

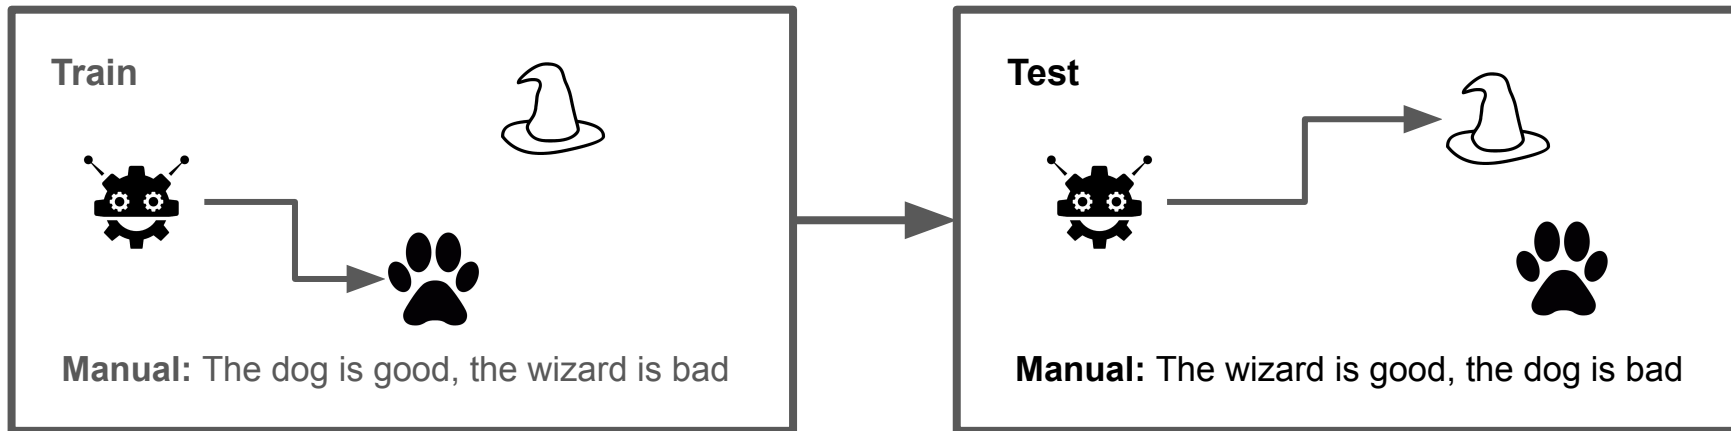
Grounding Language to Entities and Dynamics for Generalization in Reinforcement Learning

Austin W. Hanjie¹, Victor Zhong², Karthik Narasimhan¹

1. Princeton University 2. University of Washington

Introduction

Study natural language manuals to allow RL agents to generalize to new environments



Introduction

Study natural language manuals to allow RL agents to generalize to new environments



Entity Grounding: Mapping text symbols (*wizard*) to entity symbols in observation



Our Work

Prior Work

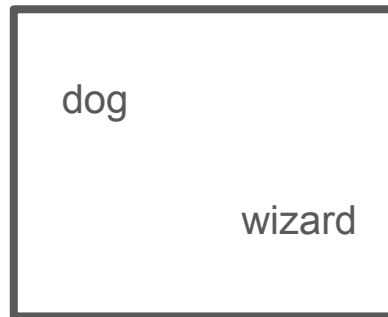
Simplified/eliminated entity grounding problem

Our Work

Prior Work

Simplified/eliminated entity grounding problem

Lexical Overlap
(Zhong et al. 2020)



1. ...dog...,
2. ...wizard...,

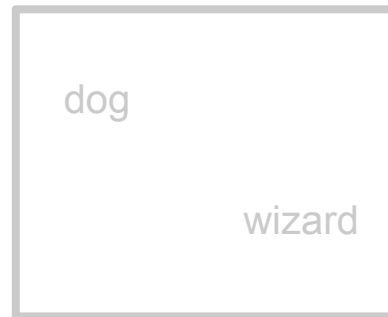
Victor Zhong, Tim Rocktäschel, & Edward Grefenstette (2020). RTFM: Generalising to New Environment Dynamics via Reading. In International Conference on Learning Representations.

Our Work

Prior Work


Simplified/eliminated entity grounding problem

Lexical Overlap
(Zhong et al. 2020)



1. ...dog...,
2. ...wizard...,

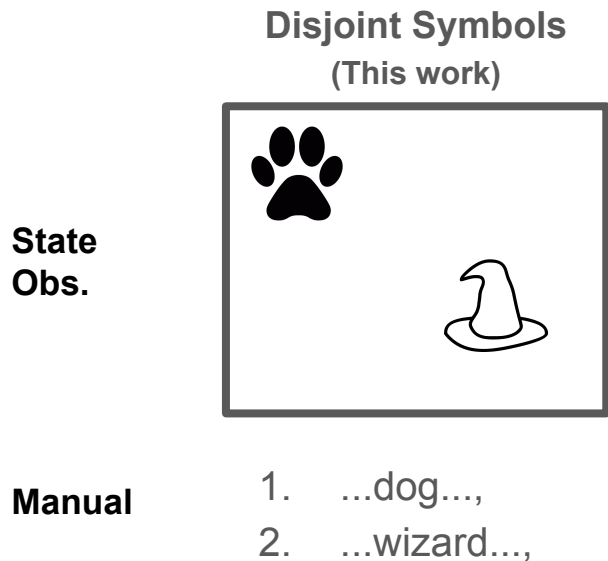
Oracle Mapping
(Narasimhan et al. 2018)



1. ...dog...,
2. ...wizard...,

Our Work

No signal connecting observation + text manual




Prior Work

Simplified/eliminated entity grounding problem




Key Novel Challenge: Learn to **ground entities** via environment interaction **without priors** connecting text to state observations.

Importance



Autonomous agents that learn to associate **text to objects** by interacting with the world **without manually specifying**



Key Contributions

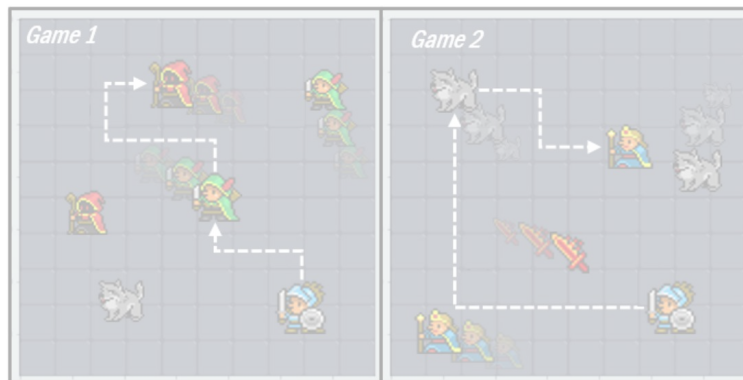
Key Contributions

Introduce the new environment *Messenger* designed to evaluate **entity grounding**

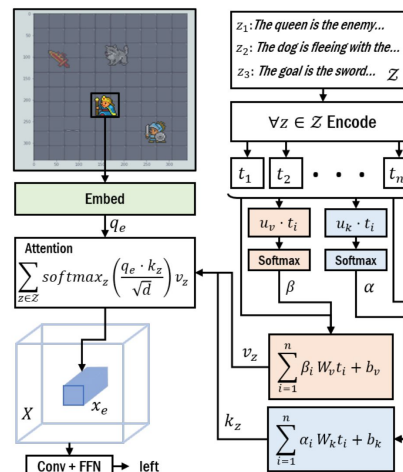


Key Contributions

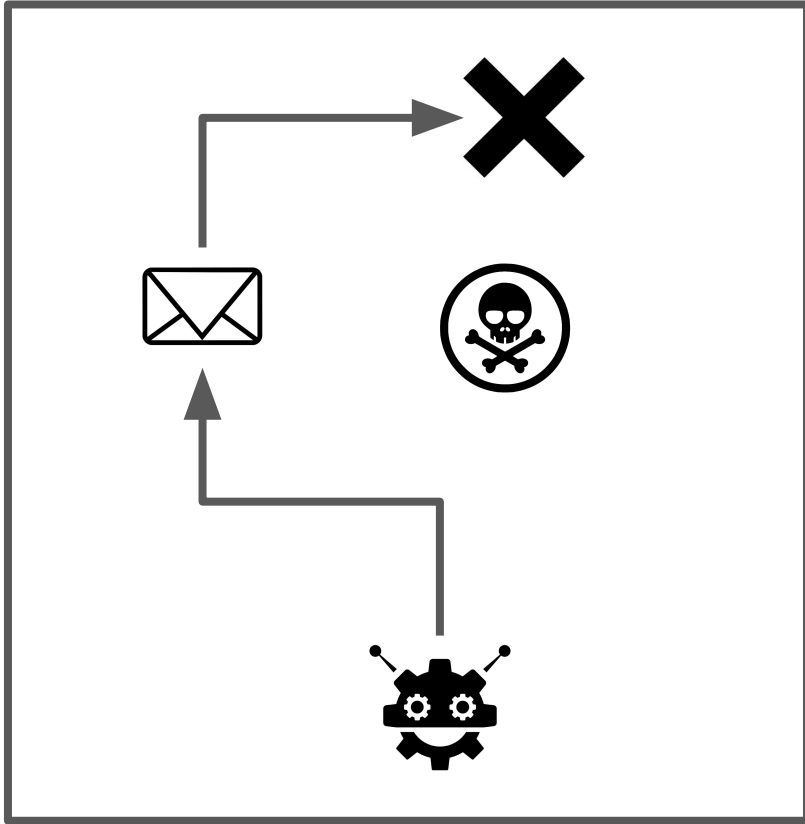
Introduce the new environment *Messenger* designed to evaluate **entity grounding**



Introduce **EMMA**, a model that outperforms other approaches on *Messenger*




Messenger



Objective: Get the message, bring it to the goal, while avoiding an enemy


Messenger



Objective: Get the message, bring it to the goal, while avoiding an enemy

The message, goal, and enemy, roles may be filled by different **entities** (e.g. wizard)

Messenger



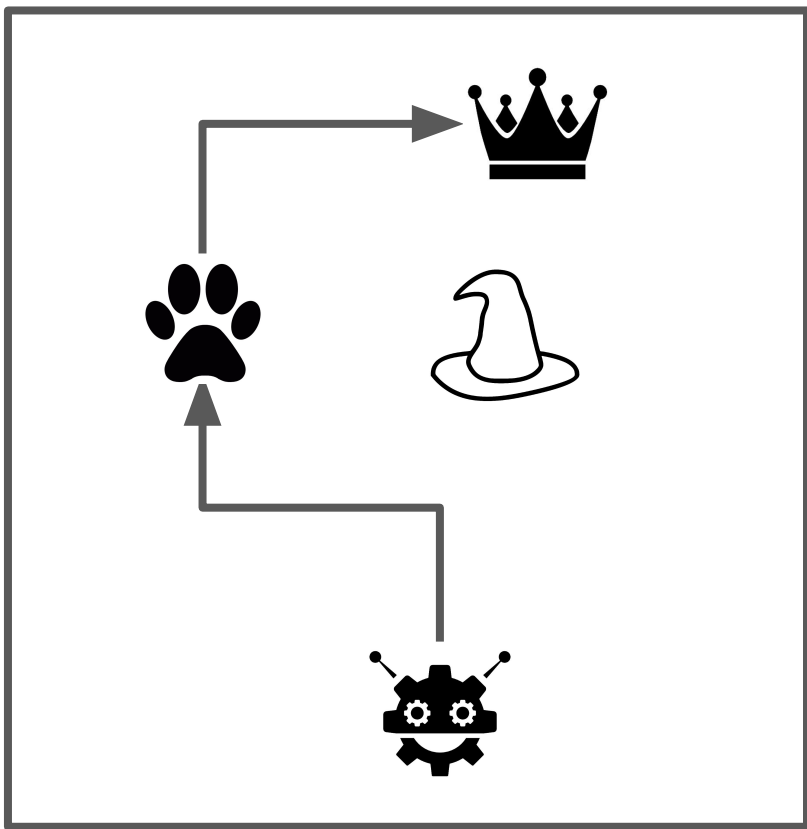
Objective: Get the message, bring it to the goal, while avoiding an enemy

The message, goal, and enemy, roles may be filled by different **entities** (e.g. wizard)

The agent must consult a **NL manual** to win:

- The wizard is an enemy fast approaching.
- The queen is the fleeing target
- The stationary dog is holding the secret message

Messenger



Objective: Get the message, bring it to the goal, while avoiding an enemy

The message, goal, and enemy, roles may be filled by different **entities** (e.g. wizard)

The agent must consult a **NL manual** to win:

- The wizard is an enemy fast approaching.
- The queen is the fleeing target
- The stationary dog is holding the secret message

44/32/32 train/val/test game variants

5000+ descriptions, **1,125** vocab size

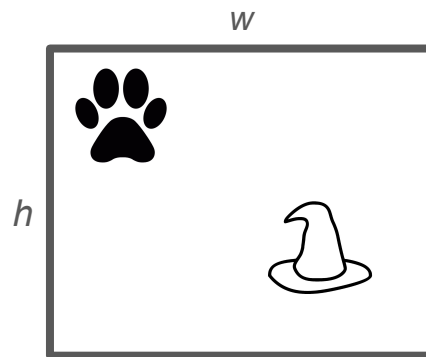
30-60 words/manual, **completely human written**

EMMA Model

Entity Mapper with Multi-modal Attn.

Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)



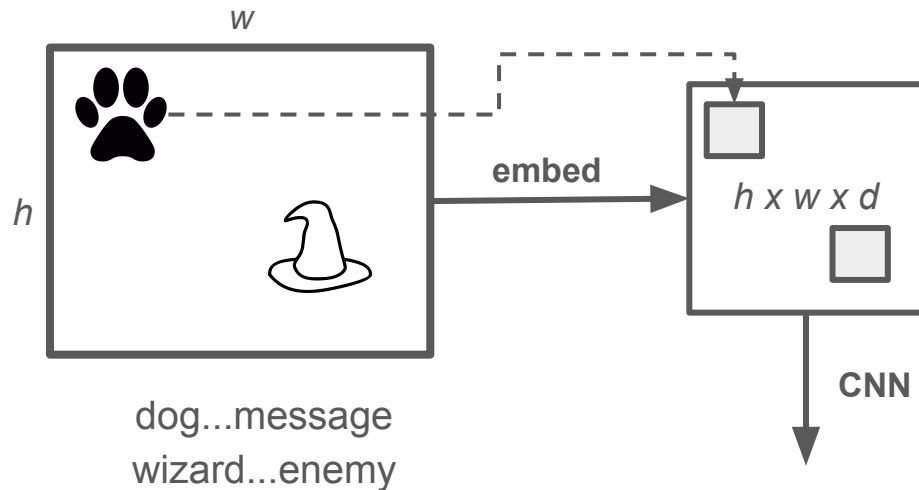
dog...message
wizard...enemy

EMMA Model

Entity Mapper with Multi-modal Attn.

Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)

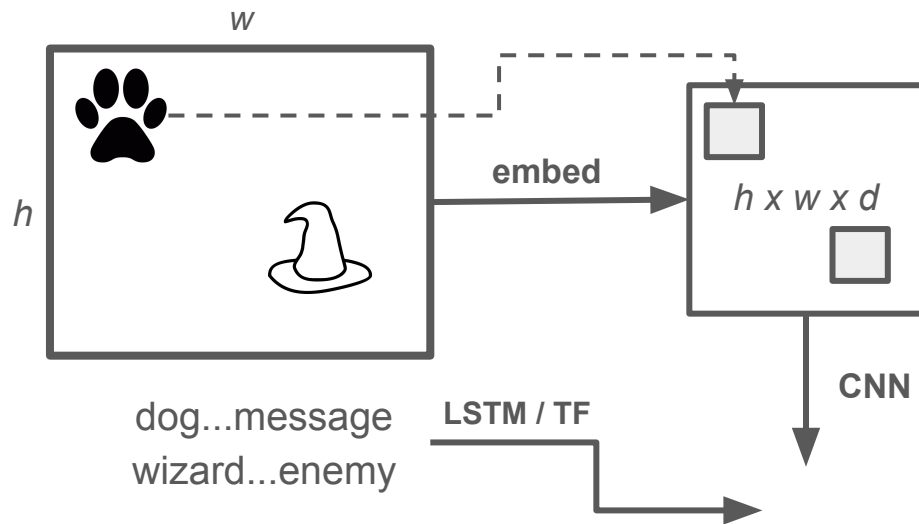


EMMA Model

Entity Mapper with Multi-modal Attn.

Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)

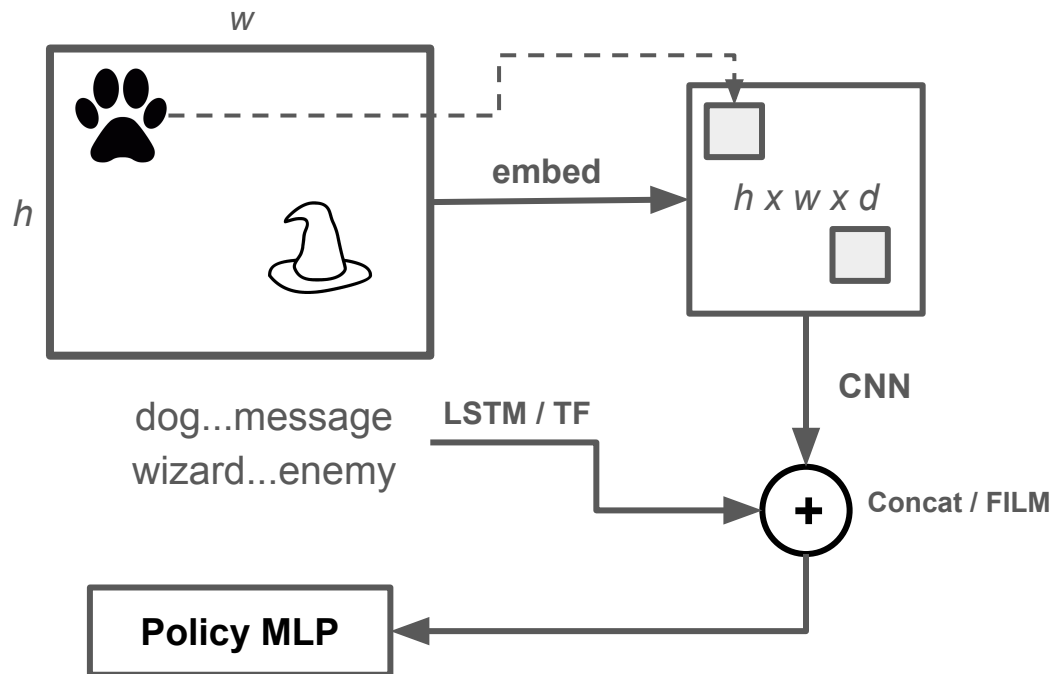


EMMA Model

Entity Mapper with Multi-modal Attn.

Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)

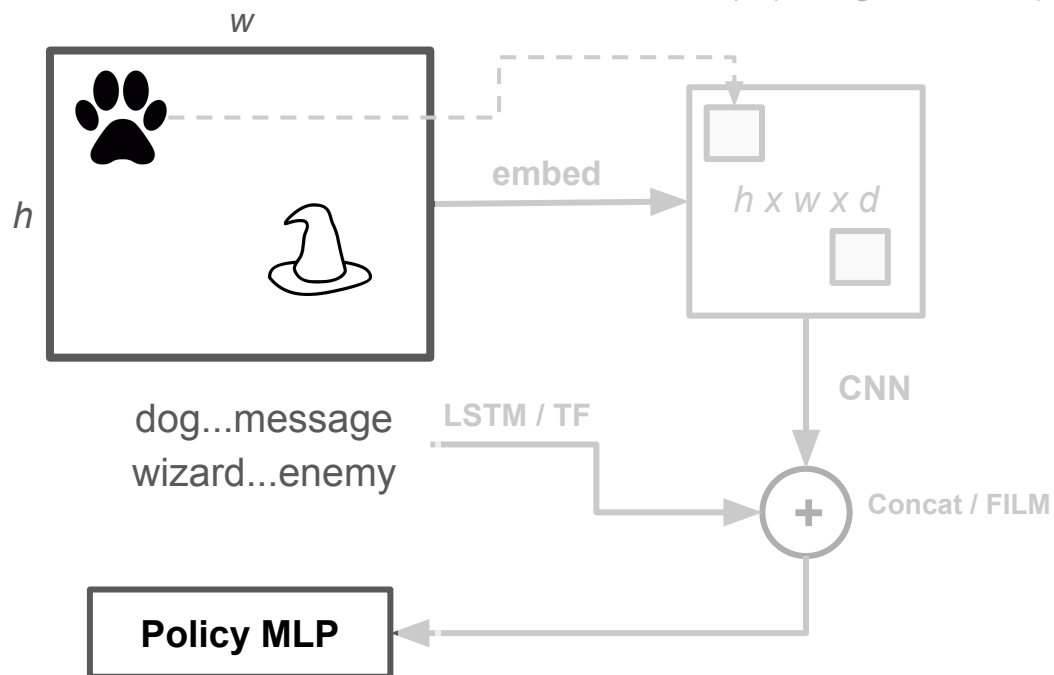


EMMA Model

Entity Mapper with Multi-modal Attn.

Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)

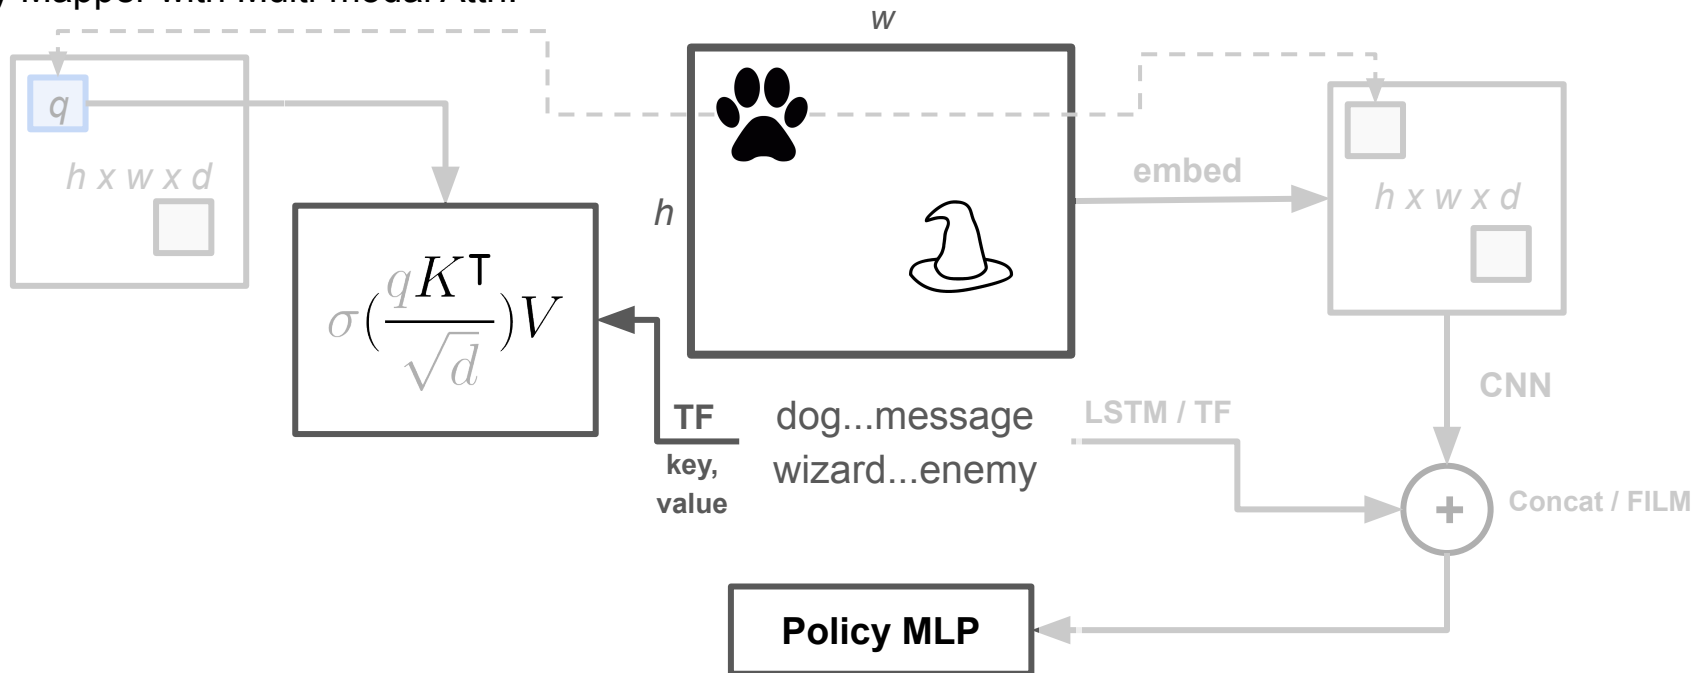


EMMA Model

Entity Mapper with Multi-modal Attn.

Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)

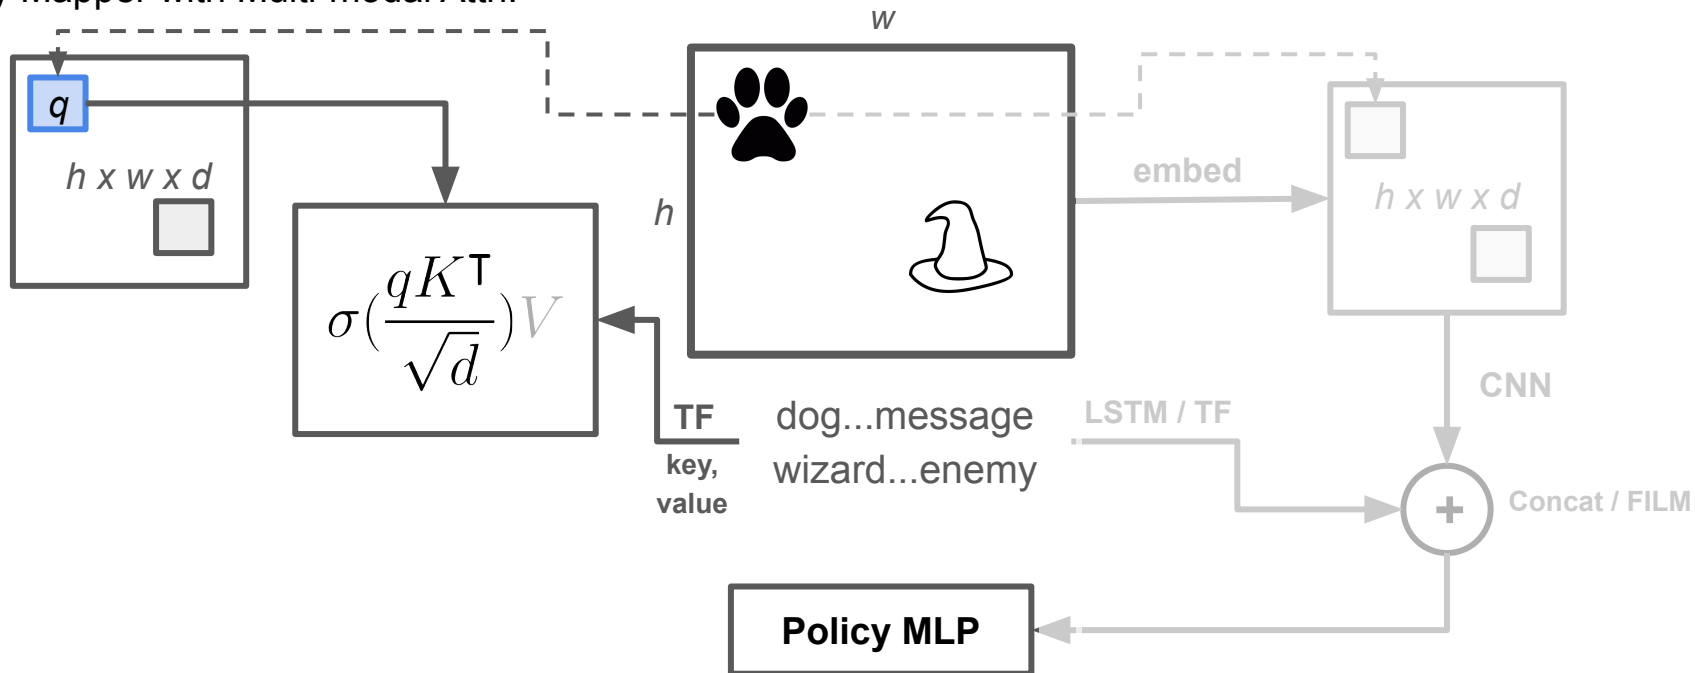


EMMA Model

Entity Mapper with Multi-modal Attn.

Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)

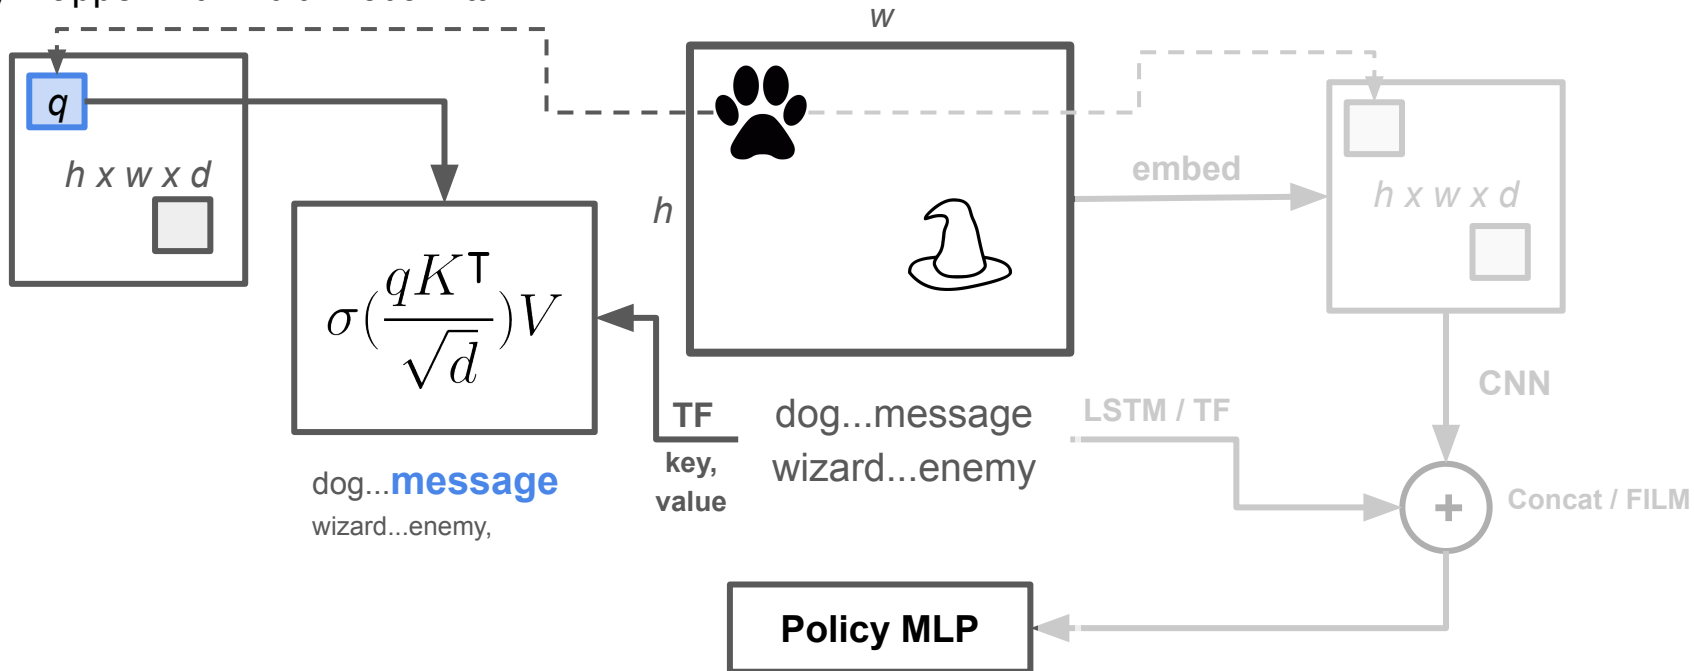


EMMA Model

Entity Mapper with Multi-modal Attn.

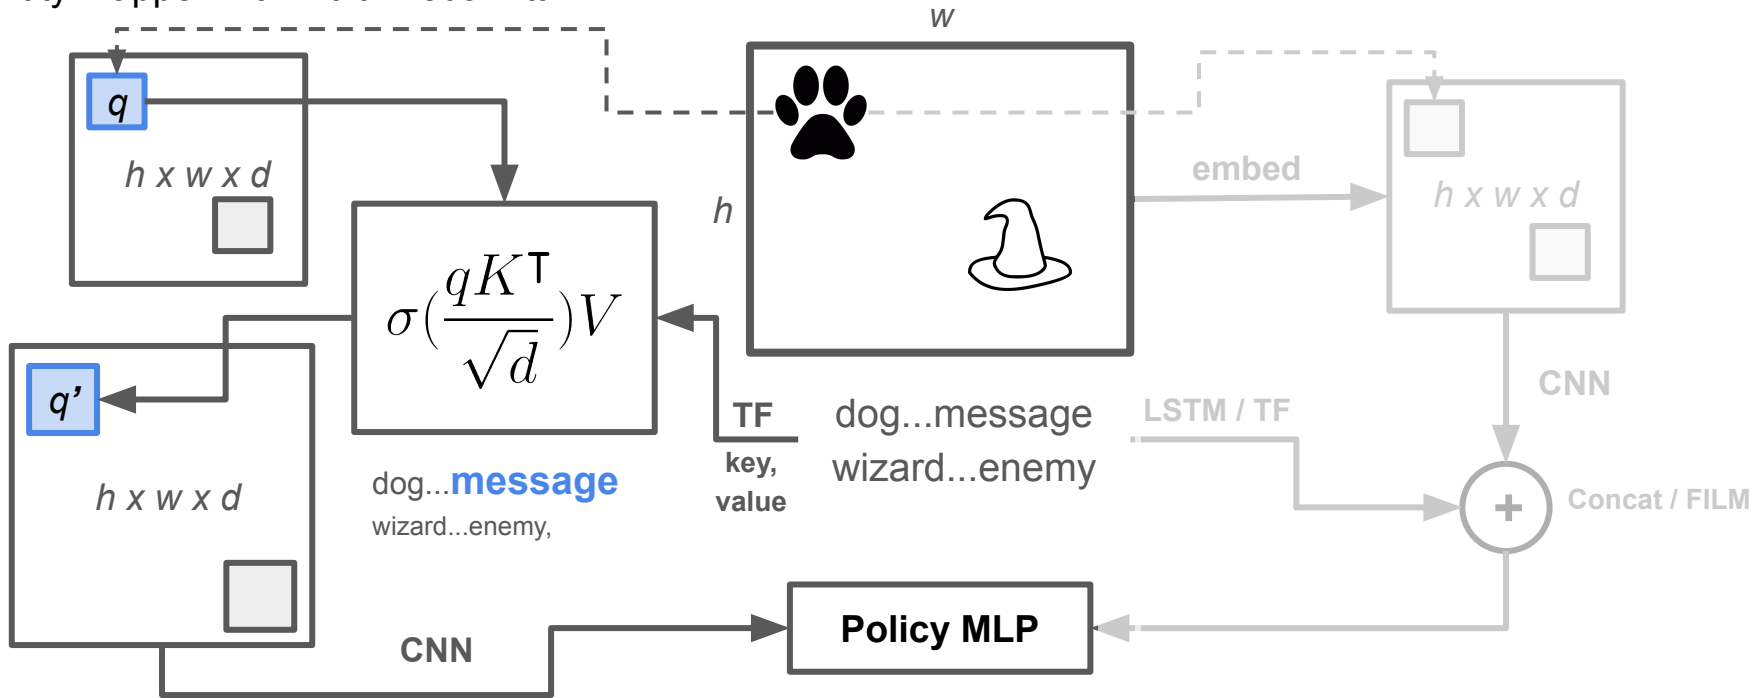
Other Approaches

G-ID, Mean-BOS, txt2pi (Zhong et al. 2020)



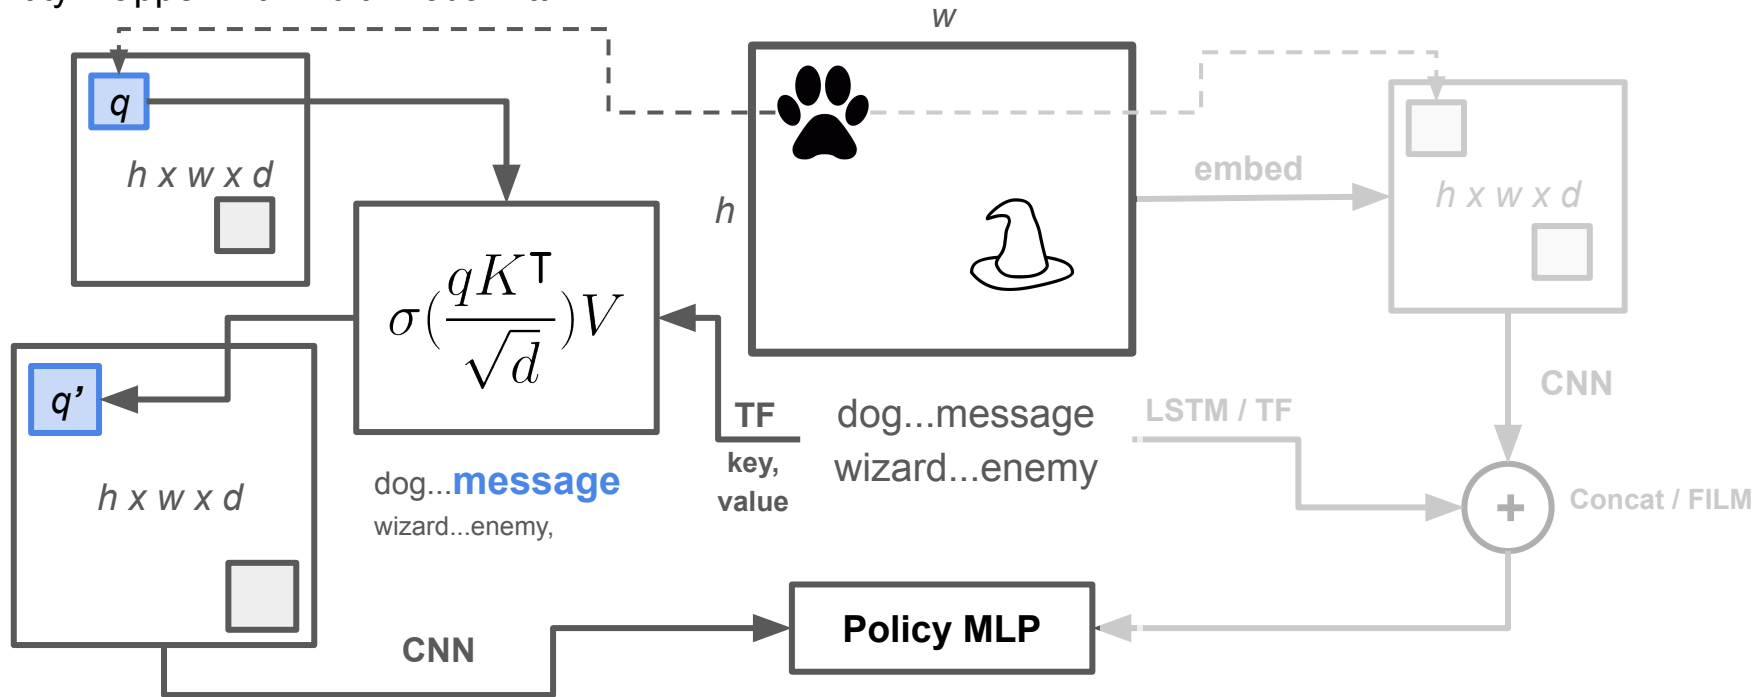
EMMA Model

Entity Mapper with Multi-modal Attn.



EMMA Model

Entity Mapper with Multi-modal Attn.



Key Idea: EMMA allows the agent to focus on **entity roles** (e.g. enemy), rather than making decisions off **entity identities** (e.g. wizard)

Results

Win rates (\pm stddev.) on stage 2 of *Messenger*

	Train	Test
Game-ID	3.6 ± 0.6	5.2 ± 0.2
Mean-Bag of Sentences	2.1 ± 0.5	4.7 ± 0.5
Bayesian Attention	69 ± 1.1	41 ± 1.7
Txt2pi (Zhong et al. 2020)	94 ± 3.5	0.3 ± 0.08
EMMA	95 ± 0.4	85 ± 0.6

Conclusions

Conclusions

- Studied use of **natural language manuals** to allow **RL agents to generalize** to new environments

Conclusions

- Studied use of **natural language manuals** to allow **RL agents to generalize** to new environments
- Introduced the ***Messenger*** environment with novel challenges in **entity grounding**

Conclusions

- Studied use of **natural language manuals** to allow **RL agents to generalize** to new environments
- Introduced the *Messenger* environment with novel challenges in **entity grounding**
- Introduced the **EMMA** model which can generalize to unseen games where prior approaches struggle.

Conclusions

- Studied use of **natural language manuals** to allow **RL agents to generalize** to new environments
- Introduced the ***Messenger*** environment with novel challenges in **entity grounding**
- Introduced the **EMMA** model which can generalize to unseen games where prior approaches struggle.
- Thank you for listening!
 - Paper at: <https://arxiv.org/abs/2101.07393>
 - Code at: <https://github.com/ahjwang/messenger-emma>