Who's Responsible? Disentangling Safety in Multi-Agent Interactions

Isaac Remy
University of Washington

Karen Leung
University of Washington

Abstract

From autonomous driving to package delivery, ensuring safe yet efficient multi-agent interaction is challenging as the interaction dynamics are influenced by hard-to-model factors such as social norms and contextual cues. Understanding these influences can aid in the design and evaluation of socially-aware autonomous agents whose behaviors are aligned with human values. In this work, we seek to codify factors governing safe multi-agent interactions via the lens of *responsibility*, i.e., an agent's willingness to deviate from their desired control to accommodate safe interaction with others. Specifically, we propose a data-driven modeling approach based on control barrier functions and differentiable optimization that efficiently learns agents' responsibility allocation from data. We demonstrate on synthetic and real-world datasets that we can obtain an interpretable and quantitative understanding of how much agents adjust their behavior to ensure the safety of others given their current environment.

Background and Intuition

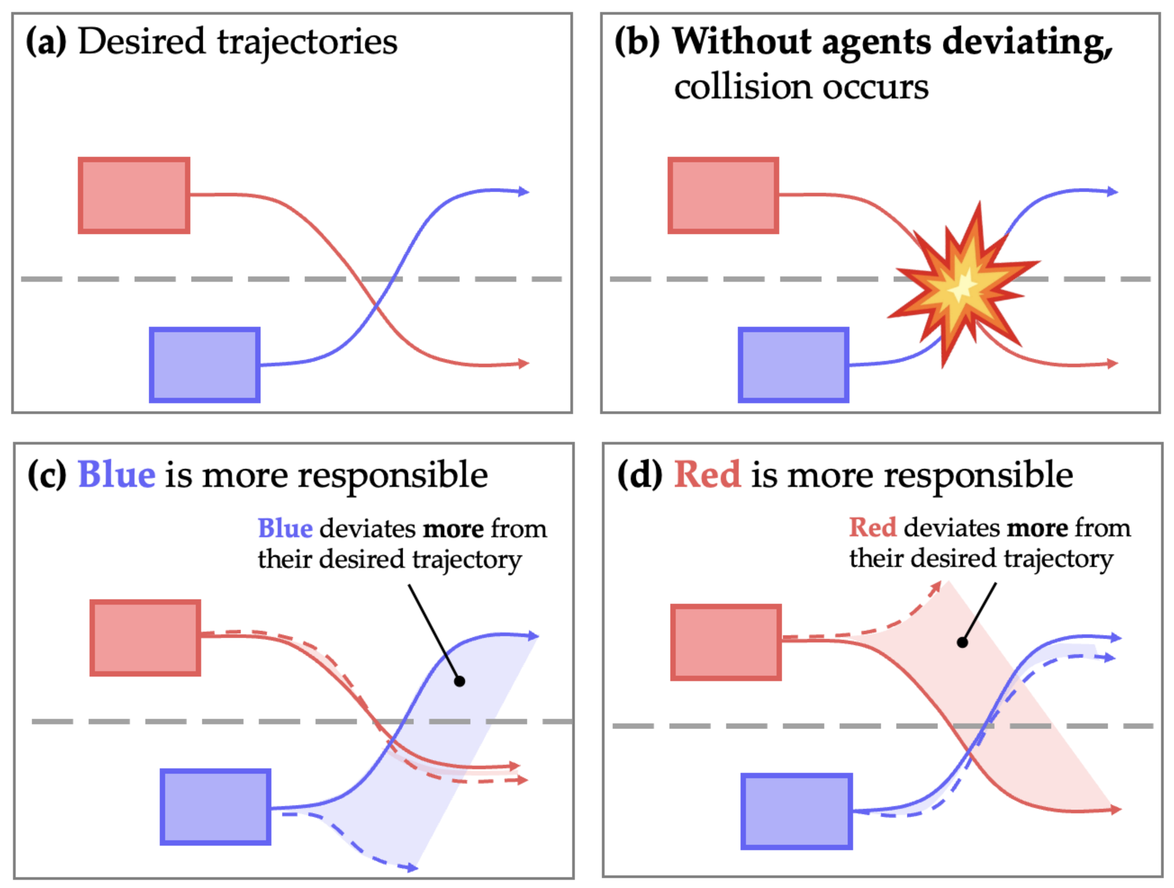

Fig. 1: Our definition of responsibility comes from the observation that people must deviate from their ideal path to ensure joint safety constraints.

Responsibility via Control Barrier Functions

CBF-QP: For some unsafe desired control that multiple agents wish to take, a *control barrier function* constraint projects to the nearest safe control, given their state and a responsibility allocation γ.

$$\mathrm{proj}(u_{1:N}^{\mathrm{desired}}; \mathbf{x}, b, \alpha, \gamma) := \underset{u_1,\ldots,u_N,\epsilon}{\mathrm{argmin}} \sum_{i=1}^{N} \left(\gamma_i \|u_i - u_i^{\mathrm{desire}}\|_2^2 + \beta_1 \|u_i\|_2^2\right) + \beta_2 \epsilon^2$$
$$\text{s.t.} \quad \nabla b(\mathbf{x})^T \left[\tilde{f}(\mathbf{x}) + \sum_{i=1}^{N} g_i(\mathbf{x}) u_i\right] + \alpha(b(\mathbf{x})) \geq -\epsilon$$
$$u_1 \in \mathcal{U}_1, \ldots, u_N \in \mathcal{U}_N$$
$$\epsilon \geq 0.$$

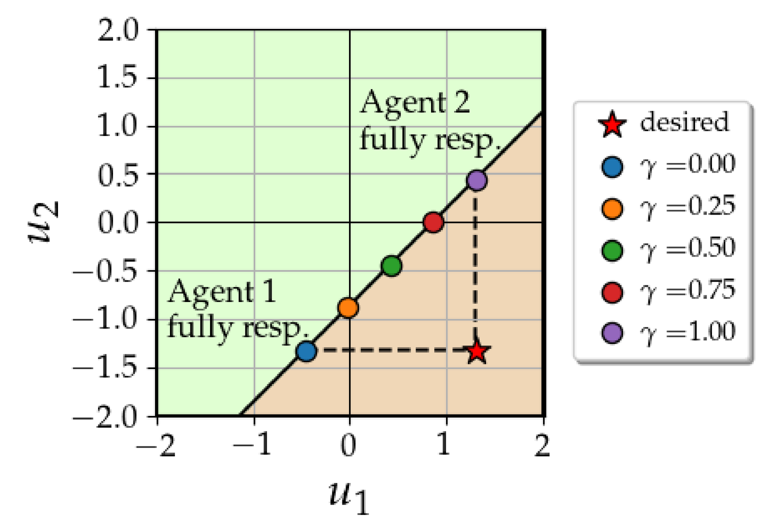

Fig. 2: A visualization of the responsibility CBF constraint in a 2-agent setting, each with a 1D control space. In the red zone (unsafe controls), the star represents the joint control the agents would like to take. The CBF-QP projects to the safe controls on the boundary.

Learning Responsibility Allocations

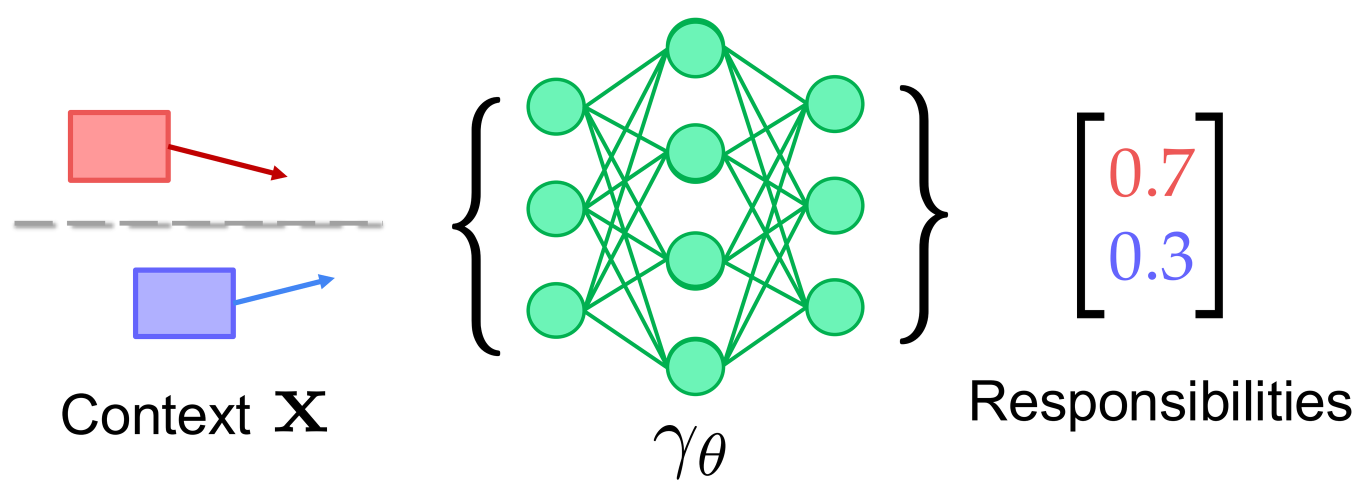

Fig. 3: Towards a contextual responsibility model. Given a set of agent states $\mathbf{X}$, we train a neural network γ_θ to perform prediction of the agents' responsibility allocations.

$$\min_{\theta} \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}^i, u_{1:N}^i) \in \mathcal{D}} \Delta(u_{1:N}^i, \tilde{u}_{1:N}^i, \gamma_\theta(\mathbf{x}^i))$$
$$\text{s.t.} \ \tilde{u}_{1:N}^i = \mathrm{proj}(\pi_{1:N}^{\mathrm{desire}}(\mathbf{x}^i); \mathbf{x}^i, b, \alpha, \gamma_\theta(\mathbf{x}^i)),$$
$$\mathbf{1}^T \gamma_\theta(\mathbf{x}^i) = 1, \quad 0 \leq \gamma_\theta(\mathbf{x}^i) \leq 1.$$

Loss function: With *differentiable optimization,* we can use the original CBF-QP to infer how good of a responsibility prediction we made for an observed datapoint $(\mathbf{x}^i, u_{1:N}^i)$.

Experiments and Future Work

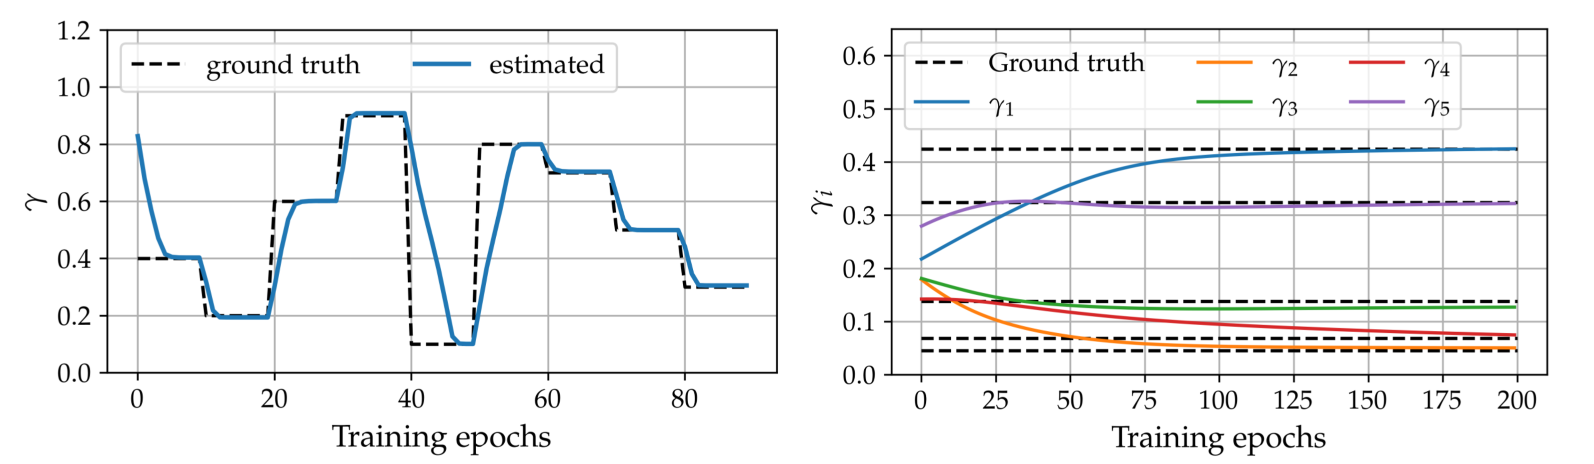

Fig. 4: Convergence to ground truth responsibility allocations using synthetically generated data for 2 and 5 agent scenarios.

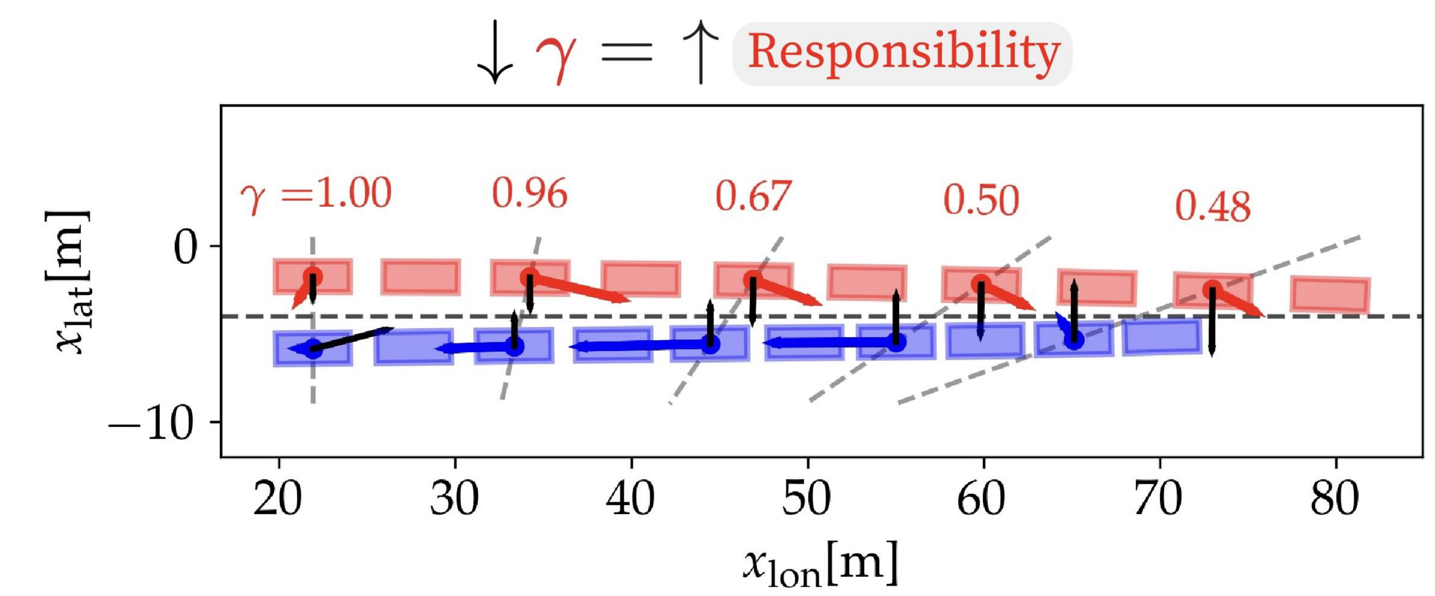

Fig. 5: Responsibility inference with a neural network trained on real-world lane-swapping data.

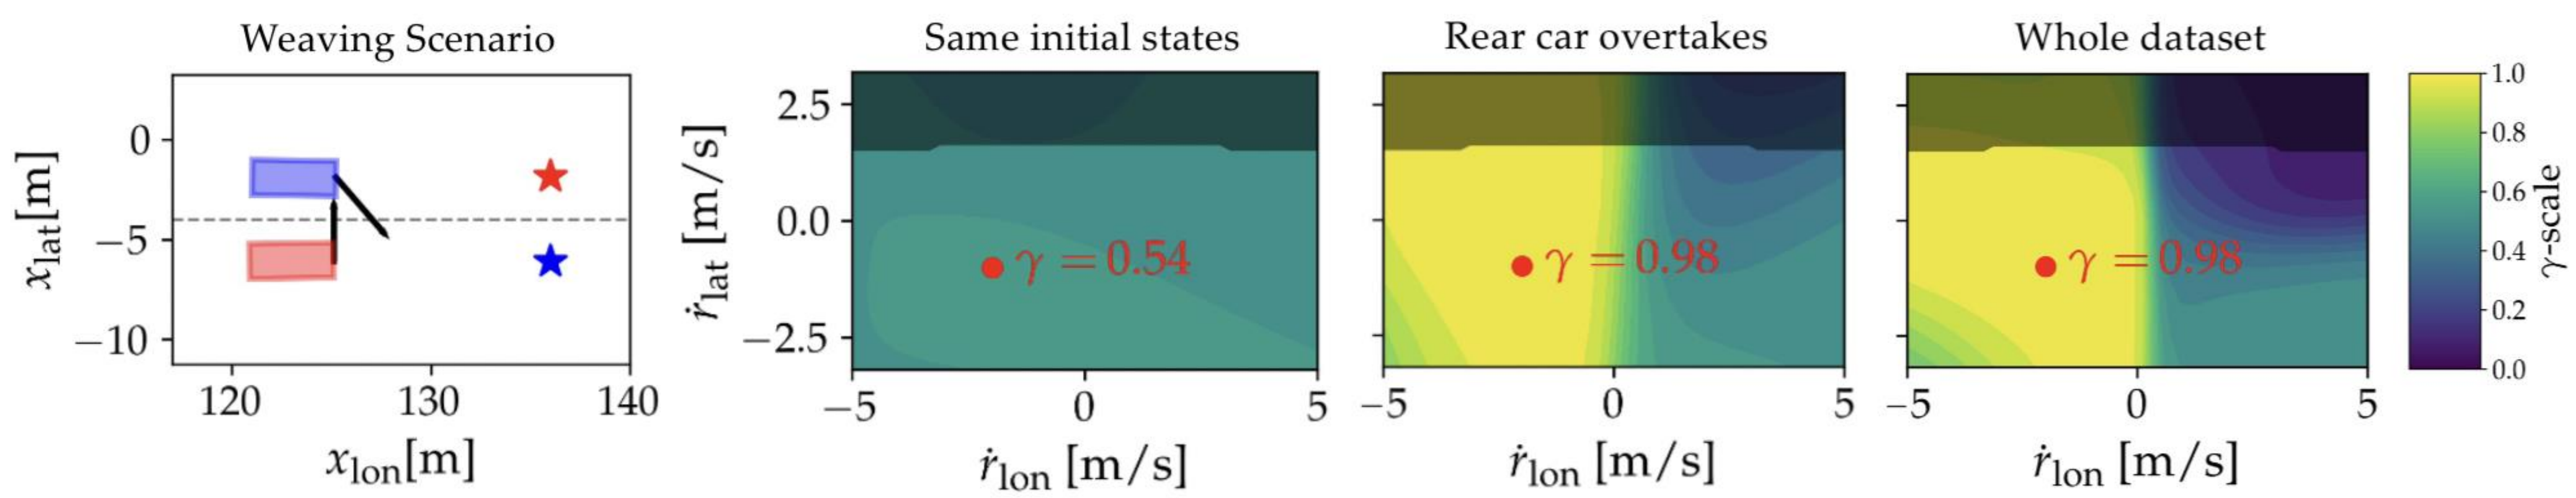

Fig. 6: Comparison of responsibility landscape after training on different sub-datasets. In the lane-swapping dataset, agents start at multiple places relative to one another. When both agents started with same initial states, our model couldn't learn a meaningful difference in responsibility allocations across different states, sparking an interest in a probabilistic approach.