DARTMOUTH

Neural Amortized Inference for Nested Multi-agent Reasoning

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

NYU

UNIVERSITY of VIRGINIA

AAAI, Vancouver, BC, Canada
February 2024

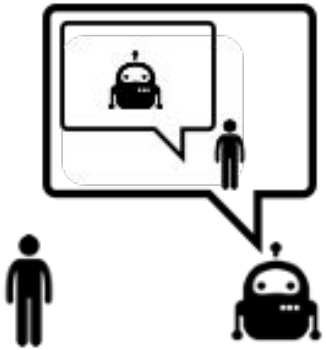

Theory of Mind

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

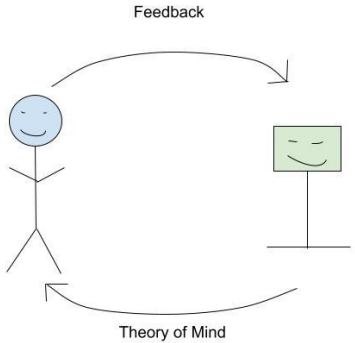

The blue car sees the red car but can’t see the pedestrian behind the building.

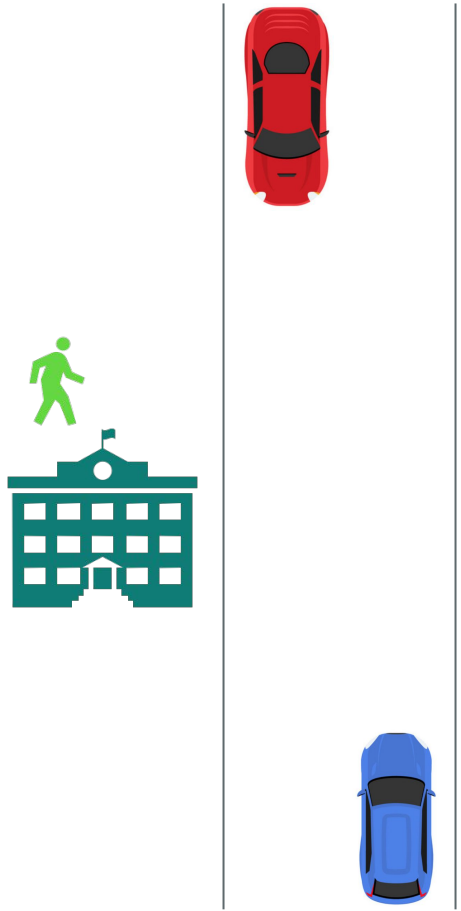

The red car sees the pedestrian
and slows down to avoid crashing.

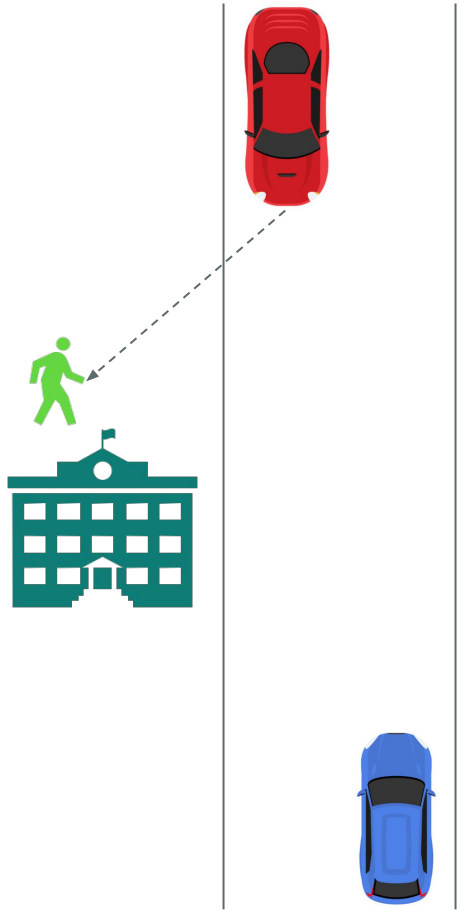

The blue car sees the red car slowing down. Not seeing any other obstacles, it infers the presence of a pedestrian behind the building and therefore slows down.

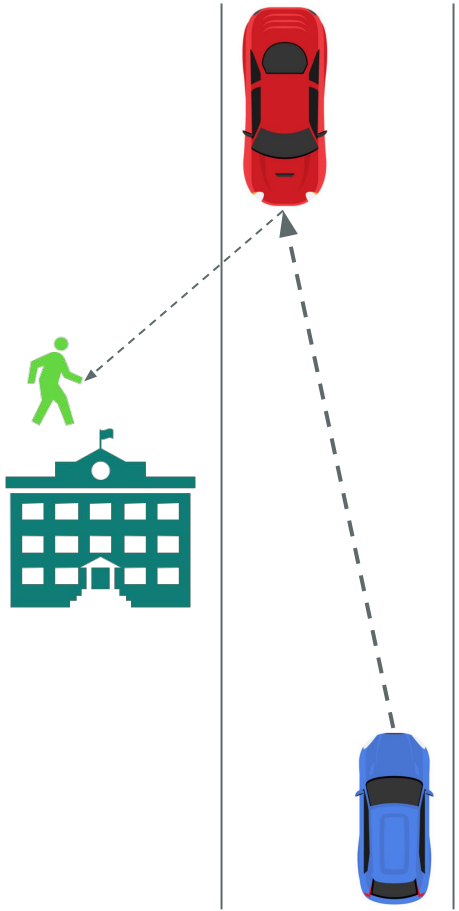

How Can We Model This?

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(mental state | plan) $\propto$ P(plan | mental state) * P(mental state)

Interactive Partially Observable Markov Decision Processes (I-POMDP)

P(Blue car existing behind the barrier) ≒ P(∃ Blue)

Interactive Partially Observable Markov Decision Processes (I-POMDP)

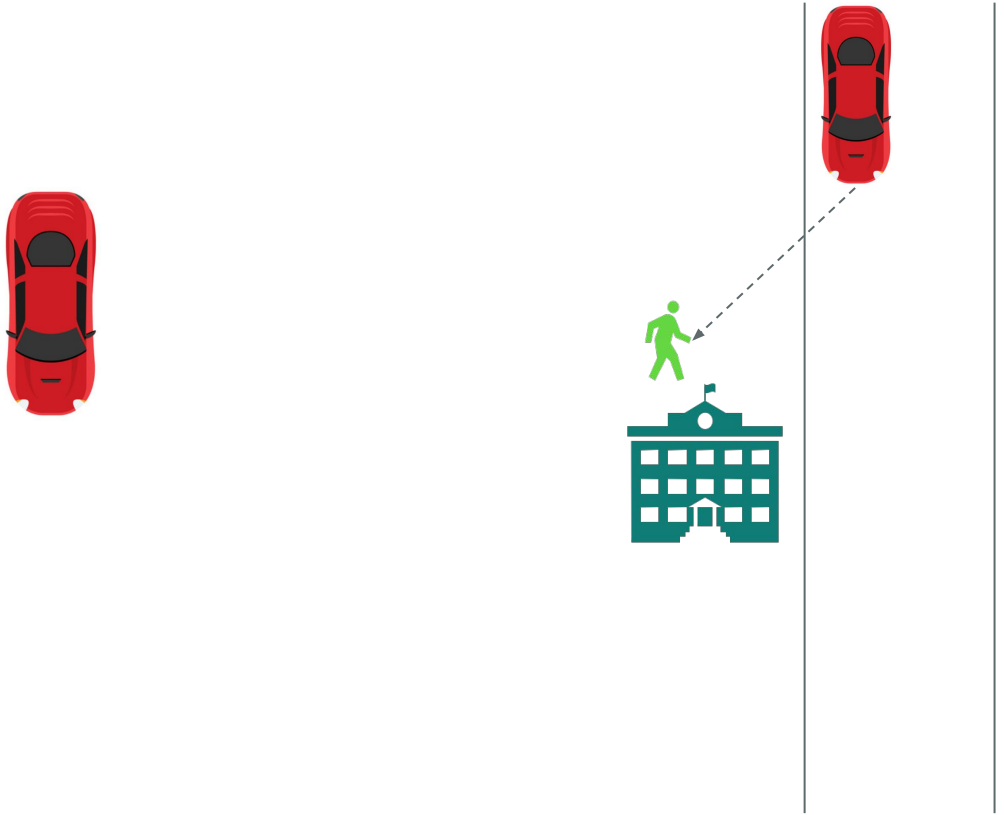

Interactive Partially Observable Markov Decision Processes (I-POMDP)

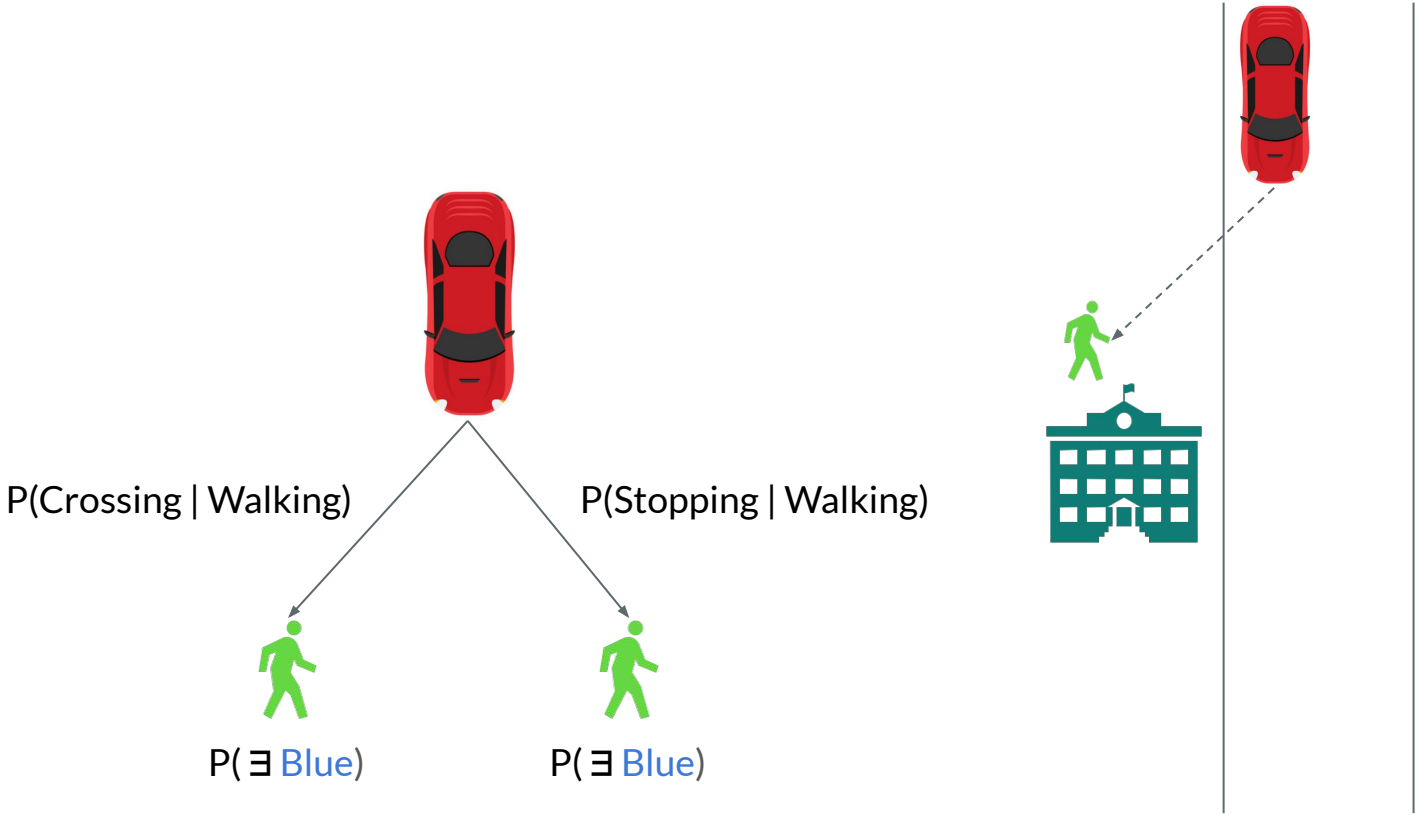

Interactive Partially Observable Markov Decision Processes (I-POMDP)

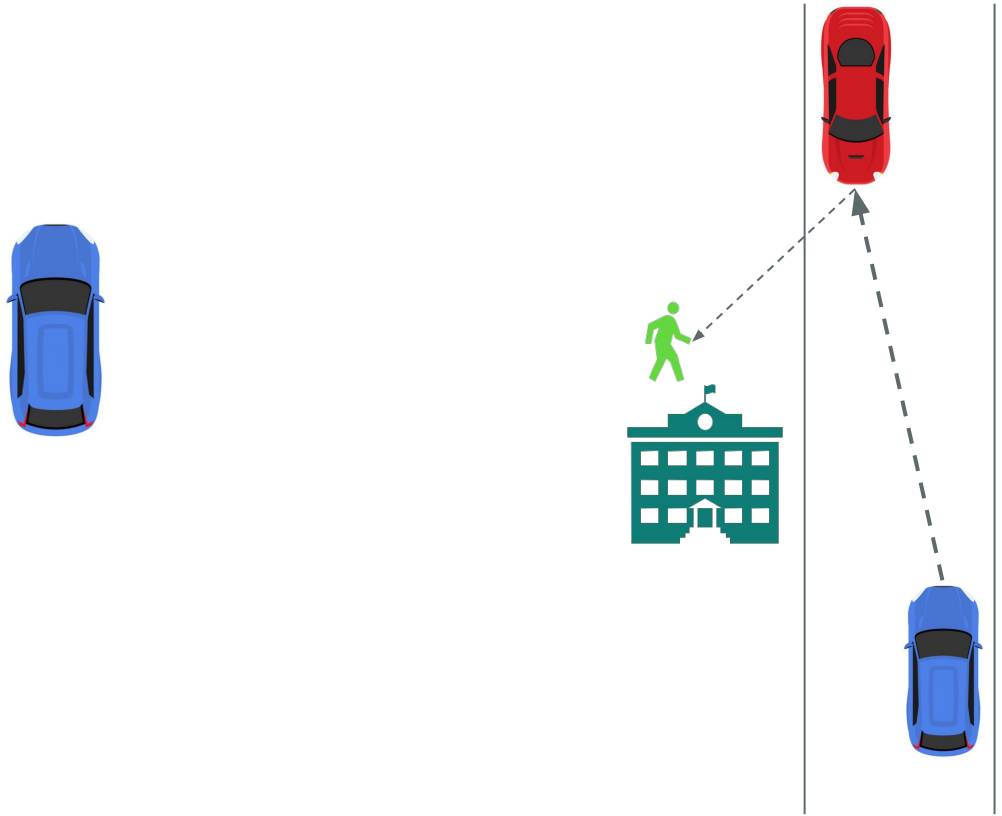

Interactive Partially Observable Markov Decision Processes (I-POMDP)

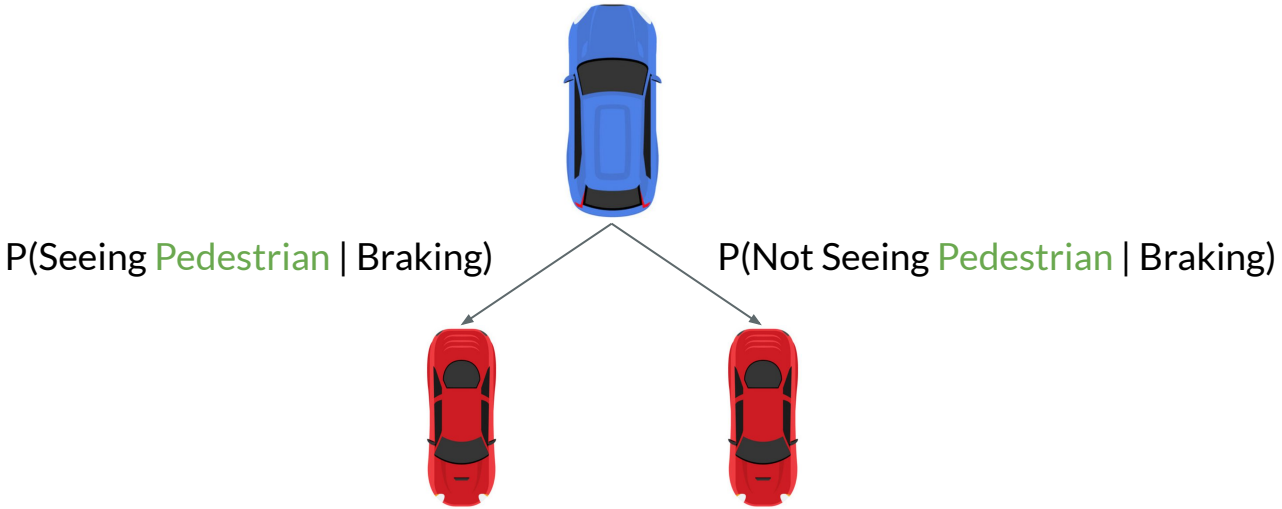

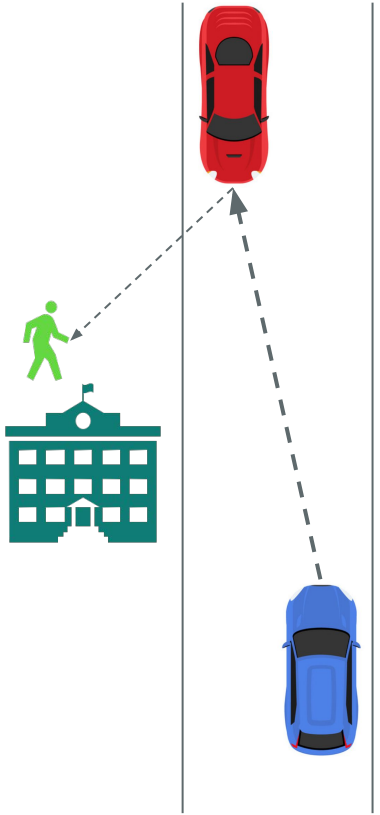

Interactive Partially Observable Markov Decision Processes (I-POMDP)

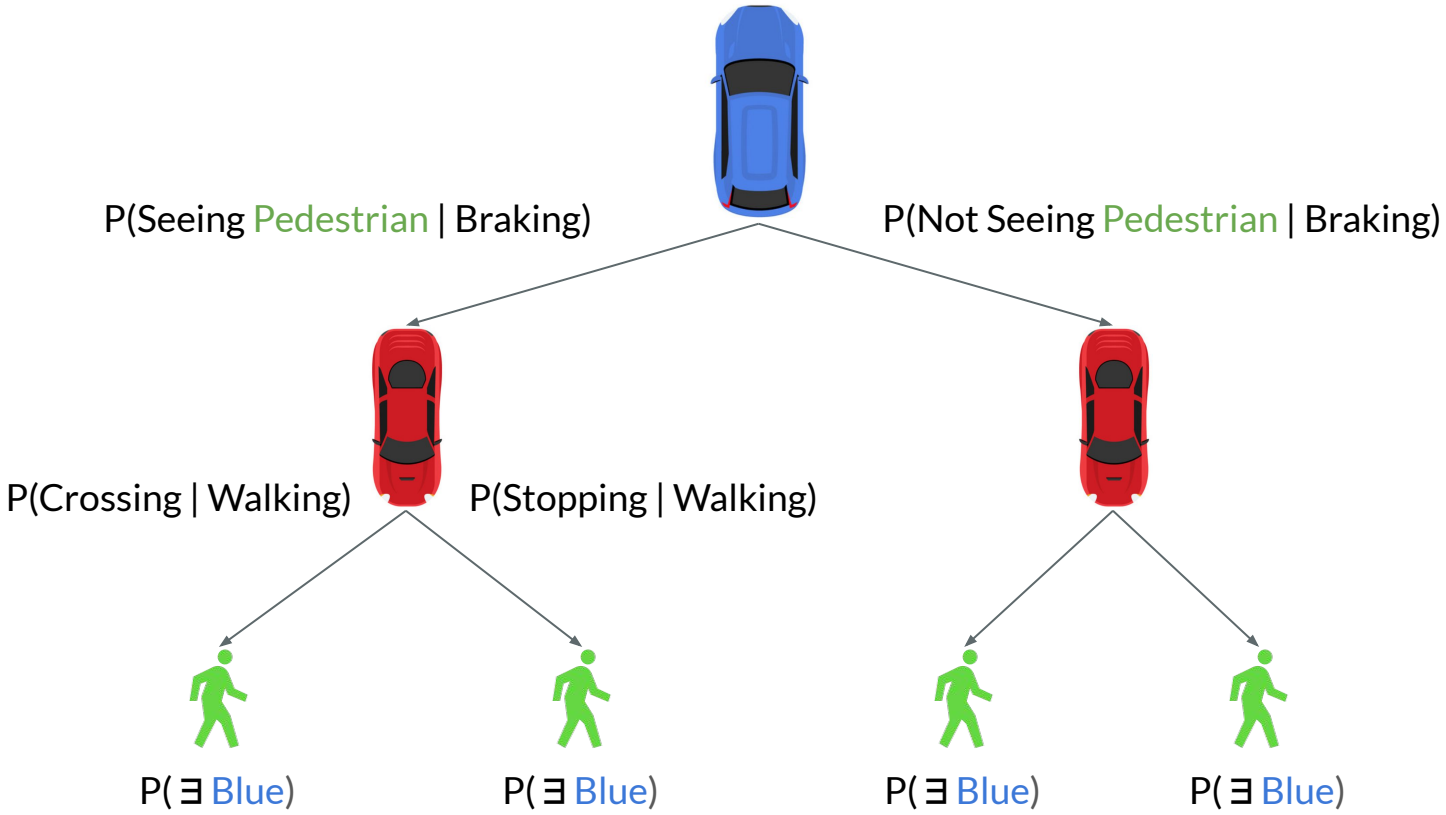

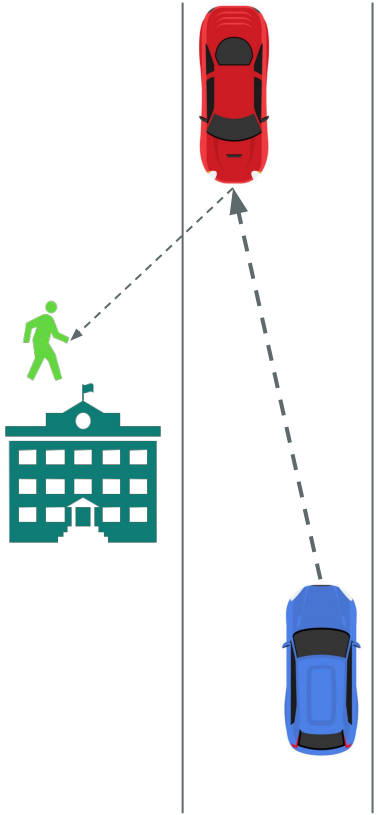

Background - Interactive Partially Observable Markov Decision Processes

State

Background - Interactive Partially Observable Markov Decision Processes

Background - Interactive Partially Observable Markov Decision Processes

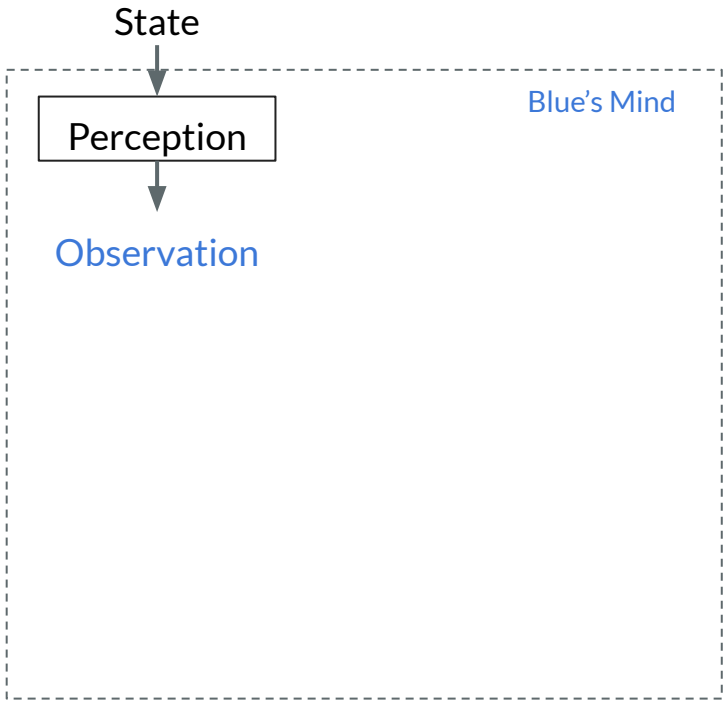

Background - Interactive Partially Observable Markov Decision Processes

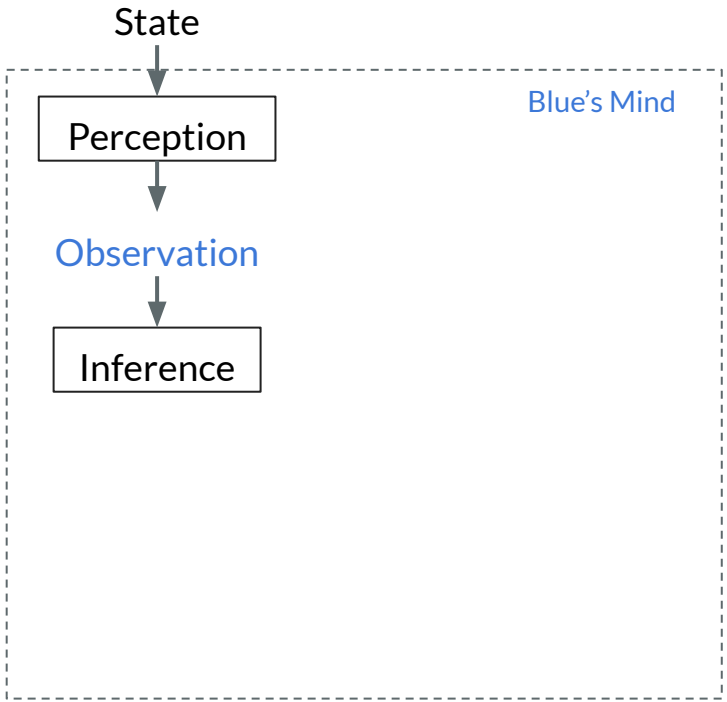

Background - Interactive Partially Observable Markov Decision Processes

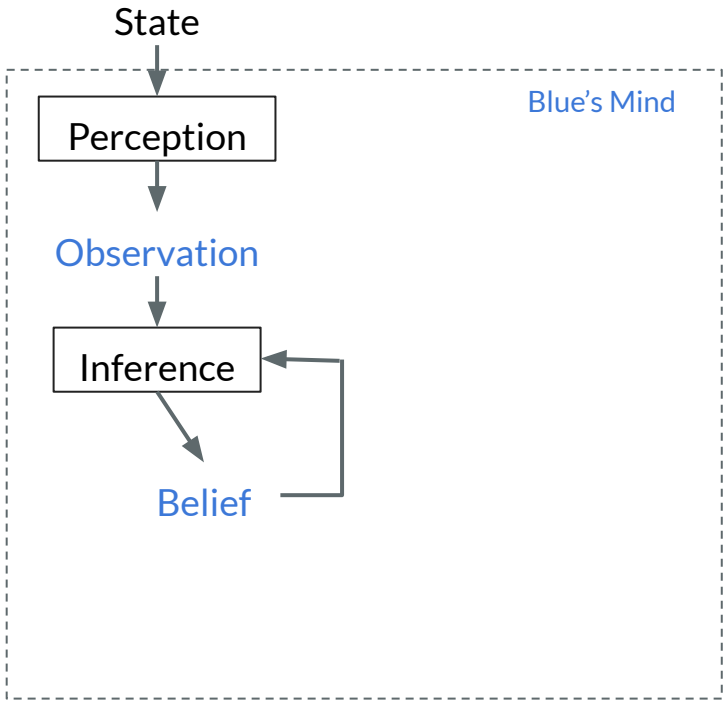

Background - Interactive Partially Observable Markov Decision Processes

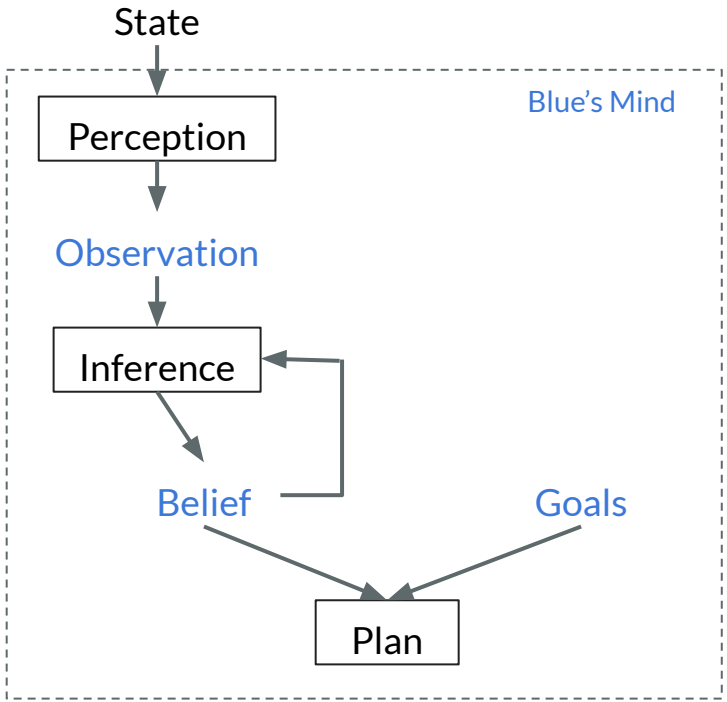

Background - Interactive Partially Observable Markov Decision Processes

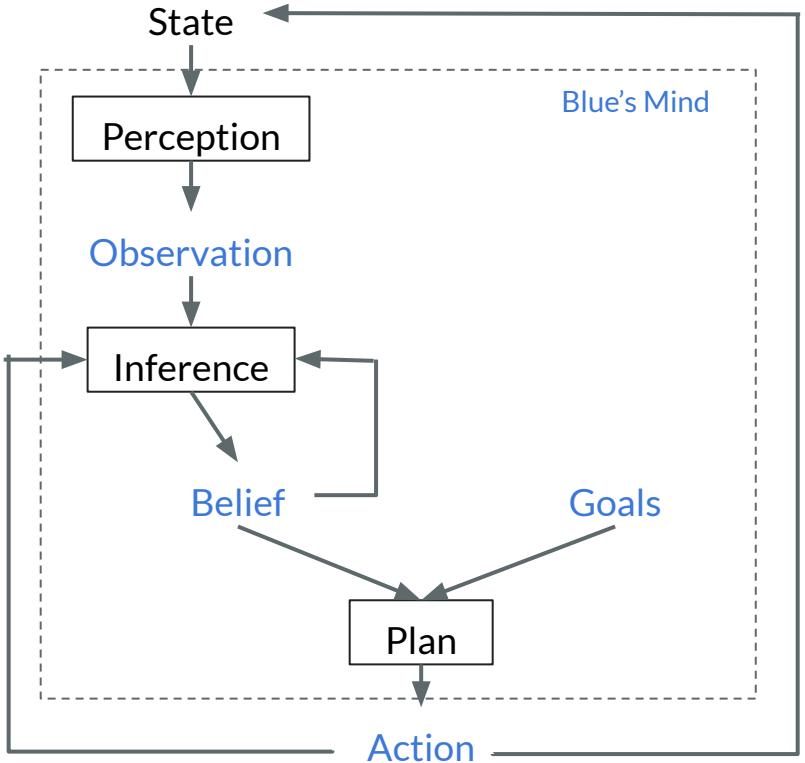

Background - Interactive Partially Observable Markov Decision Processes

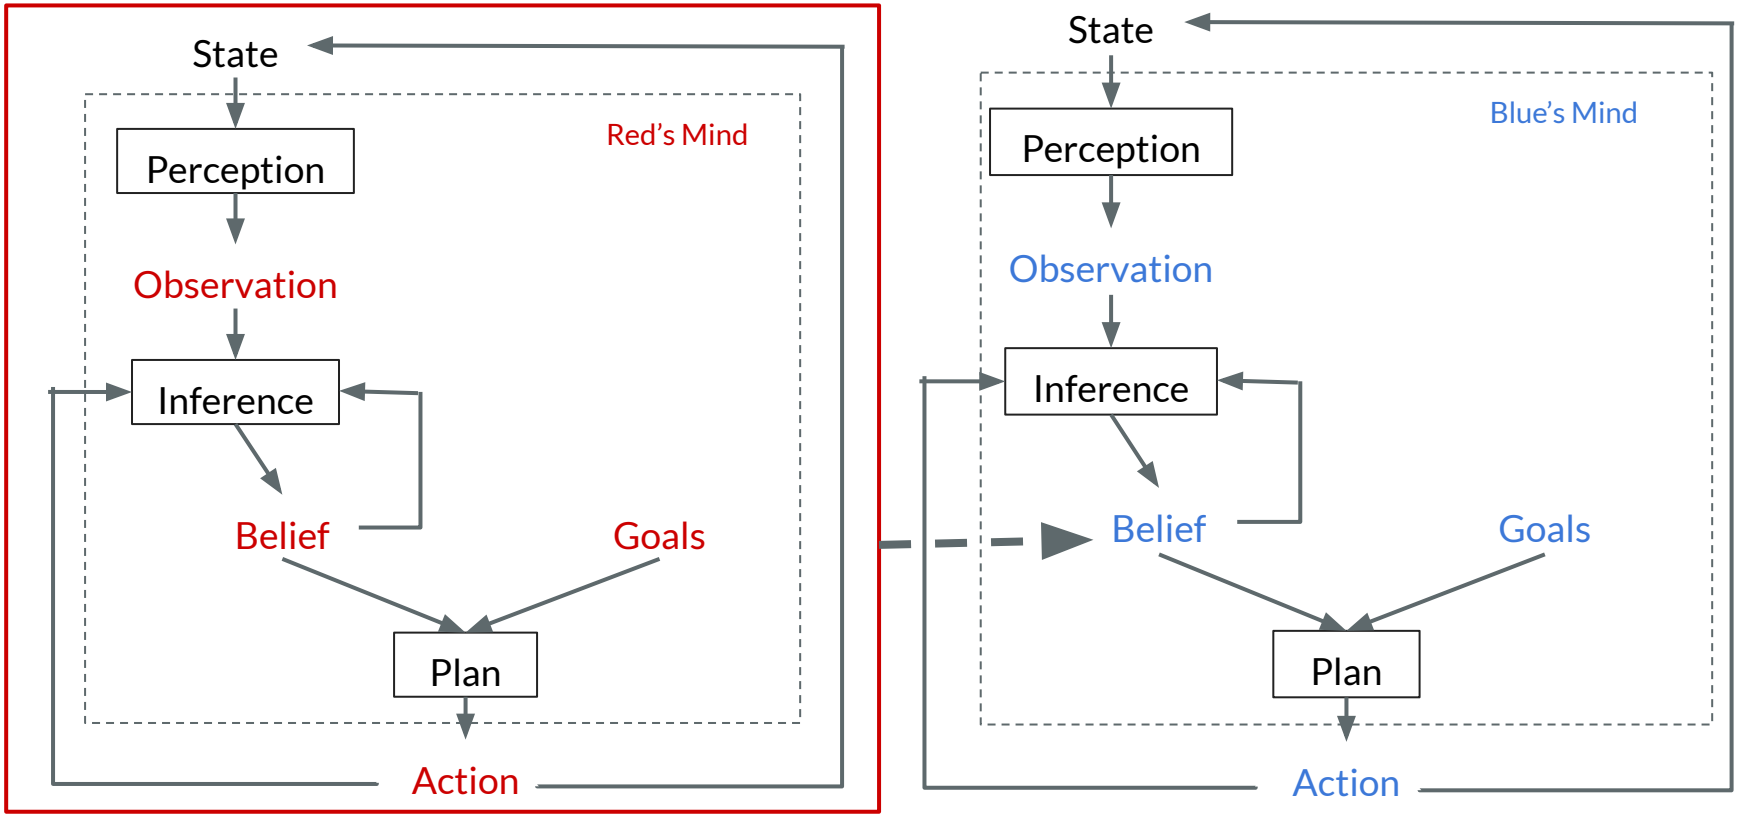

Our Approach: Amortize Lower Levels and Update Beliefs through Importance Sampling

Our Approach – Amortized Inference

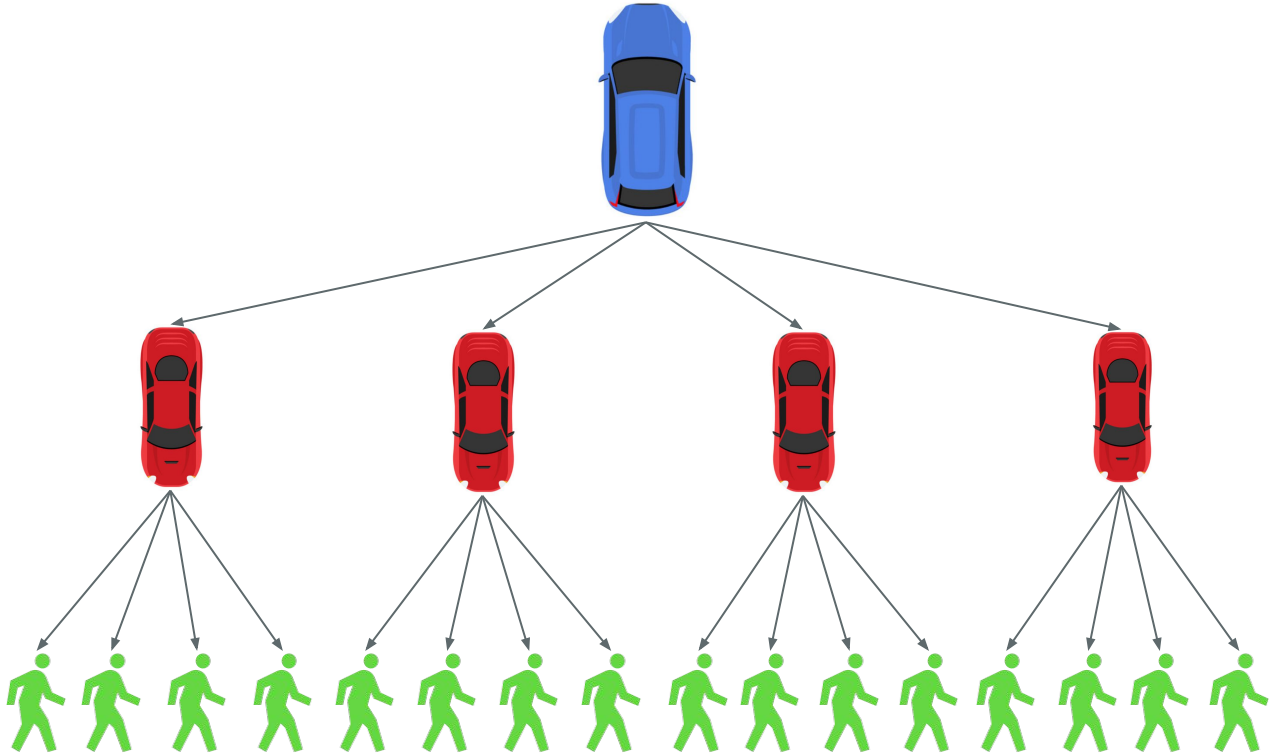

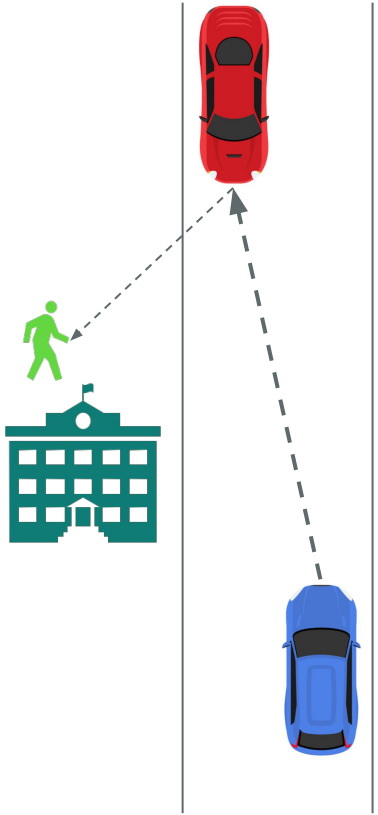

Our Approach – Amortized Inference

Learn Distribution of Pedestrian goals called NN0

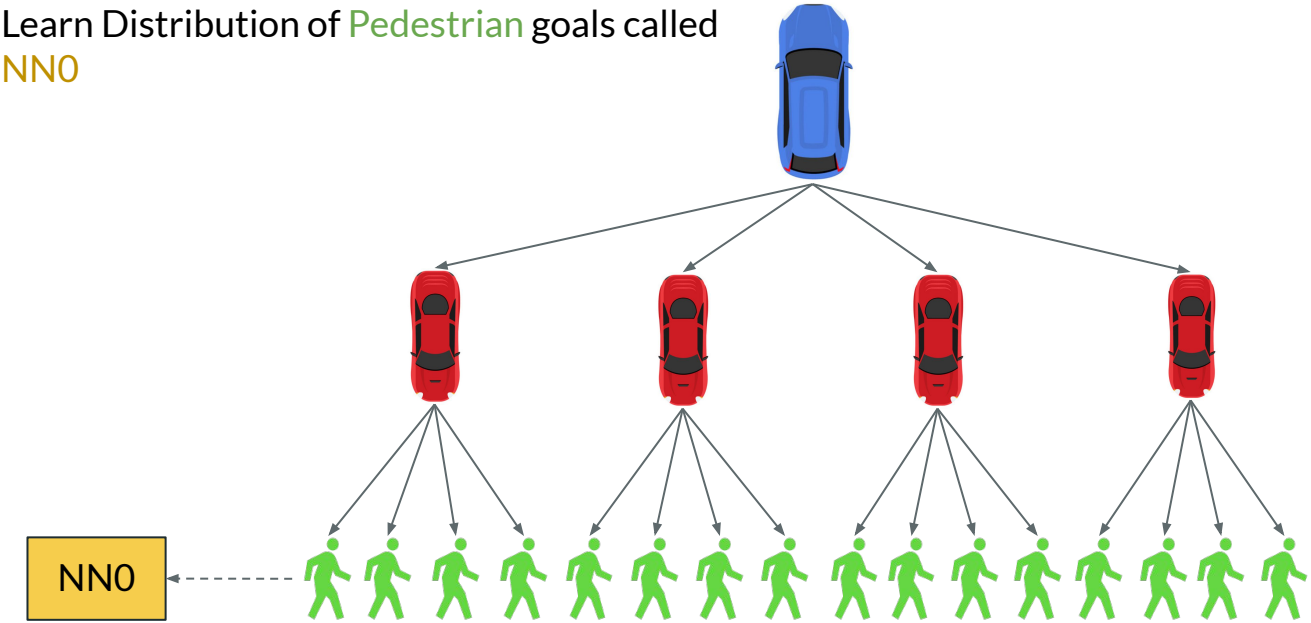

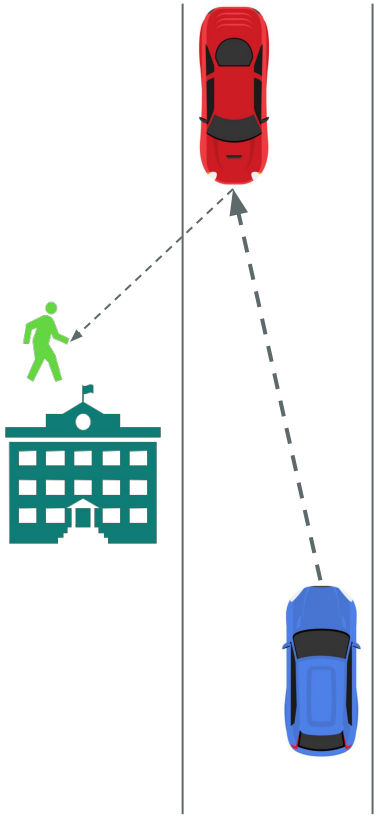

Our Approach – Amortized Inference

Sample most likely Pedestrian goals using Learned Distribution from NN0

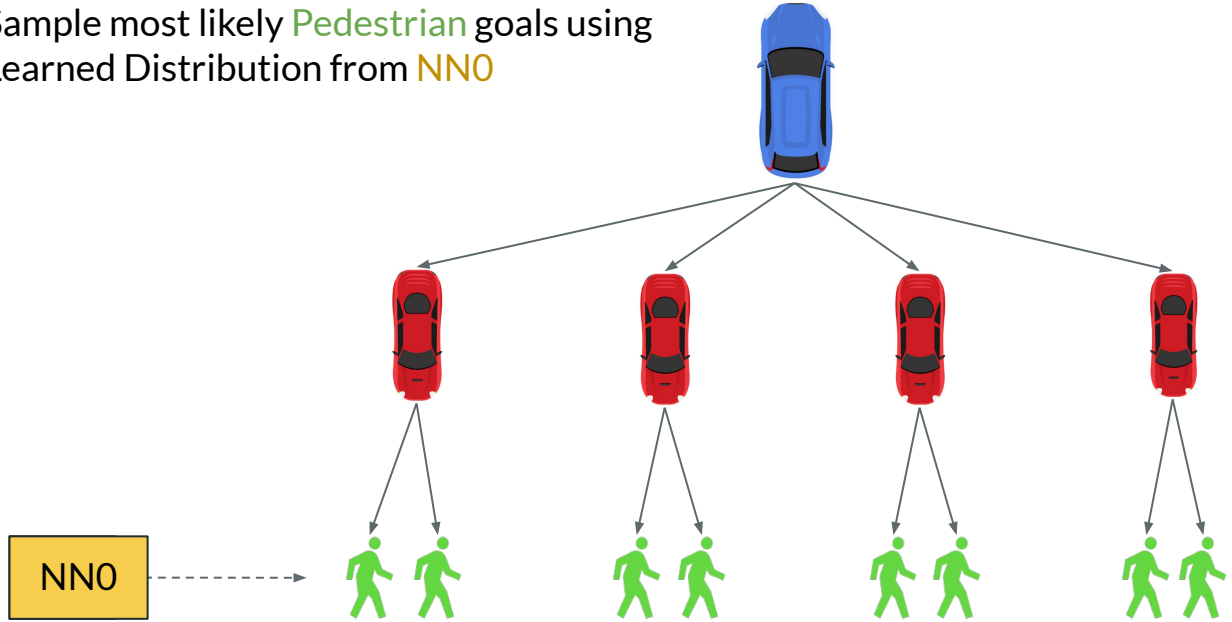

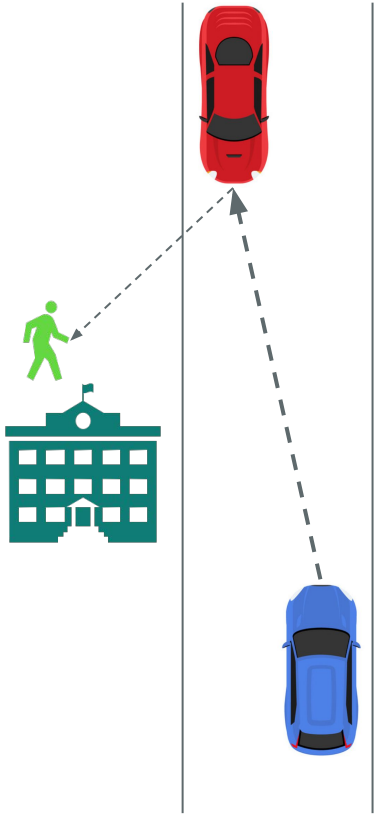

Our Approach – Amortized Inference

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

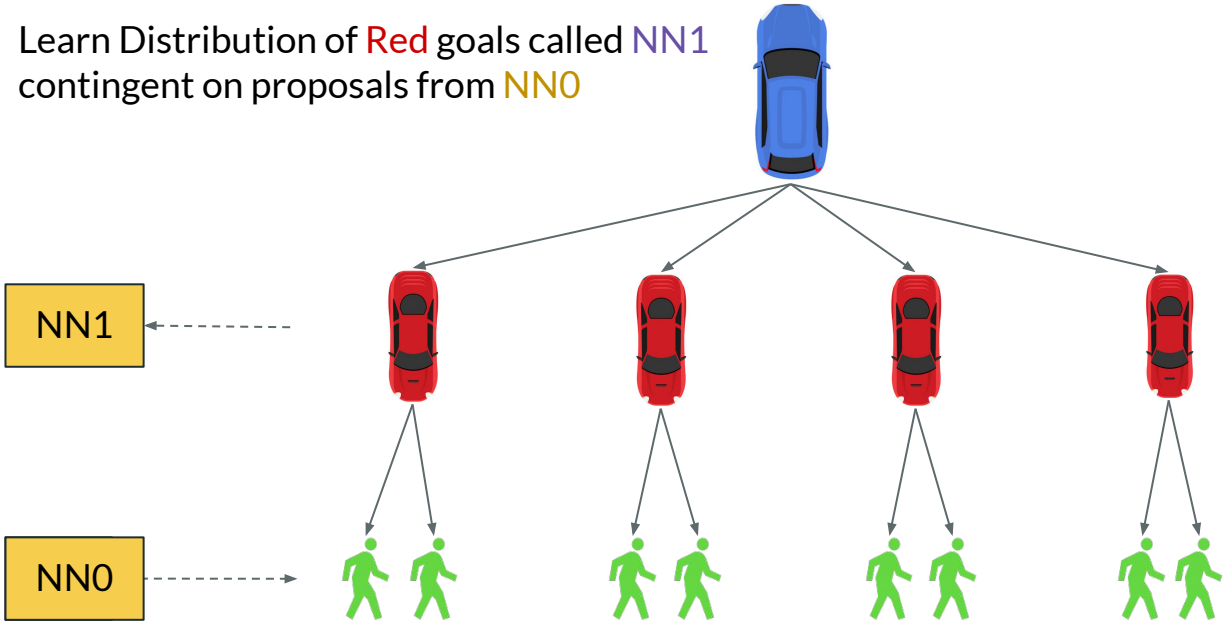

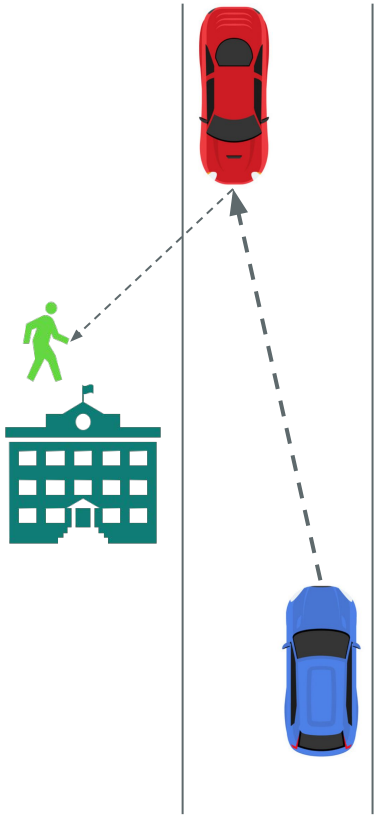

Our Approach – Amortized Inference

Sample most likely Red goals using Learned Distribution from NN1

NN1

NN0

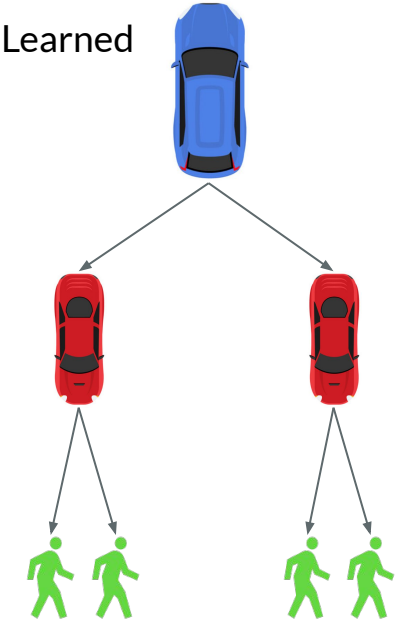

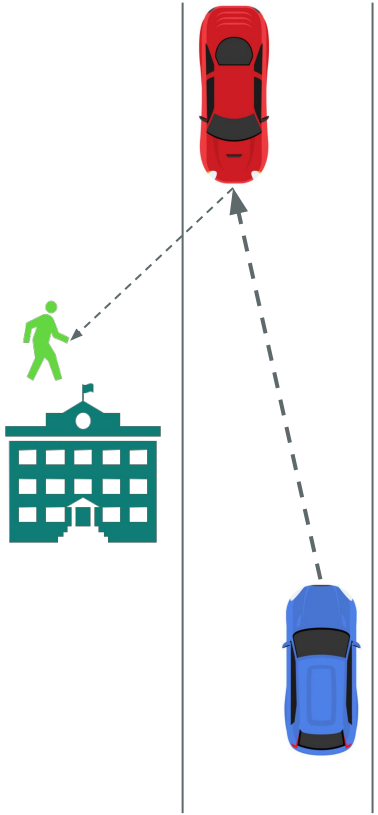

Our Approach – Building a Dataset for Recursive Reasoning

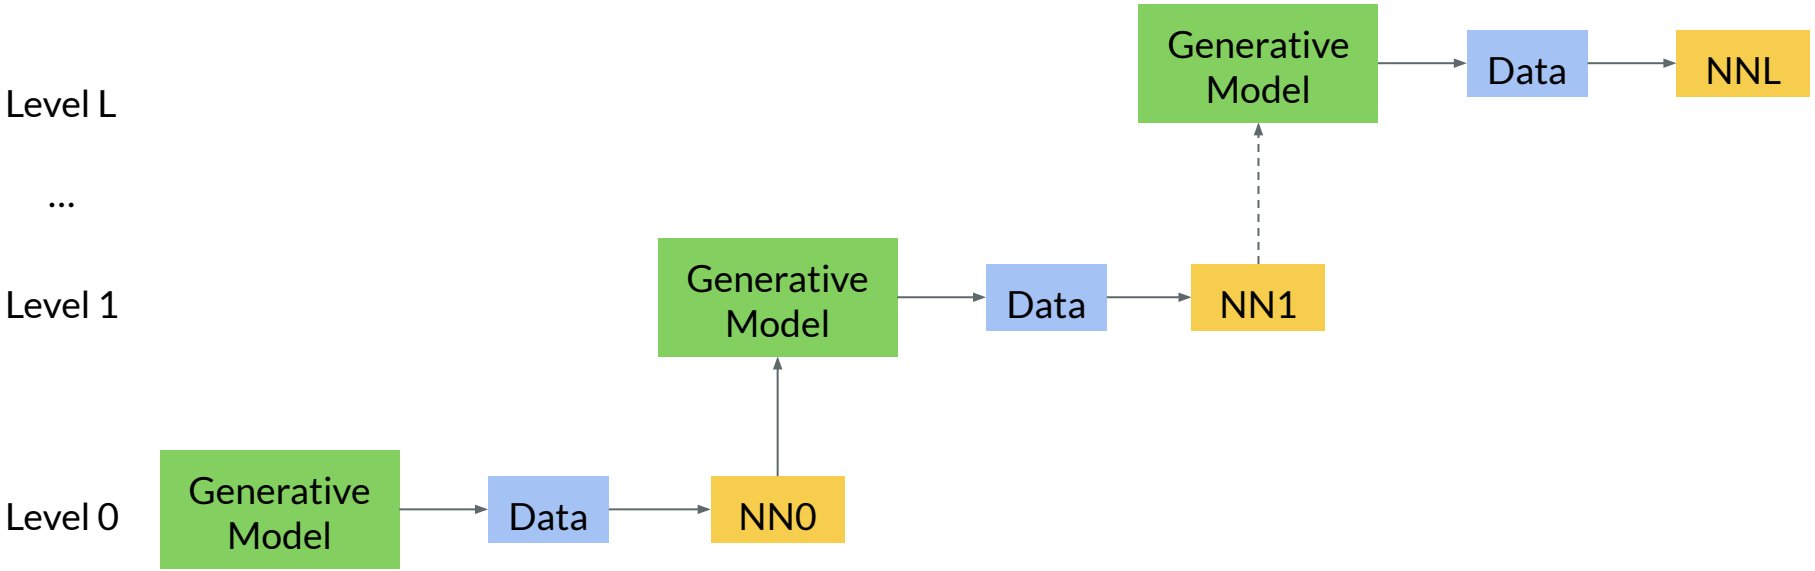

Experiment - Setup

- 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 **22 times larger** than prior work

Alice

Bob

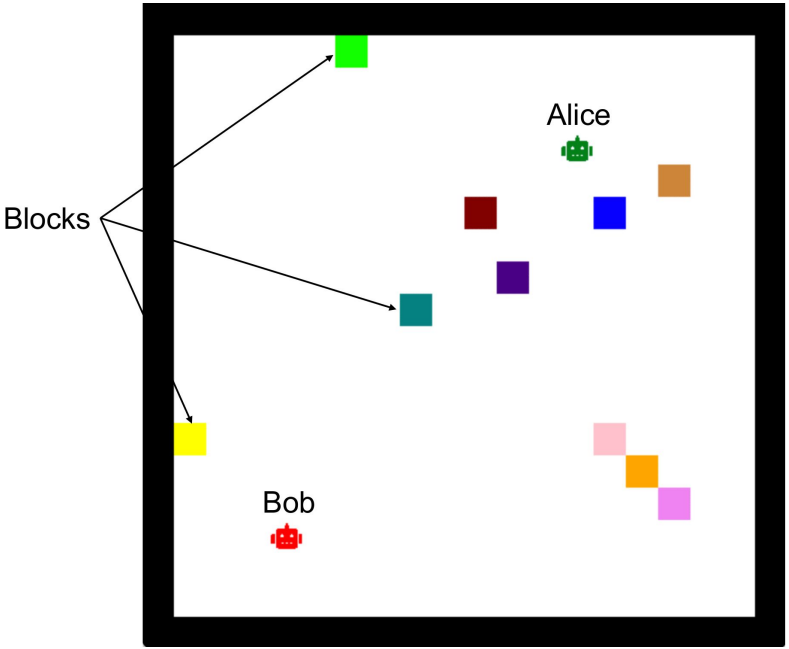

Experiment - Data Generation

- Procedurally Generate Data by sampling *random States and Goals*

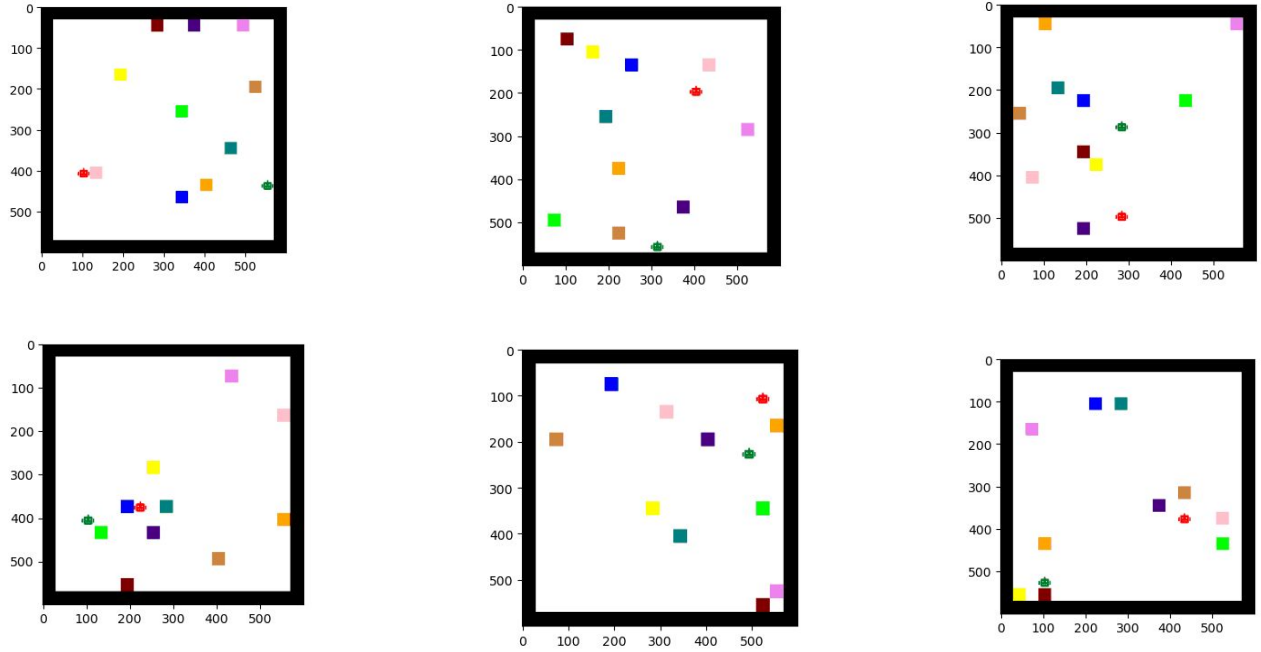

Experiment - Results on Larger Dataset

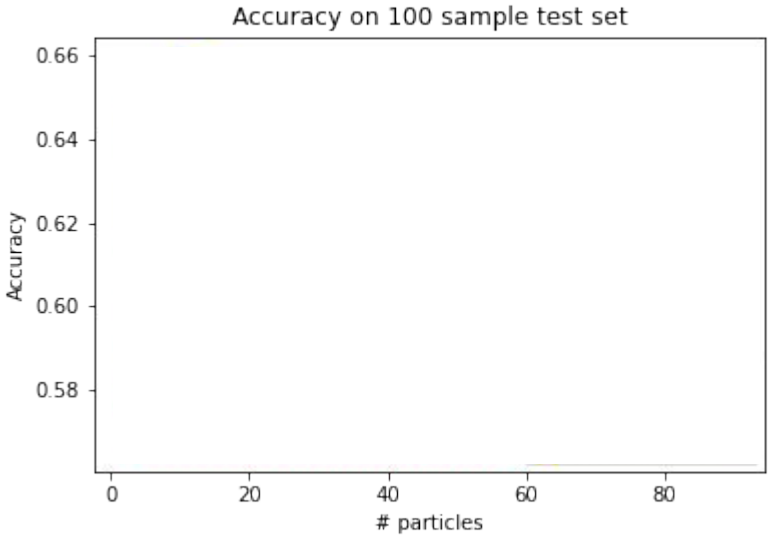

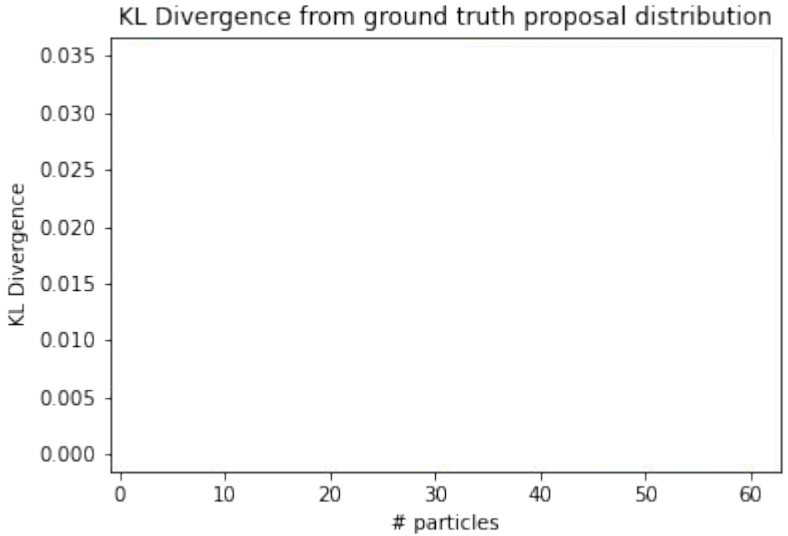

Experiment - Results on Larger Dataset

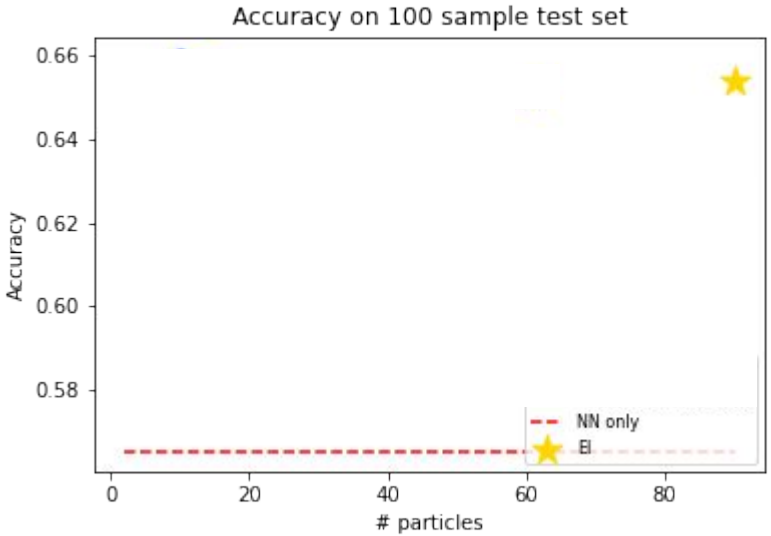

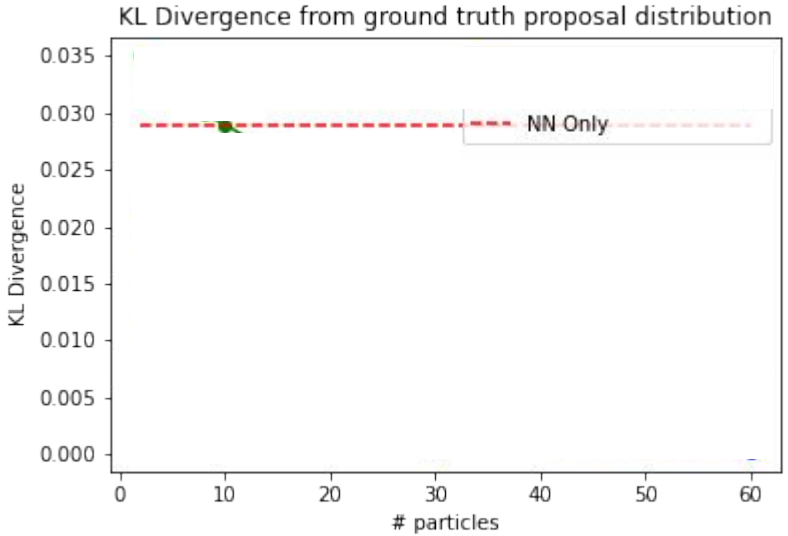

Experiment - Results on Larger Dataset

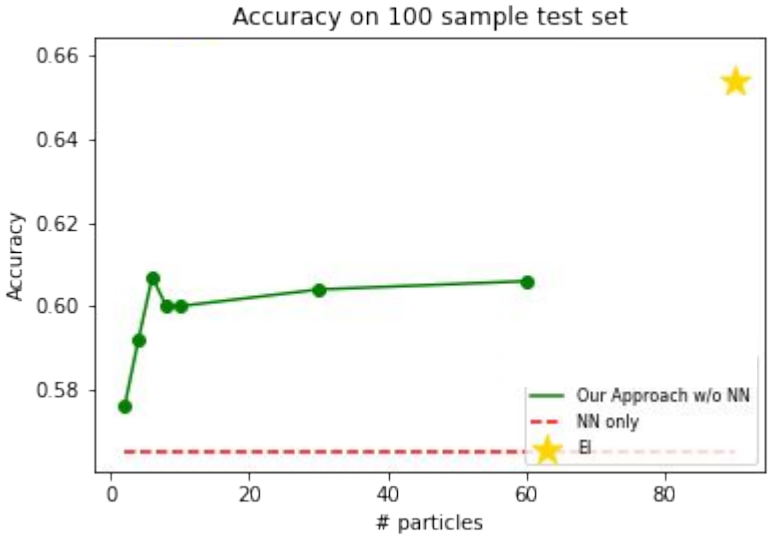

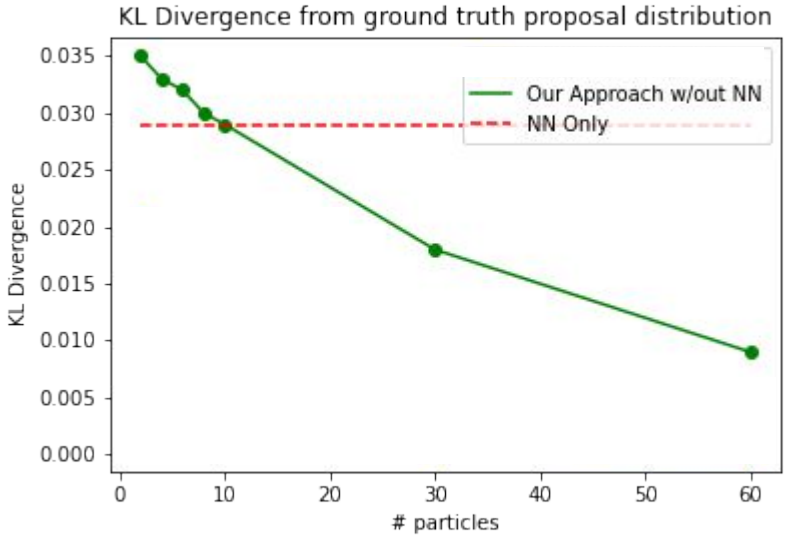

Experiment - Results on Larger Dataset

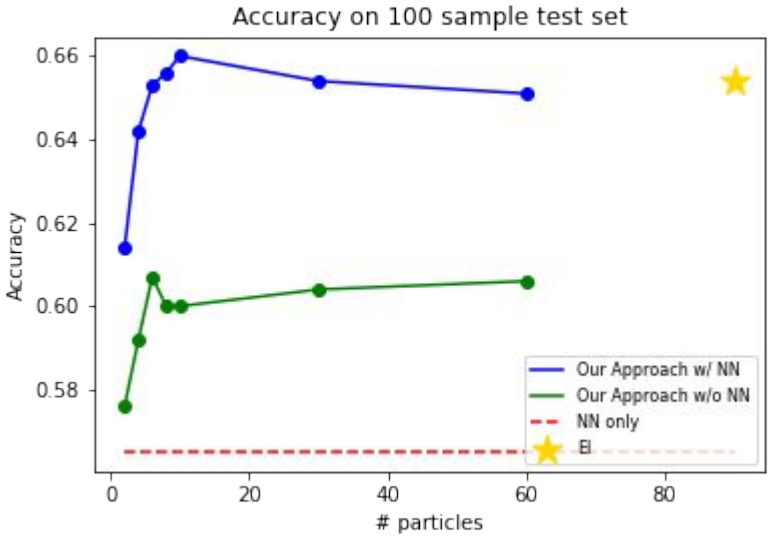

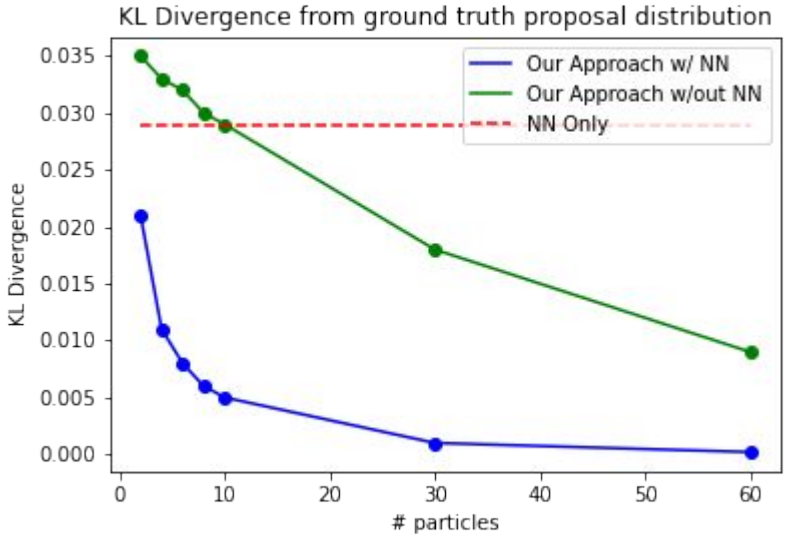

Experiment - Setup

- 3 Agents, Alice (Green), Bob (Red), Charlie (Blue)
  - All cars are randomly assigned to turn left, turn right, or continue moving straight
- No agent knows which direction others are turning
- At each timestep, we ask the model to predict the actions of every car
  - Action space is Accelerate, Brake, Rotate Left, Rotate Right, Signal Danger
- Goal Space **18 times larger** than prior work

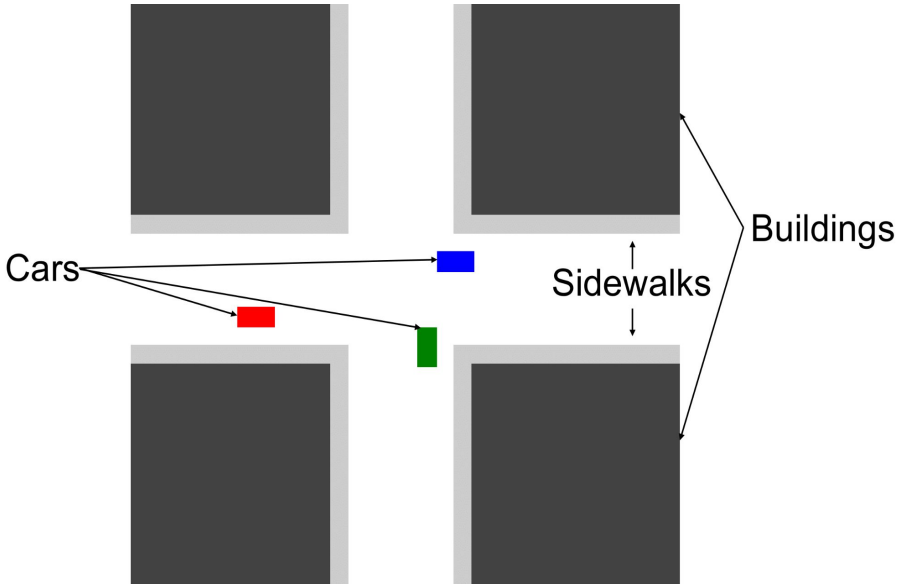

Experiment - Data Generation

- Procedurally Generate Data by sampling *random States and Goals*

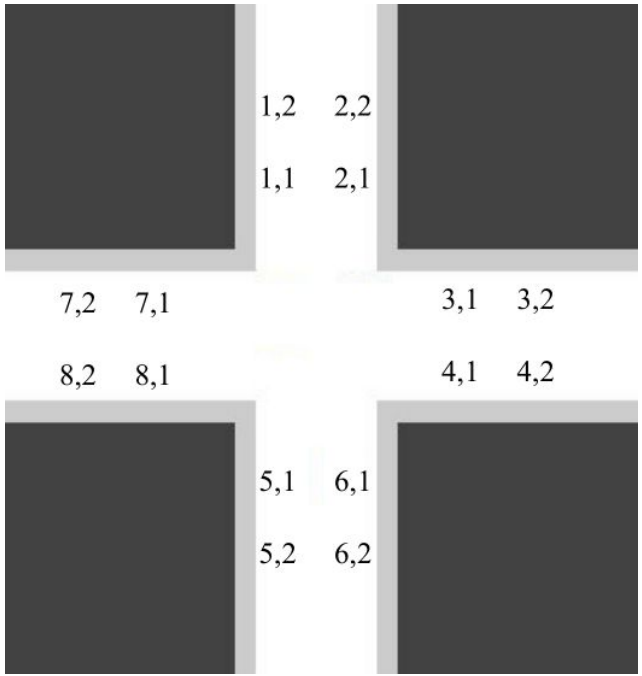

Experiment - Results on Large Dataset

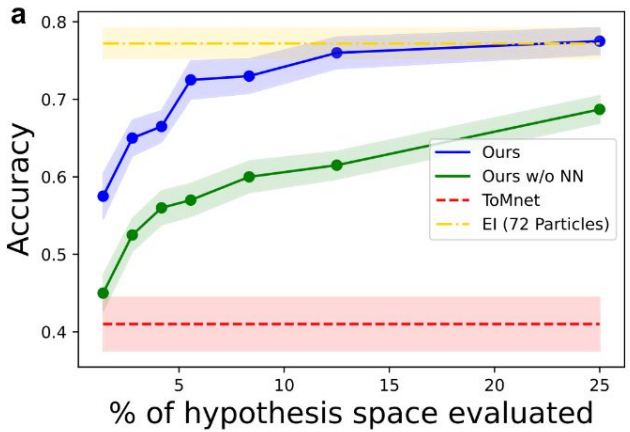

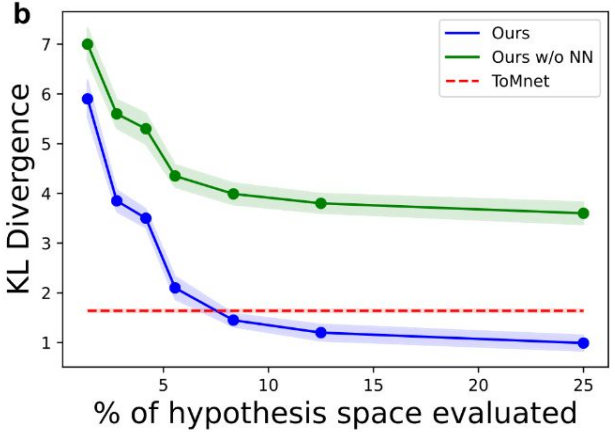

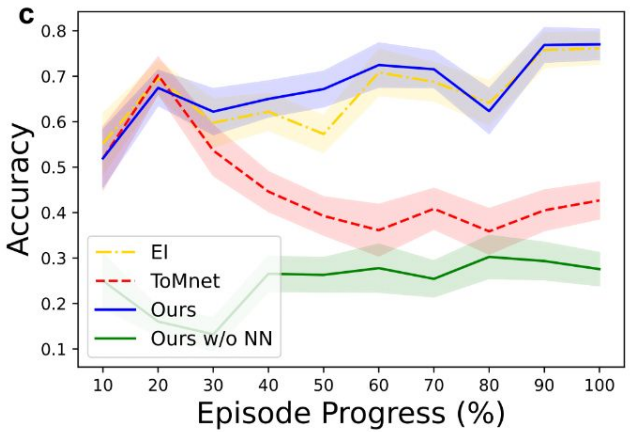

Experiment - Generalization Results on Large Dataset

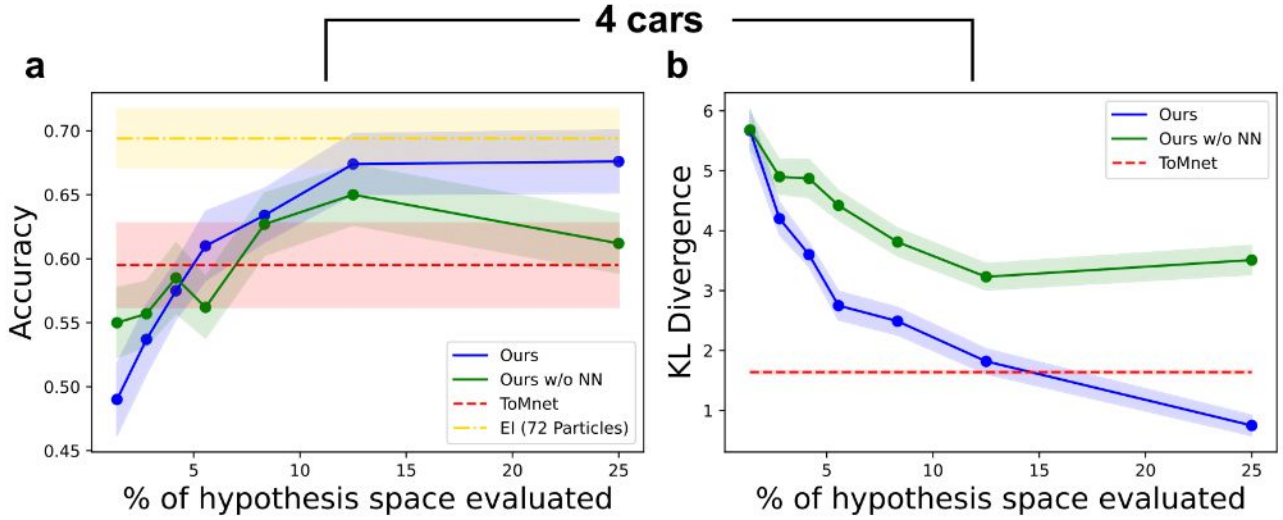

Experiment - Generalization Results on Large Dataset

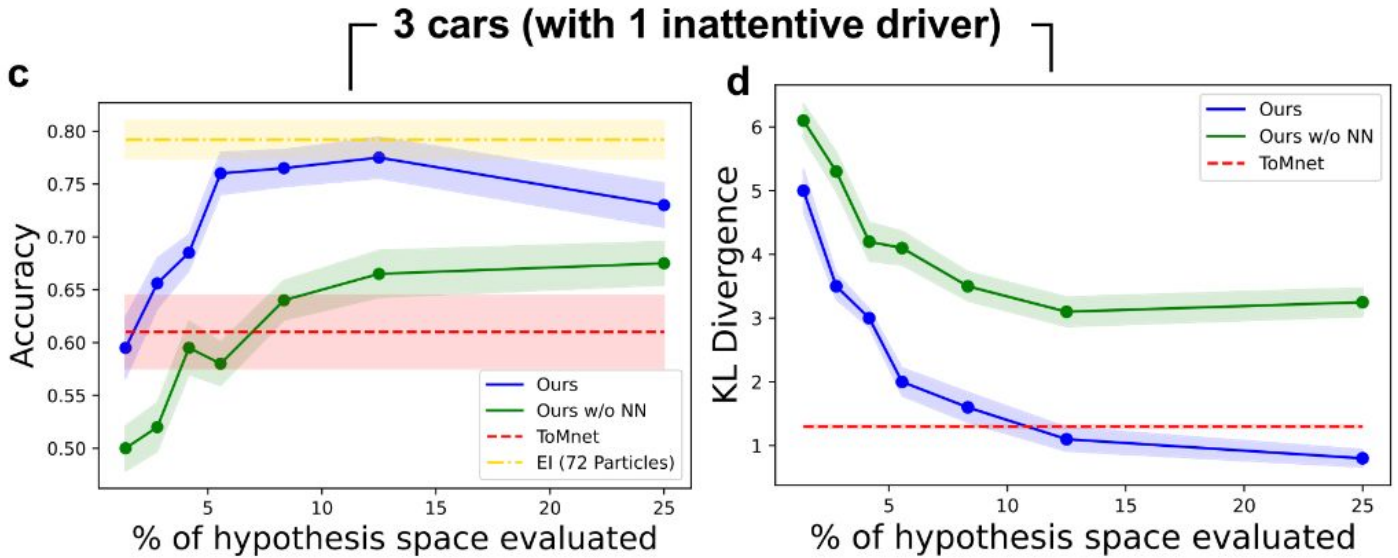

Conclusion

- Neural amortized inference **drastically accelerates** nested social reasoning while preserving robustness and uncertainty estimation
- This approach has been demonstrated successfully on *goal* inference **and** *physical state* inference

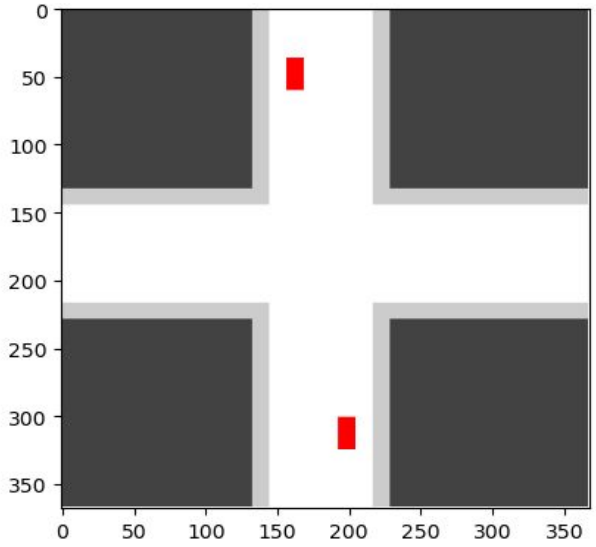

Code

https://github.com/KJha02/AmortizedNestedReasoning