



Neural Amortized Inference for Nested Multi-agent Reasoning

Kunal Jha, Tuan Anh Le, Chuanyang
Jin, Yen-Ling Kuo, Joshua
Tenenbaum, Tianmin Shu




DARTMOUTH

(July 2023)

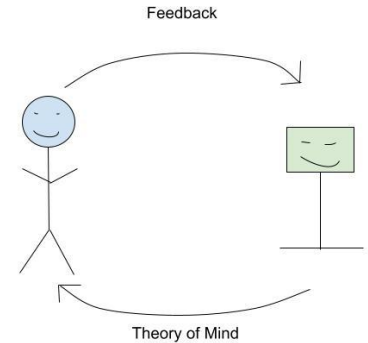


NYU




Theory of Mind


- Theory of Mind: Reasoning about the **hidden mental variables** that drive **observable actions**




The blue car sees the red car but can't see the pedestrian behind the building.



The red car sees the pedestrian and slows down to avoid crashing.



The blue car sees the red car slowing down. Not seeing any other obstacles, it infers the presence of a pedestrian behind the building and therefore slows down.



How Can We Model This?

Partially Observable Markov Decision Processes (POMDP)

- The entire state of the world is not visible
- Each agent stores (*mental states, actions, observations*)

Bayesian Inverse Planning

$$P(\text{mental state} \mid \text{plan}) \propto P(\text{plan} \mid \text{mental state}) * P(\text{mental state})$$




Interactive Partially Observable Markov Decision Processes (I-POMDP)

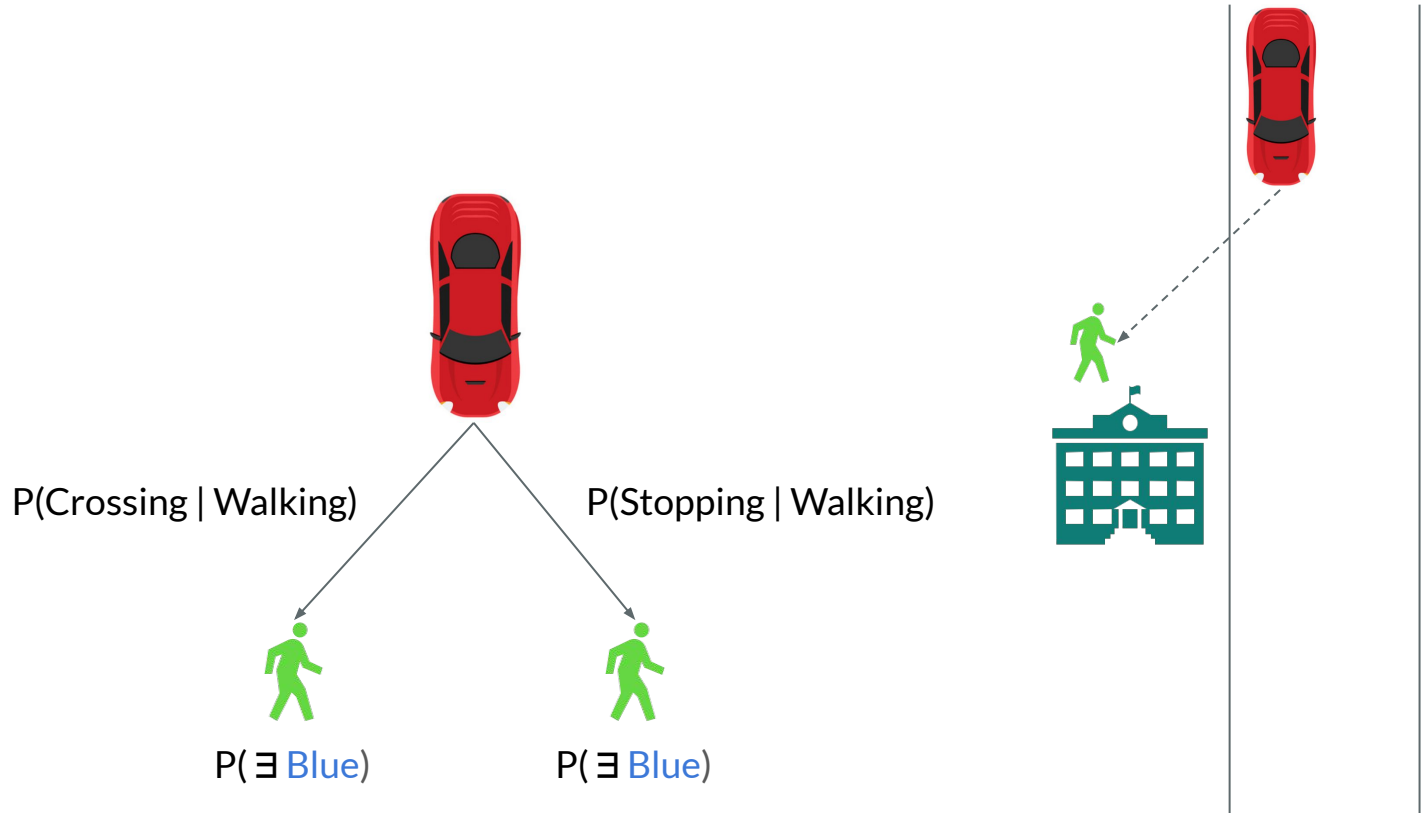


$P(\text{Blue car existing behind the barrier}) = P(\exists \text{Blue})$

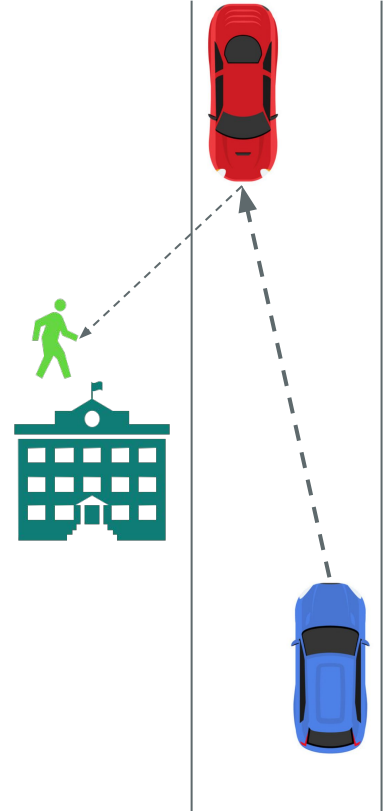
Interactive Partially Observable Markov Decision Processes (I-POMDP)



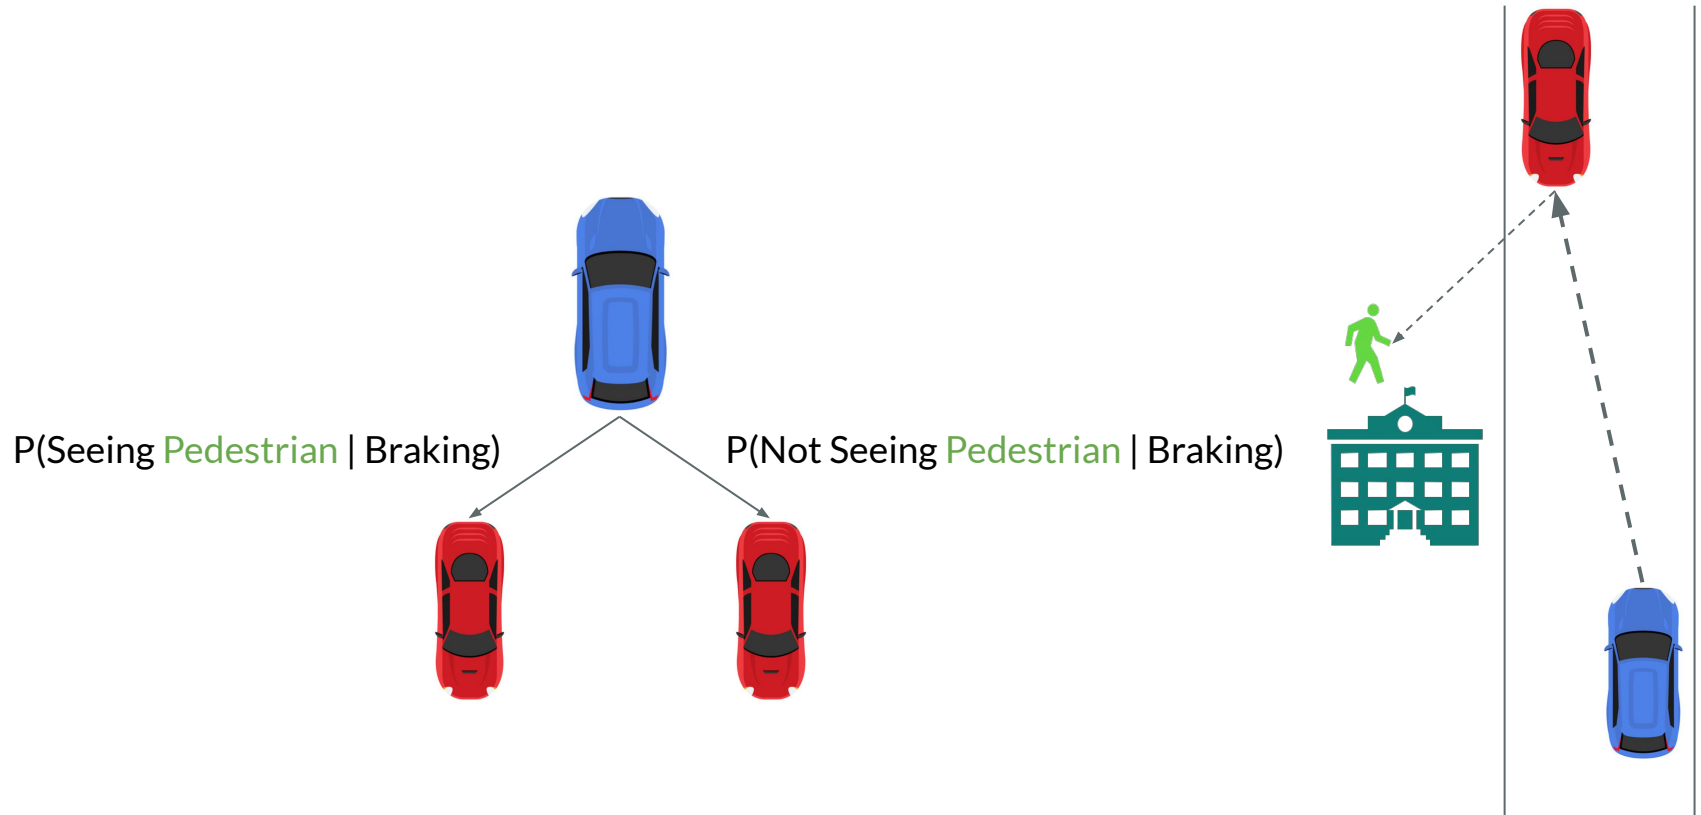
Interactive Partially Observable Markov Decision Processes (I-POMDP)



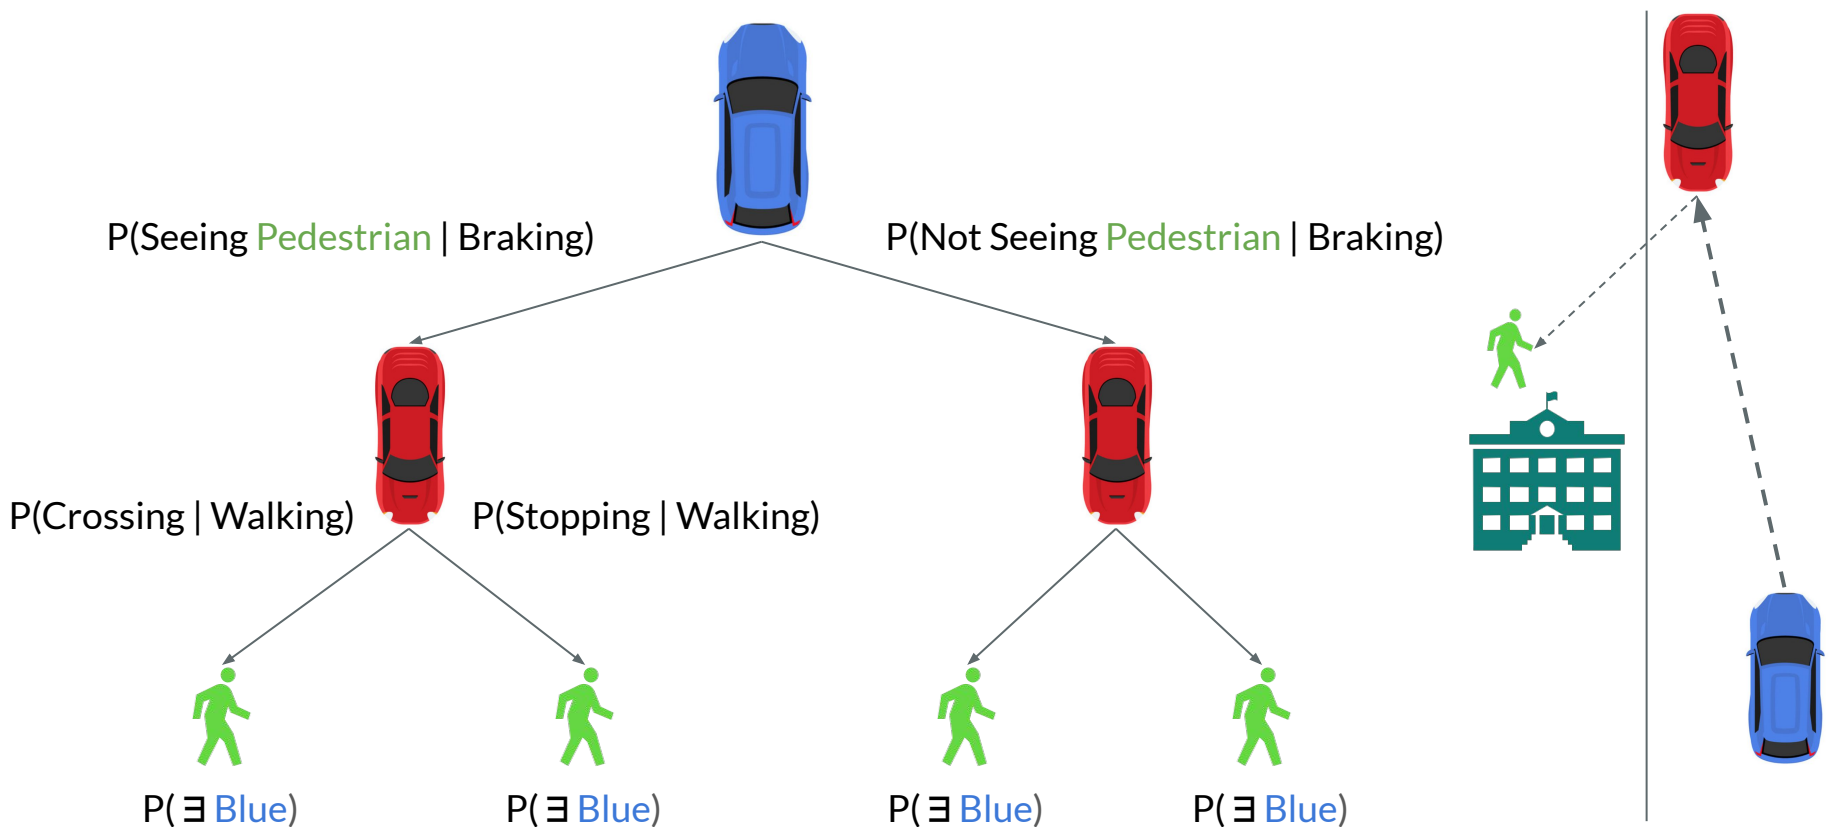
Interactive Partially Observable Markov Decision Processes (I-POMDP)



Interactive Partially Observable Markov Decision Processes (I-POMDP)



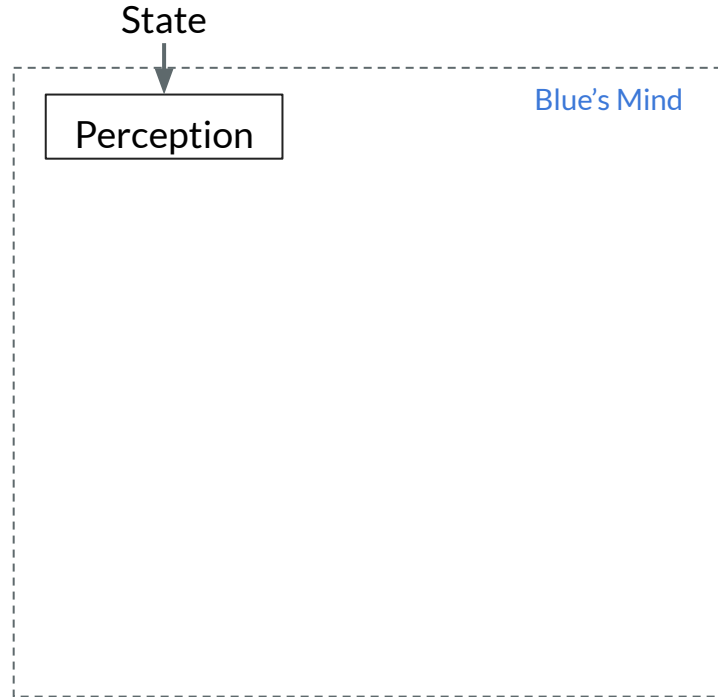
Interactive Partially Observable Markov Decision Processes (I-POMDP)



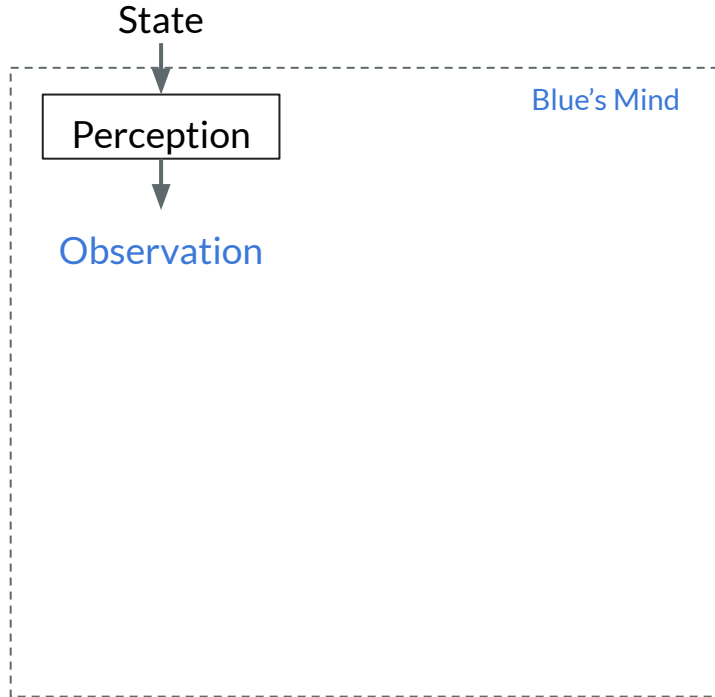
Background - Interactive Partially Observable Markov Decision Processes

State

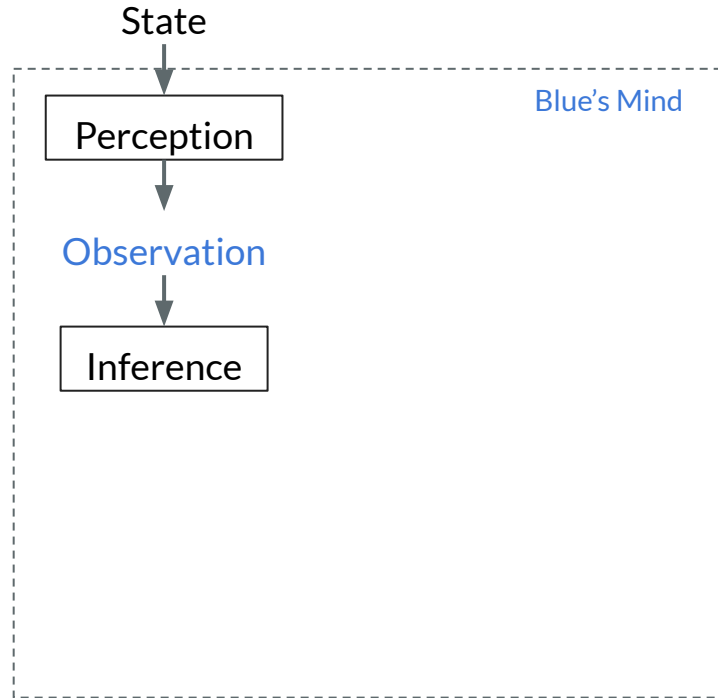
Background - Interactive Partially Observable Markov Decision Processes



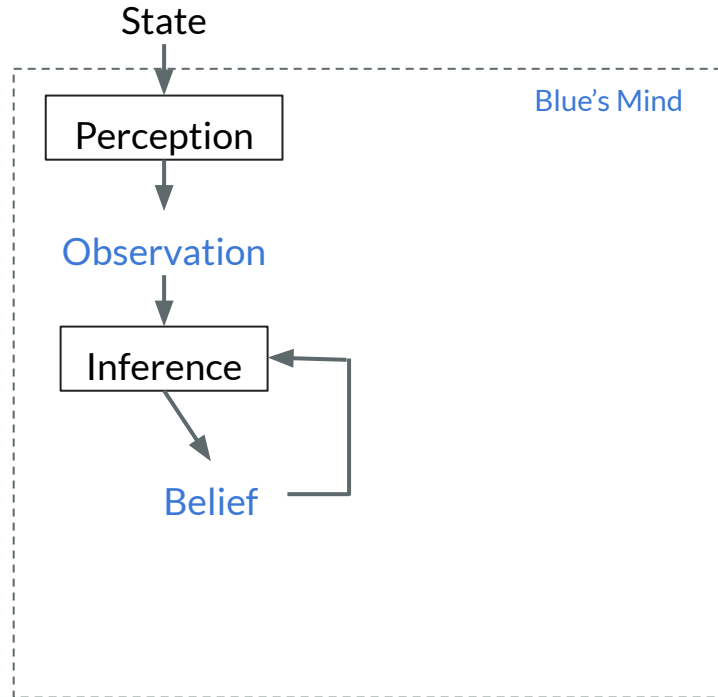
Background - Interactive Partially Observable Markov Decision Processes



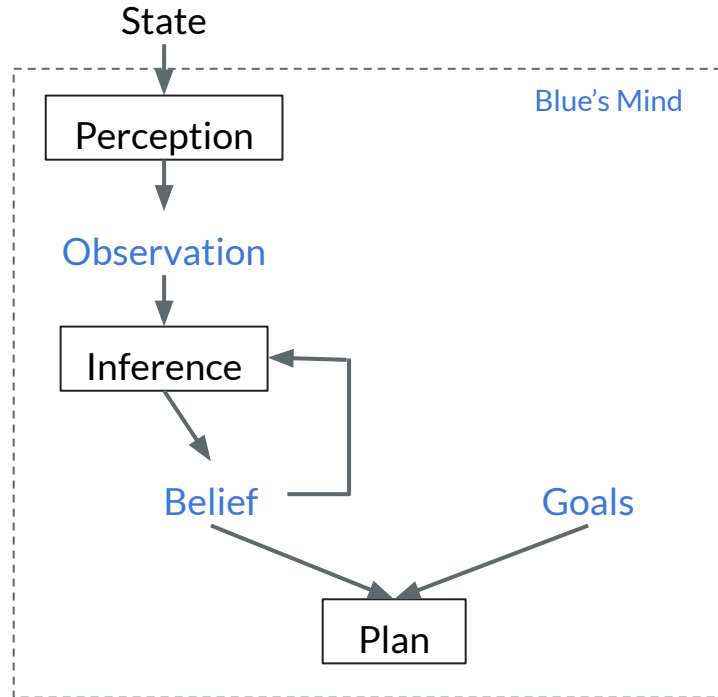
Background - Interactive Partially Observable Markov Decision Processes



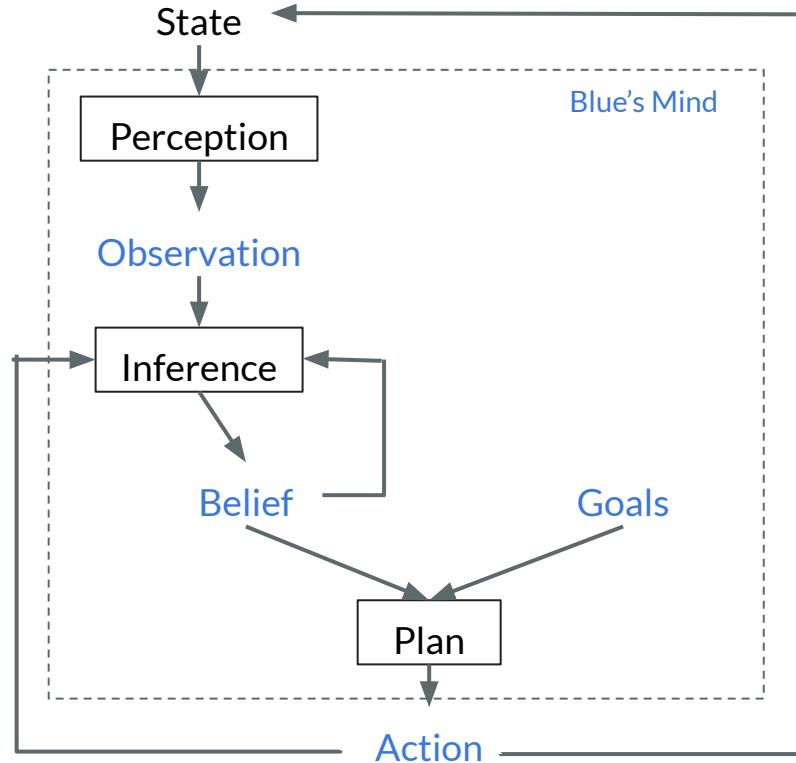
Background - Interactive Partially Observable Markov Decision Processes



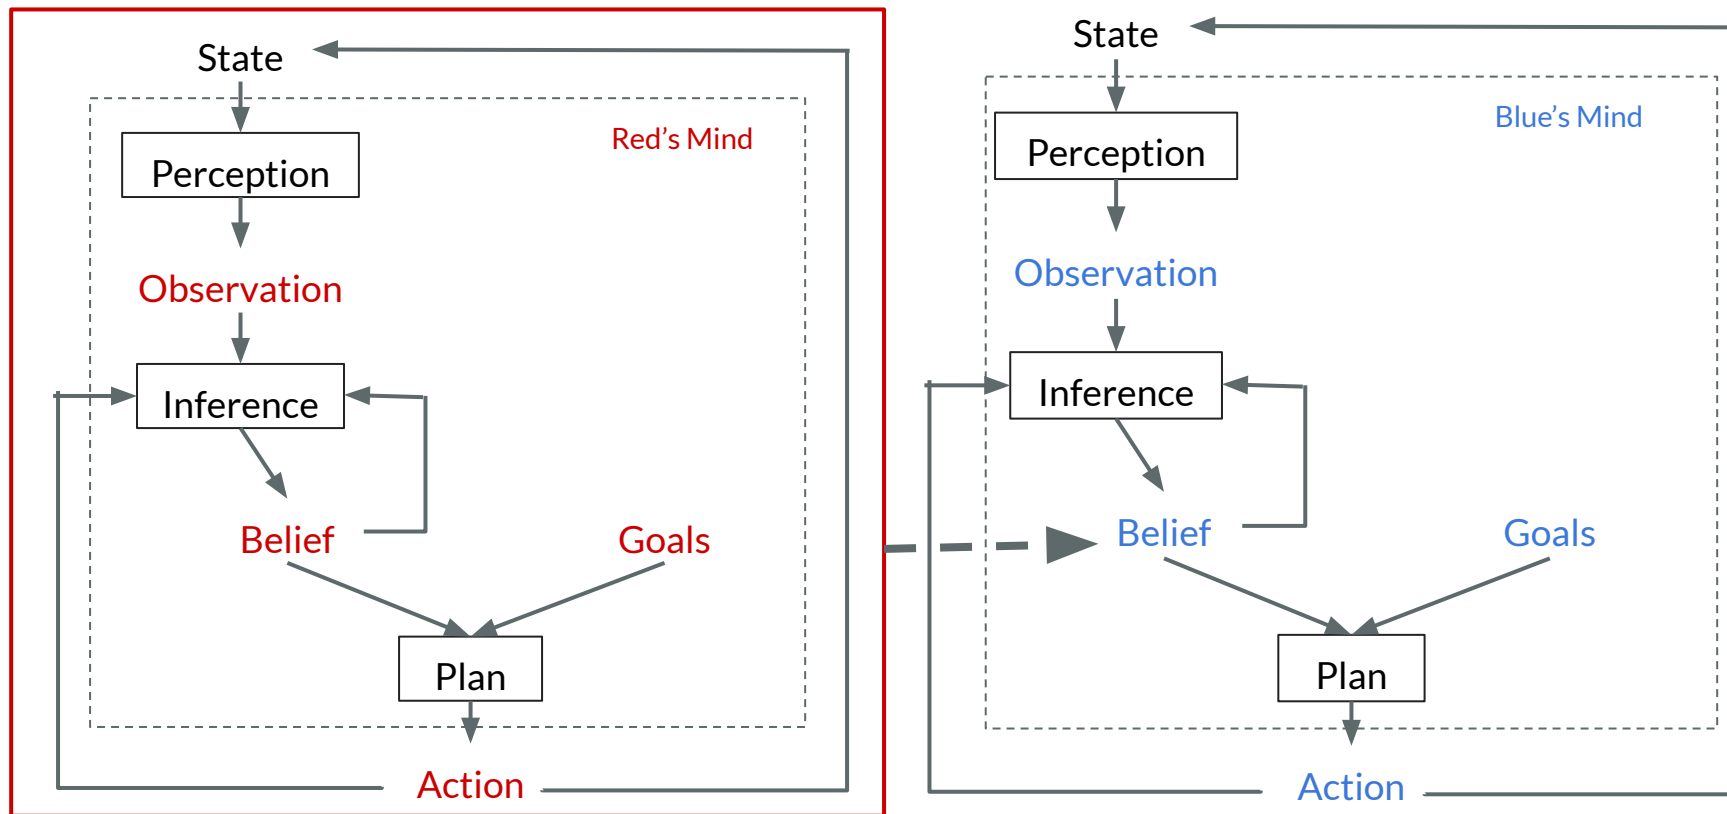
Background - Interactive Partially Observable Markov Decision Processes



Background - Interactive Partially Observable Markov Decision Processes

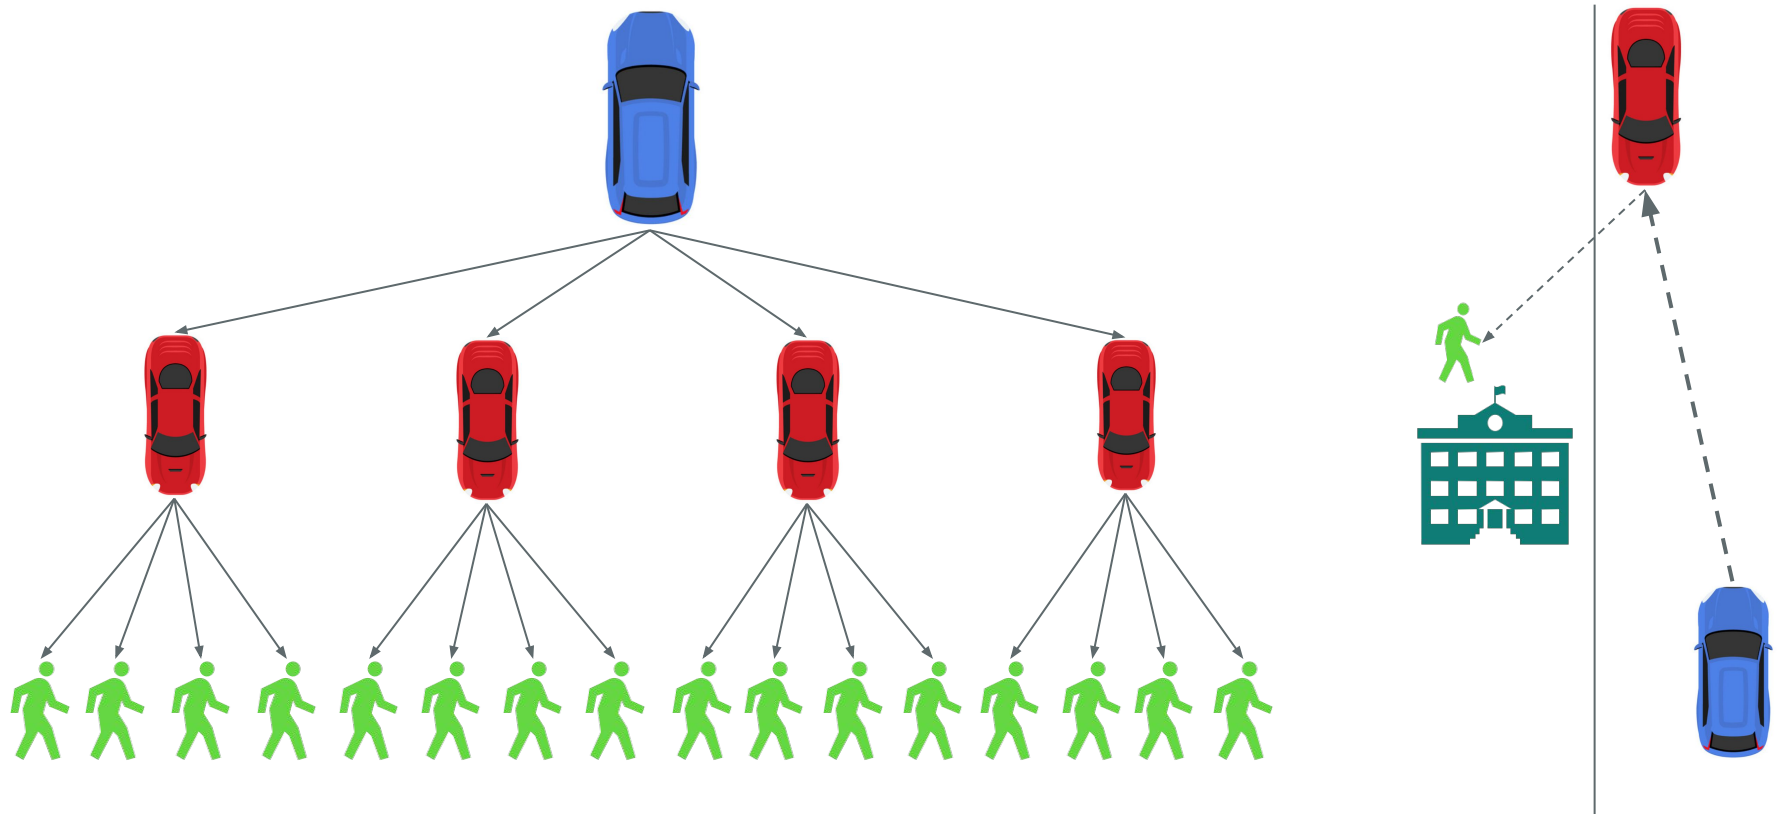


Background - Interactive Partially Observable Markov Decision Processes



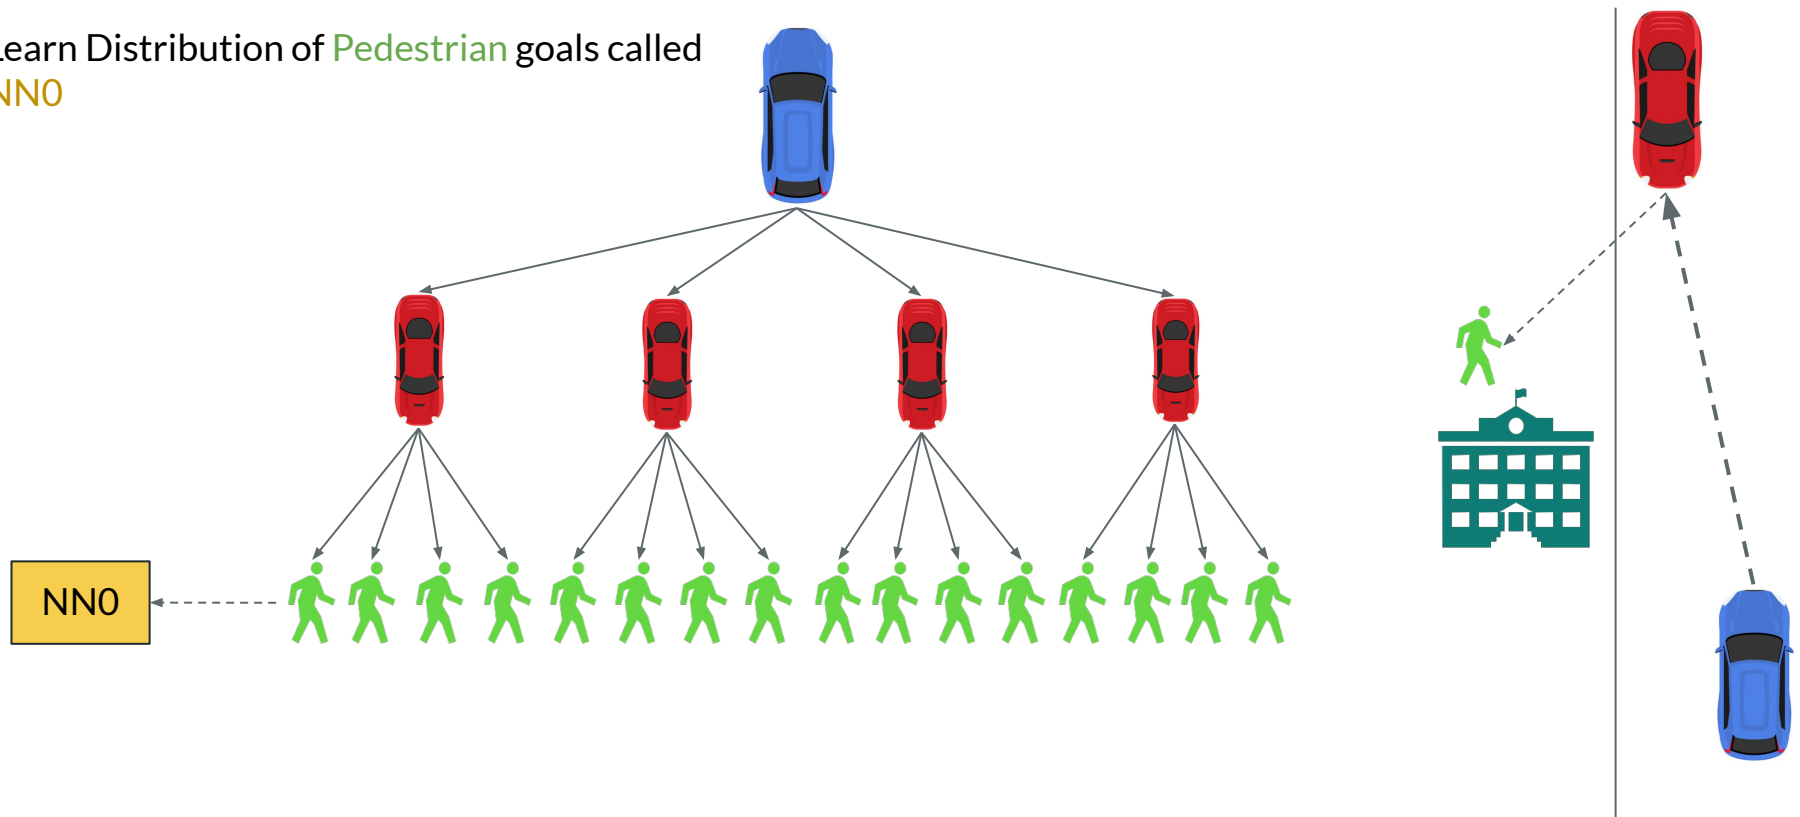
**Our Approach: Amortize Lower Levels and
Update Beliefs through Importance Sampling**

Our Approach - Amortized Inference



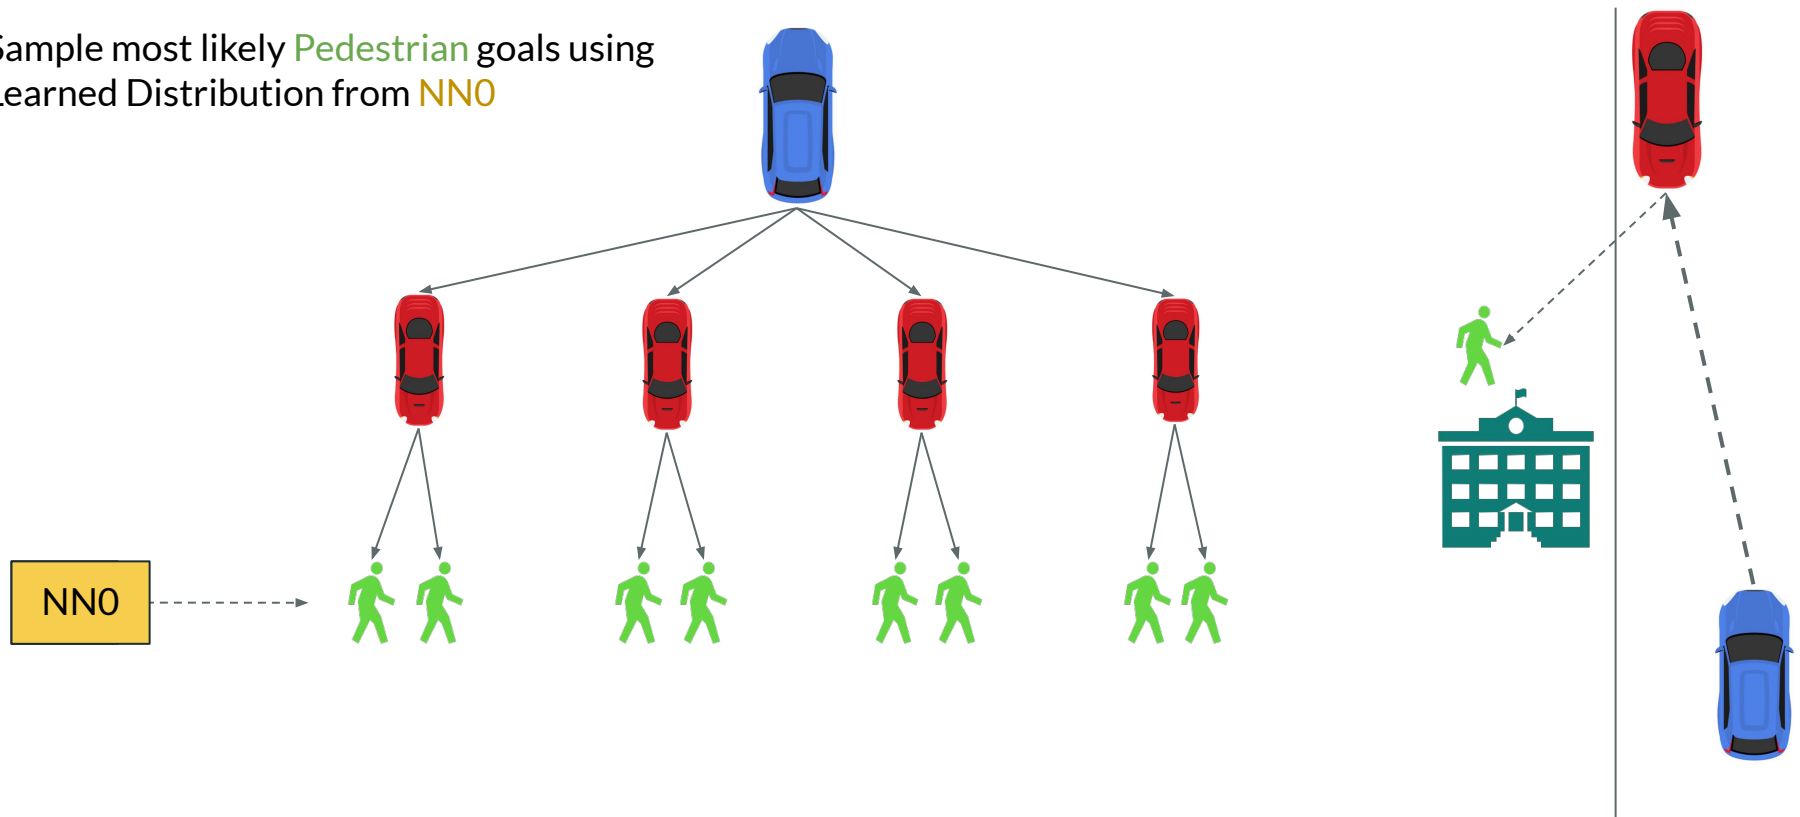
Our Approach - Amortized Inference

Learn Distribution of Pedestrian goals called
NNO



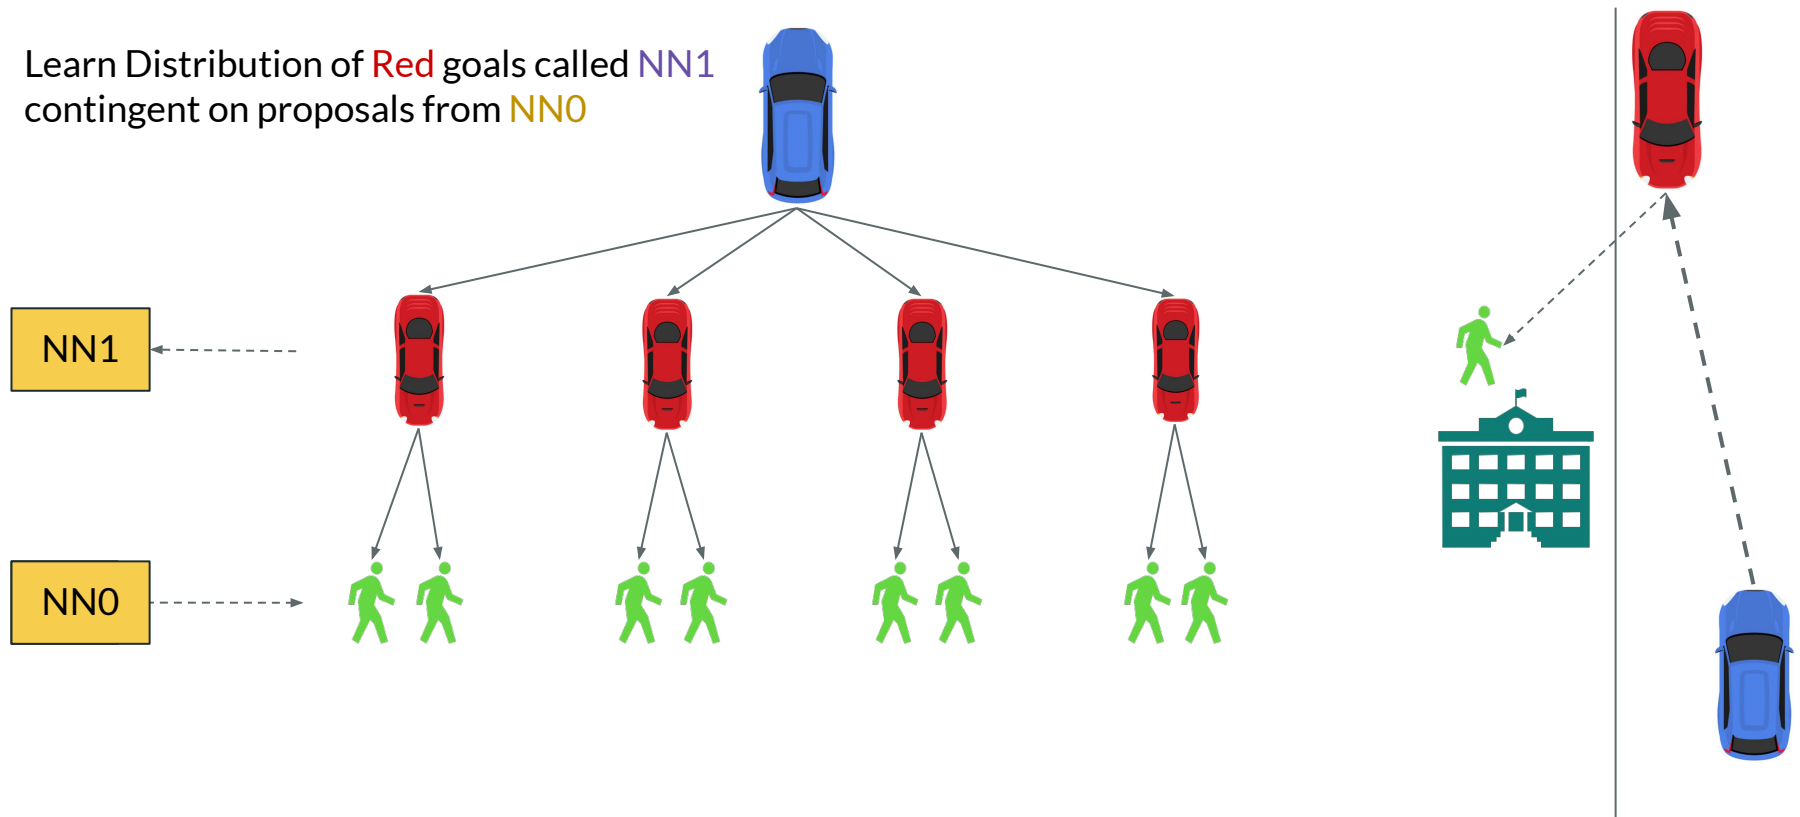
Our Approach - Amortized Inference

Sample most likely Pedestrian goals using
Learned Distribution from NNO



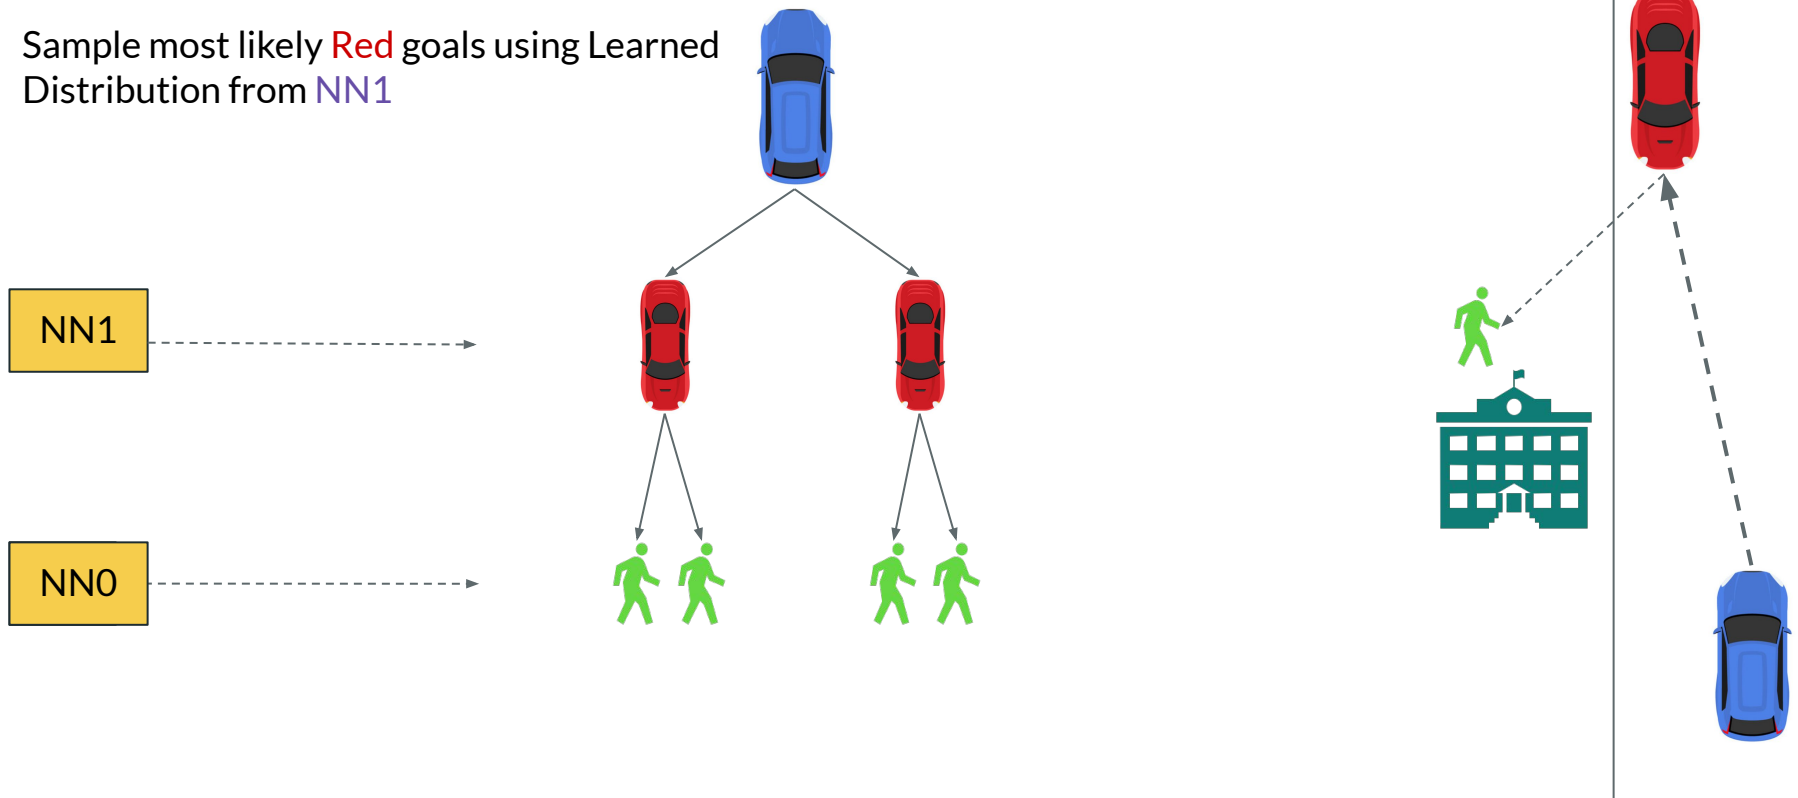
Our Approach - Amortized Inference

Learn Distribution of **Red** goals called **NN1**
contingent on proposals from **NN0**

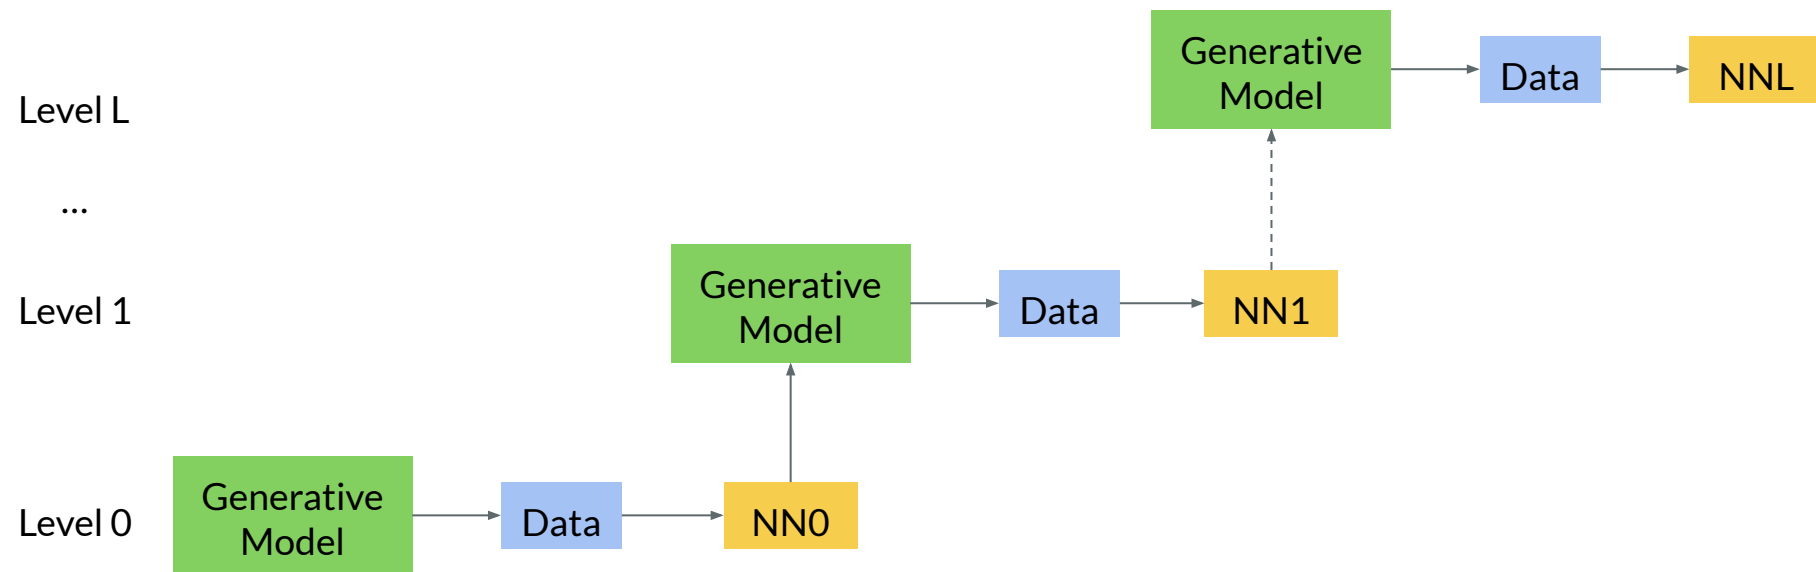


Our Approach - Amortized Inference

Sample most likely **Red** goals using Learned Distribution from **NN1**



Our Approach - Building a Dataset for Recursive Reasoning



Experiment - Setup

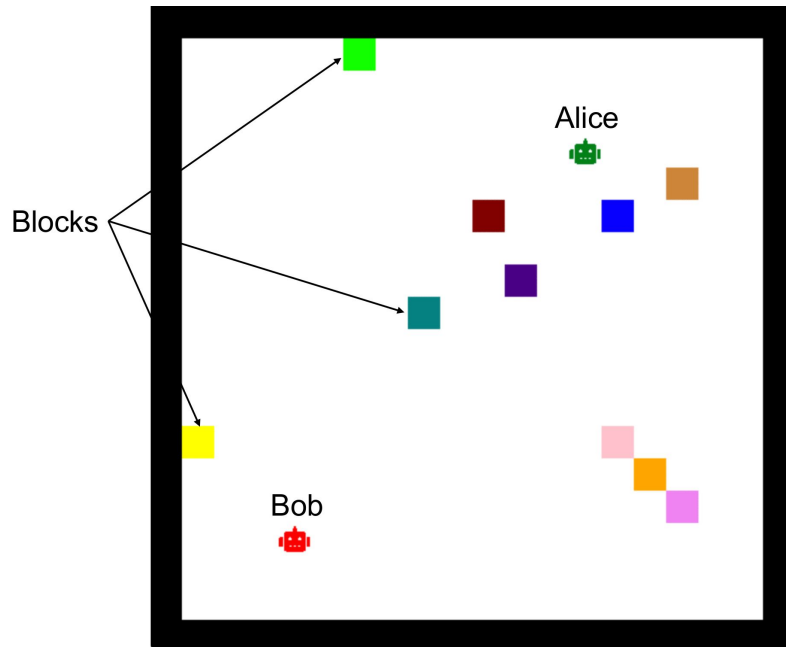
- 2 Agents, **Alice (Green)**, **Bob (Red)**
 - **Alice** wants to move two colored blocks together
 - **Bob** wants to help or hinder Alice
- Neither agent knows each other's goal
- At each timestep, we ask a model to infer the intentions of **Bob**
- Goal Space 20 times larger than prior work



Alice









Bob

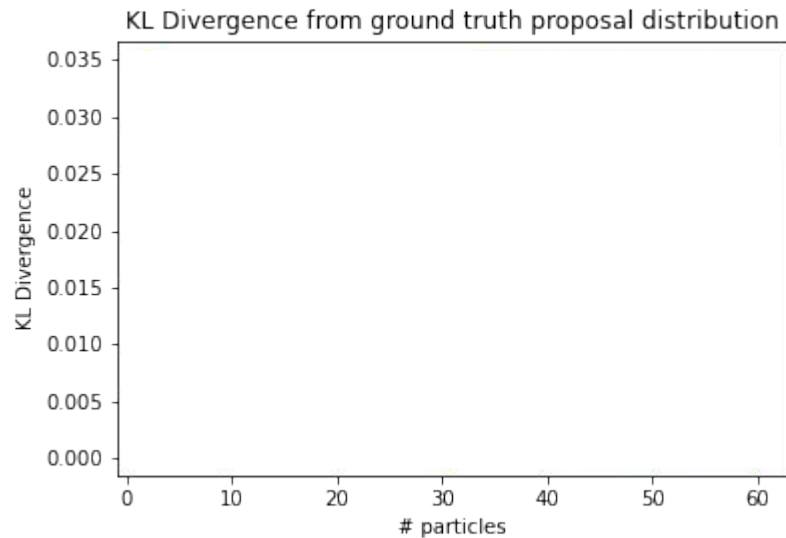
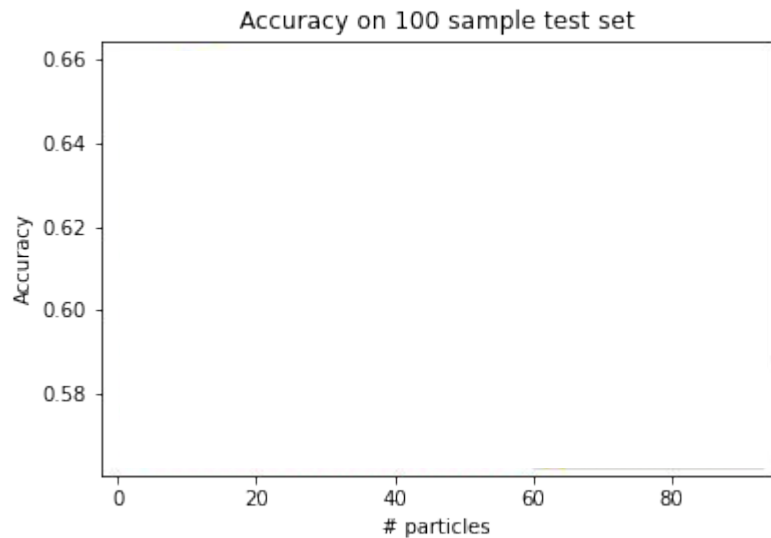


Experiment - Data Generation



- Procedurally Generate Data by sampling random States and Goals





Experiment - Results on Larger Dataset





Experiment - Results on Larger Dataset





Experiment - Results on Larger Dataset




Experiment - Results on Larger Dataset




Experiment - Results from **Bob** Helping **Alice**


Alice's Goals:  and 



Bob's belief about Alice's desires (N/A)

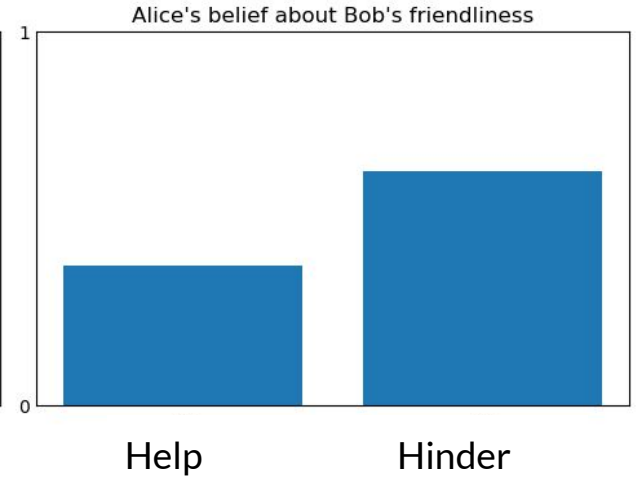
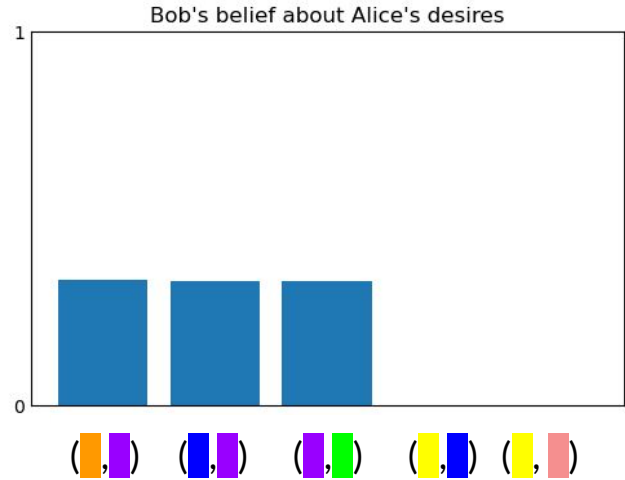
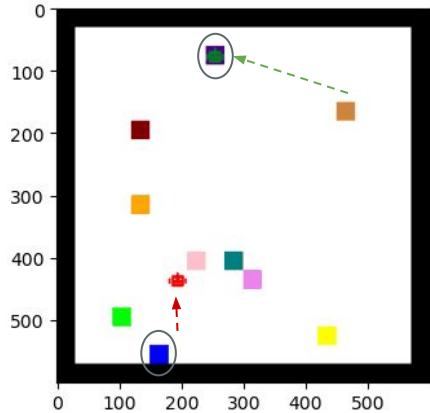


Alice's belief about Bob's friendliness






Experiment - Results from **Bob** Helping **Alice**

Alice's Goals: ■ and ■
Bob's Social Goal: Help






Experiment - Results from **Bob** Helping **Alice**

Alice's Goals: ■ and ■
Bob's Social Goal: Help






Experiment - Results from **Bob** Helping **Alice**

Alice's Goals: ■ and ■
Bob's Social Goal: Help








Experiment - Results from **Bob** Helping **Alice**

Alice's Goals: ■ and ■
Bob's Social Goal: Help




Experiment - Results from **Bob** Hindering Alice



Alice's Goals:  and 



Experiment - Results from **Bob** Hindering **Alice**


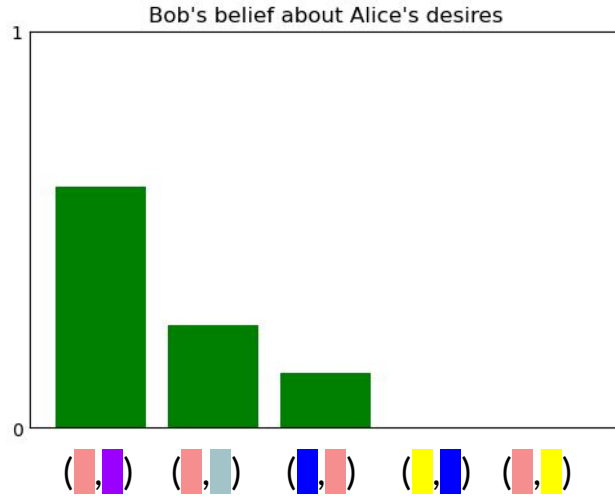



Alice's Goals: ■ and ■
Bob's Social Goal: Hinder






Experiment - Results from **Bob** Hindering **Alice**

Alice's Goals: ■ and ■
Bob's Social Goal: Hinder




Experiment - Results from **Bob** Hindering Alice

Alice's Goals: ■ and ■
Bob's Social Goal: Hinder




Experiment - Results from **Bob** Hindering **Alice**


Alice's Goals: ■ and ■
Bob's Social Goal: Hinder



Bob's belief about Alice's desires



Alice's belief about Bob's friendliness



Conclusion

- Neural amortized inference **drastically accelerates** nested social reasoning while preserving robustness and uncertainty estimation
- Current experiments focus on *goal* inference, but our framework extends to *physical state* inference

