Auto-Encoding Bayesian Inverse Games

Xinjie Liu*, Lasse Peters*, Javier Alonso-Mora, Ufuk Topcu, David Fridovich-Keil
*equal contribution

TEXAS TUDelft

Abstract

- **Game theory** naturally models the coupling of agents' decisions in **multi-agent interaction**. However, **complete game models** are often **unavailable** in real-world scenarios, e.g., due to unknown agents' objectives.
- **Main Contribution:** We propose a tractable approach for approximate **Bayesian inference** of **posterior distributions** of unknown game parameters.
- The method embeds a **differentiable game solver** into a **variational autoencoder** (VAE), naturally handling **continuous** and **multi-modal** distributions.

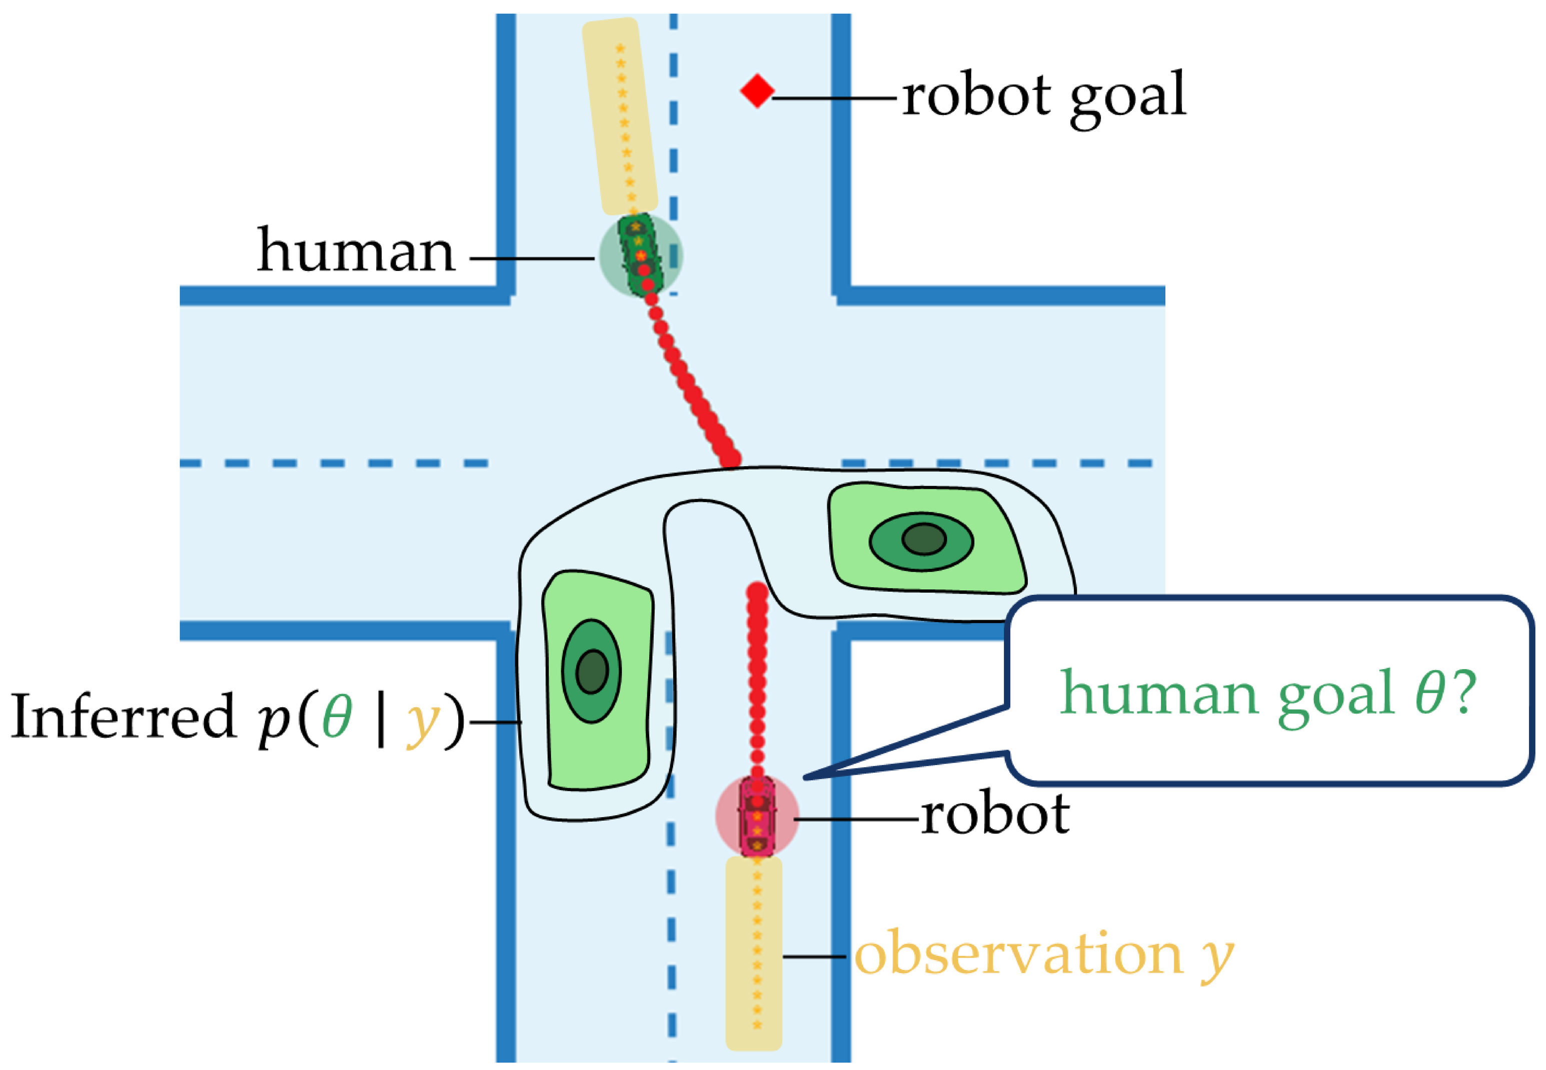

Preliminaries: Generalized Nash Games

$$\text{robot:} \quad \mathcal{S}^r_\theta(\tau^h) := \arg\min_{\tau_r} \; J^r_\theta(\tau^r, \tau^h) \quad \text{s.t.} \quad g^r_\theta(\tau^r, \tau^h) \geq 0$$

$$\text{human:} \quad \mathcal{S}^h_\theta(\tau^r) := \arg\min_{\tau_h} \; J^h_\theta(\tau^h, \tau^r) \quad \text{s.t.} \quad g^h_\theta(\tau^h, \tau^r) \geq 0$$

- **Coupled** trajectory **optimization** problems.
- Solution: Generalized Nash equilibrium (GNE).
- Parameter $\boldsymbol{\theta}$: Unknown aspects of the game, e.g., agents' goal position, desired driving speed, lane preference, etc.

Formalizing Bayesian Inverse Games

$$p(\theta \mid y) = \frac{p(y \mid \theta)p(\theta)}{p(y)} \quad \text{s.t. Nash equilibrium conditions}$$

Challenges:

- **Unknown** prior $p(\theta)$.
- Observation model $p(y \mid \theta)$ involves **game solve**, and the posterior $p(\theta \mid y)$ is in general **non-Gaussian** or even **multi-modal**.
- The computation of the normalizing constant is **intractable** due to the **marginalization** $p(y) = \int p(y \mid \theta)p(\theta)\mathrm{d}\theta$.

Auto-Encoding Bayesian Inverse Games

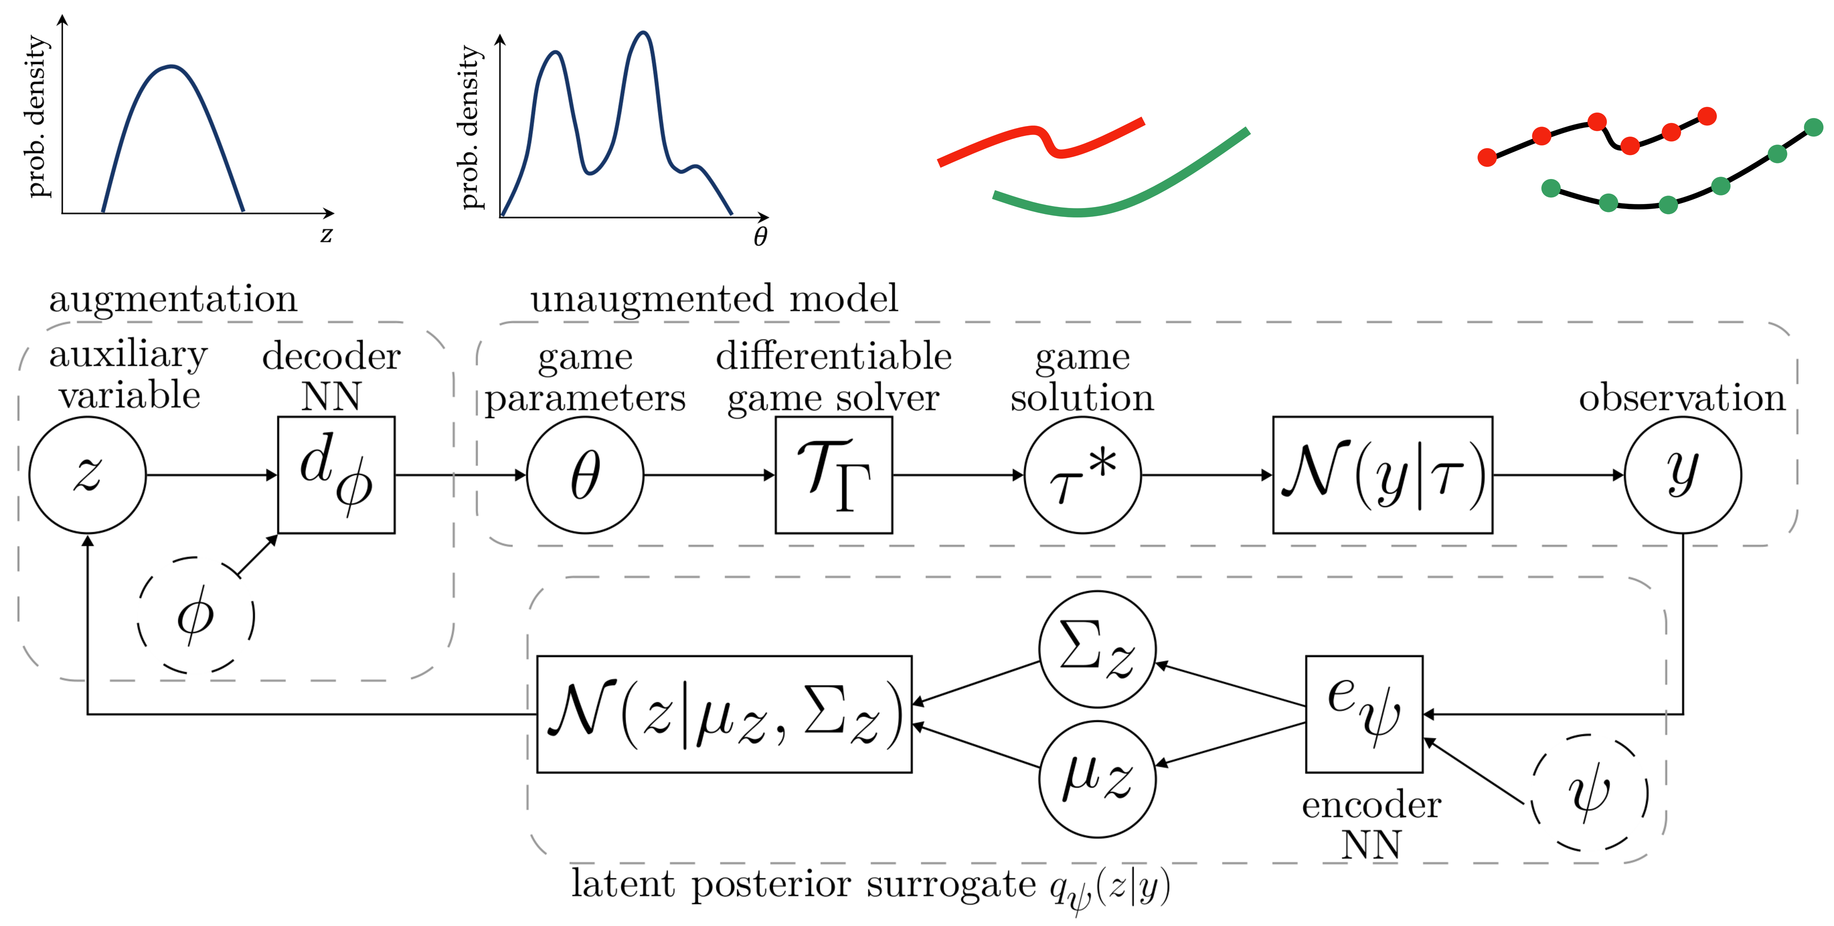

- Solves a **variational inference** problem to approximate Bayesian inference.
- Results in **a structured VAE** framework, where the **differentiable game solver** encodes the game structure and constraints.
- Naturally handles **continuous, multi-modal** distributions. The pipeline supports efficient sampling from the inferred posteriors and does **not** require **game solve** at **runtime**.
- The structured VAE can be trained from an **unlabeled** dataset of observed interactions.

Baseline: maximum likelihood estimation (MLE)

Ours: $p(\theta \mid y) = \frac{p(y \mid \theta)p(\theta)}{p(y)}$ s.t. Nash equilibrium conditions

Baseline: $\cancel{p(\theta \mid y)} = \frac{p(y \mid \theta)\cancel{p(\theta)}}{\cancel{p(y)}}$

$$\hat{\theta} \in \arg\max_\theta \; p(y \mid \theta) \quad \text{s.t. Nash equilibrium conditions}$$

- **Ignores** prior.
- Only provides **point estimates** without uncertainty quantification and performs **poorly** in case of **uninformative observations**.

Experimental Results

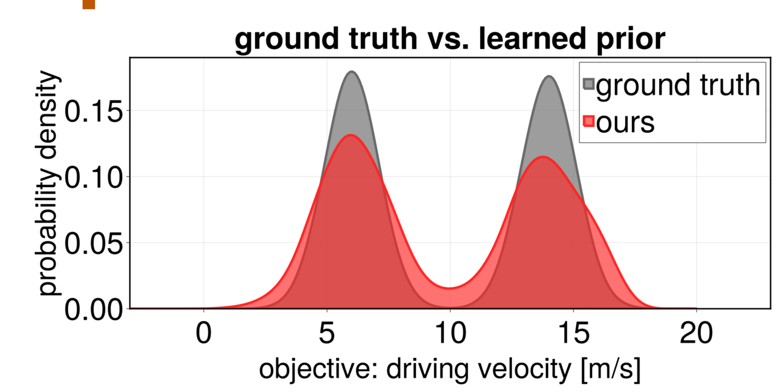

Fig. 1

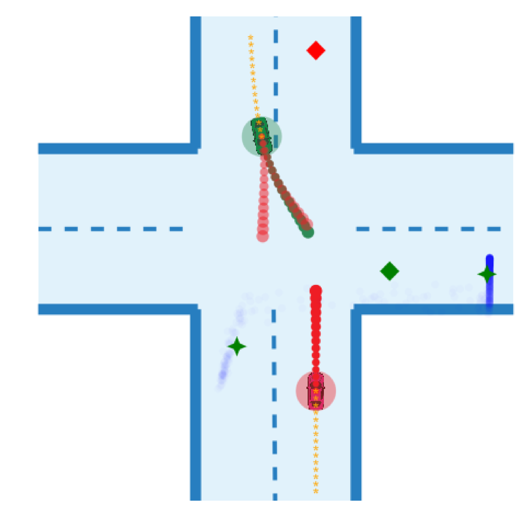

Fig. 2

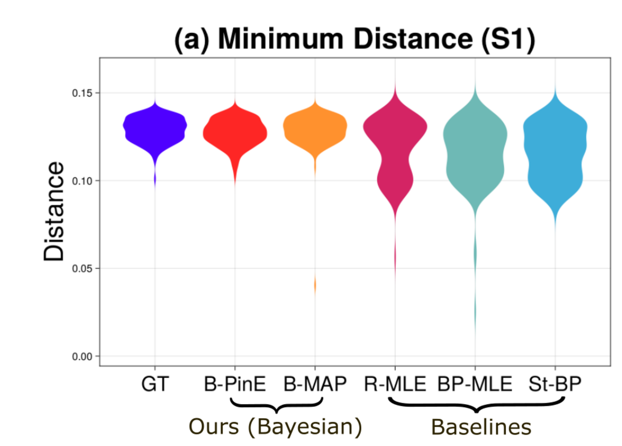

Fig. 3

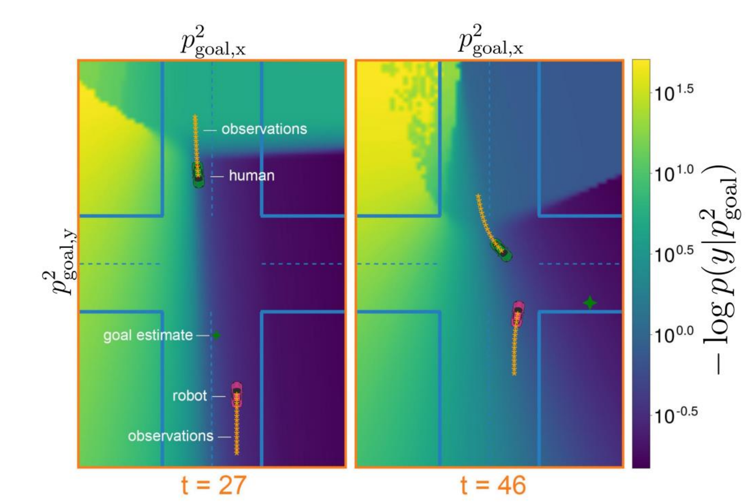

Fig. 4

The proposed Bayesian inverse game approach **learns unknown priors** from **unlabeled** interactions (Fig. 1), captures **multi-modality** of unknown game parameter posteriors (Fig. 2) and gives **improved** downstream motion planning **safety** (Fig. 3). The MLE baseline performs poorly in case of uninformative observations (Fig. 4).

Project Website

xinjie-liu.github.io/projects/bayesian-inverse-games/

Contact

Xinjie Liu
https://xinjie-liu.github.io