi-Agent System Using Mixed Integer Linear Programming. In: Proceedings of the IEEE Conference on Decision and Control (CDC). vol. 1, pp. 107–111. Las Vegas, NV, USA (2002). <https://doi.org/10.1109/CDC.2002.1184476>
7. Esterle, K., Hart, P., Bernhard, J., Knoll, A.: Spatiotemporal Motion Planning with Combinatorial Reasoning for Autonomous Driving. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC). pp. 1053–1060 (2018). <https://doi.org/10.1109/ITSC.2018.8570003>
8. Fabiani, F., Grammatico, S.: Multi-Vehicle Automated Driving as a Generalized Mixed-Integer Potential Game. *IEEE Transactions on Intelligent Transportation Systems* **21**(3), 1064–1073 (2020). <https://doi.org/10.1109/TITS.2019.2901505>
9. Hang, P., Lv, C., Xing, Y., Huang, C., Hu, Z.: Human-Like Decision Making for Autonomous Driving: A Noncooperative Game Theoretic Approach. *IEEE Transactions on Intelligent Transportation Systems* **22**(4), 2076–2087 (2021). <https://doi.org/10.1109/TITS.2020.3036984>
10. Kavuncu, T., Yaraneri, A., Mehr, N.: Potential iLQR: A Potential-Minimizing Controller for Planning Multi-Agent Interactive Trajectories. In: Robotics. Robotics: Science and Systems, Massachusetts Institute of Technology, United States (2021). <https://doi.org/10.15607/RSS.2021.XVII.084>
11. Kurtz, V., Lin, H.: Temporal Logic Motion Planning With Convex Optimization via Graphs of Convex Sets. *IEEE Transactions on Robotics* **39**(5), 3791–3804 (2023). <https://doi.org/10.1109/TRO.2023.3291463>
12. Lee, Y., Bakolas, E.: Relay Pursuit of an Evader by a Heterogeneous Group of Pursuers using Potential Games. In: 2021 American Control Conference (ACC). pp. 3182–3187 (2021). <https://doi.org/10.23919/ACC50511.2021.9482912>
13. Liu, M., Kolmanovsky, I., Tseng, H.E., Huang, S., Filev, D., Girard, A.: Potential Game-Based Decision-Making for Autonomous Driving. *IEEE Transactions on Intelligent Transportation Systems* **24**(8), 8014–8027 (2023). <https://doi.org/10.1109/TITS.2023.3264665>
14. Lopez, V.G., Lewis, F.L., Liu, M., Wan, Y., Nagesh Rao, S., Filev, D.: Game-Theoretic Lane-Changing Decision Making and Payoff Learning for Autonomous Vehicles. *IEEE Transactions on Vehicular Technology* **71**(4), 3609–3620 (2022). <https://doi.org/10.1109/TVT.2022.3148972>
15. Marcucci, T.: Graphs of Convex Sets with Applications to Optimal Control and Motion Planning. Massachusetts Institute of Technology (2024)
16. Marcucci, T., Petersen, M., von Wrangel, D., Tedrake, R.: Motion planning around obstacles with convex optimization. *Science Robotics* **8**(84), eadf7843 (2023). <https://doi.org/10.1126/scirobotics.adf7843>, <https://www.science.org/doi/abs/10.1126/scirobotics.adf7843>
17. Marcucci, T., Umenberger, J., Parrilo, P., Tedrake, R.: Shortest Paths in Graphs of Convex Sets. *SIAM Journal on Optimization* **34**(1), 507–532 (2024). <https://doi.org/10.1137/22M1523790>, <https://doi.org/10.1137/22M1523790>
18. Monderer, D., Shapley, L.S.: Potential Games. *Games and Economic Behavior* **14**(1), 124–143 (1996). <https://doi.org/https://doi.org/10.1006/game.1996.0044>, <https://www.sciencedirect.com/science/article/pii/S0899825696900445>

19. Paden, B., Cáp, M., Yong, S.Z., Yershov, D.S., Frazzoli, E.: A Survey of Motion Planning and Control Techniques for Self-driving Urban Vehicles. *CoRR abs/1604.07446* (2016), <http://arxiv.org/abs/1604.07446>
20. Richards, A., How, J.P.: Aircraft Trajectory Planning with Collision Avoidance Using Mixed Integer Linear Programming. In: Proceedings of the American Control Conference (ACC). vol. 3, pp. 1936–1941. Anchorage, AK, USA (2002). <https://doi.org/10.1109/ACC.2002.1023918>
21. Shoham, Y., Leyton-Brown, K.: Multiagent systems: Algorithmic, game-theoretic, and logical foundations. Cambridge University Press (2008)
22. Tang, J., Mao, Z., Yang, L., Ma, H.: Space-Time Graphs of Convex Sets for Multi-Robot Motion Planning. arXiv preprint arXiv:2503.00583 (2025)
23. Tobia Marcucci: gcsopt: Graphs of Convex Sets optimization library. <https://github.com/TobiaMarcucci/gcsopt> (2026), gitHub repository, accessed March 2026
24. Vicini, M., Albut, S., Gindullina, E., Badia, L.: Decision Making via Game Theory for Autonomous Vehicles in the Presence of a Moving Obstacle. 2022 IEEE International Conference on Communication, Networks and Satellite (COMNETSAT) pp. 393–398 (2022), <https://api.semanticscholar.org/CorpusID:255417477>
25. 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). <https://doi.org/10.1109/TRO.2020.3047521>
26. Wang, Z., Taubner, T., Schwager, M.: Multi-agent sensitivity enhanced iterative best response: A real-time game theoretic planner for drone racing in 3D environments. *Robotics and Autonomous Systems* **125**, 103410 (2020). <https://doi.org/https://doi.org/10.1016/j.robot.2019.103410>, <https://www.sciencedirect.com/science/article/pii/S0921889019301939>
27. Williams, Z., Chen, J., Mehr, N.: Distributed Potential iLQR: Scalable Game-Theoretic Trajectory Planning for Multi-Agent Interactions. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 01–07 (2023). <https://doi.org/10.1109/ICRA48891.2023.10161176>
28. Yan, Y., Peng, L., Shen, T., Wang, J., Pi, D., Cao, D., Yin, G.: A Multi-Vehicle Game-Theoretic Framework for Decision Making and Planning of Autonomous Vehicles in Mixed Traffic. *IEEE Transactions on Intelligent Vehicles* **8**(11), 4572–4587 (2023). <https://doi.org/10.1109/TIV.2023.3321346>
29. Zanardi, A., Mion, E., Bruschetta, M., Bolognani, S., Censi, A., Frazzoli, E.: Urban Driving Games With Lexicographic Preferences and Socially Efficient Nash Equilibria. *IEEE Robotics and Automation Letters* **6**(3), 4978–4985 (2021). <https://doi.org/10.1109/LRA.2021.3068657>