DARTMOUTH ENGINEERING

Misspecification in Games as Models
The Game-to-Real Gap and Game Theoretic MPC

Bryce L. Ferguson

Assistant Professor

Thayer School of Engineering at Dartmouth College

Game Theory Meets MPC: Advances in Multi-Agent Control Workshop at CDC

December 9th, 2025

A Motivating Example: Autonomous Racing

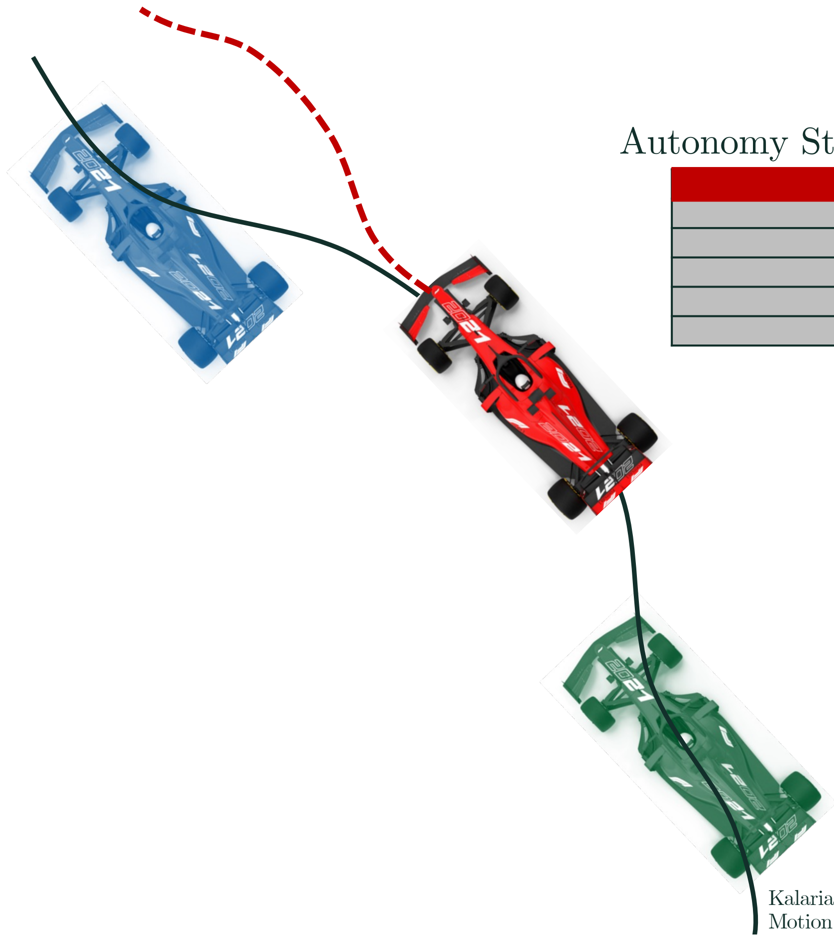

Autonomy Stack

Follow the optimal race line

Intervehicle strategy
- Passing
- Blocking
- Using limited fuel

Approach:

1. Design rewards
   $r_i(u, \theta) =$ dist ahead or behind
2. Cast as a large Markov Game
   $J_i(\pi, s_0) = \mathbb{E}_{u^t \sim \pi(s^t)} \left[\sum_{t=0}^{\infty} \gamma^t r_i(u^t, s^t)\right]$
3. Solve for a Nash equilibrium joint policy
   π^{NE} via Actor-Critic MARL
4. Take ego policy and implement in the autonomy stack
   $\pi^{\mathrm{NE}} = (\pi_1^{\mathrm{NE}}, \ldots, \pi_{ego}^{\mathrm{NE}}, \ldots, \pi_n^{\mathrm{NE}})$

Kalaria, Dvij, Chinmay Maheshwari, and Shankar Sastry. "{\alpha}-RACER: Real-Time Algorithm for Game-Theoretic Motion Planning and Control in Autonomous Racing using Near-Potential Function." *arXiv preprint arXiv:2412.08855*(2024).

A Motivating Example: Autonomous Racing

Simulated Multi-Agent Interaction

Realized Multi-Agent Interaction

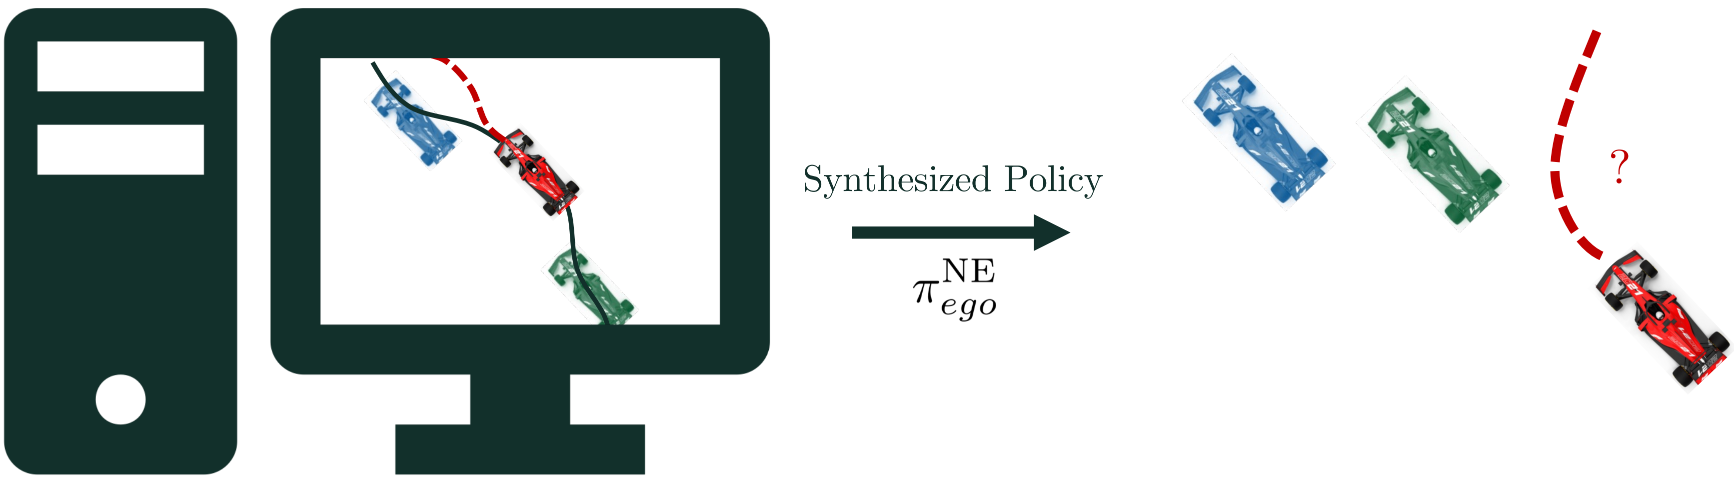

Difference between ***simulated*** and ***realized*** interaction can cause degraded performance of policy

How do we quantify and overcome this ***'Strategic' sim-to-real gap?***

Realizations of Game Theoretic Planning and Prediction

Dynamic Zero-Sum Games

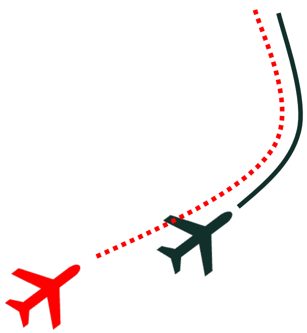

Iterative LQ

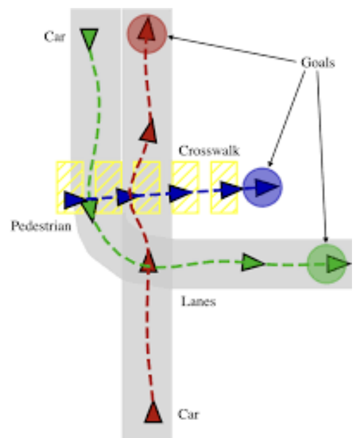

Model Predictive Games

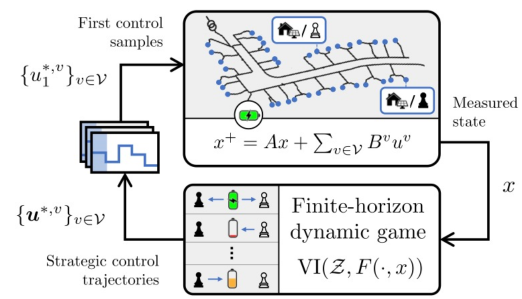

1. J. Sprinkle, J. M. Eklund, H. J. Kim and S. Sastry, "Encoding aerial pursuit/evasion games with fixed wing aircraft into a nonlinear model predictive tracking controller," 2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601), Nassau, Bahamas, 2004, pp. 2609-2614 Vol.3, doi: 10.1109/CDC.2004.1428851.
2. Başar, Tamer, ed. *Dynamic games and applications in economics*. Vol. 265. Springer Science & Business Media, 1986.
3. H. Chen, C. W. Scherer and F. Allgower, "A game theoretic approach to nonlinear robust receding horizon control of constrained systems," *Proceedings of the 1997 American Control Conference (Cat. No.97CH36041)*, Albuquerque, NM, USA, 1997, pp. 3073-3077 vol.5, doi: 10.1109/ACC.1997.612023.
4. D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan and C. J. Tomlin, "Efficient Iterative Linear-Quadratic Approximations for Nonlinear Multi-Player General-Sum Differential Games," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, Paris, France, 2020
5. S. Hall, G. Belgioioso, D. Liao-McPherson and F. Dorfler, "Receding Horizon Games with Coupling Constraints for Demand-Side Management," 2022 IEEE 61st Conference on Decision and Control (CDC), Cancun, Mexico, 2022, pp. 3795-3800, doi: 10.1109/CDC51059.2022.9992497.

Games as Models

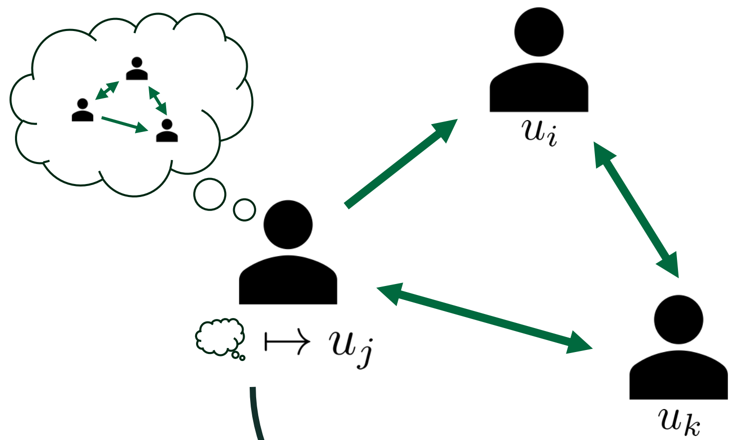

- Each agent possesses a model for how the collective group will interact
- Use conjectures about other players' intents and capabilities to predict the outcome and your role in it
- Opt for models that possess efficiently computed solutions

<u>Games and Prediction</u>

$$G = \langle N, \{\mathcal{U}_i\}_{i \in N}, \{J_i\}_{i \in N}, \texttt{Soln} \rangle$$

Players — Control/Action set — Objective — Solution Concept

$$\texttt{Soln} : \{J : \mathcal{U} \to \mathbb{R}\}^n \to 2^{\mathcal{U}} \qquad \texttt{Soln}\left(\{J_i\}_{i \in N}\right) \subseteq \mathcal{U}$$

Predictive Capabilities of Games

$$G^{(j)} = \langle N, \{\mathcal{U}_i^{(j)}\}_{i \in N}, \{J_i^{(j)}\}_{i \in N}, \texttt{Soln}^{(j)} \rangle$$

Conjectured game of player j

Player j predicted behavior: $u^{(j)} \in \texttt{Soln}^{(j)}(J^{(j)})$

Predictive Capabilities of Games

$$G^{(j)} = \langle N, \{\mathcal{U}_i^{(j)}\}_{i \in N}, \{J_i^{(j)}\}_{i \in N}, \texttt{Soln}^{(j)} \rangle$$

Conjectured game of player j

Player j predicted behavior: $u^{(j)} \in \texttt{Soln}^{(j)}(J^{(j)})$

Each player predicts behavior individually

$G^{(1)} \quad G^{(2)} \quad \cdots \quad G^{(j)} \quad \cdots \quad G^{(n)}$

Players act in response to their own predictions

Realized Behavior: $u^{\circ} = (u_1^{(1)}, u_2^{(2)}, \ldots, u_j^{(j)}, \ldots, u_n^{(n)})$

- Need not be an equilibrium/solution
- Gap between simulated and real behavior
- Exaggerated by greater mischaracterizations

Misspecified Game-theoretic Planning

Players conjecture different games

$$G^{(i)} \neq G^{(j)}$$

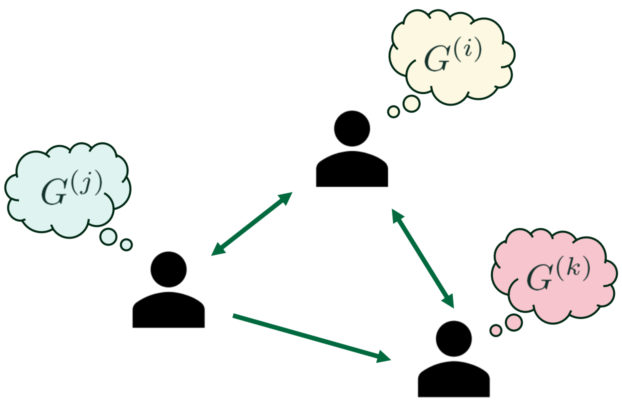

Misconjectured cost functions $J_i \neq J_i^{(j)}$

Misconjectured constraints $\mathcal{U}_i \neq \mathcal{U}_i^{(j)}$

Misaligned solution concepts $\texttt{Soln}^{(i)} \neq \texttt{Soln}^{(j)}$

Multiplicity of equilibria/predictions $|\texttt{Soln}(G)| > 1$

Game-to-Real Gap

$$J_i\big(u^{\circ}\big) - J_i\big(u^{(i)}\big)$$

Quantify magnitude

Dynamics with misspecification

Sampling, learning, and overcoming

Related Perspectives

- Sim-to-real gap in RL
- Sensitivity of equilibria in games
- Bayesian Games
- ‘Robust’ game theory
- Strategically robust equilibria
- K-level rationality
- Inverse learning in games

<u>Our Perspective:</u>
Adopt classic, computationally friendly games and equilibria as interaction models, even though they may not be entirely accurate (as in practice, none will)

How much do we lose from our model being incorrect, and how do we overcome?

Recent submission: *Game-to-Real Gap: Quantifying the Effect of Model Misspecification in Network Games*

- Focus on well-studied case of Nash equilibria in strongly monotone network games
- Characterize and bound the game-to-real gap conditioned on game-defining parameters
- Identify forecast and network specific properties

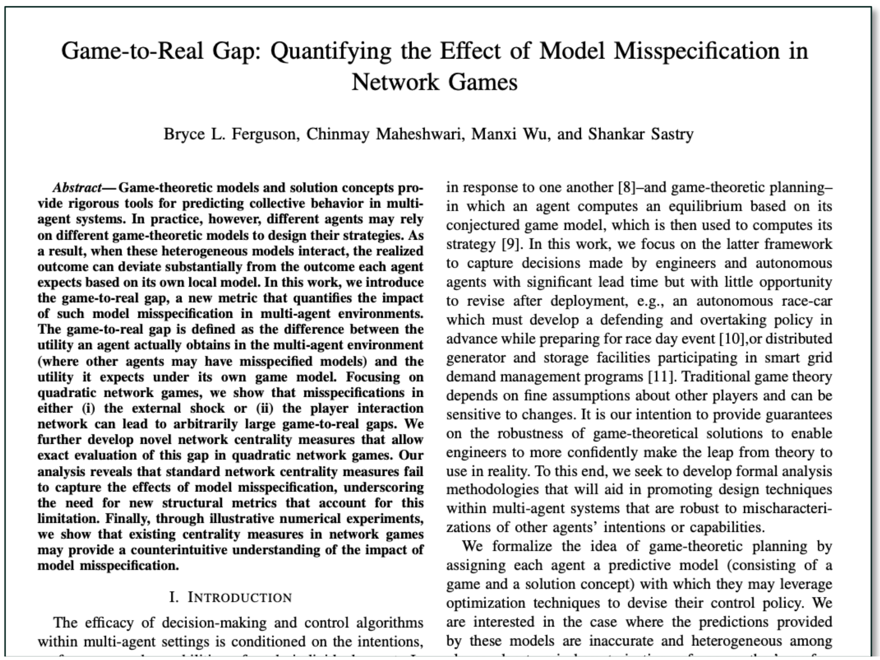

Game-to-Real Gap: Quantifying the Effect of Model Misspecification in Network Games

Bryce L. Ferguson, Chinmay Maheshwari, Manxi Wu, and Shankar Sastry

***Abstract*— Game-theoretic models and solution concepts provide rigorous tools for predicting collective behavior in multi-agent systems. In practice, however, different agents may rely on different game-theoretic models to design their strategies. As a result, when these heterogeneous models interact, the realized outcome can deviate substantially from the outcome each agent expects based on its own local model. In this work, we introduce the game-to-real gap, a new metric that quantifies the impact of such model misspecification in multi-agent environments. The game-to-real gap is defined as the difference between the utility an agent actually obtains in the multi-agent environment (where other agents may have misspecified models) and the utility it expects under its own game model. Focusing on quadratic network games, we show that misspecifications in either (i) the external shock or (ii) the player interaction network can lead to arbitrarily large game-to-real gaps. We further develop novel network centrality measures that allow exact evaluation of this gap in quadratic network games. Our analysis reveals that standard network centrality measures fail to capture the effects of model misspecification, underscoring the need for new structural metrics that account for this limitation. Finally, through illustrative numerical experiments, we show that existing centrality measures in network games may provide a counterintuitive understanding of the impact of model misspecification.**

I. Introduction

The efficacy of decision-making and control algorithms within multi-agent settings is conditioned on the intentions,

in response to one another [8]–and game-theoretic planning–in which an agent computes an equilibrium based on its conjectured game model, which is then used to computes its strategy [9]. In this work, we focus on the latter framework to capture decisions made by engineers and autonomous agents with significant lead time but with little opportunity to revise after deployment, e.g., an autonomous race-car which must develop a defending and overtaking policy in advance while preparing for race day event [10],or distributed generator and storage facilities participating in smart grid demand management programs [11]. Traditional game theory depends on fine assumptions about other players and can be sensitive to changes. It is our intention to provide guarantees on the robustness of game-theoretical solutions to enable engineers to more confidently make the leap from theory to use in reality. To this end, we seek to develop formal analysis methodologies that will aid in promoting design techniques within multi-agent systems that are robust to mischaracterizations of other agents' intentions or capabilities.

We formalize the idea of game-theoretic planning by assigning each agent a predictive model (consisting of a game and a solution concept) with which they may leverage optimization techniques to devise their control policy. We are interested in the case where the predictions provided by these models are inaccurate and heterogeneous among

Joint work with...

Chinmay Maheshwari
Assistant Professor
Johns Hopkins University

Manxi Wu
Assistant Professor
UC Berkeley

Shankar Sastry
Professor
UC Berkeley

Network Games with Misspecification

$u_i \in \mathbb{R}^{m_i}, \ \forall i \in N$

$J_i{}^{(j)}(u) = \frac{1}{2} u_i^\top u_i - u_i^\top \left(P_{i,-}^{(j)} u + \epsilon_i{}^{(j)} \right)$

$u^{(j)} \in \texttt{Soln}(G^{(j)})$ if it is a Nash equilibrium, i.e.,
$J_i^{(j)}(u) \leq J_i^{(j)}(u_i', u_{-i}), \ \forall u_i \in \mathcal{U}^{(j)}, \ i \in N$

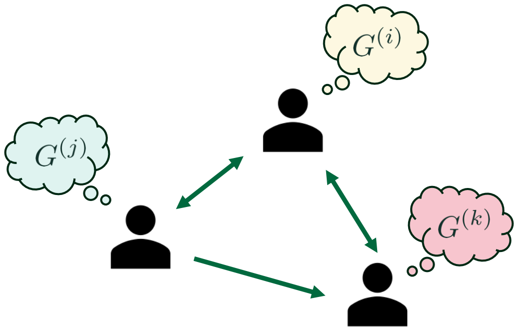

Forecast Misspecification

$\epsilon^{(i)} \neq \epsilon^{(j)}$

Interaction Misspecification

$P^{(i)} \neq P^{(j)}$

Network Games with Misspecification

Proposition 1: In quadratic network aggregative games, for any $\delta, M > 0$, there exits some $\{\epsilon^{(i)}\}_{i \in N}$ and interaction matrix P such that

$$\|\epsilon^{(i)} - \epsilon^{(j)}\|_2 \leq \delta \quad \forall\, i, j \in N,$$

but that for each $i \in N$,

$$J_i(u^\circ) - J_i(u^{(i)}) > M.$$

Definition 1: The *Shock Misspecification Centrality* between player i and j is denoted by

$$\mathcal{B}_{i,j} = [(I-P)^{-1}]_{i,-}^\top P_{i,j} [(I-P)^{-1}]_{j,-} \in \mathbb{R}^{m \times m}.$$

Theorem 2 (abrv.): The game-to-real gap is characterized by

$$J_i(u^\circ) - J_i(u^{(i)}) = \sum_{j \neq i} \epsilon^{(i)\top} \mathcal{B}_{i,j} (\epsilon^{(i)} - \epsilon^{(j)}).$$

Graphs centrality measures are insufficient for characterizing game-to-real gap, need pair-wise centrality

$$P = \begin{bmatrix} 0 & 0 & \gamma \\ \gamma & 0 & 0 \\ 0 & \gamma & 0 \end{bmatrix},$$

$$\epsilon^{(1)} = \begin{bmatrix} 1 \\ \beta \\ \beta \end{bmatrix},\ \epsilon^{(2)} = \begin{bmatrix} \beta \\ 1 \\ \beta \end{bmatrix},\ \epsilon^{(3)} = \begin{bmatrix} \beta \\ \beta \\ 1 \end{bmatrix}$$

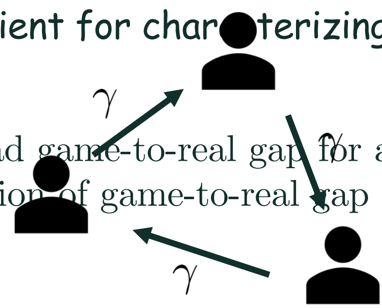

$$\lim_{\gamma \to 1} J_i(u^\circ) - J_i(u^{(i)}) = \infty$$

Observation: embedded cycles cause a larger game-to-real gap

Interaction Misspecification:

- Proposition 3: arbitrarily bad game-to-real gap for any bounded norm of graph difference
- Proposition 4: characterization of game-to-real gap by weighted pair-wise centrality measure

Simultaneous Misspecification:

- Corollary 1: characterization of game-to-real gap

Ongoing work: more general game-to-real gap by parametric VI analysis, apply to LQ games, network routing games, and zero-sum games

Designing Around Misspecification

- Interaction Sample complexity

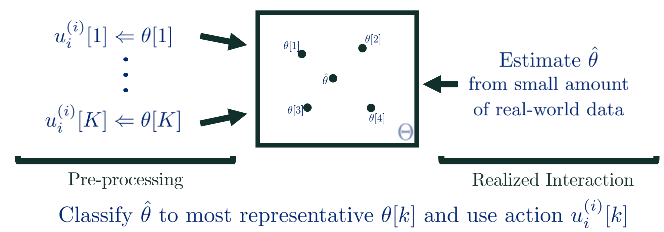

- Robust model optimization

$$\underset{\theta^{(i)} \in \Theta}{\text{minimize}} \quad \max_{\theta^{(-i)} \sim \Theta_{-i}} \left[J_i \left(u^\circ \left(\theta^{(i)}, \theta^{(-i)} \right) \right) \right]$$

$$\underset{\theta^{(i)} \in \Theta}{\text{minimize}} \quad \mathbb{E}_{\theta^{(-i)} \sim \mu_i} \left[J_i \left(u^\circ \left(\theta^{(i)}, \theta^{(-i)} \right) \right) \right]$$

- Online planning
  - Model predictive games

Ada Yildirim
PhD Student
Dartmouth College

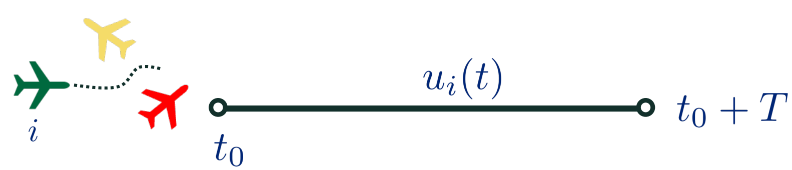

A Motivating Example: Autonomous Fleet Drifting

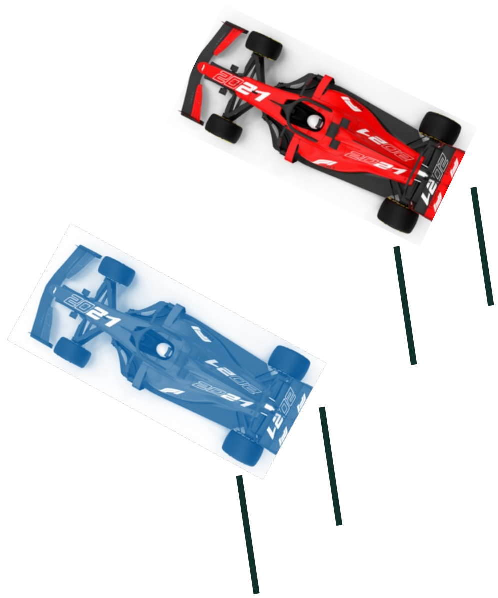

Single MPC controller deployed on each vehicle

Finely tuned based on vehicle characteristics and state

Over time, characteristics of each vehicle change,
these changes are only observed locally

Cars are now solving different optimization problems

If agents start optimizing different objective functions, how does MPC controller fare?

Can we design a controller where agents consider each others' objectives?

What if the models of each others' objectives is not completely accurate?

Model Predictive Control (MPC)

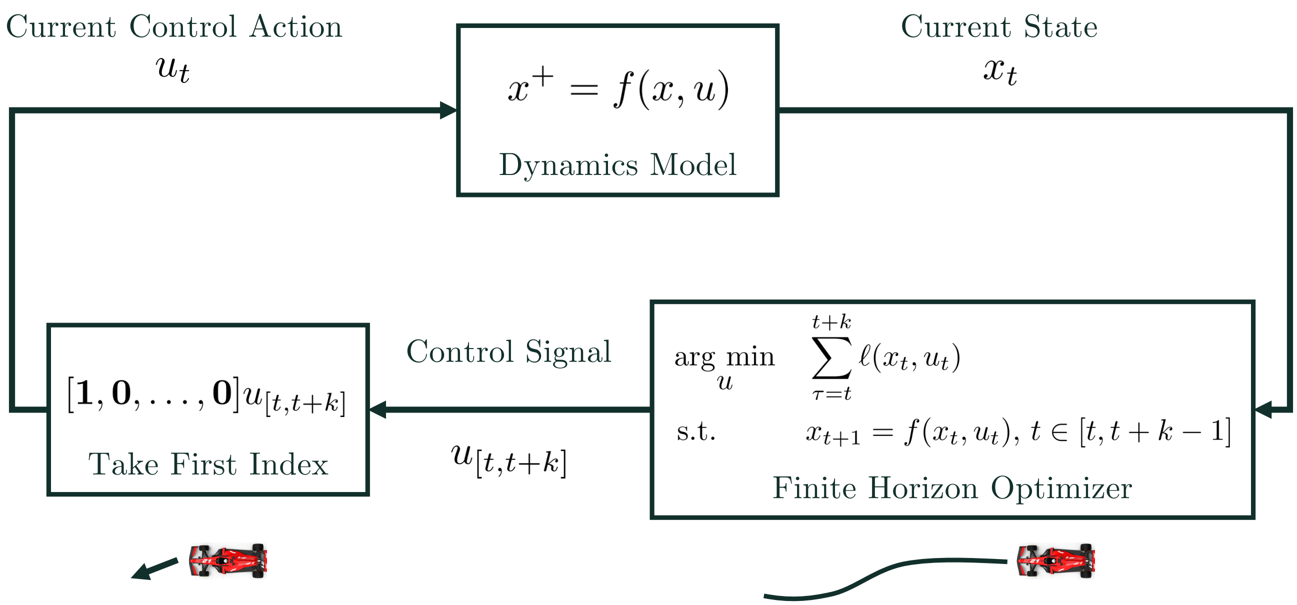

Solve for optimal control signal over finite horizon and deploy first timestep

Game-Theoretic MPC

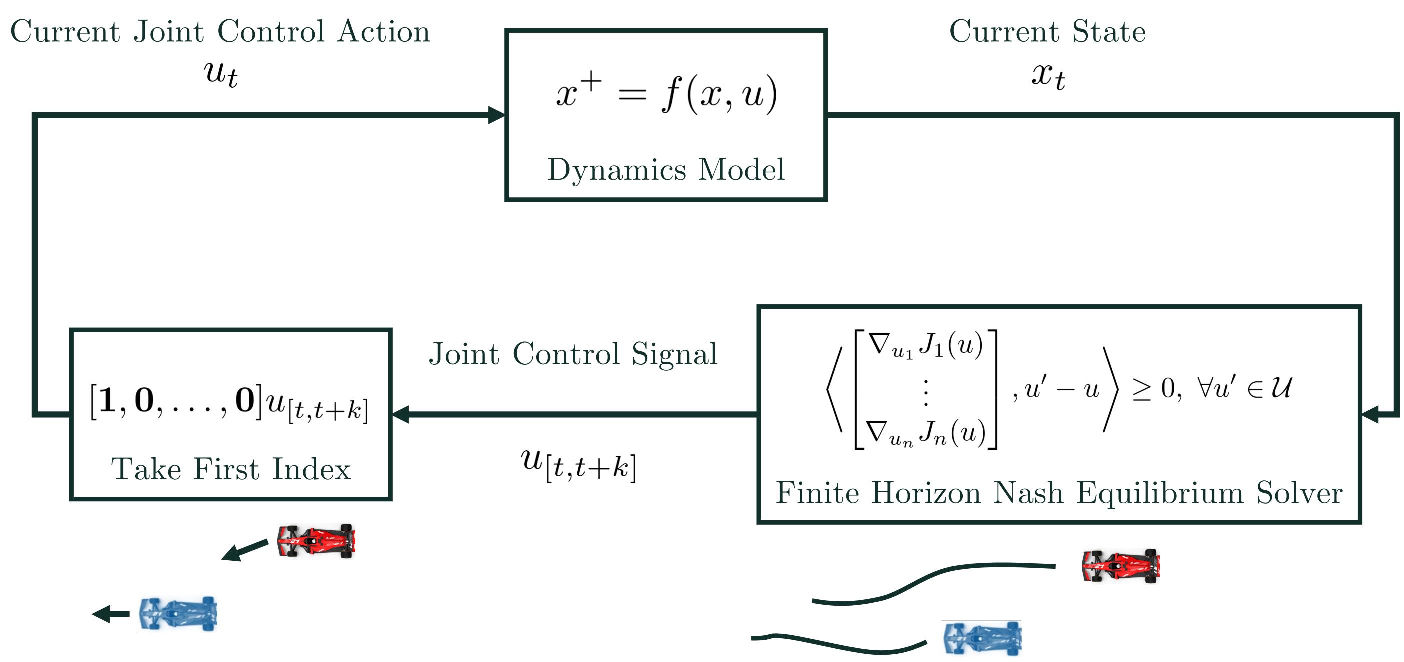

Solve for ***Nash Equilibrium*** control signal over finite horizon and deploy first timestep

Hall, Sophie, Dominic Liao-McPherson, Giuseppe Belgioioso, and Florian Dörfler. "Stability Certificates for Receding Horizon Games." *arXiv preprint arXiv:2404.12165* (2024).

Misspecificed Game-Theoretic MPC

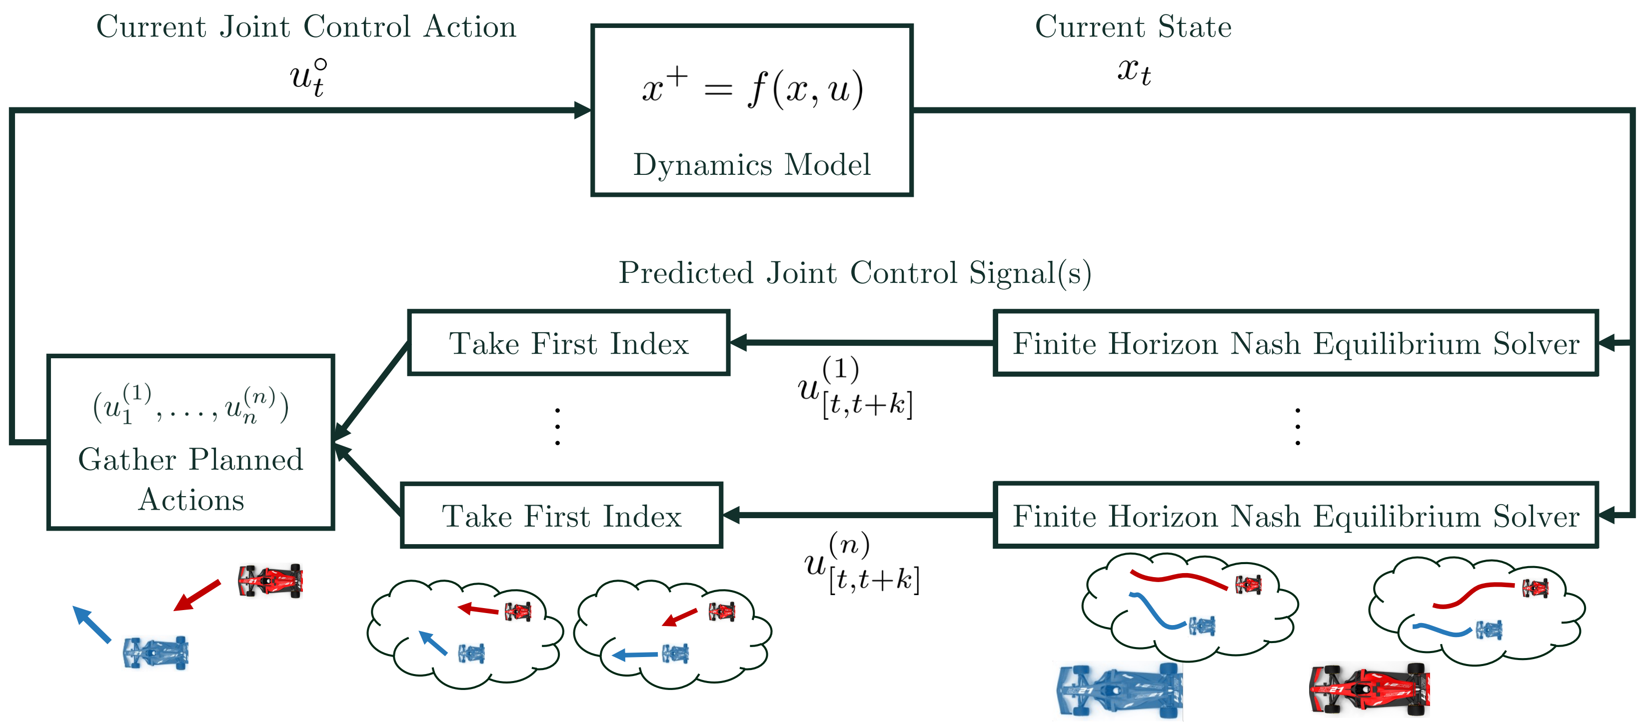

Each player solves for a ***Nash Equilibrium*** control signal over finite horizon and deploys ***respective*** first timestep

Stability Criteria with Misspecifications (n=2)

$$\begin{bmatrix} A^\top PA - P & A^\top P\hat{B}_1 & A^\top P\hat{B}_2 \\ \star & \hat{B}_1^\top P\hat{B}_1 & \hat{B}_1^\top P\hat{B}_2 \\ \star & \star & \hat{B}_2^\top P\hat{B}_2 \end{bmatrix} + \begin{bmatrix} \frac{\lambda_2}{\mu_1^2}F_{1,x}^\top F_{1,x} & -\frac{\lambda_1}{2}F_{1,x}^\top & 0 \\ \star & -(\lambda_1\mu_1+\lambda_2)I & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} \frac{\lambda_4}{\mu_2^2}F_{2,x}^\top F_{2,x} & 0 & -\frac{\lambda_3}{2}F_{2,x}^\top \\ 0 & 0 & 0 \\ \star & 0 & -(\lambda_3\mu_2+\lambda_4)I \end{bmatrix} \preceq -\epsilon I \quad \text{(A)}$$

Result 1. *If each player's game is strongly monotone and has a convex joint action space; the open loop dynamics are stable, and there exists a positive-definite matrix P and non-negative scalars $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ (where $\lambda_1+\lambda_2>0, \lambda_3+\lambda_4>0$) that make (A) true for some $\epsilon>0$, then the following conditions hold:*

(i) *there exists a globally asymptotically stable equilibrium point, $\bar{x}\in\mathbb{R}^n$, of the closed-loop system;*

(ii) *the optimal control problem (OCP) is recursively feasible, $\forall\,\mathbf{x}\in\mathbb{R}^{n_x}$;*

(iii) *the control inputs satisfy constraints for all times, $u_t\in\mathcal{Z}\ \forall t$.*

Stability criteria depends on a non-trivial combination of each player's game

Dynamics with Misspecifications

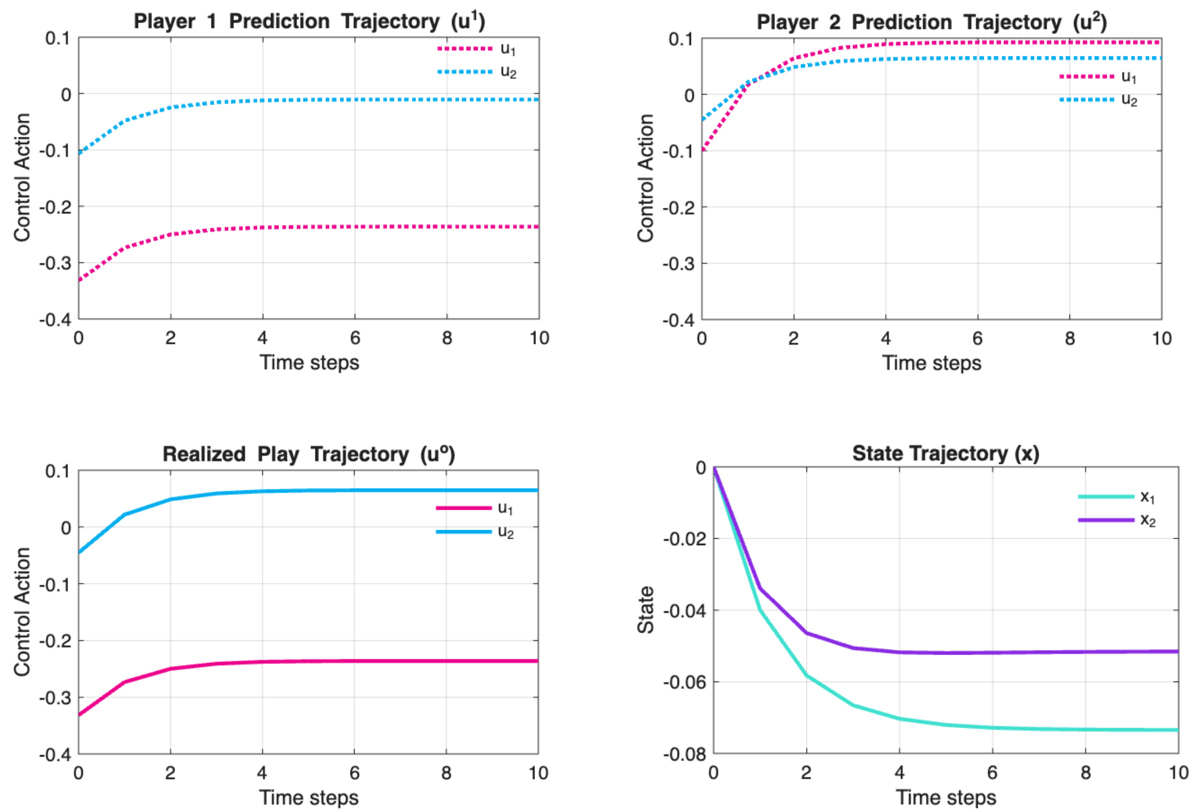

Equilibria with Misspecifications

How does the equilibrium of misspecified GT MPC change with greater misalignment?

$$G^{(1)} = G^A \qquad G^{(2)} = \lambda G^A + (1-\lambda) G^B$$

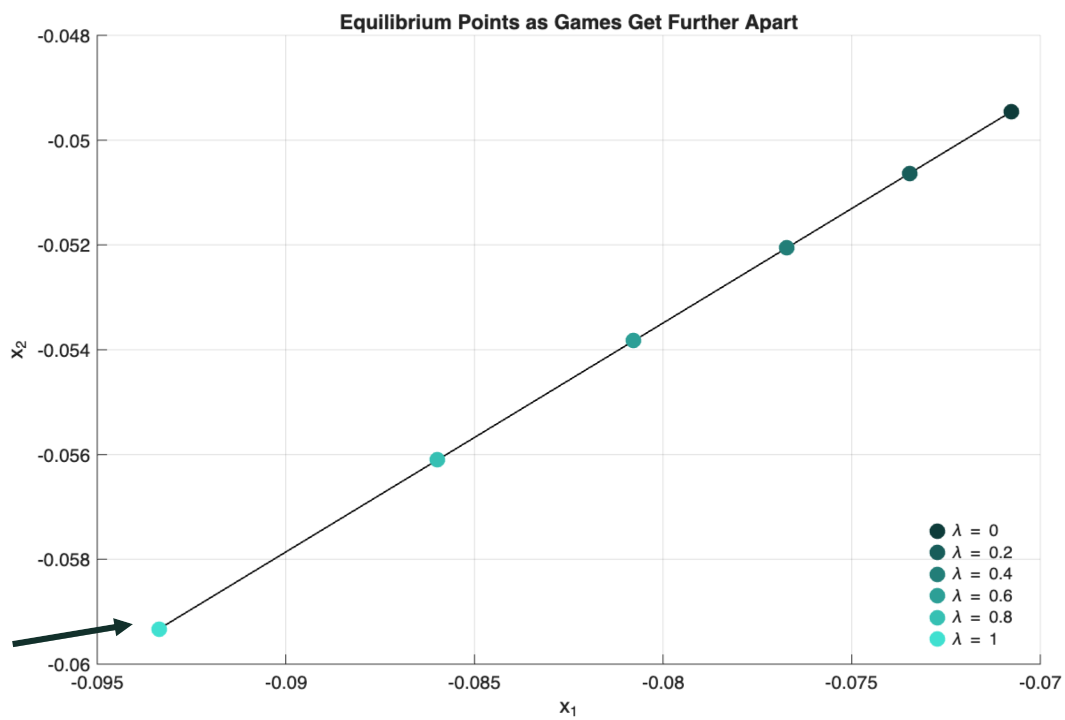

No misspecification

Equilibria with Misspecifications

How does the equilibrium of misspecified GT MPC change with greater misalignment?

$$G^{(1)} = G^{A} \qquad\qquad G^{(2)} = \lambda G^{A} + (1-\lambda) G^{B}$$

No misspecification

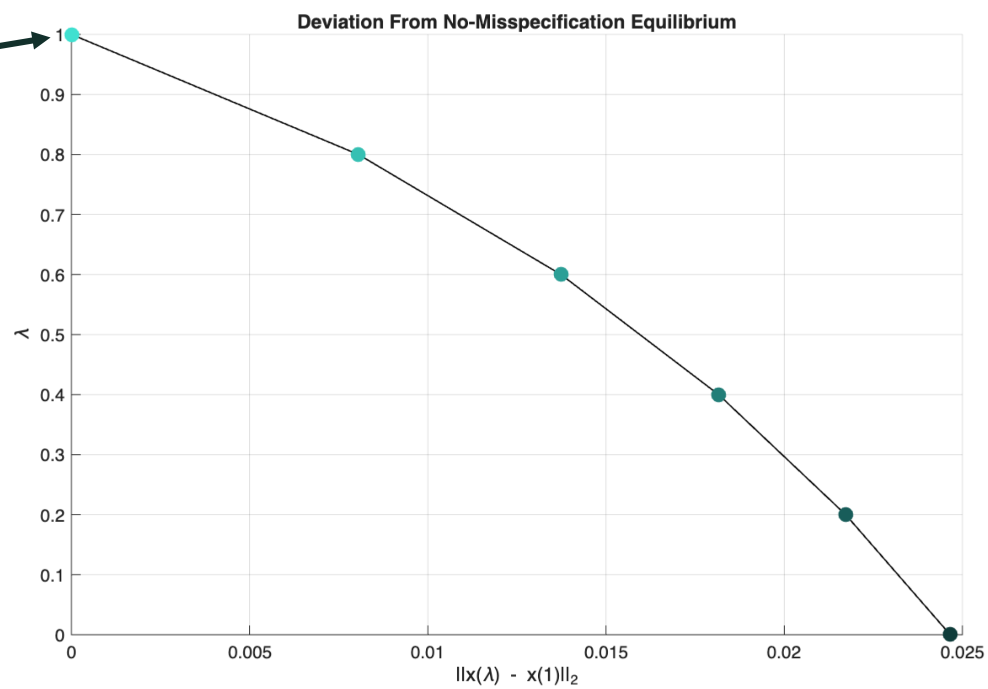

Conclusion

- Introduction of the Game-to-Real gap to capture strategic misspecifications a la sim-to-real gap
- Characterization of the Game-to-Real gap in network games
- Extension of strategic misspecification to Game Theoretic Model Predictive Control

Interested in talking more?

Ada Yildirim
PhD Student
Dartmouth College
ada.yildirim.th@dartmouth.edu