

# Two Body Problem: Collaborative Visual Task Completion CVPR 2019

<https://prior.allenai.org/projects/two-body-problem>

Speaker: Unnat Jain, UIUC [[webpage](#)]




# Collaborations




Adam (lou-Jen) Liu  
UIUC



Alex Schwing  
UIUC



Ali Farhadi  
AI2, UW, XNOR.ai



Ani Kembhavi  
AI2



Eric Kolve  
AI2



Luca Weihs  
AI2




Mohammad Rastegari  
AI2, XNOR.ai



Svetlana Lazebnik  
UIUC



# Typical visual tasks




Object detection  
[COCO]




What is the mustache  
made of?

Question Answering  
VQA

# Visual embodied tasks



# Visual embodied tasks





We live in a *multi-agent* world

# Multi-agent world



# Multi-agent world






# Multi-agent world



# Multi-agent virtual world: Existing work



# Single-agent virtual world: Existing work



# Multi-agent virtual world: Ours



# Task



Agent 1



Agent 2

# Task



Agent 1



Agent 2

+ Navigate to TV

# Task



Agent 1



Agent 2

+ Navigate to TV

# Collaborative task



Agent 1




Agent 2

+ Navigate to TV

+ Collaboration  
[Joint pickup]



# Find and Lift Furniture



Top-down view (for illustration)

**Agent 1** and **Agent 2**

1. Navigate to TV
2. Collaborative pickup

**Agent 1** quickly finds it

**Agent 2** is on the wrong Side

Need for communication

# Task Definition

Step 1: Navigation

Step 2: Collaborative pickup


Three conditions for successful pickup, both agents need to:

1. Be near the TV and looking at it
2. Simultaneously execute *PICKUP*
3. Be on different sides of the TV

# Agent – Environment interaction

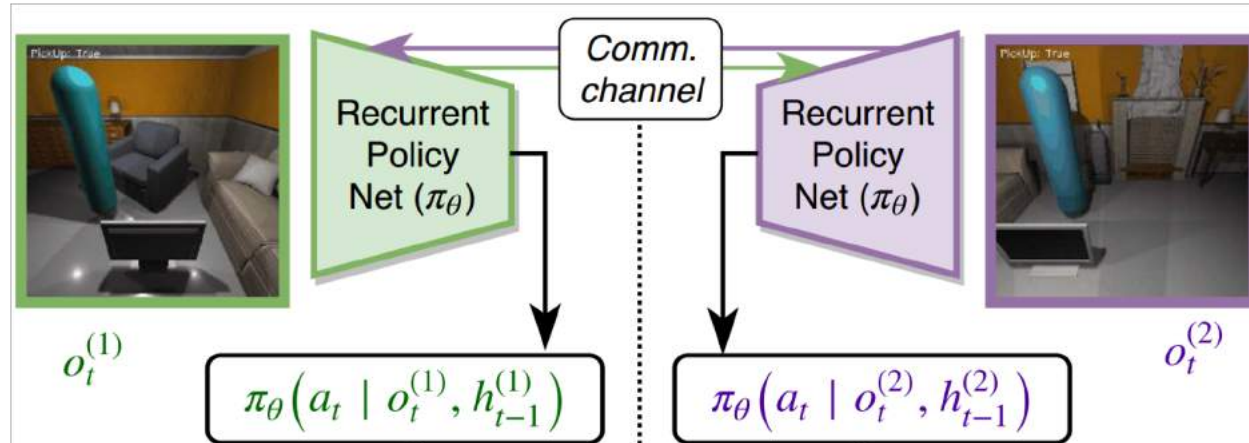


$o_t^{(1)}$

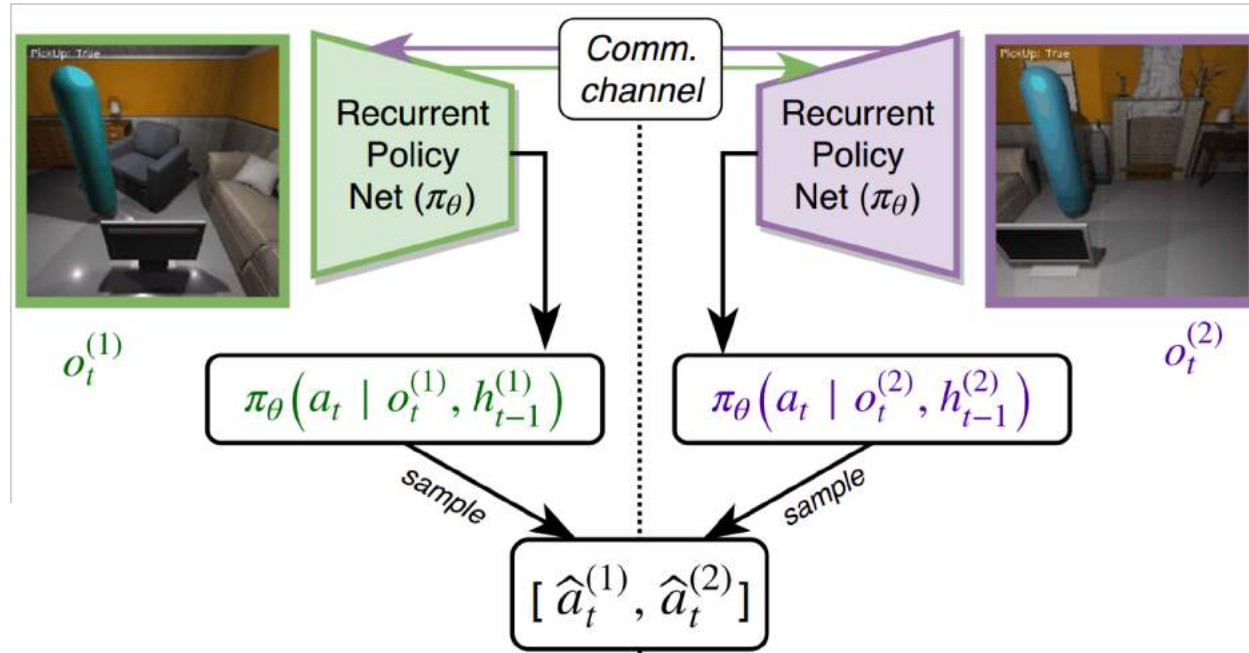


$o_t^{(2)}$

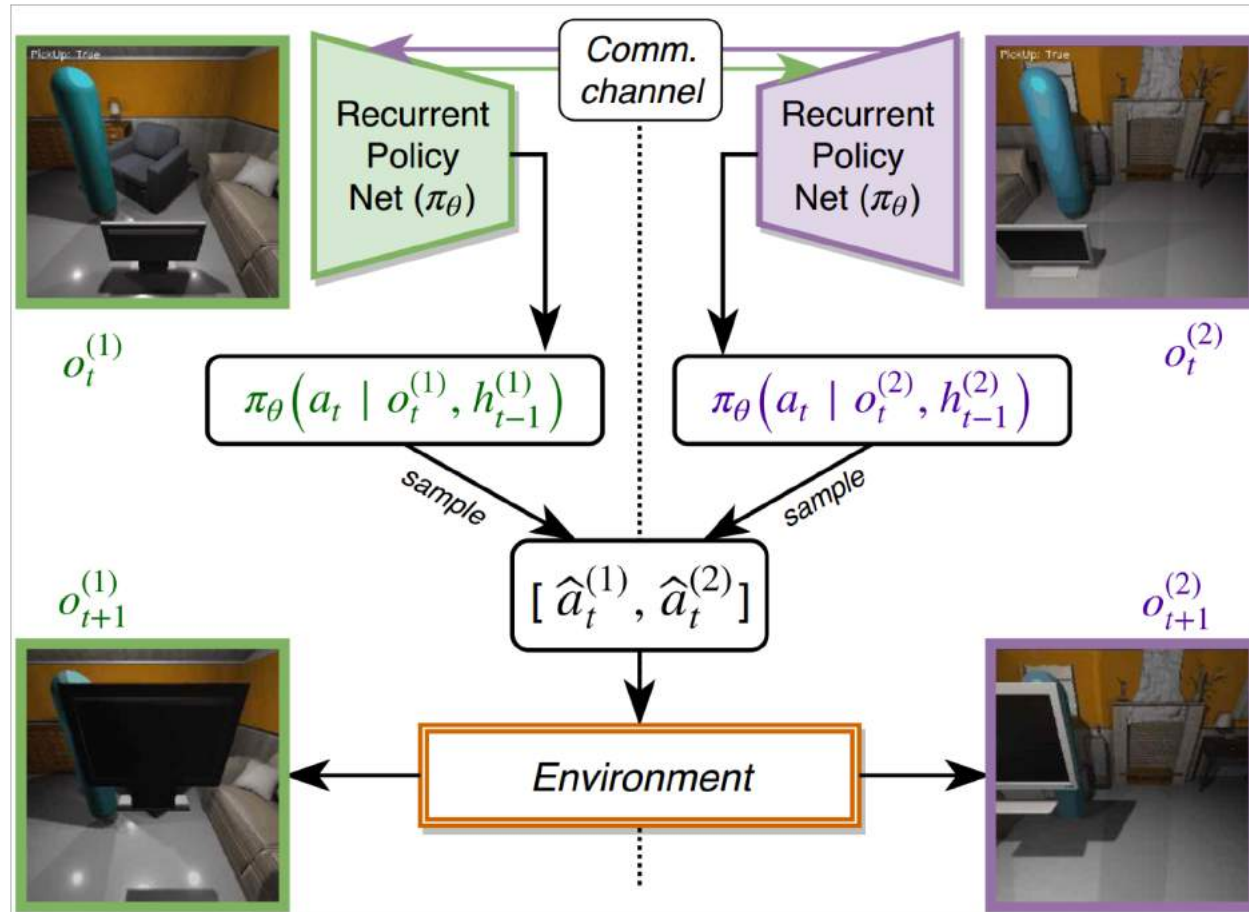
# Agent – Environment interaction



# Agent – Environment interaction




# Agent – Environment interaction



# Reward structure and Loss

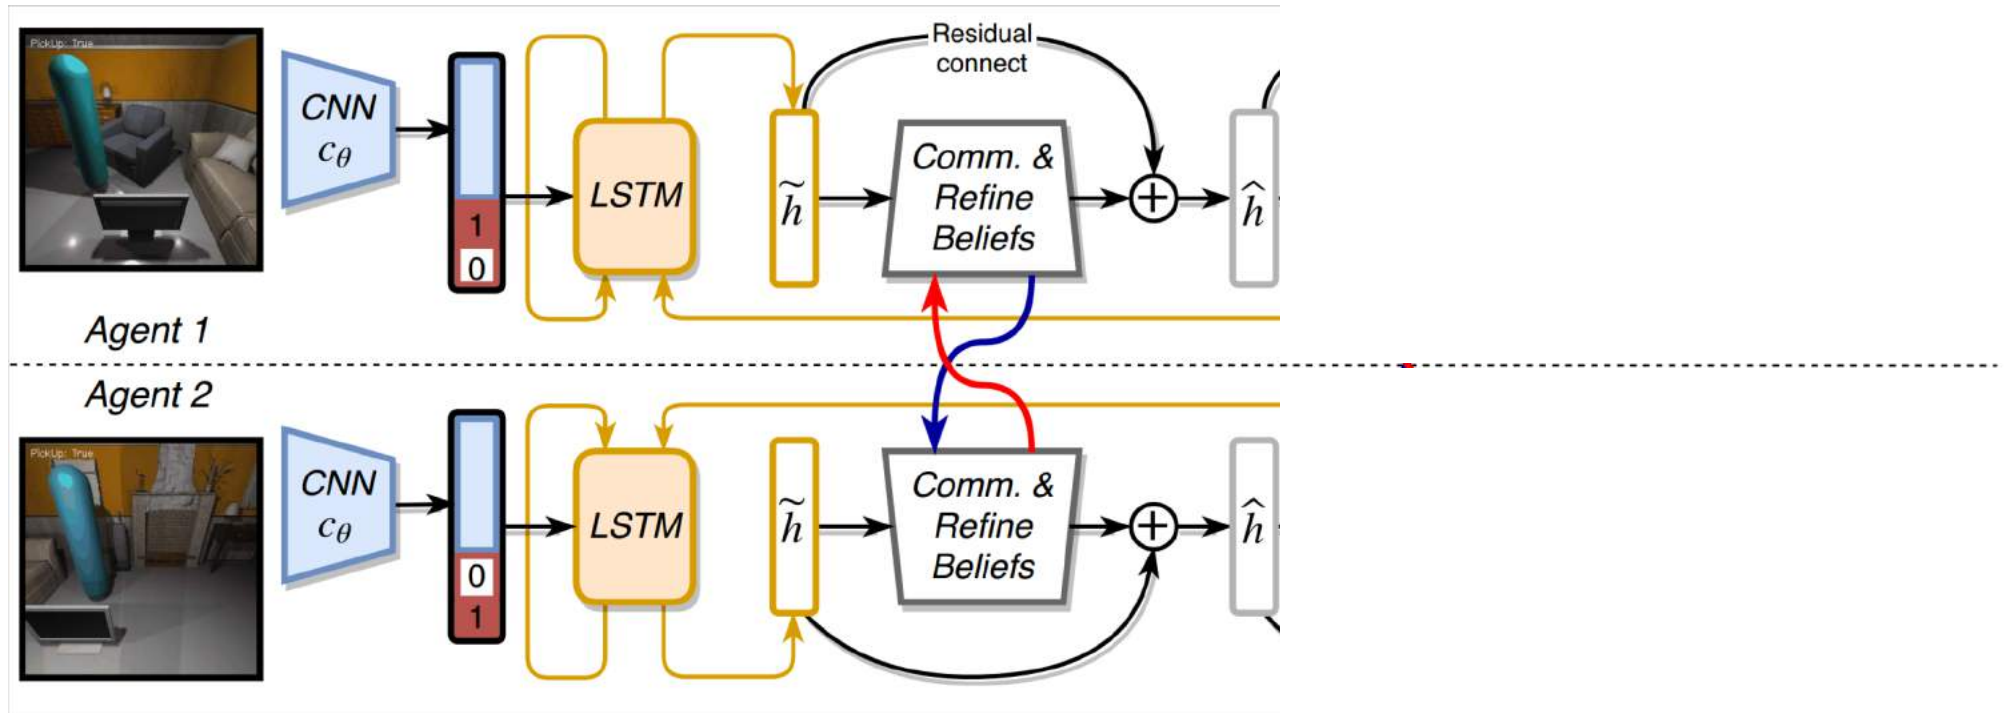
- Rewards
  - +1 for performing a successful joint pickup
  - 0.01 for every time step
  - 0.1 for a failed pickup action
  - 0.02 for any other failed action
- Asynchronous advantage actor critic
  - Value loss  $L = \sum (R - V(s))^2$
  - Policy loss  $L = -\log(\pi(s)) * A(s)$

# Two Body Network

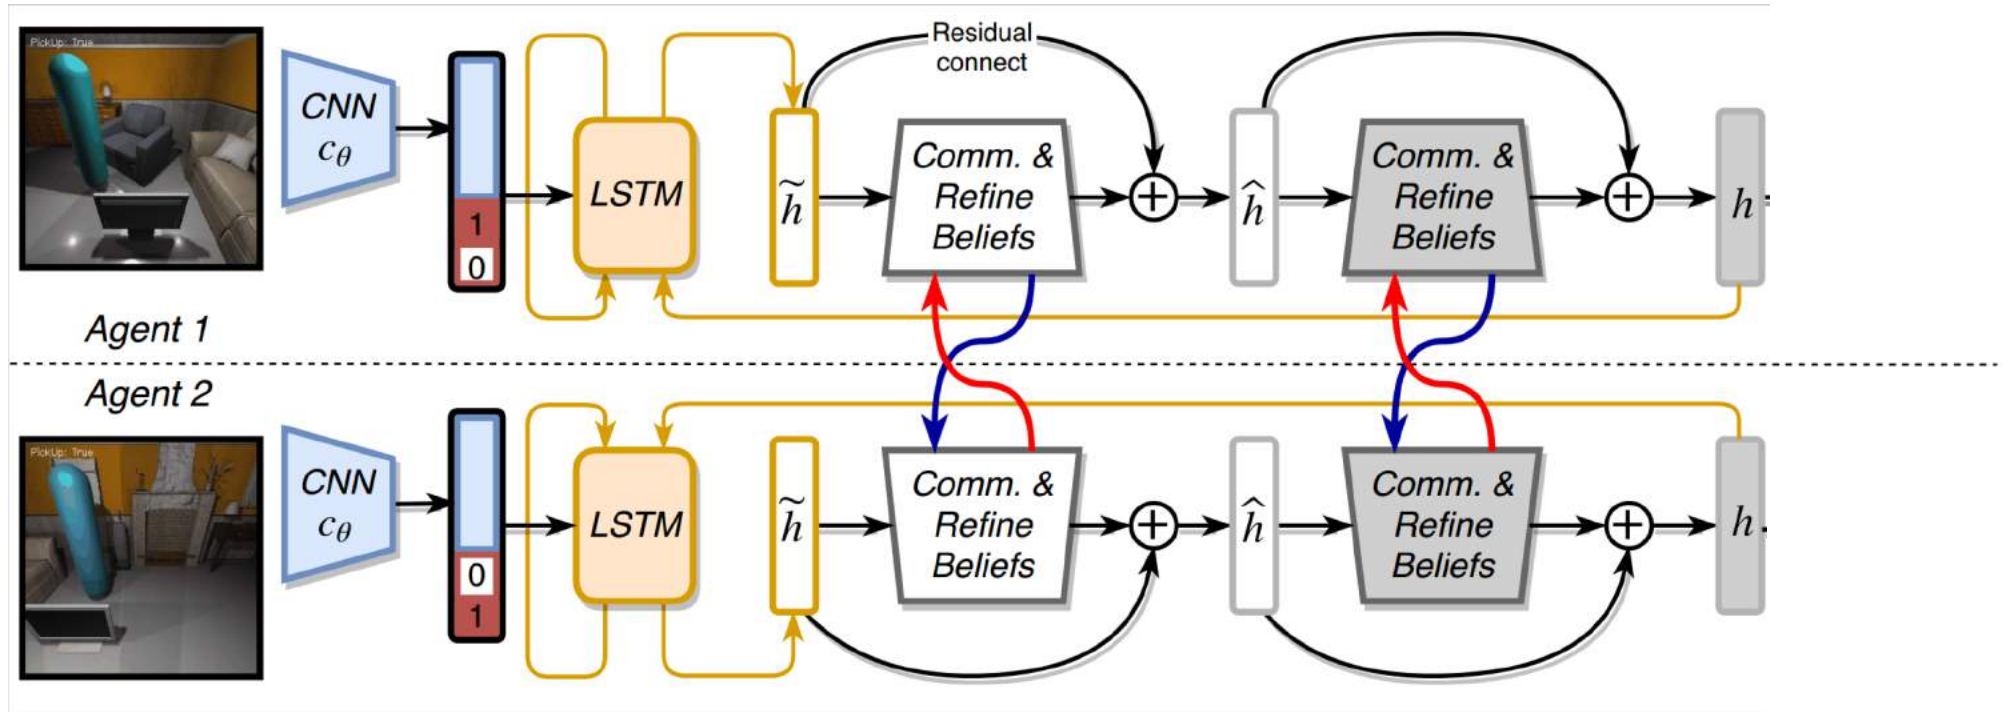




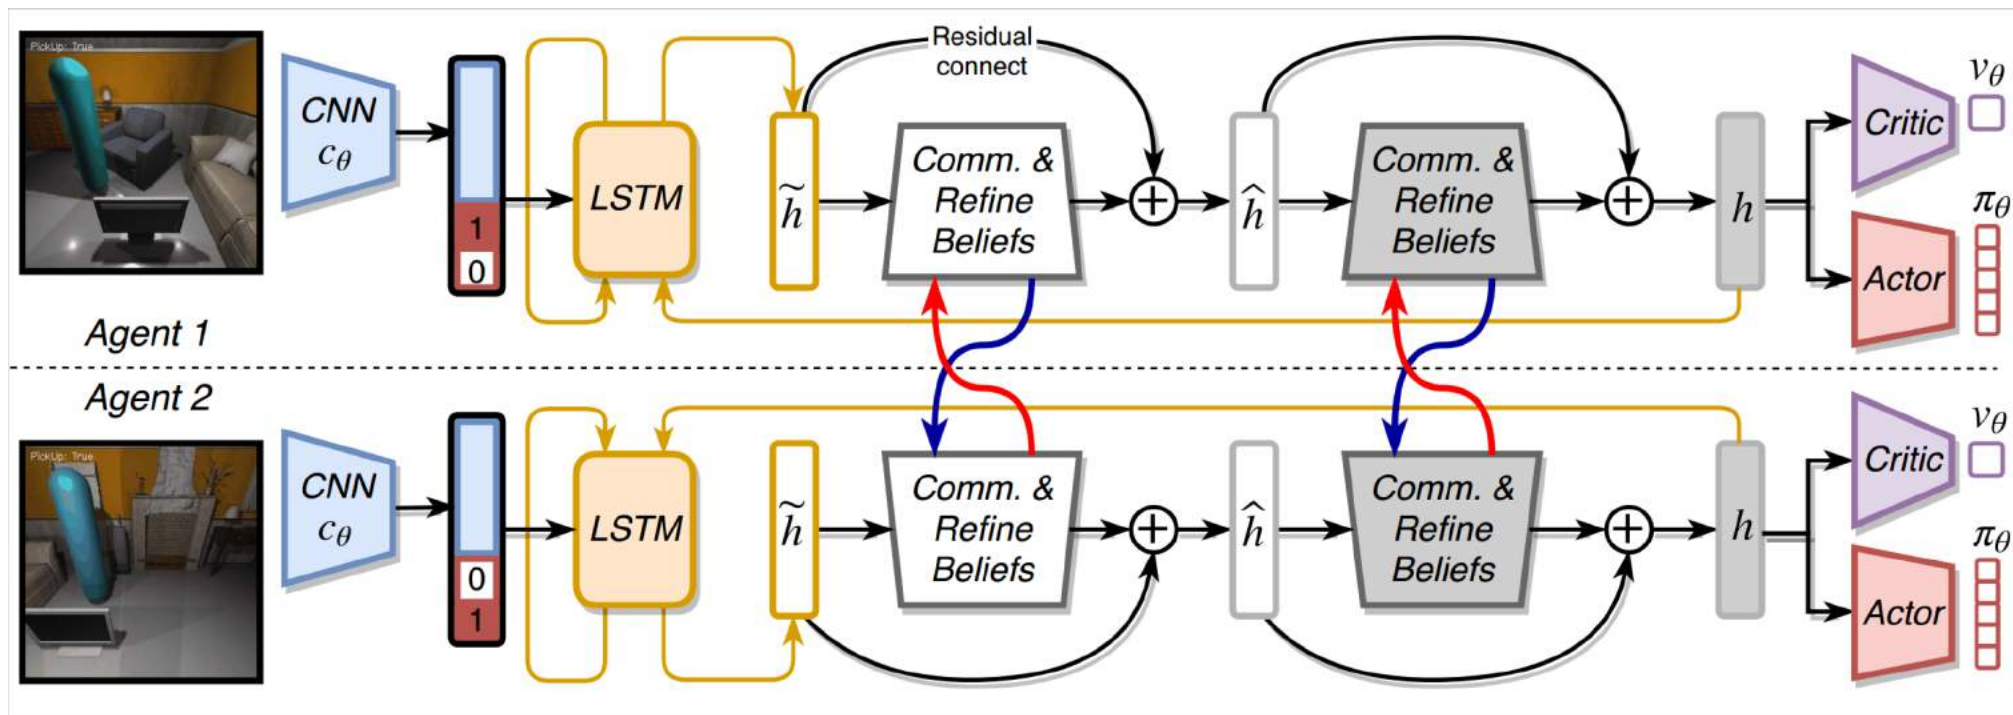
# Two Body Network



# Two Body Network

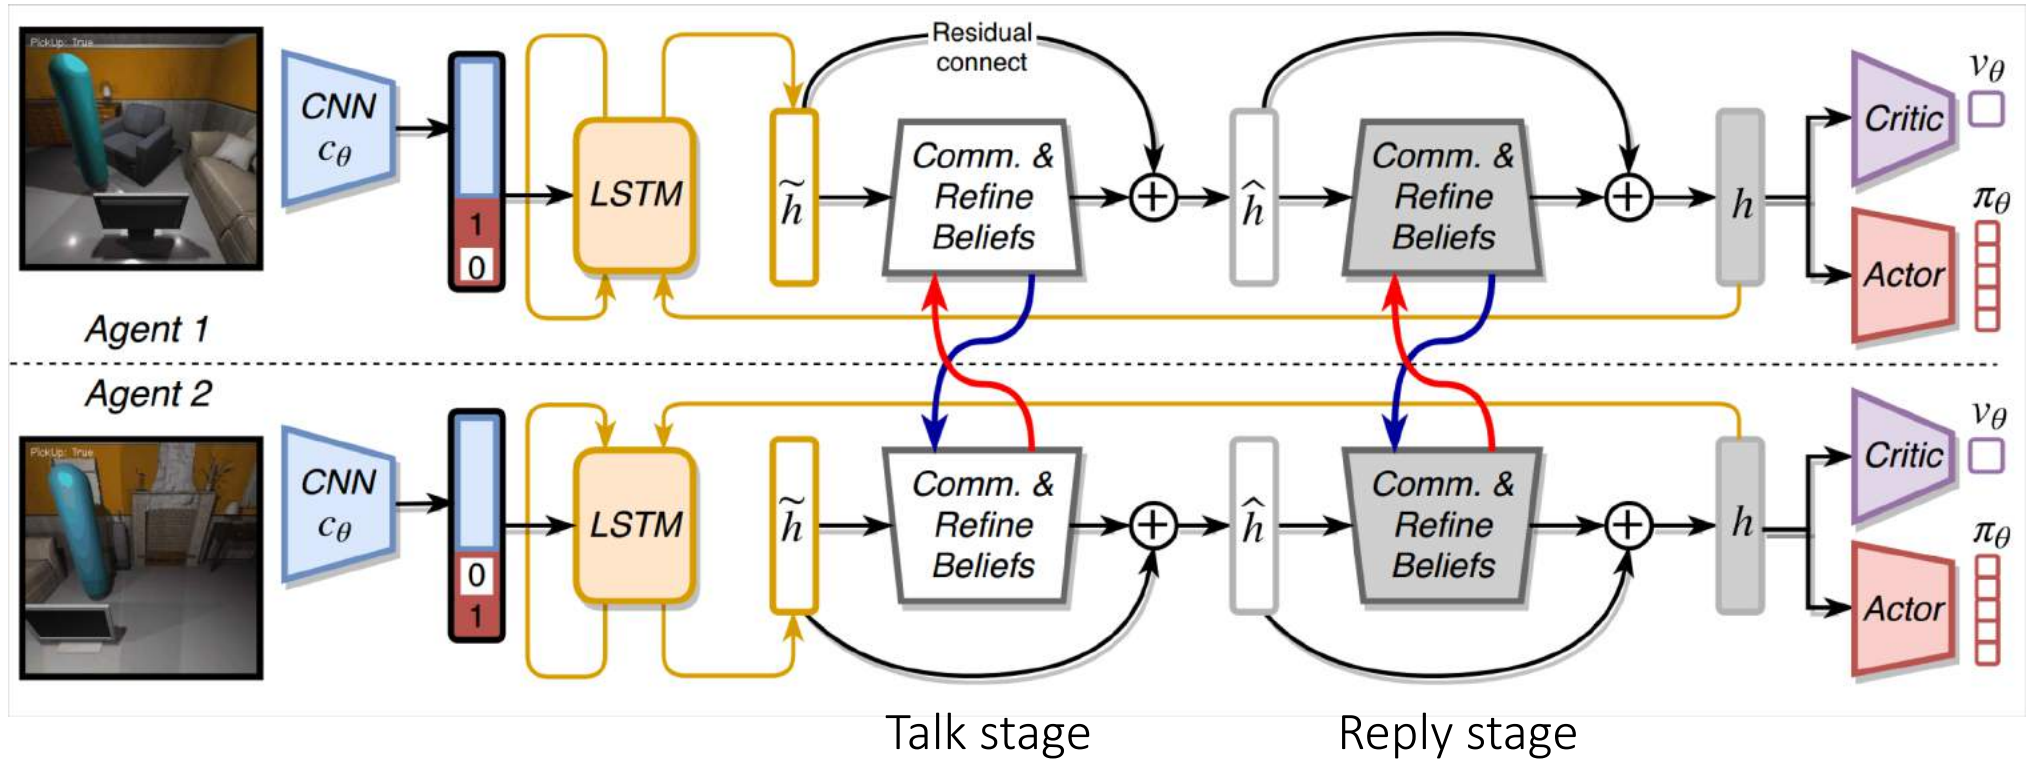


# Two Body Network

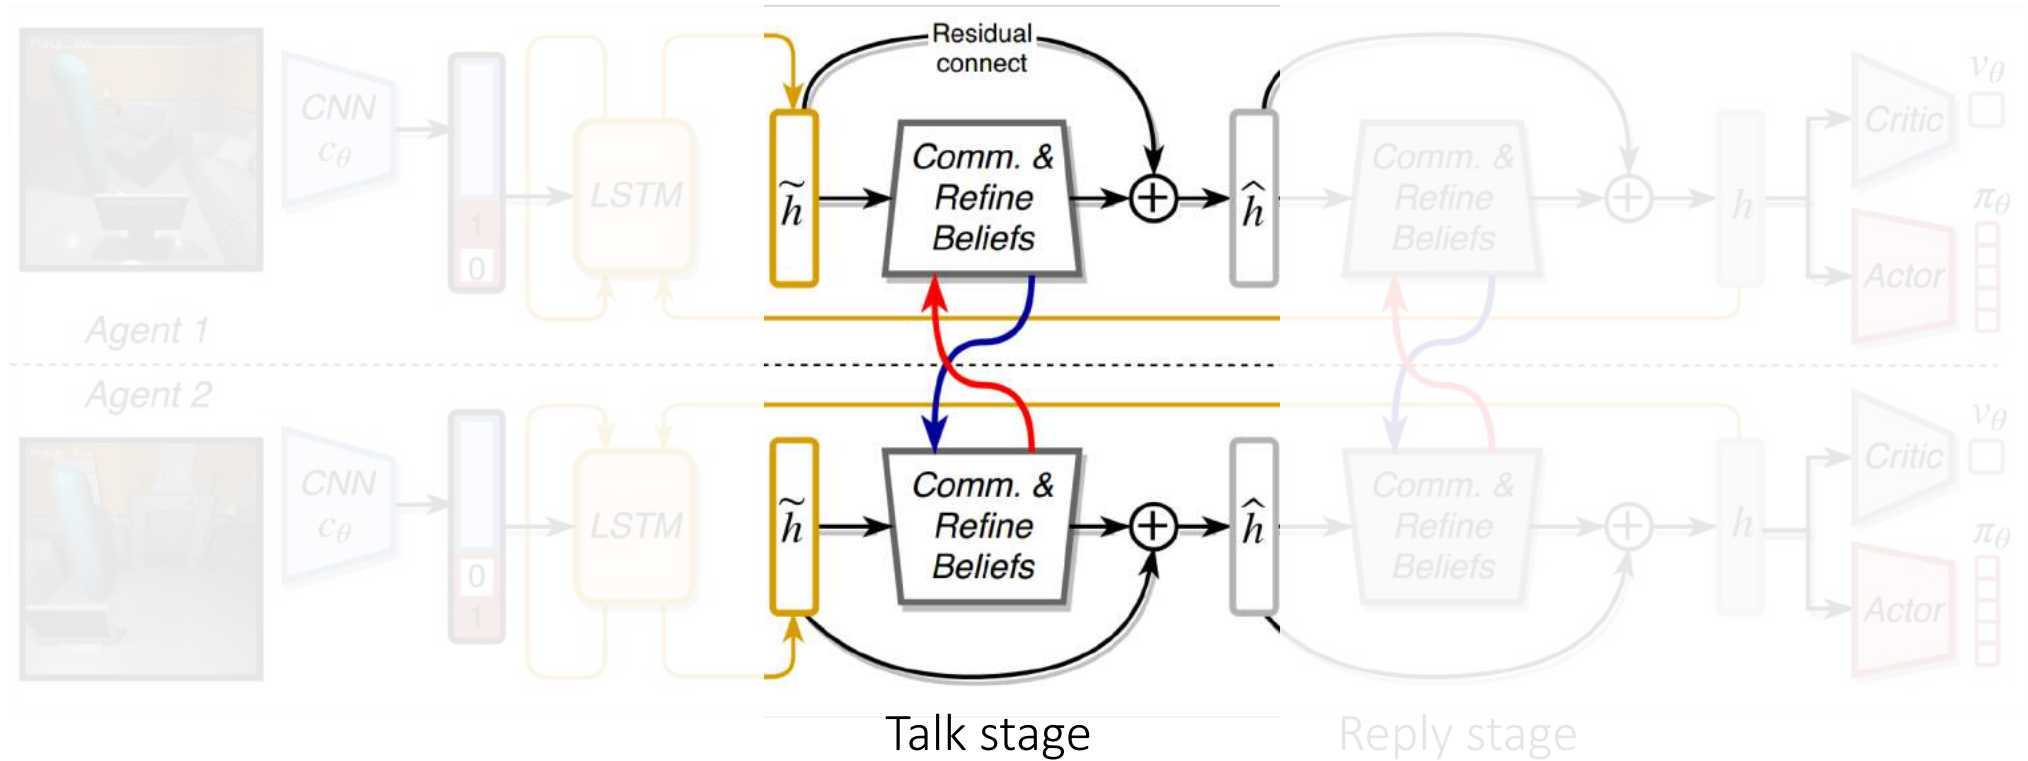


Asynchronous advantage actor critic (A3C) algorithm  
+ Imitation learning (vs shortest path oracle agent)

# Two Body Network

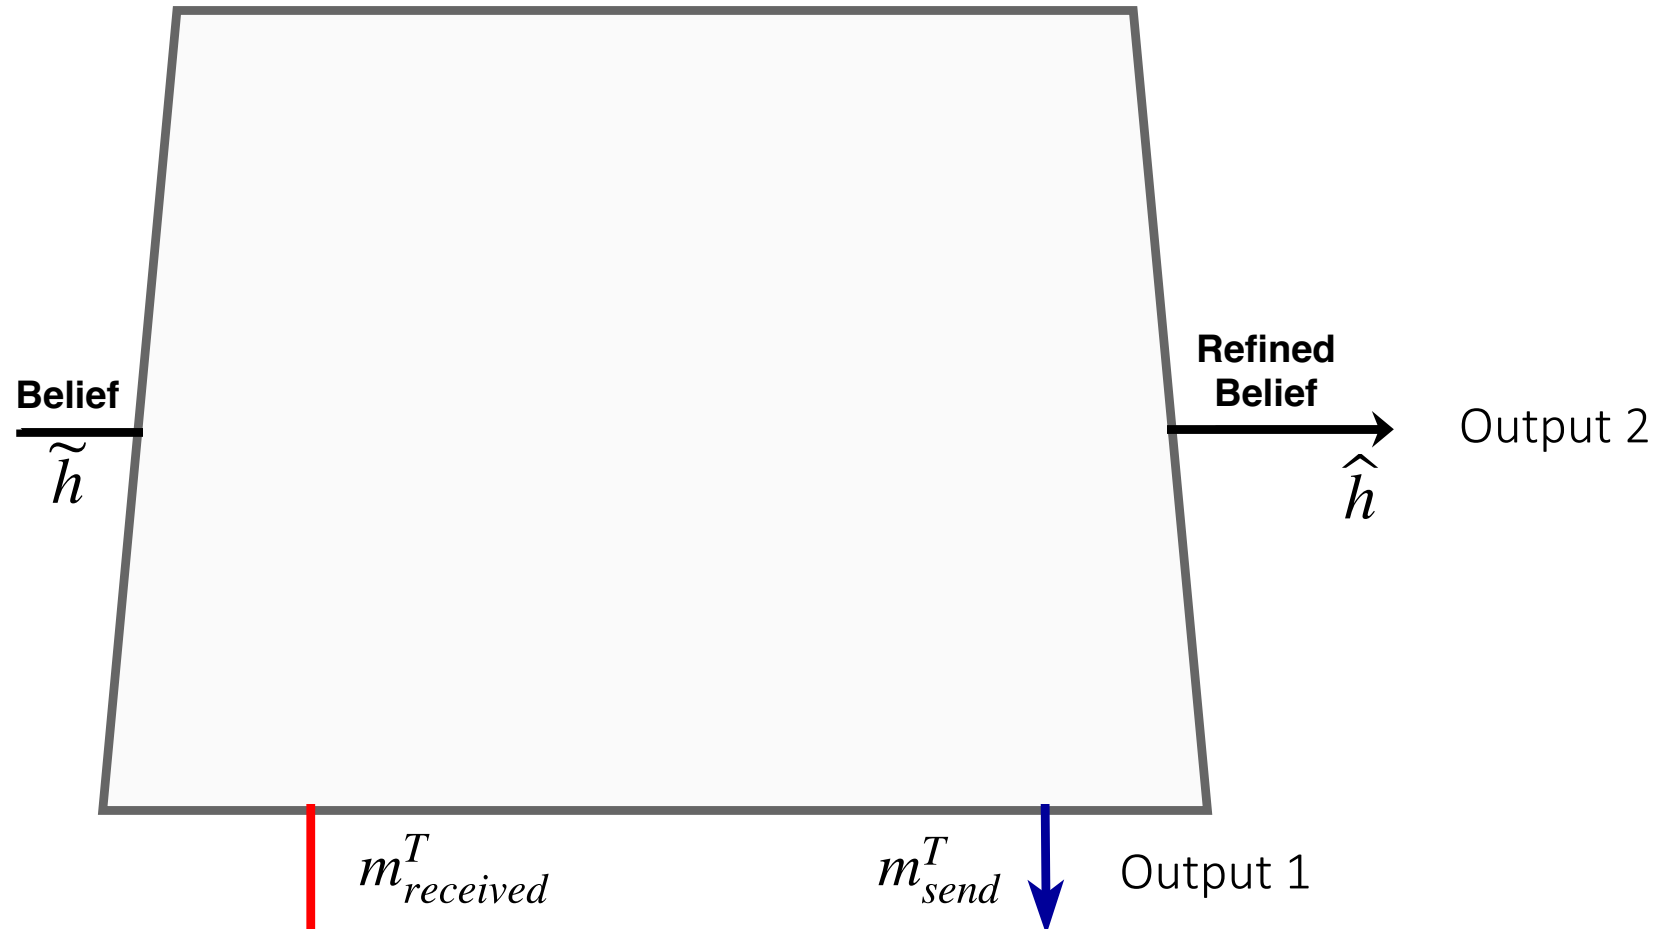


# Two Body Network



# Communication and Belief Refinement

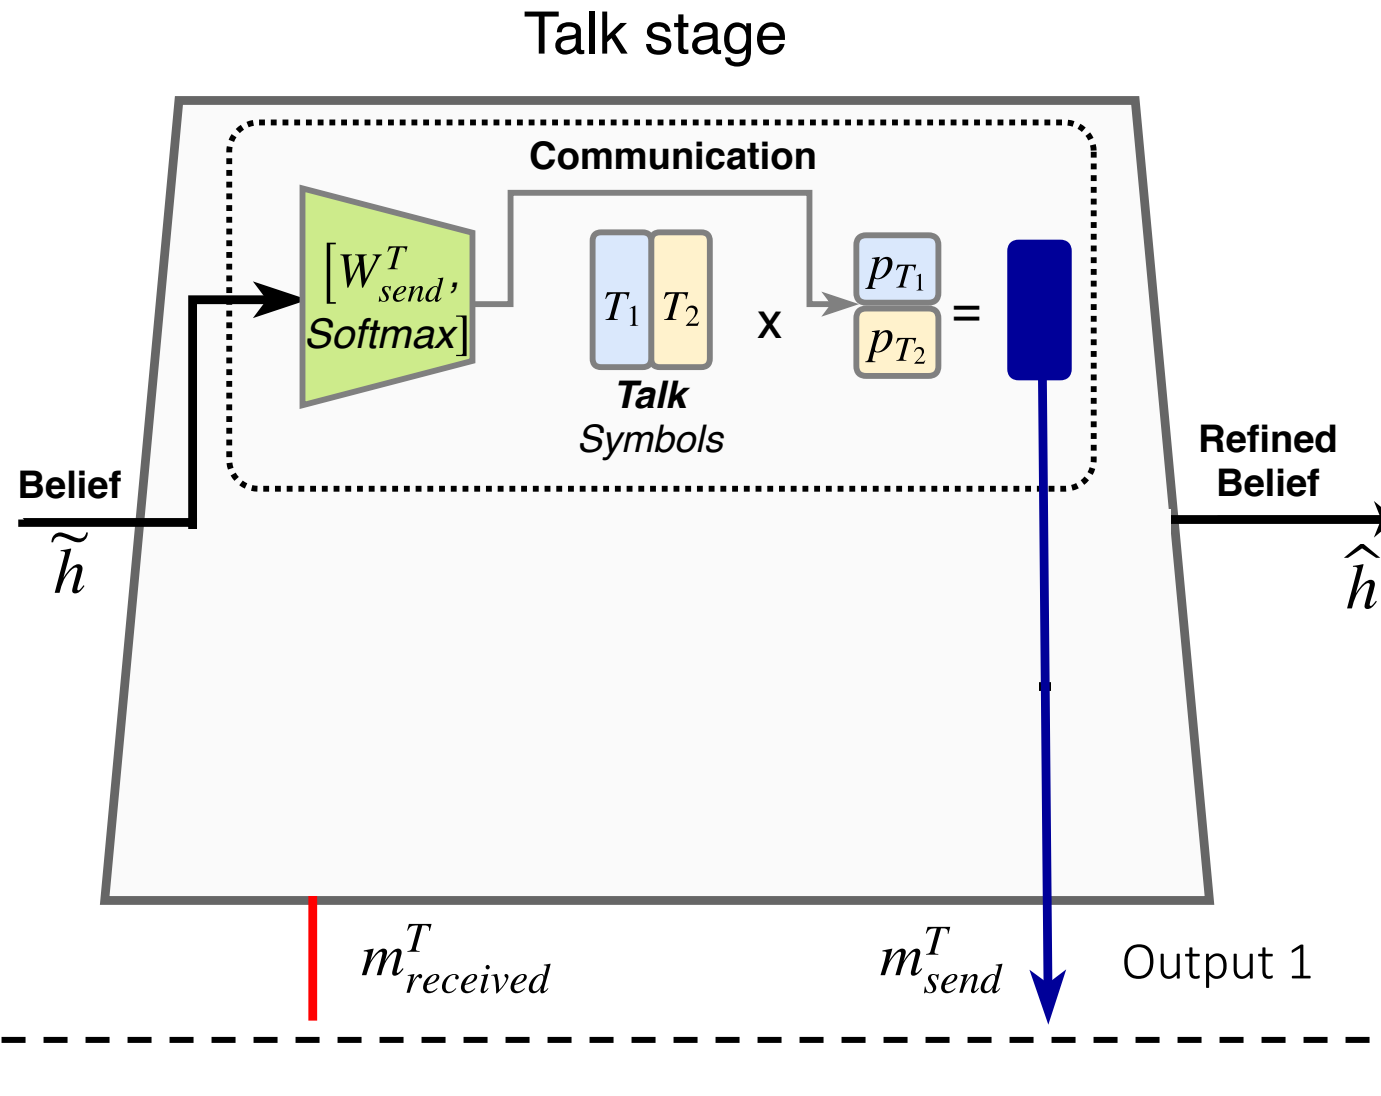
Talk stage



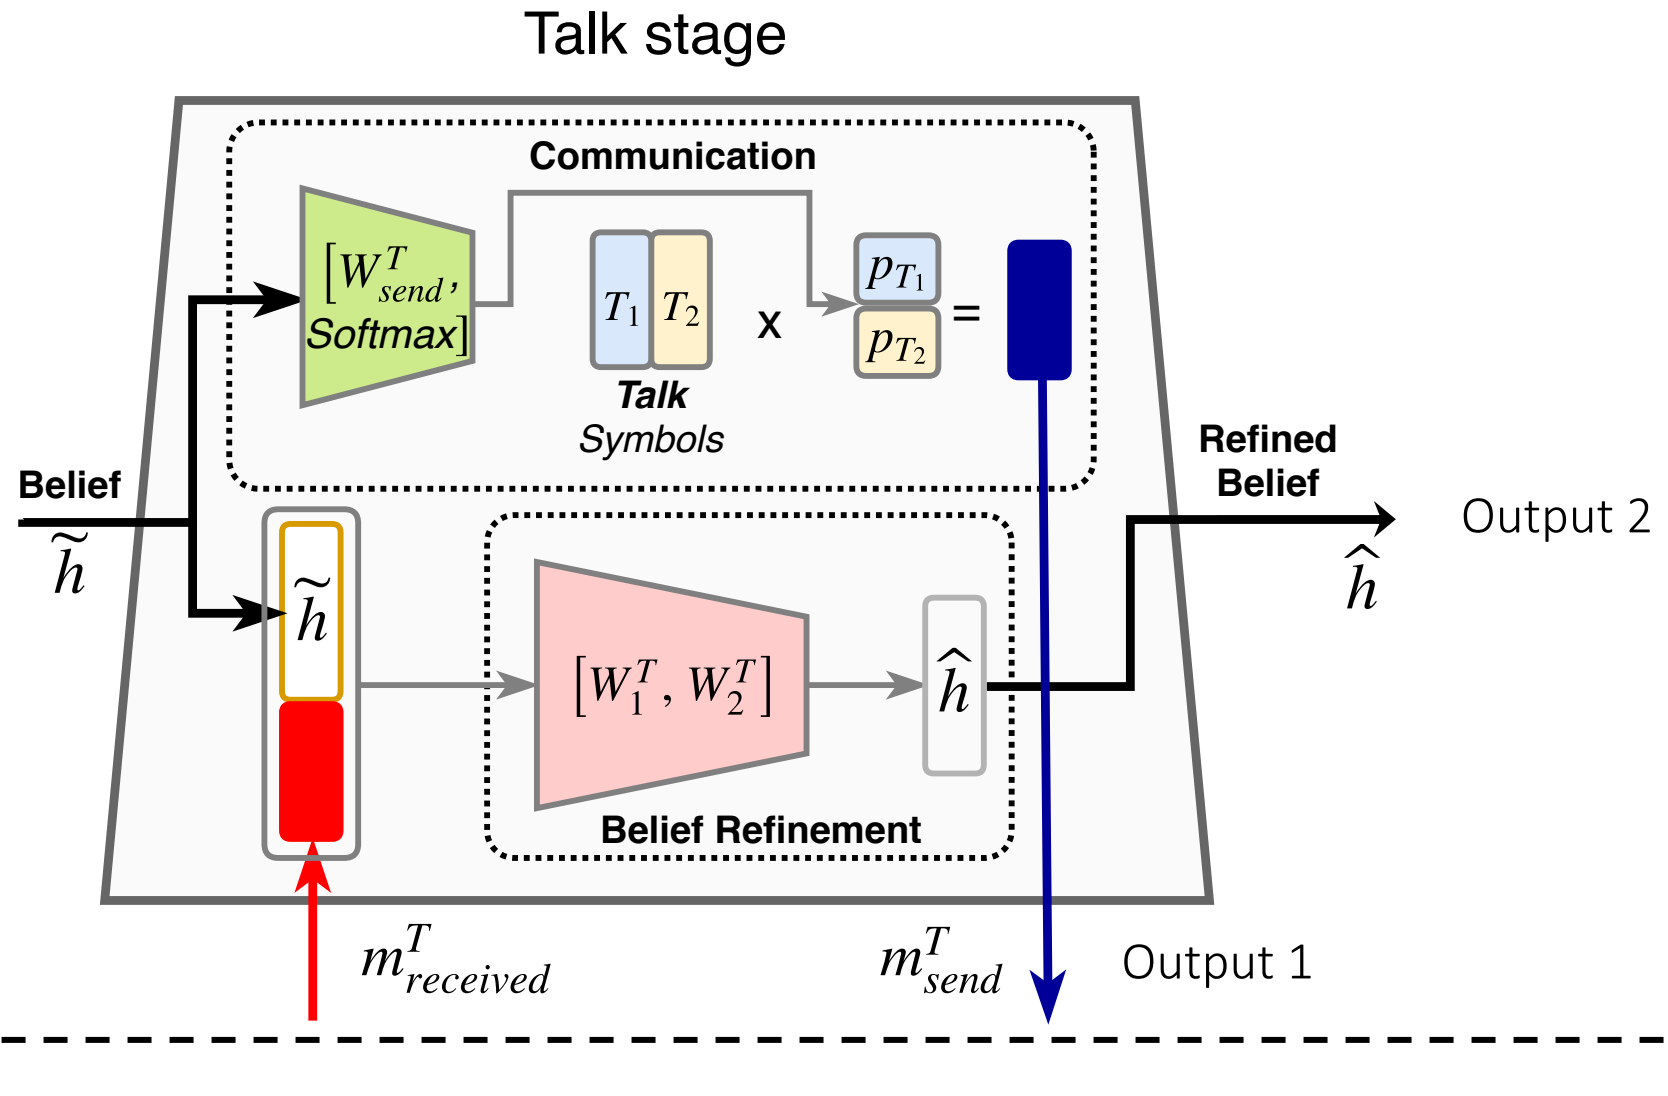
Agent 1

Agent 2

# Communication and Belief Refinement

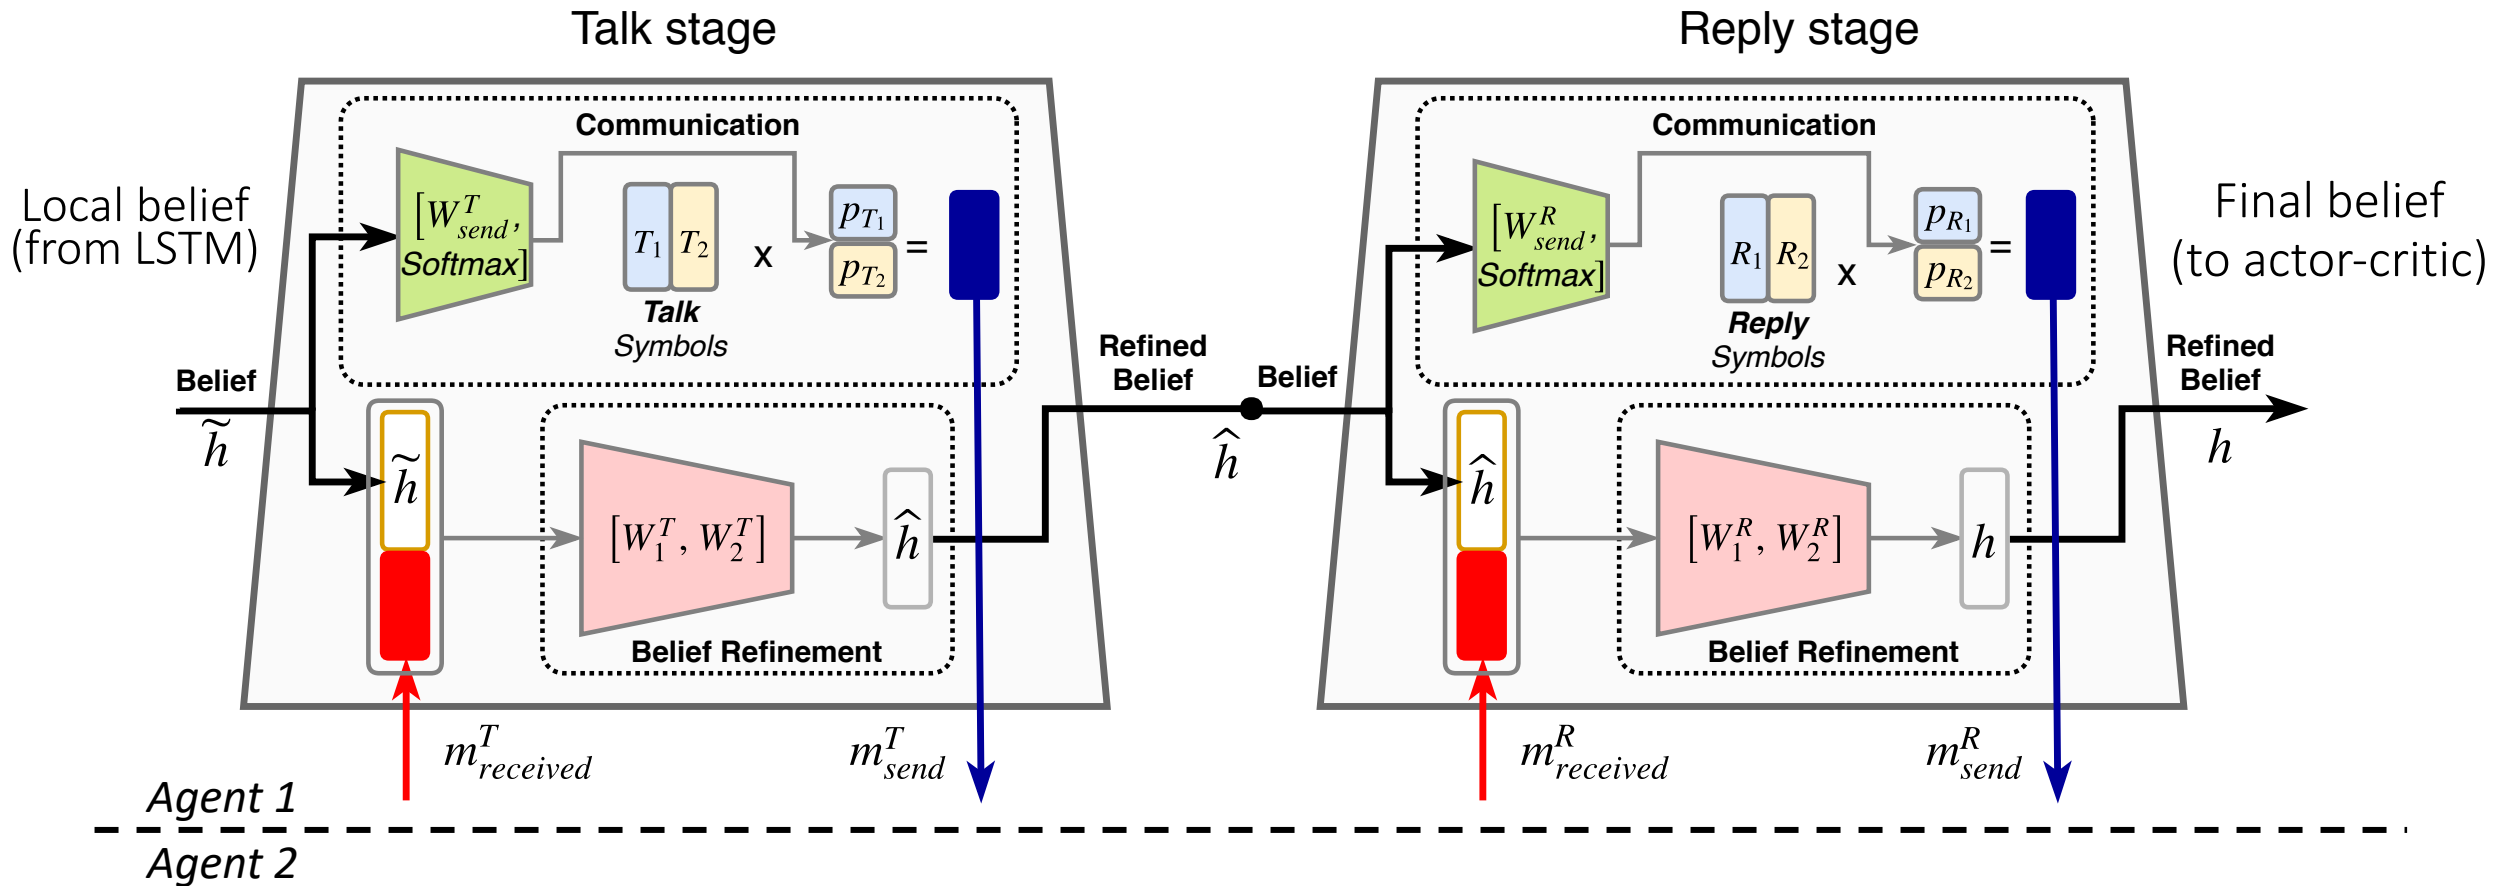


# Communication and Belief Refinement






# Talk and reply modules




# Without explicit communication



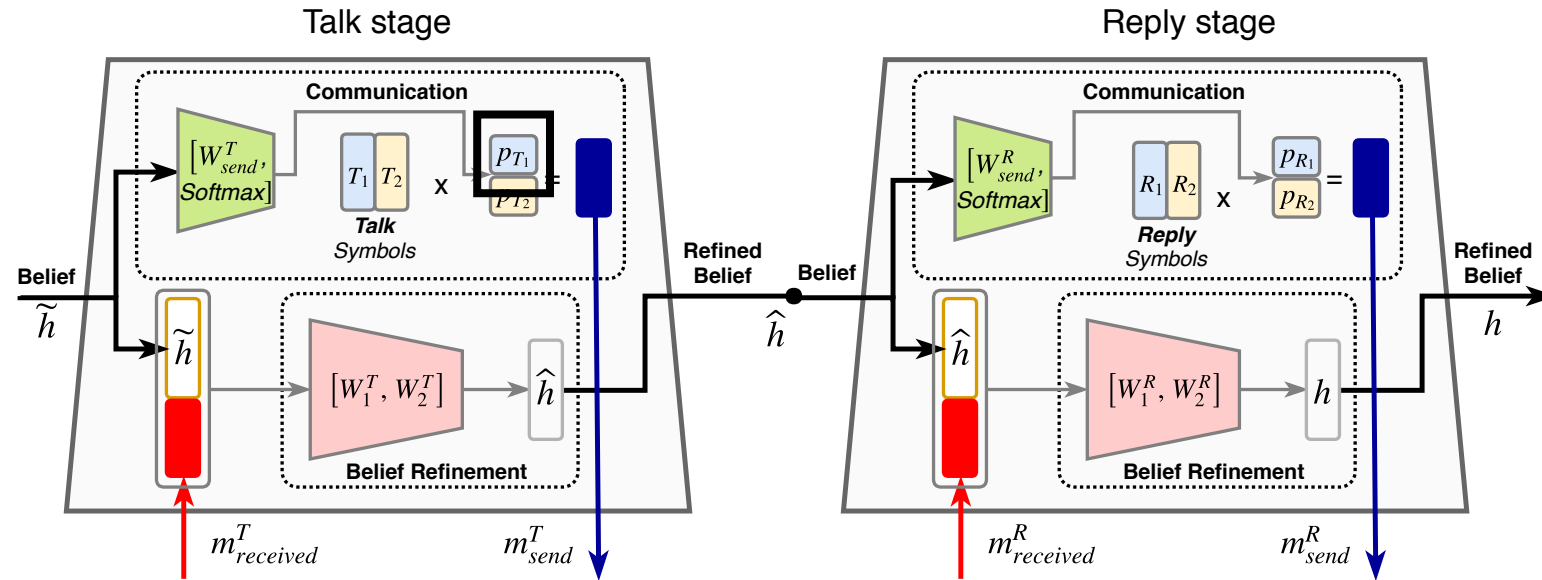
Total steps: 165  
Unsuccessful pickups: 6

# With explicit communication

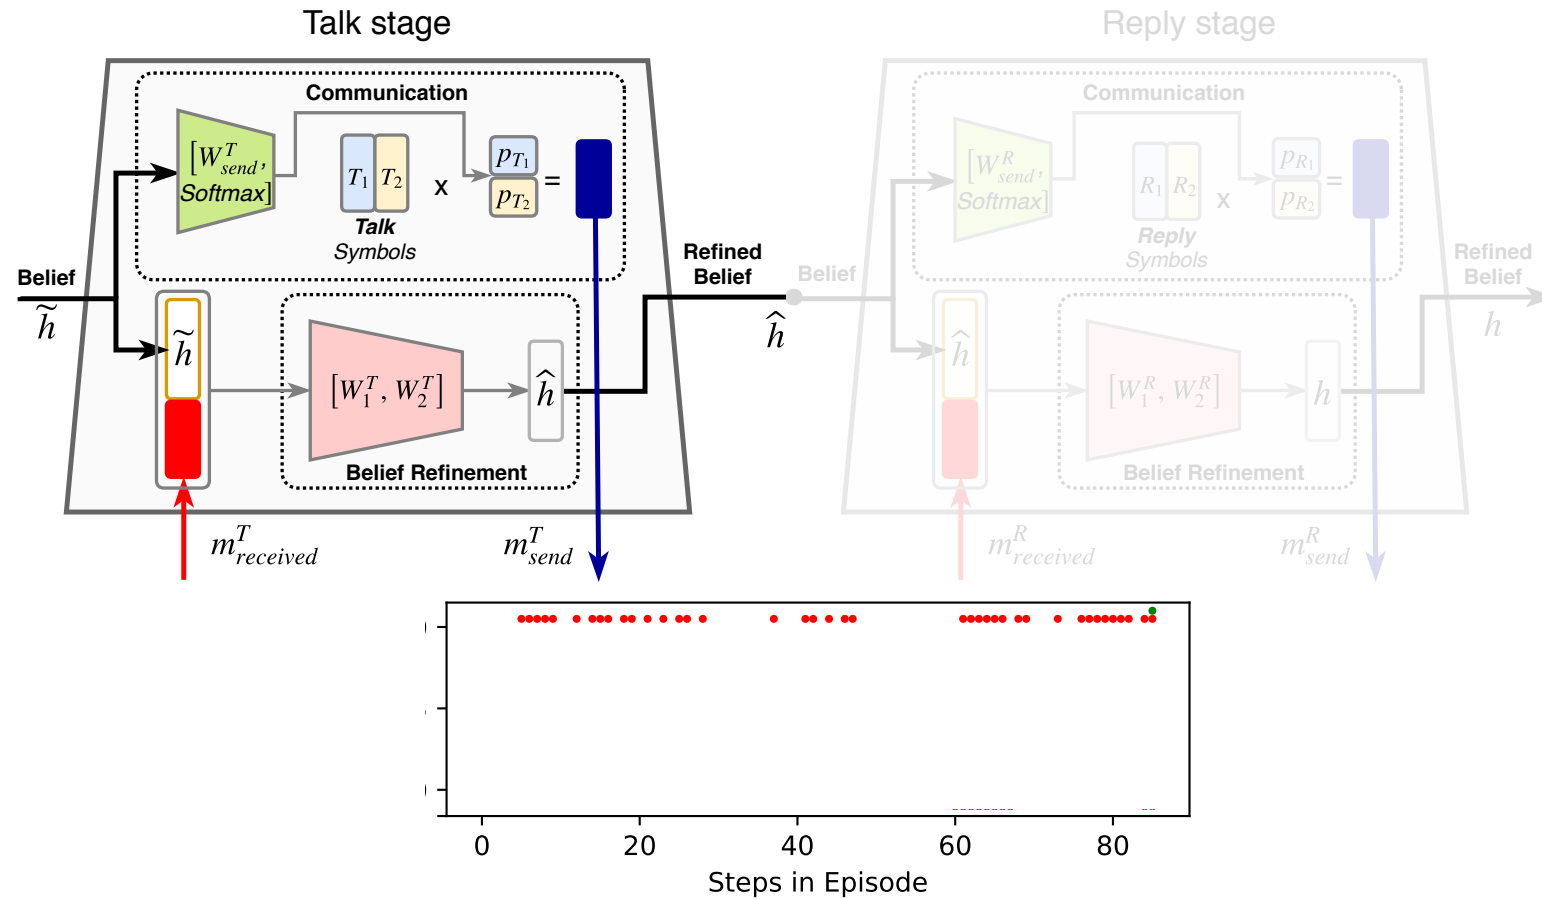


Total steps: 86  
Unsuccessful pickups: 0


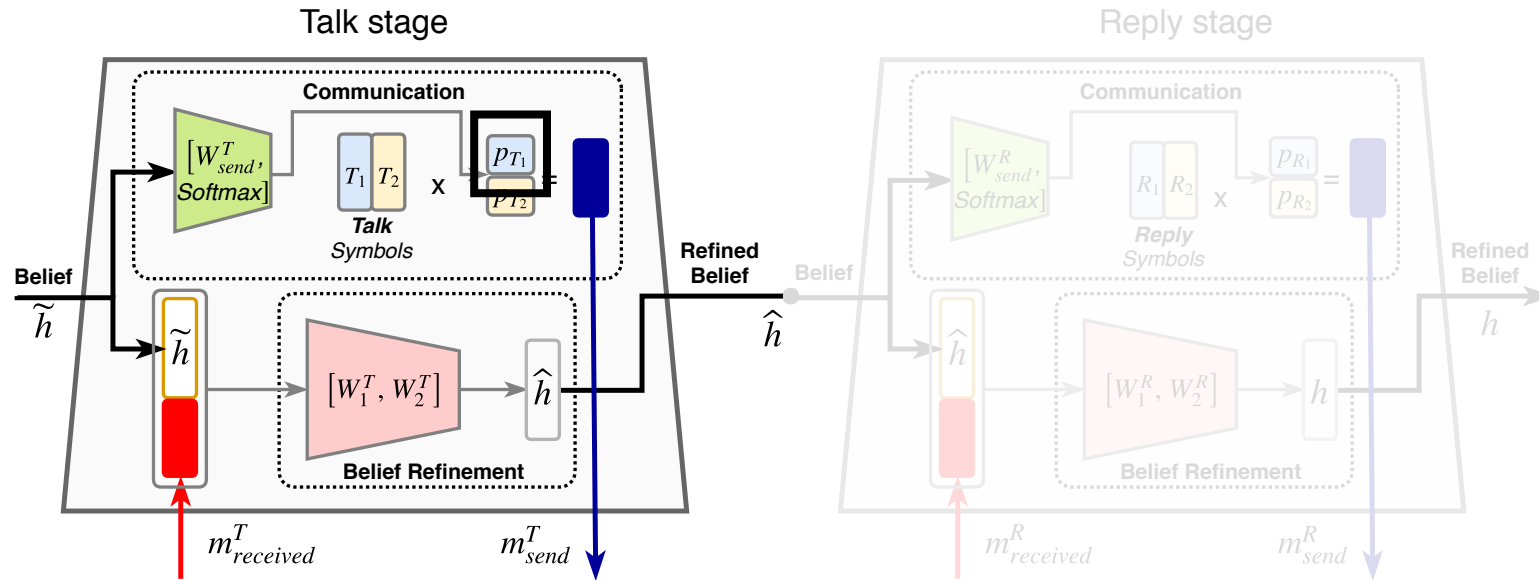
# Interpretation of messages




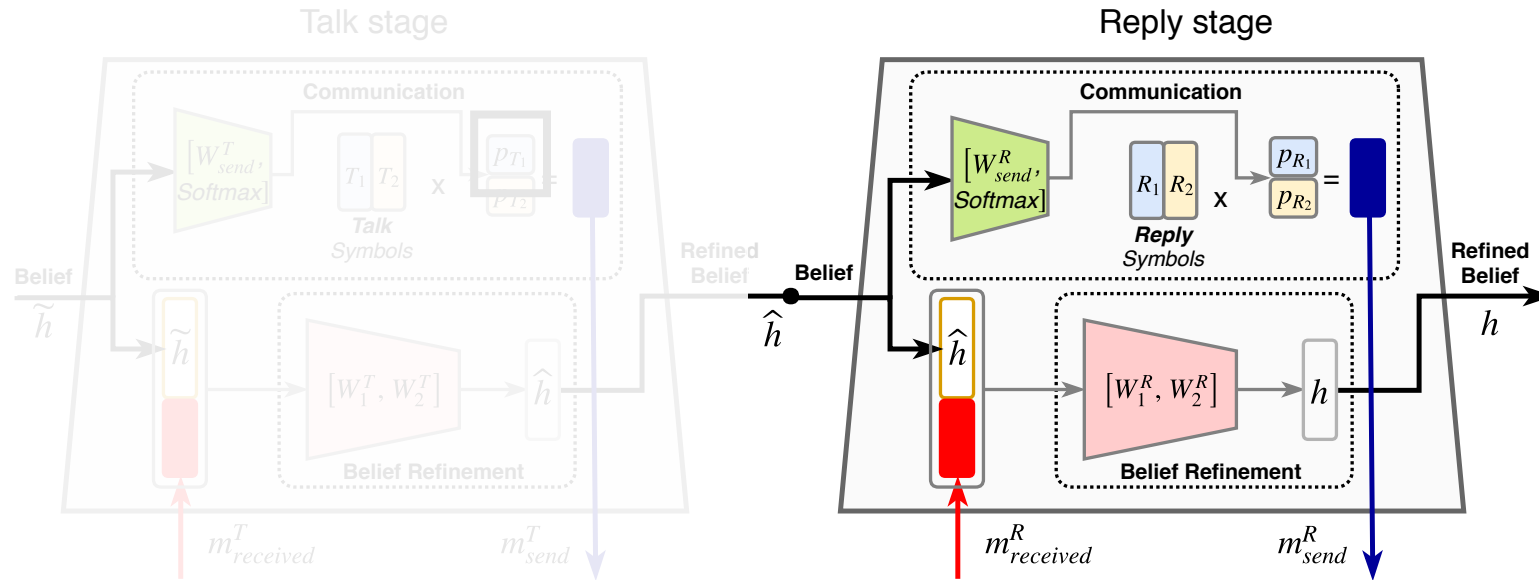
# Interpretation of messages




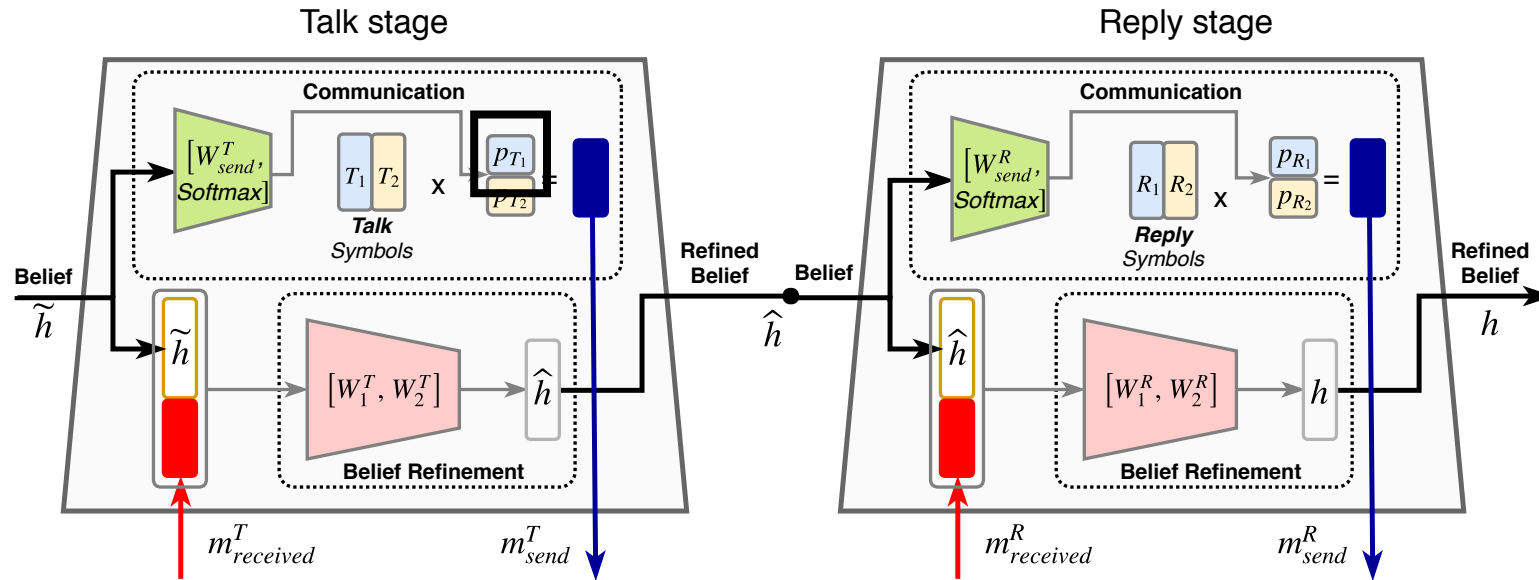
# Interpretation of messages



# Interpretation of messages



# Interpretation of messages





# Without explicit communication



Total steps: 72


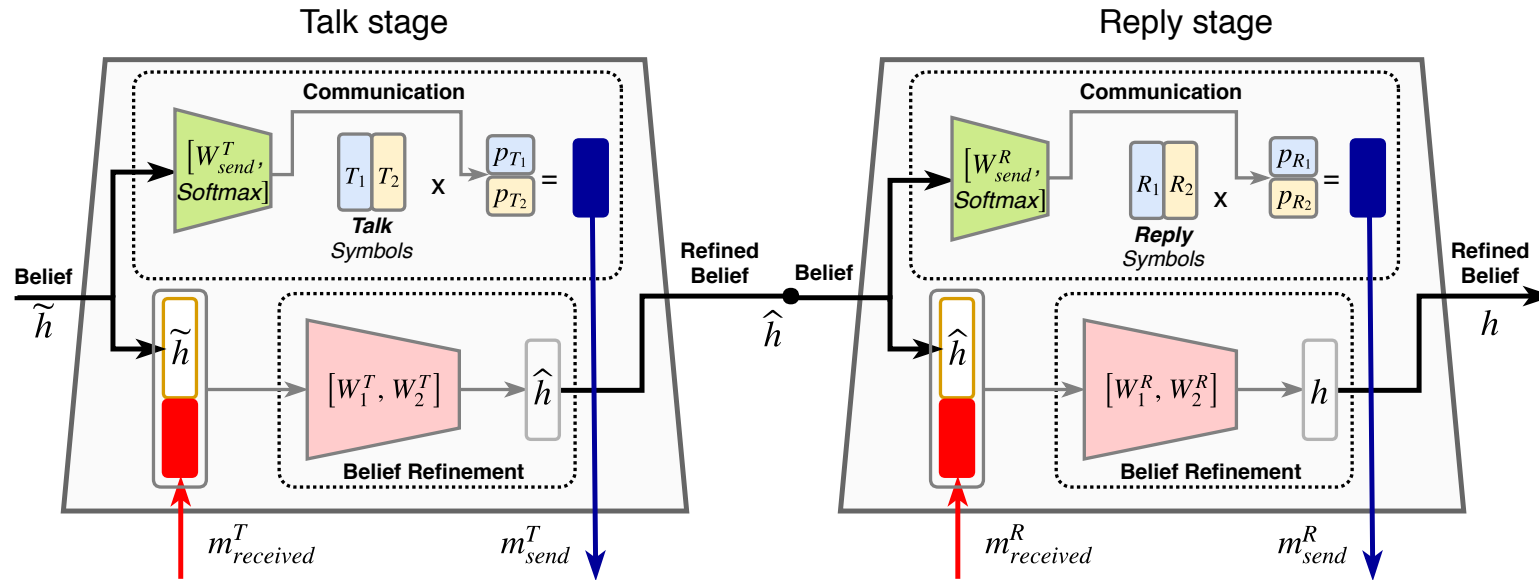
Unsuccessful pickups: 4

# With explicit communication



Total steps: 17  
Unsuccessful pickups: 0

# Communication symbols



# Implicit communication

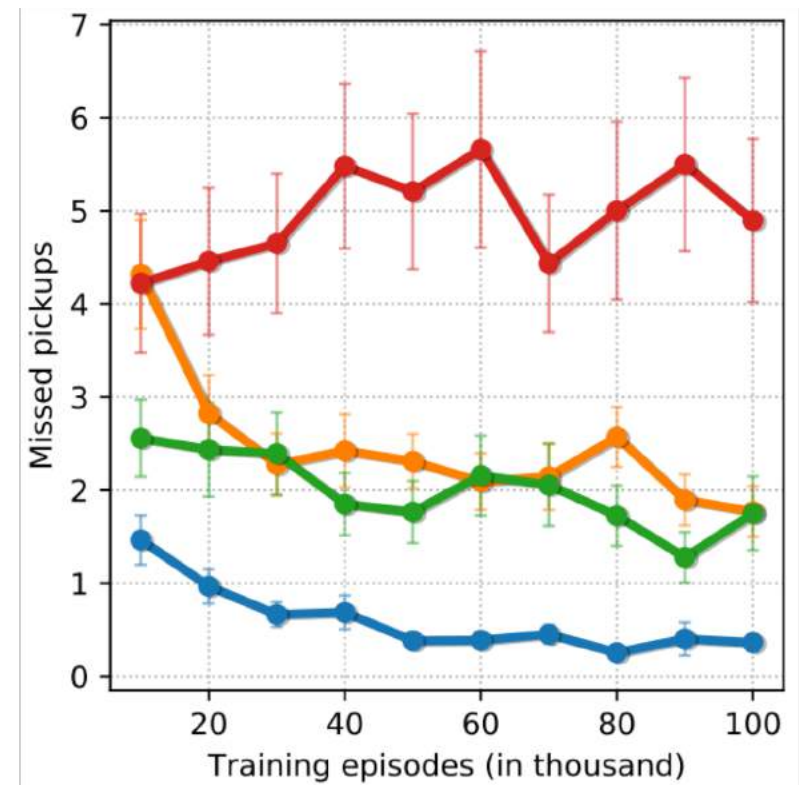


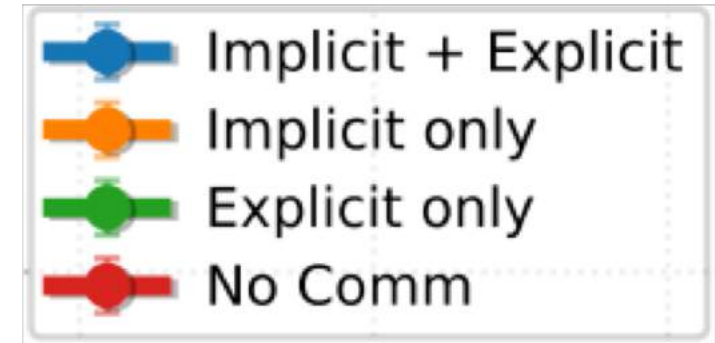
# Without Implicit communication




Total steps: 217

Unsuccessful pickups: 23


# Results: Effect of communication



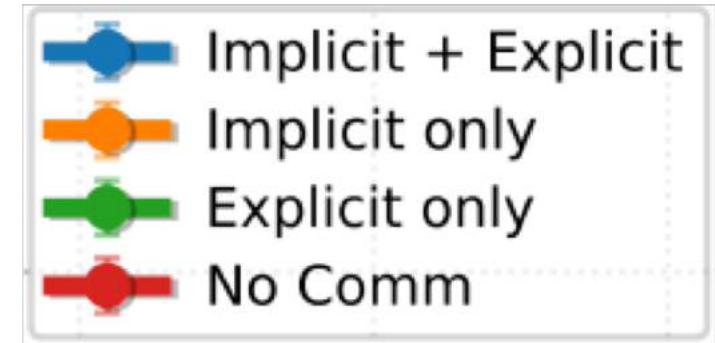
# Result: Generalization



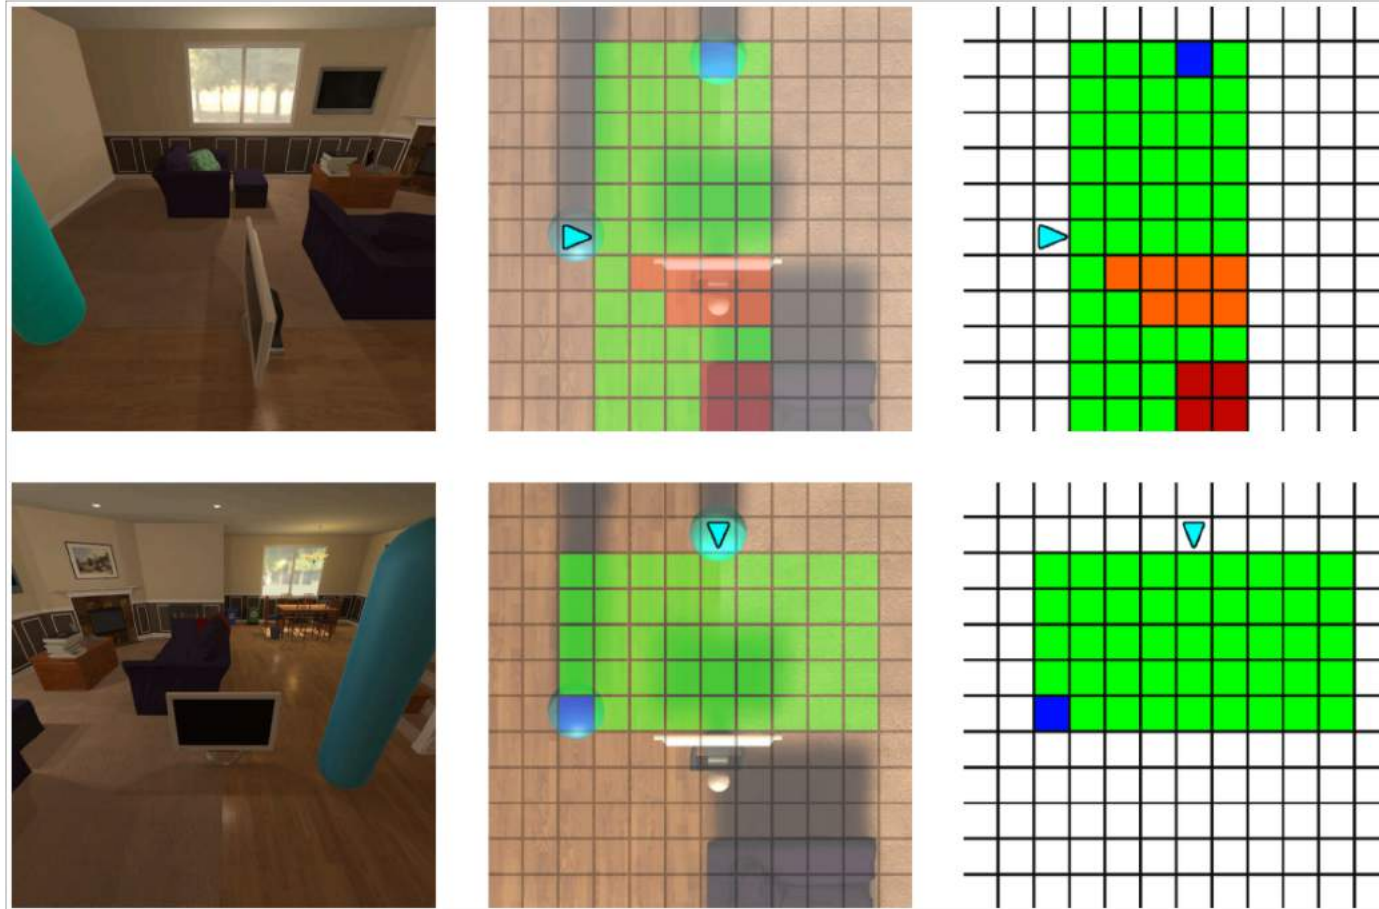
Seen environments



Unseen environments





# Visual vs Gridworld






# Visual vs Gridworld



# Visual vs Gridworld



Data	Accuracy	Missed pickups	Unsuccess. pickups
Visual	$59.0 \pm 4.0$	$0.3 \pm 0.09$	$2.9 \pm 0.8$
Visual+depth	$65.7 \pm 3.9$	$0.4 \pm 0.1$	$3.2 \pm 0.9$
Grid world	<b><math>78.2 \pm 3.4</math></b>	<b><math>0.1 \pm 0.05</math></b>	<b><math>0.7 \pm 0.1</math></b>

# Summary

## 1. Multi-agent simulation framework



# Summary

1. Multi-agent simulation framework
2. Collaborative find and lift task



Agent 1

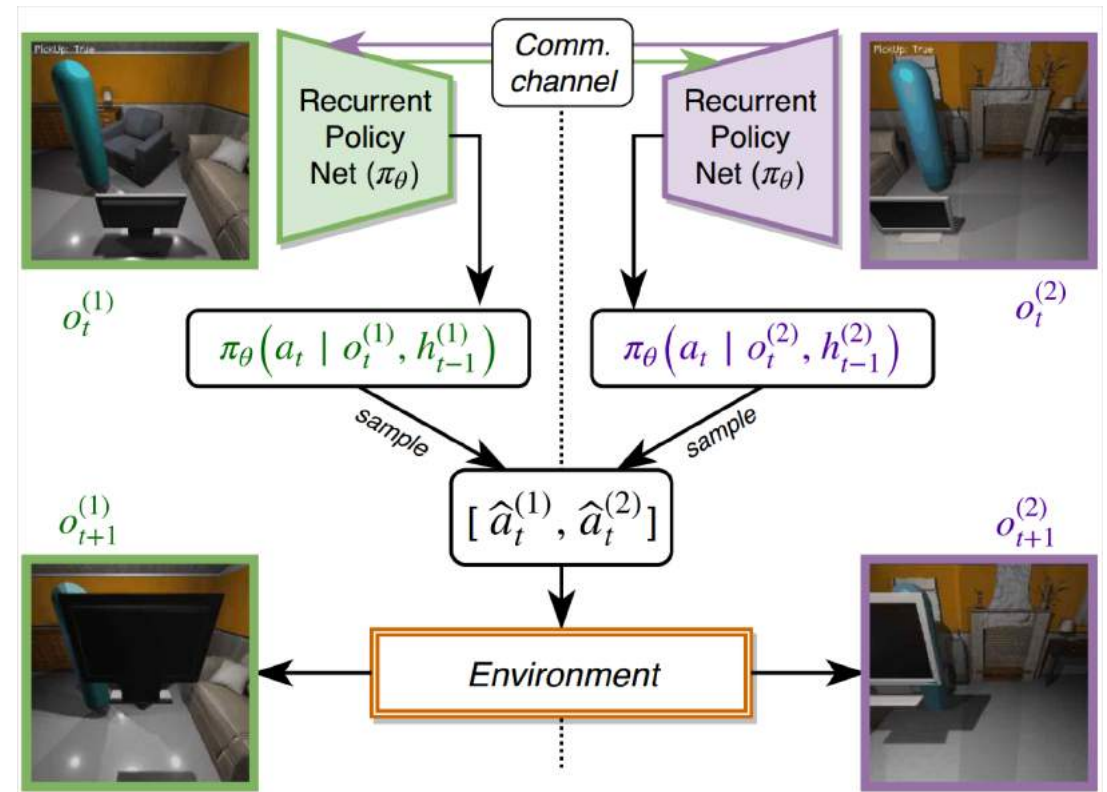


Agent 2

- + Navigate to TV
- + Collaboration  
[Joint pickup]

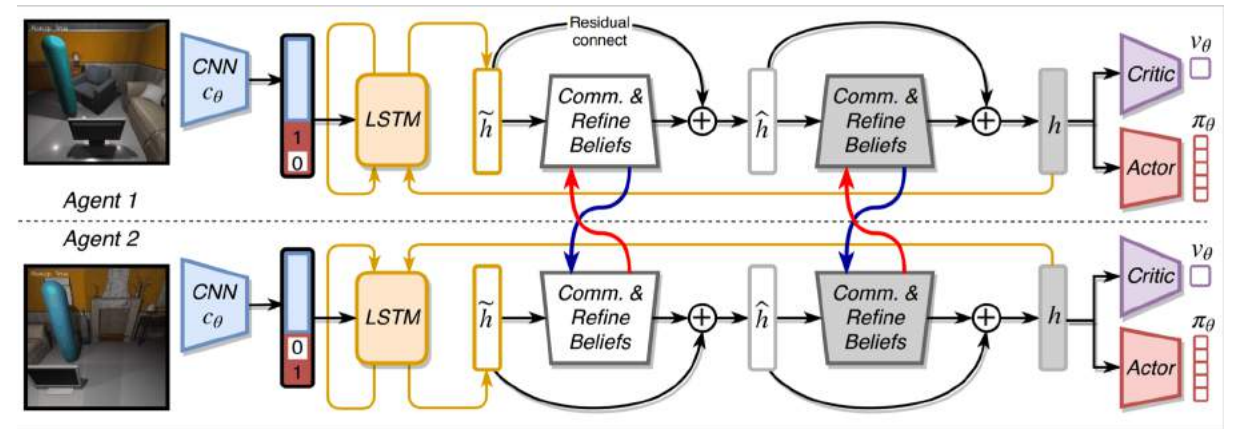
# Summary

1. Multi-agent simulation framework
2. Collaborative find and lift task
3. Reinforcement + Imitation Learning



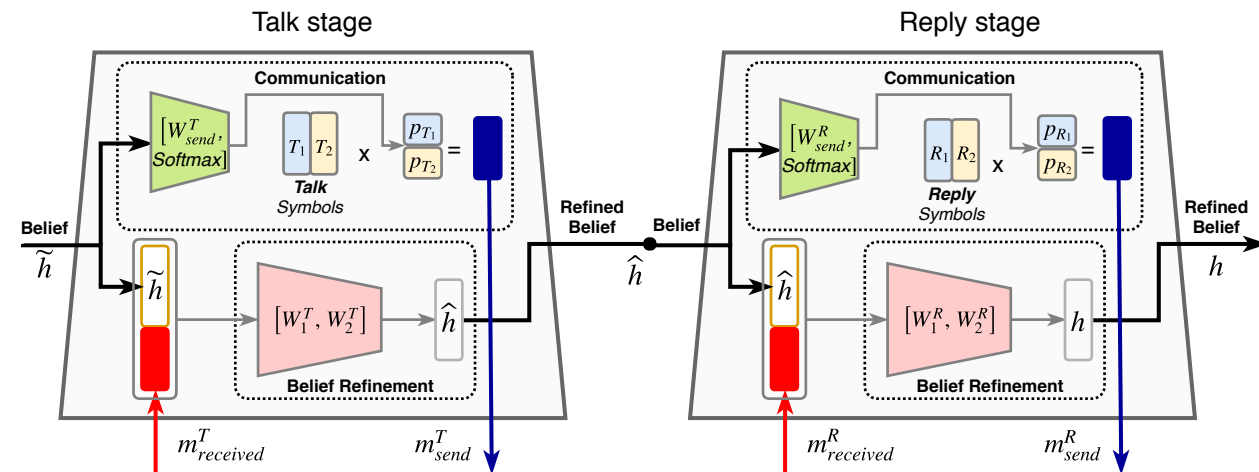
# Summary

1. Multi-agent simulation framework
2. Collaborative find and lift task
3. Reinforcement + Imitation Learning
4. Two Body Network model




# Summary

1. Multi-agent simulation framework
2. Collaborative find and lift task
3. Reinforcement + Imitation Learning
4. Two Body Network model
5. Comm. and BR module



# Summary

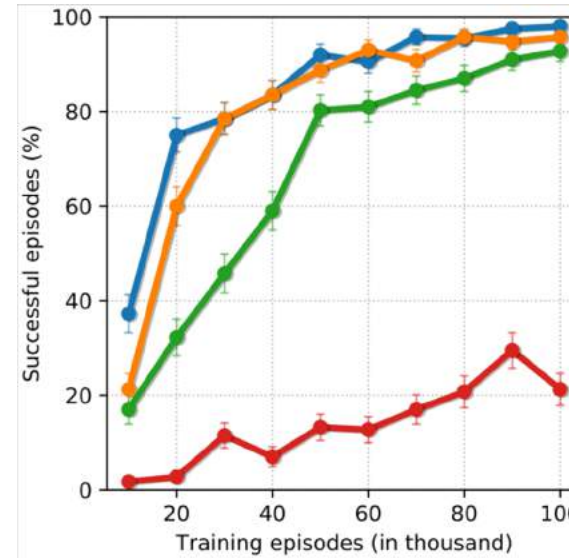
1. Multi-agent simulation framework
2. Collaborative find and lift task
3. Reinforcement + Imitation Learning
4. Two Body Network model
5. Comm. and BR module
6. Interpretation of messages






# Summary

1. Multi-agent simulation framework
2. Collaborative find and lift task
3. Reinforcement + Imitation Learning
4. Two Body Network model
5. Comm. and BR module
6. Interpretation of messages
7. Results on communication and generalization




Seen environments



Unseen environments

# Summary

1. Multi-agent simulation framework
2. Collaborative find and lift task
3. Reinforcement + Imitation Learning
4. Two Body Network model
5. Comm. and BR module
6. Interpretation of messages
7. Results on communication and generalization
8. Visual vs Gridworld



# Two Body Problem: Collaborative Visual Task Completion CVPR 2019

Speaker: Unnat Jain, UIUC [[webpage](#)]




Adam (lou-Jen) Liu  
UIUC



Alex Schwing  
UIUC



Ali Farhadi  
AI2, UW, XNOR.ai



Ani Kembhavi  
AI2



Eric Kolve  
AI2



Luca Weihs  
AI2



M. Rastegari  
AI2, XNOR.ai



Svetlana Lazebnik  
UIUC

**I** ILLINOIS

AI2 PRIOR  
Perceptual Reasoning & Interaction Research

**W**

xnor.ai