COS 484: Natural Language Processing

L19: Language Grounding - 2

Spring 2022

Logistics

- Sign up for project meetings on April 19
  - Mandatory for every team to meet with your staff guide
- Fill up preference form for poster session on April 21

Some grounding tasks

- **Vision**
  - Captioning
  - Visual question answering (VQA)
  - Spatial reasoning
- **Interaction**
  - Instruction following
  - Text-based games
  - RL for NLP

Instruction Following

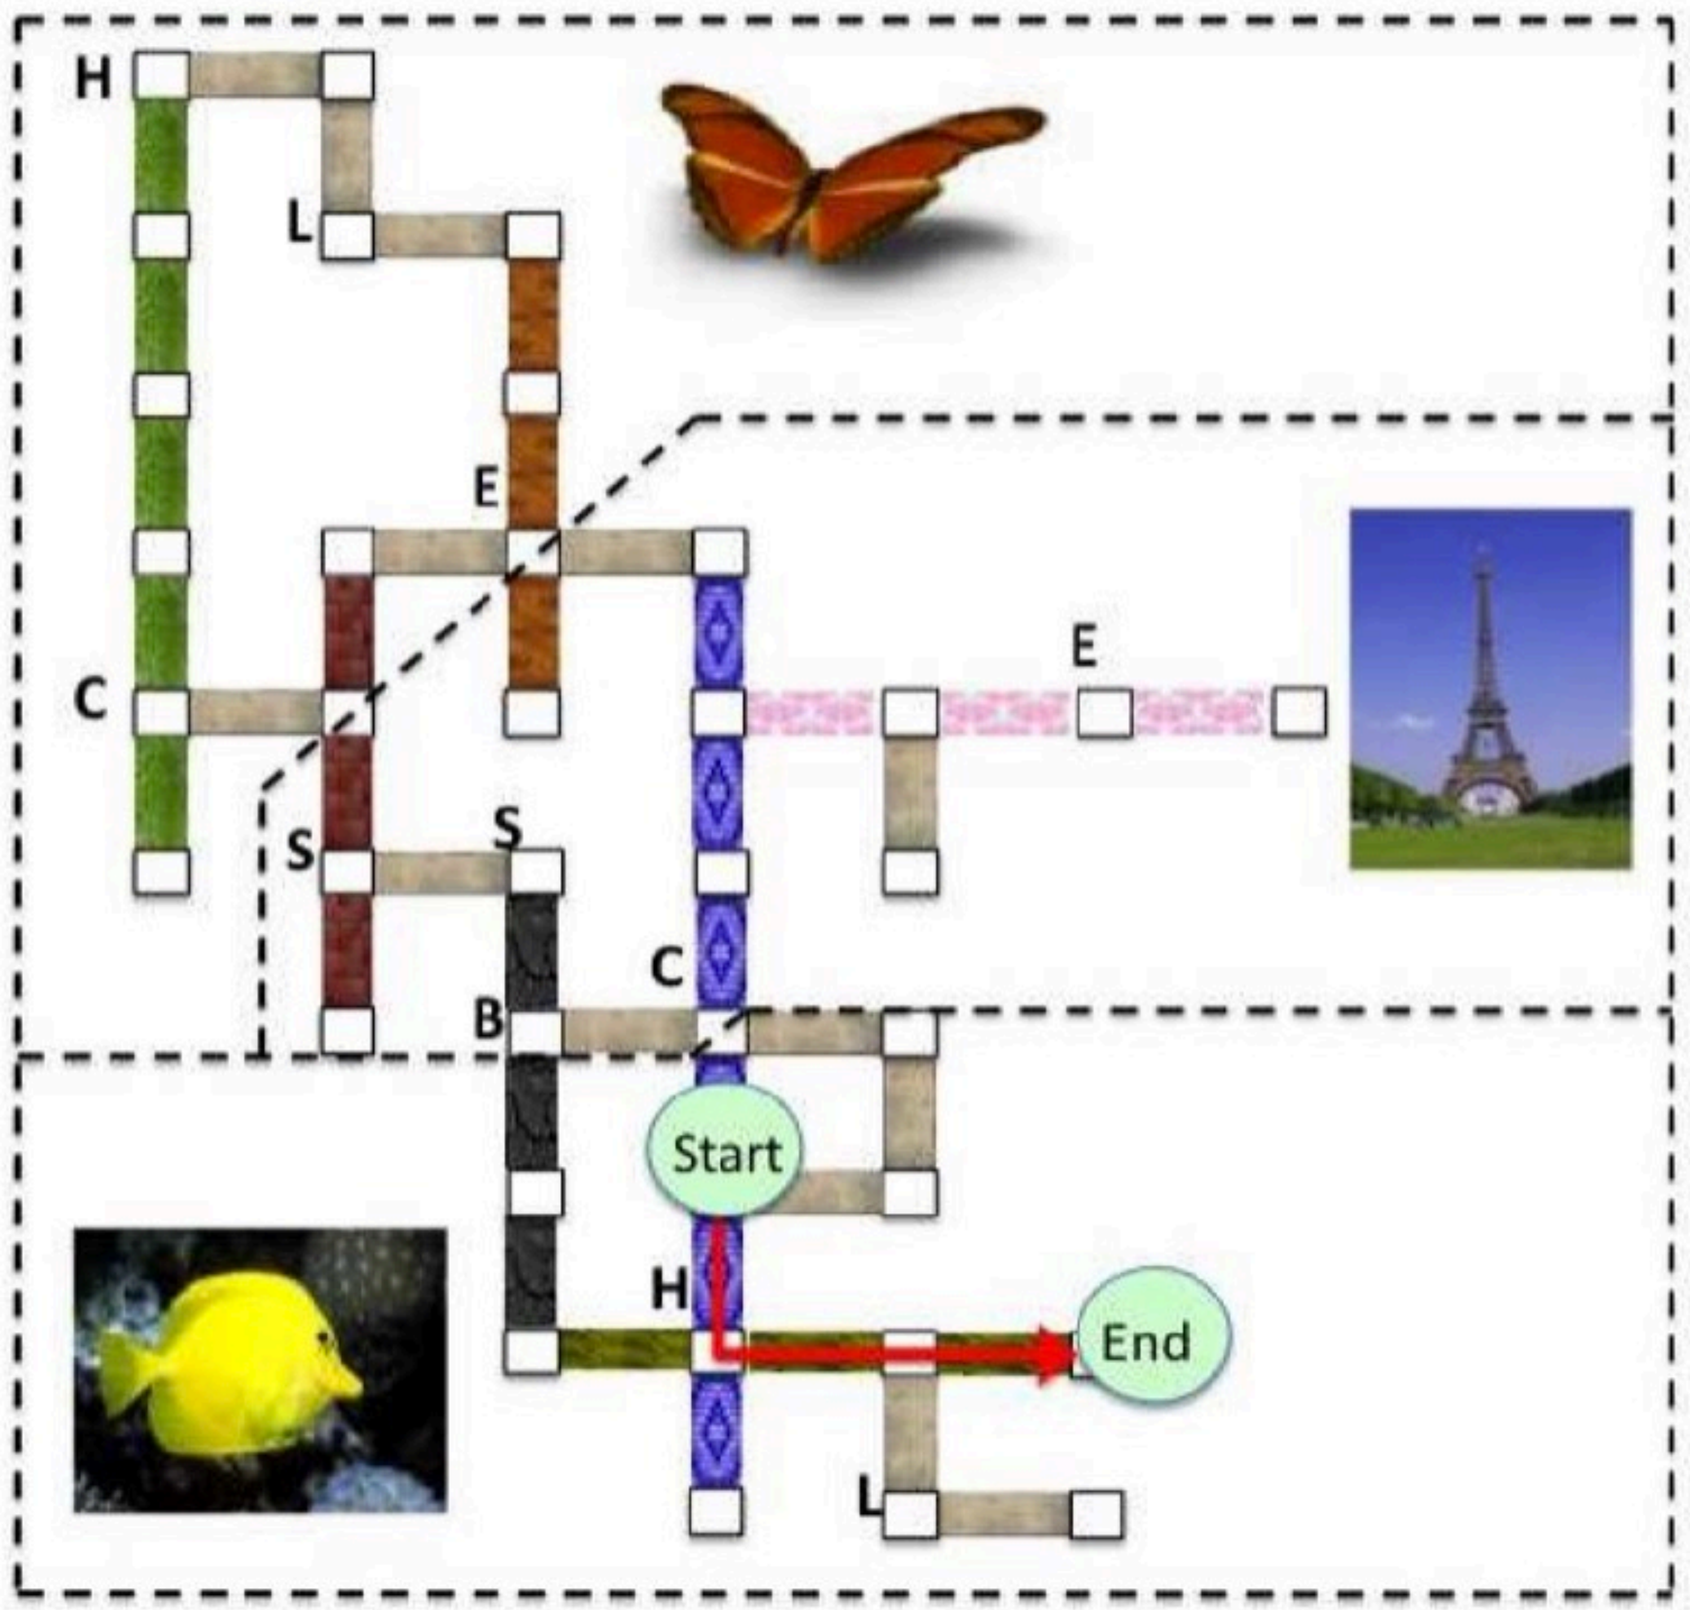

- Want to be able to follow instructions in a virtual environment
- "Go along the blue hall, then turn left away from the fish painting and walk to the end of the hallway"

Grounding language to actions

(MacMahon et al., 2006)

Instruction Following

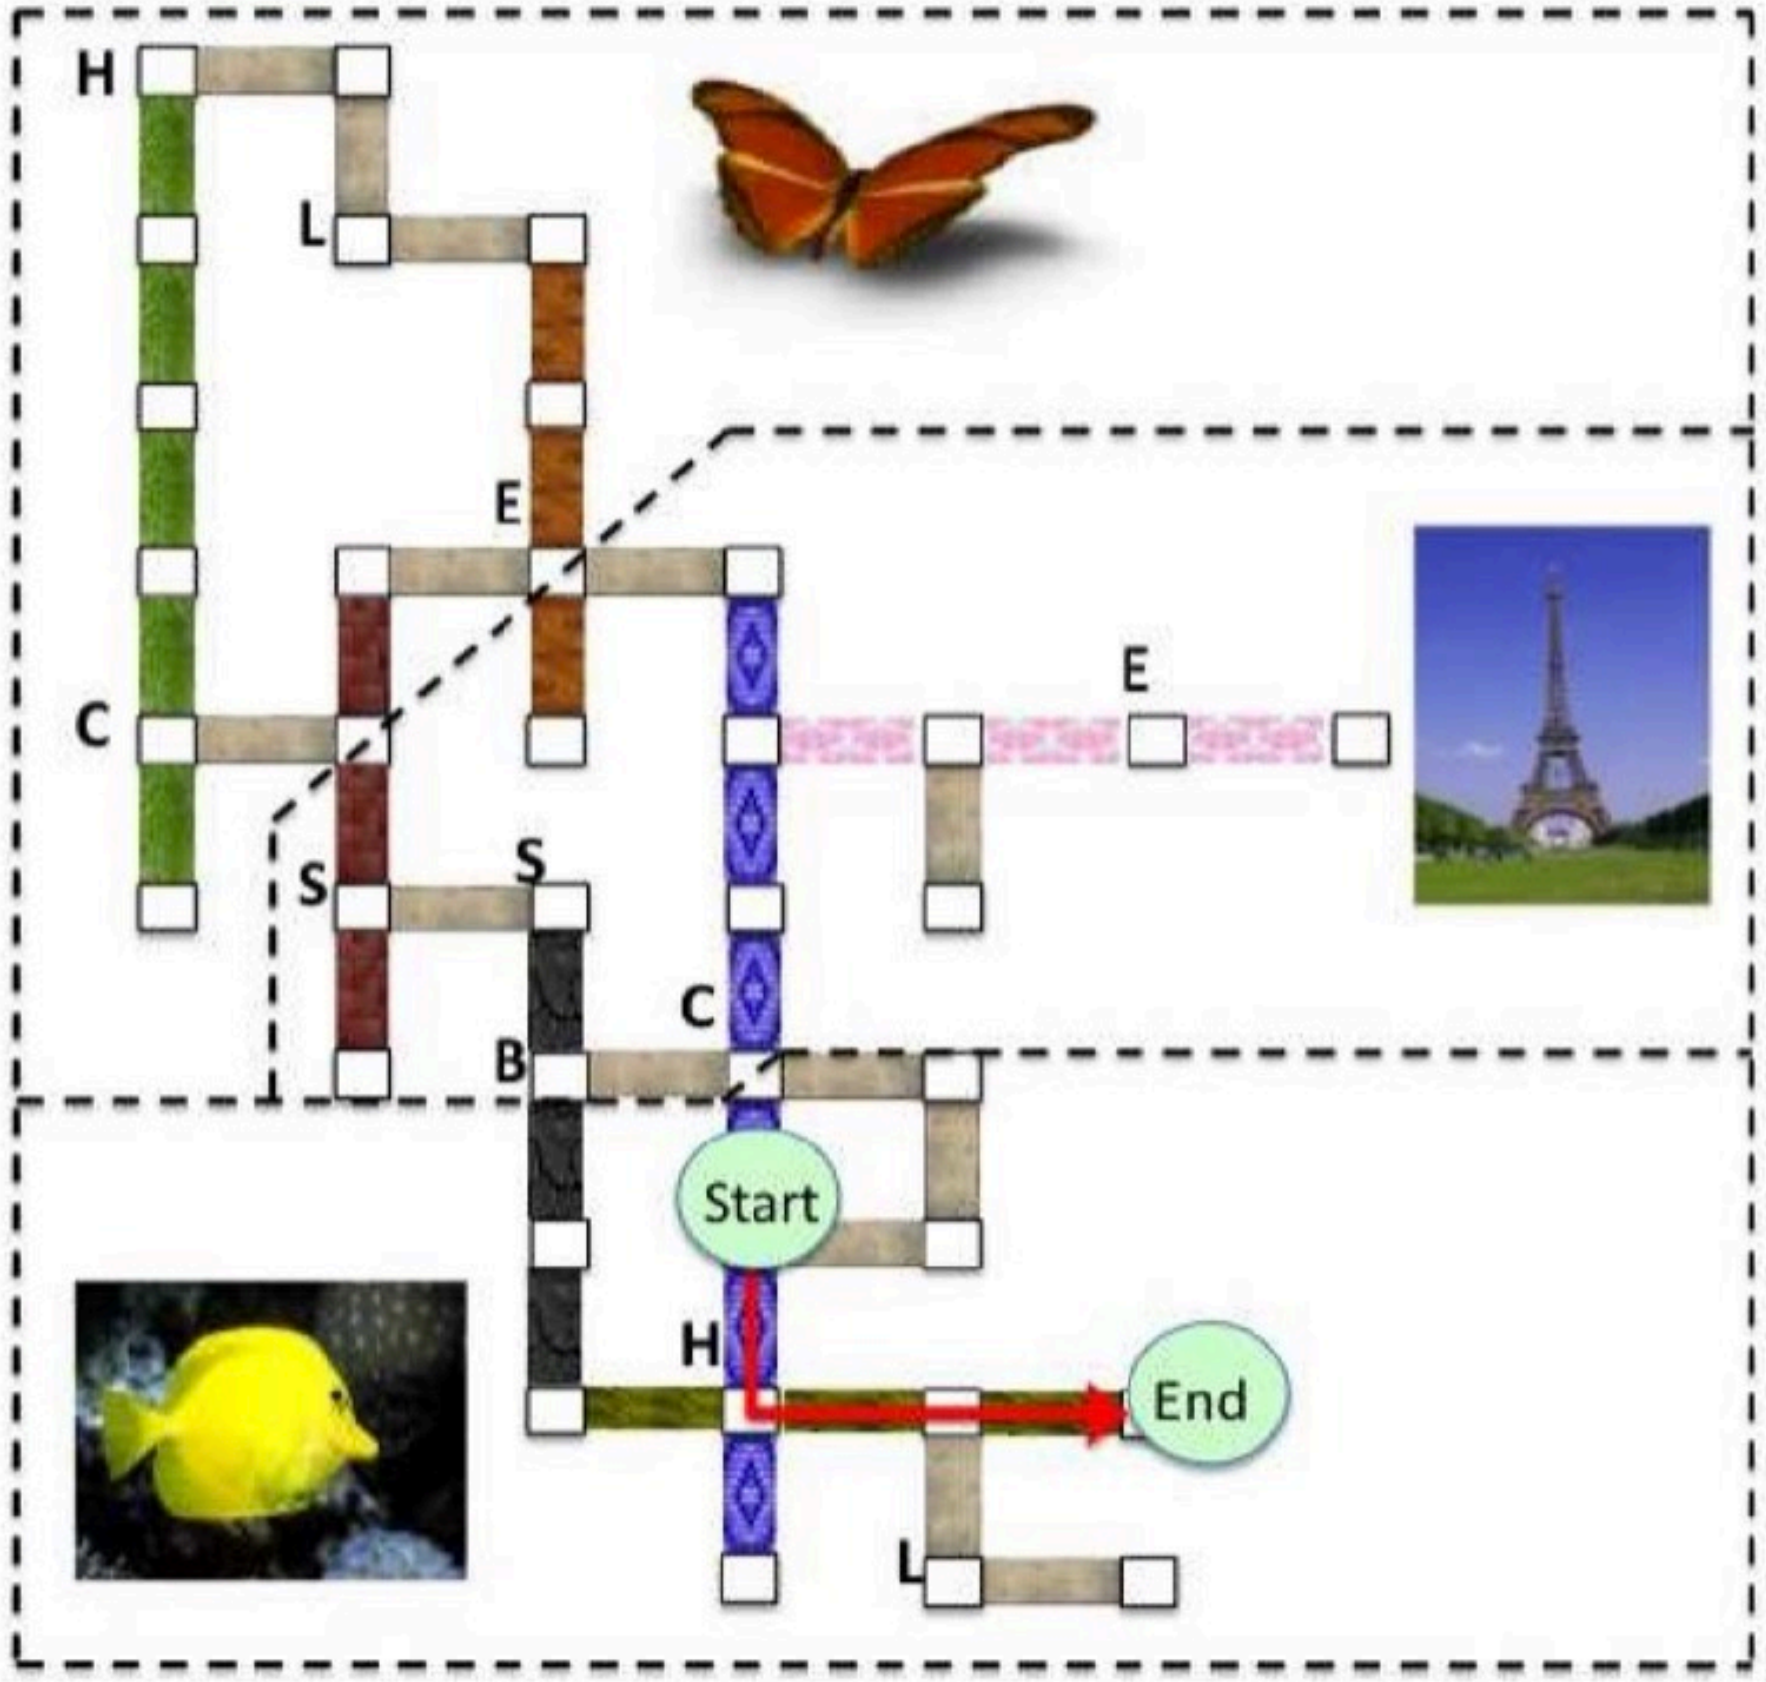

Instruction: "Go away from the lamp to the intersection of the red brick and wood"

Basic: Turn ( ) ,
Travel (steps: 1)

Landmarks: Turn ( ) ,
Verify (left: WALL , back: LAMP , back: HATRACK , front: BRICK HALL) ,
Travel (steps: 1) ,
Verify (side: WOOD HALL)

- Train semantic parser on (utterance, action) pairs

(Chen and Mooney, 2011)

Grounding semantics in control applications

1. Use feedback from task to understand language

Walk across the bridge

Reward
+1

Alleviate dependence on supervised annotation

2. Use language to improve performance in control applications

Score: 7

+

1. **Ghosts** chase and try to kill you
2. Collect all the **pellets**
3. …

Score: 107

Reinforcement learning

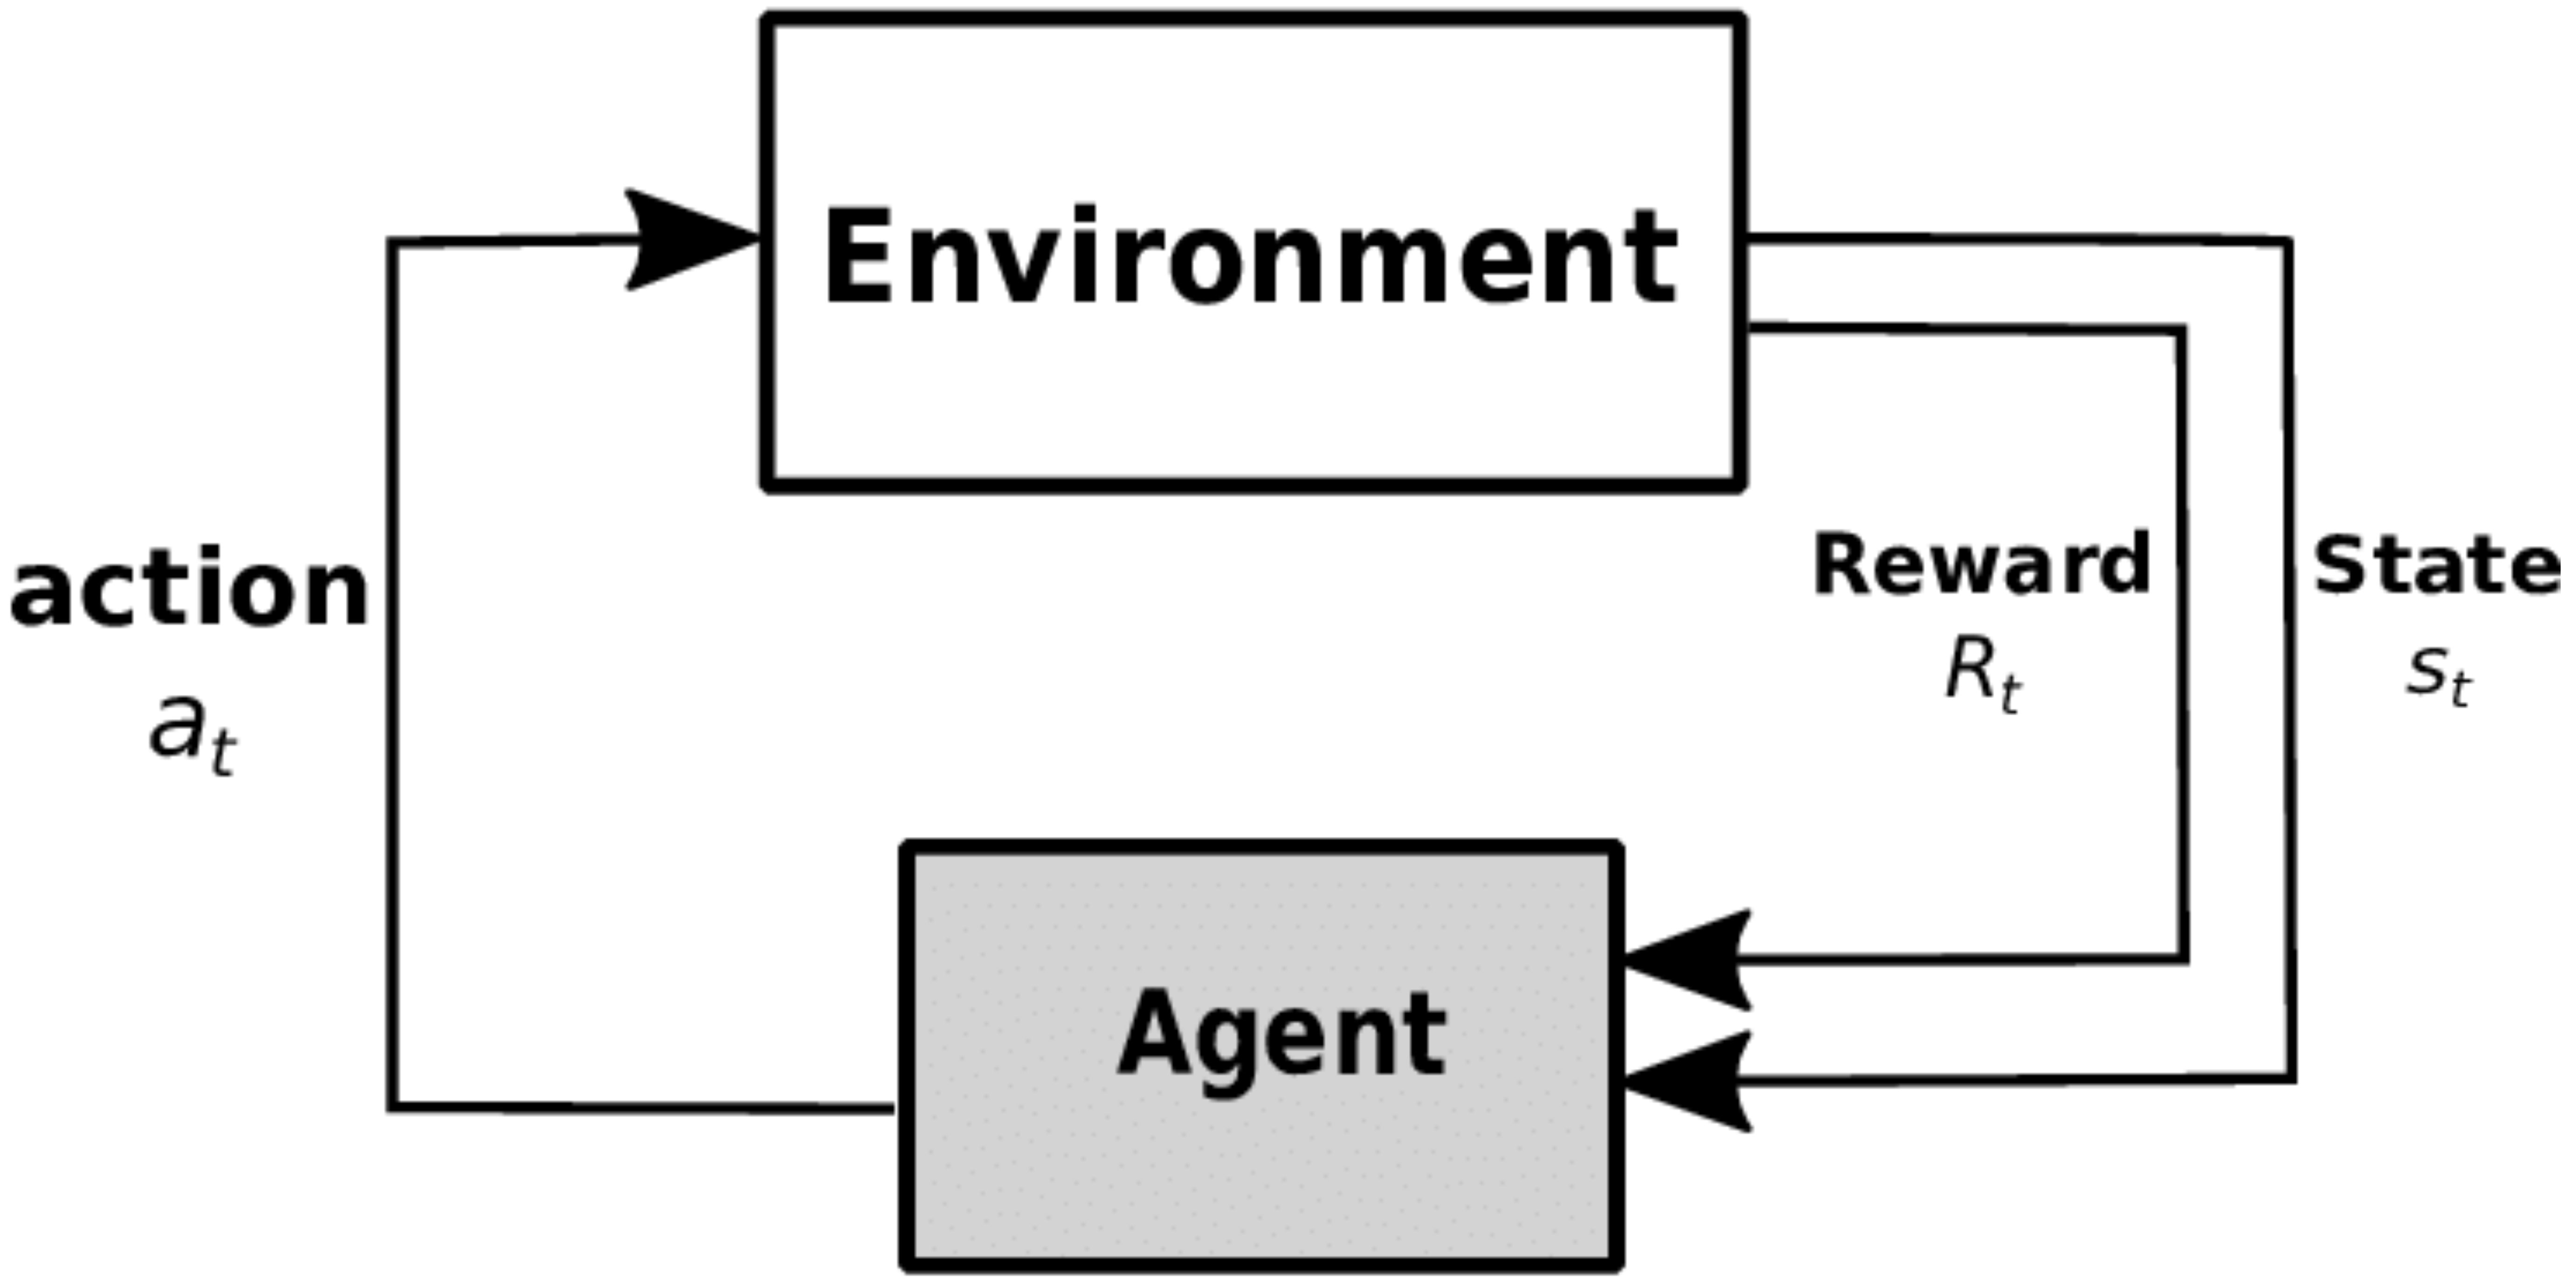

Reinforcement Learning

- Delayed feedback

action 1

○○○○

action n

Reward +10

$\Rightarrow$ *How to perform credit assignment for individual actions*

- Large number of possible action sequences

$\Rightarrow$ *Need for effective exploration*

Improved language understanding translates
to improved task performance

Playing Civilization by reading game manuals

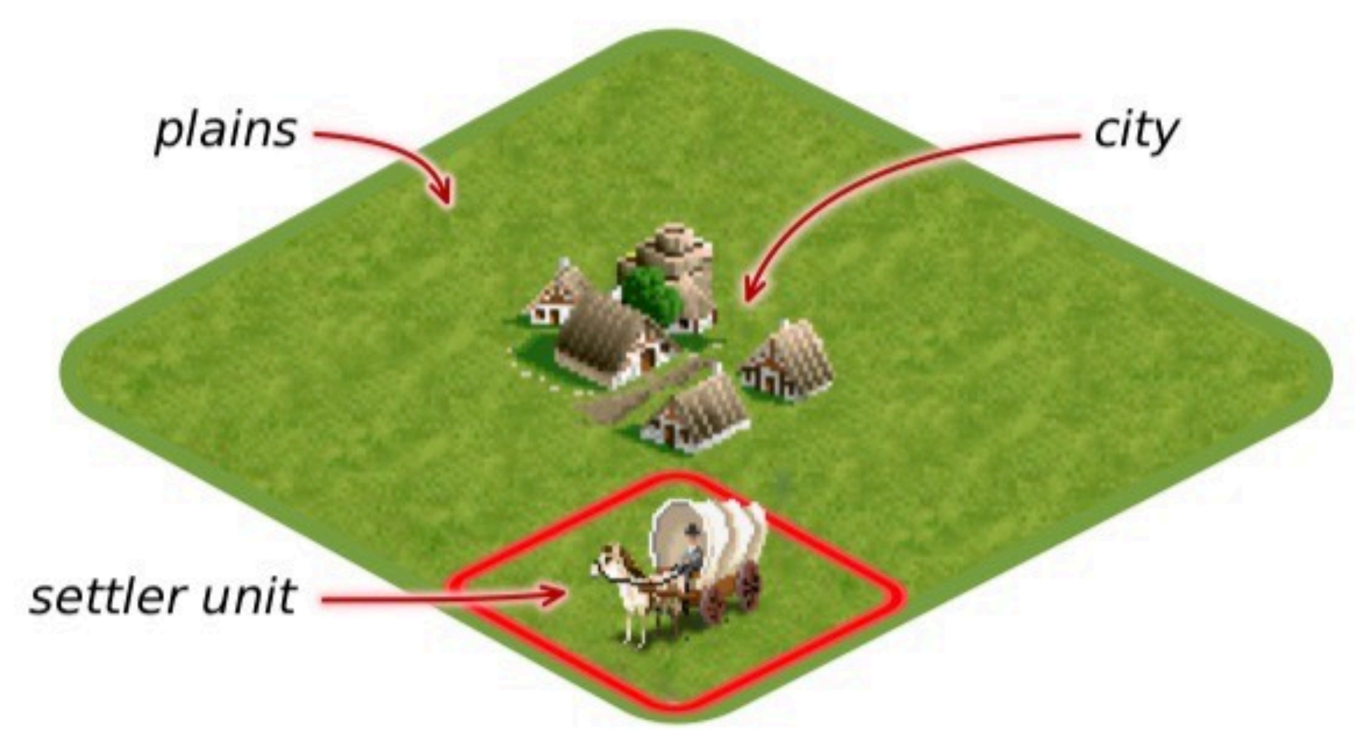

Relevant text: *"Use settlers to irrigate land near your city"*

Predicted action words: *"irrigate"*, *"settler"*

Predicted state words: *"land"*, *"near"*, *"city"*

Settlers unit, candidate action 1: ***irrigate***

Features:
action = **irrigate** and action-word = "*irrigate*"
action = **irrigate** and state-word = "*land*"
action = **irrigate** and terrain = plains
action = **irrigate** and unit-type = settler
state-word = "*city*" and near-city = true

Settlers unit, candidate action 2: ***build-city***

Features:
action = **build-city** and action-word = "*irrigate*"
action = **build-city** and state-word = "*land*"
action = **build-city** and terrain = plains
action = **build-city** and unit-type = settler
state-word = "*city*" and near-city = true

Neural network for policy

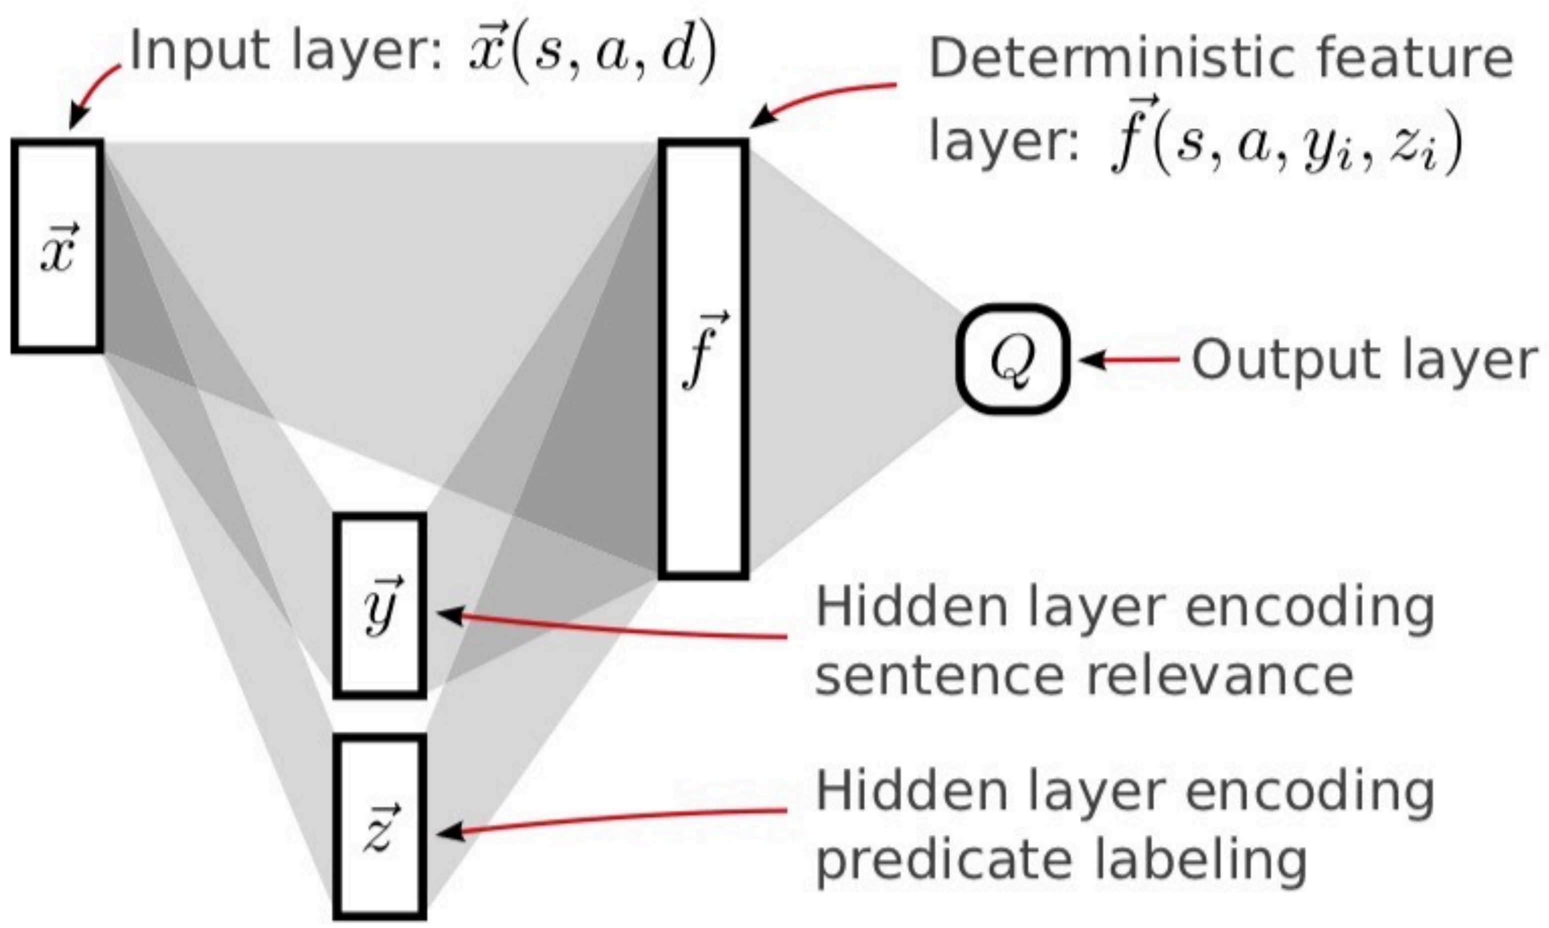

Method	% Win	% Loss	Std. Err.
Random	0	100	—
Built-in AI	0	0	—
Game only	17.3	5.3	± 2.7
Latent variable	26.1	3.7	± 3.1
Full model	**53.7**	5.9	± 3.5
Randomized text	40.3	4.3	± 3.4

(Branavan et al., 2012)

Learning a grounding

*The **dangerous enemy** is the **alien** that is inching near you.*

*The **wolf** is running from you while holding a **secret message**.*

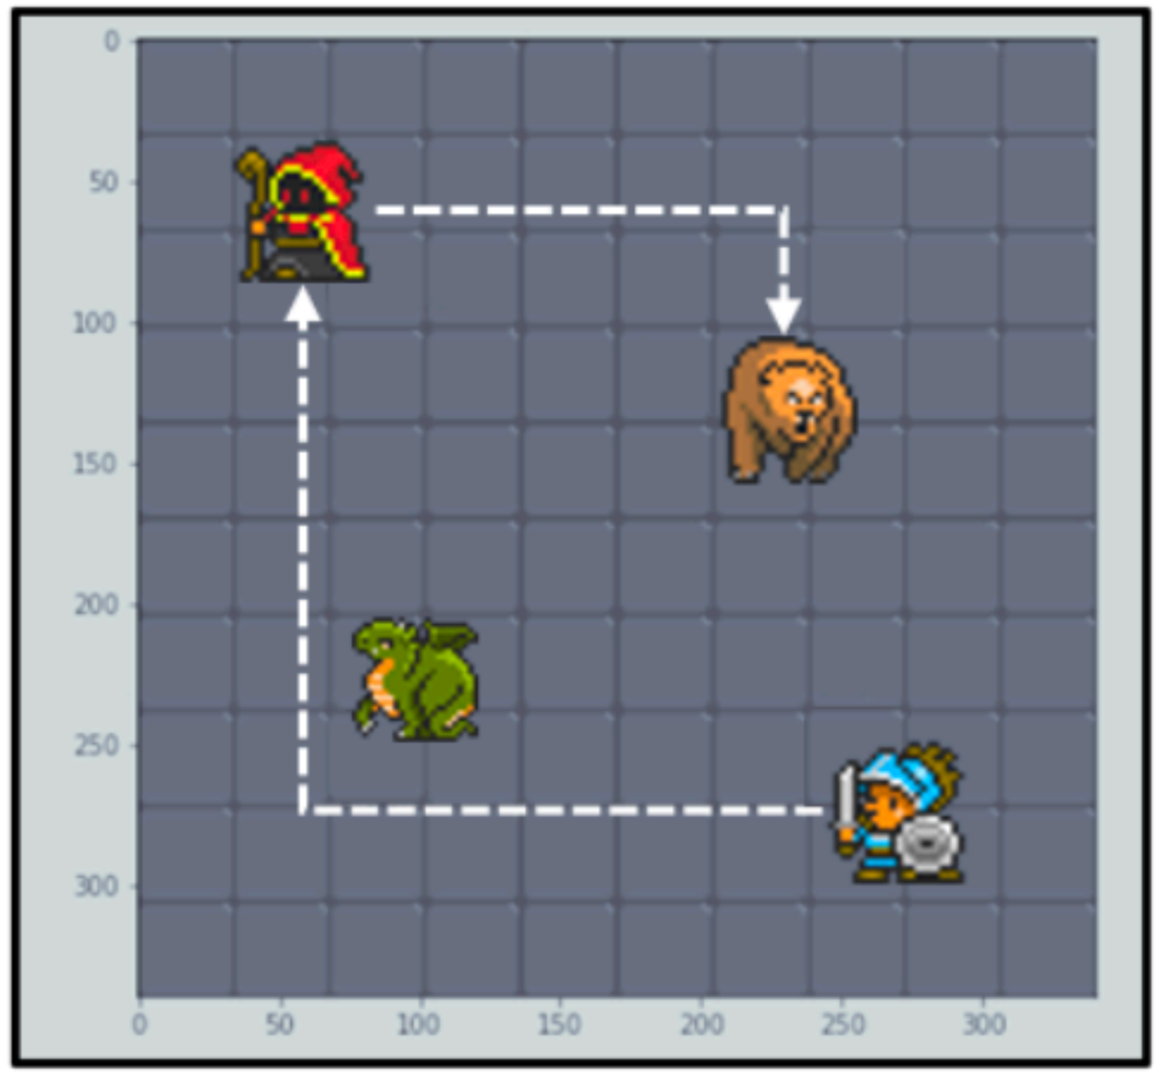

*The **bear** that is coming near you is the **crucial goal**.*

*The **dragon** which is running away is a **adversary** and the **adversary** is **deadly**.*

- How do we map symbols in language (i.e. words) to entities and concepts in the world?
- Can an agent learn grounding through interaction

[*Grounding Language to Entities and Dynamics for Generalization in Reinforcement Learning.* Austin W. Hanjie, Victor Zhong, Karthik Narasimhan; ICML 2021]

GAME 1 MANUAL

1. at a particular locale, there exists a motionless mongrel that is a formidable adversary.
2. the top-secret paperwork is in the crook's possession, and he's heading closer and closer to where you are.
3. the crucial target is held by the wizard and the wizard is fleeing from you.
4. the mugger rushing away is the opposition posing a serious threat.
5. the thing that is not able to move is the mage who possesses the enemy that is deadly.
6. *the vital goal is found with the canine, but it is running away from you.*

Messenger

- Agent can move around and interact in a simulated environment
- Receives global state observations, rewards
- Has access to a text "manual" describing entities and dynamics, throughout an episode
- Agent is not provided any prior mapping between the observations () and symbols in text (*wizard, mage*) to help it "read" the manual.

[*Grounding Language to Entities and Dynamics for Generalization in Reinforcement Learning.* Austin W. Hanjie, Victor Zhong, Karthik Narasimhan; ICML 2021]

Messenger

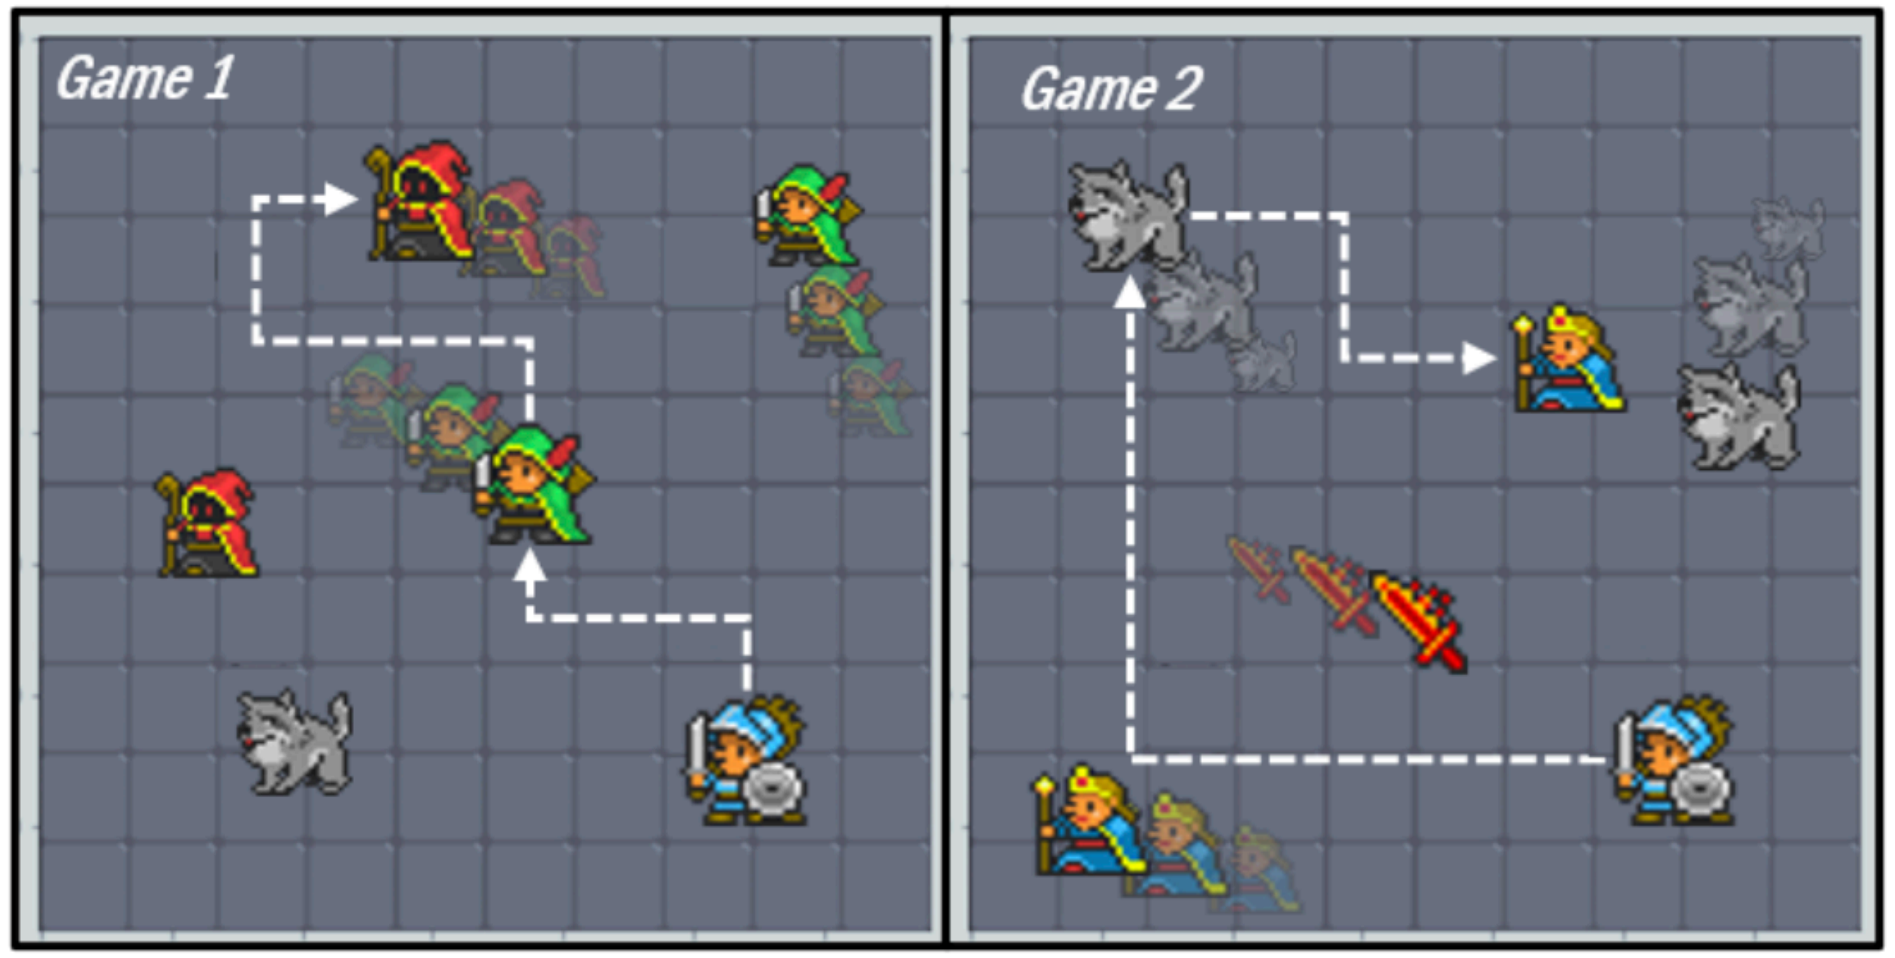

GAME 1 MANUAL

1. at a particular locale, there exists a motionless mongrel that is a formidable adversary.
2. the top-secret paperwork is in the crook's possession, and he's heading closer and closer to where you are.
3. the crucial target is held by the wizard and the wizard is fleeing from you.
4. the mugger rushing away is the opposition posing a serious threat.
5. the thing that is not able to move is the mage who possesses the enemy that is deadly.
6. *the vital goal is found with the canine, but it is running away from you.*

- Multi-game benchmark with separate train and test splits
- In each game, agent has to first pick up a *message*, and deliver it to *goal* entity, while avoiding an *enemy*
- Each game has different entities, each with different roles and different dynamics
    - *There may be multiple entities of the same type! (e.g. mage in game 1)*
- The agent must consult a natural language manual in order to consistently win
    - *Manual may contain extraneous/incorrect information (e.g. point 6 here).*

Messenger: Statistics

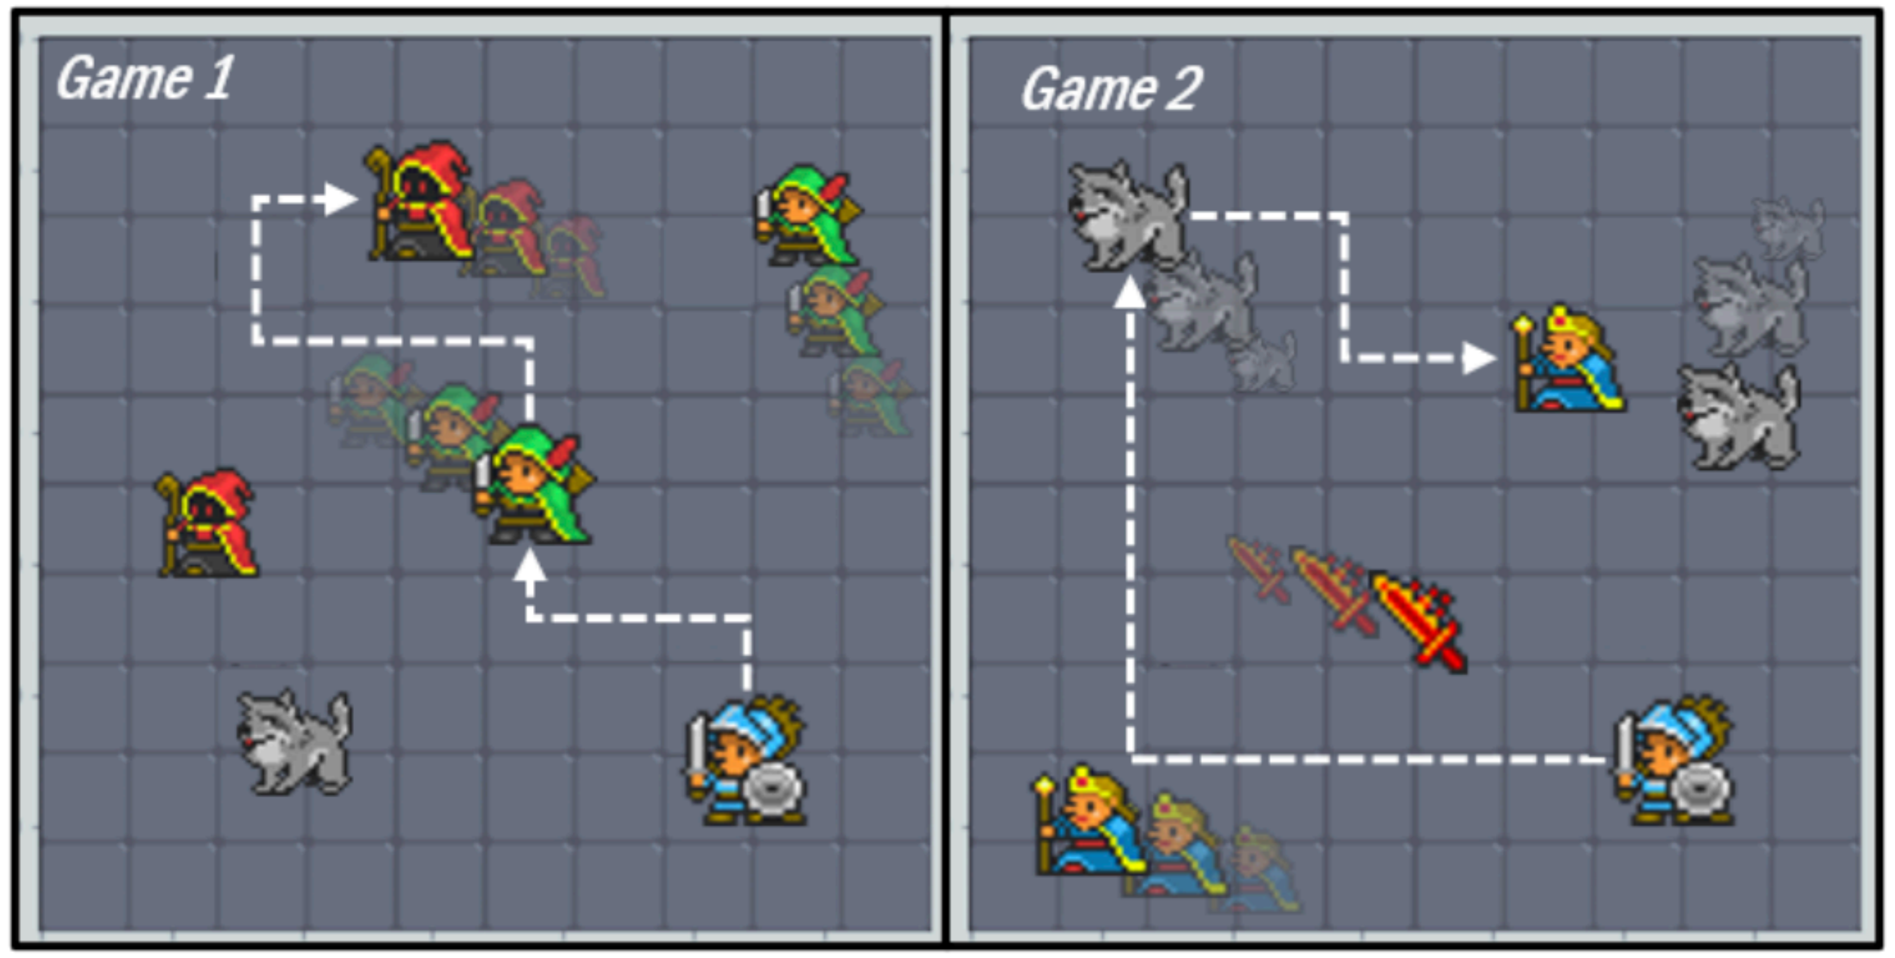

GAME 1 MANUAL

1. at a particular locale, there exists a motionless mongrel that is a formidable adversary.
2. the top-secret paperwork is in the crook's possession, and he's heading closer and closer to where you are.
3. the crucial target is held by the wizard and the wizard is fleeing from you.
4. the mugger rushing away is the opposition posing a serious threat.
5. the thing that is not able to move is the mage who possesses the enemy that is deadly.
6. *the vital goal is found with the canine, but it is running away from you.*

- Random instantiation of roles each time
- 44/32/32 train/val/test game variants
- 5000+ textual descriptions, vocabulary size of 1125
- 30-60 words/manual, completely human written (crowdsourced)

Why is Messenger challenging?

> *"The top-secret paperwork is in the crook's possession, and he's heading closer and closer to where you are"*

- Agent has to learn an accurate grounding purely through interaction
- Wide variation in how an entity is described - e.g. use of multiple synonyms (*crook, thief*), non-templated freeform text
- No overlap in terms of entity-role-dynamics combinations between train and test games

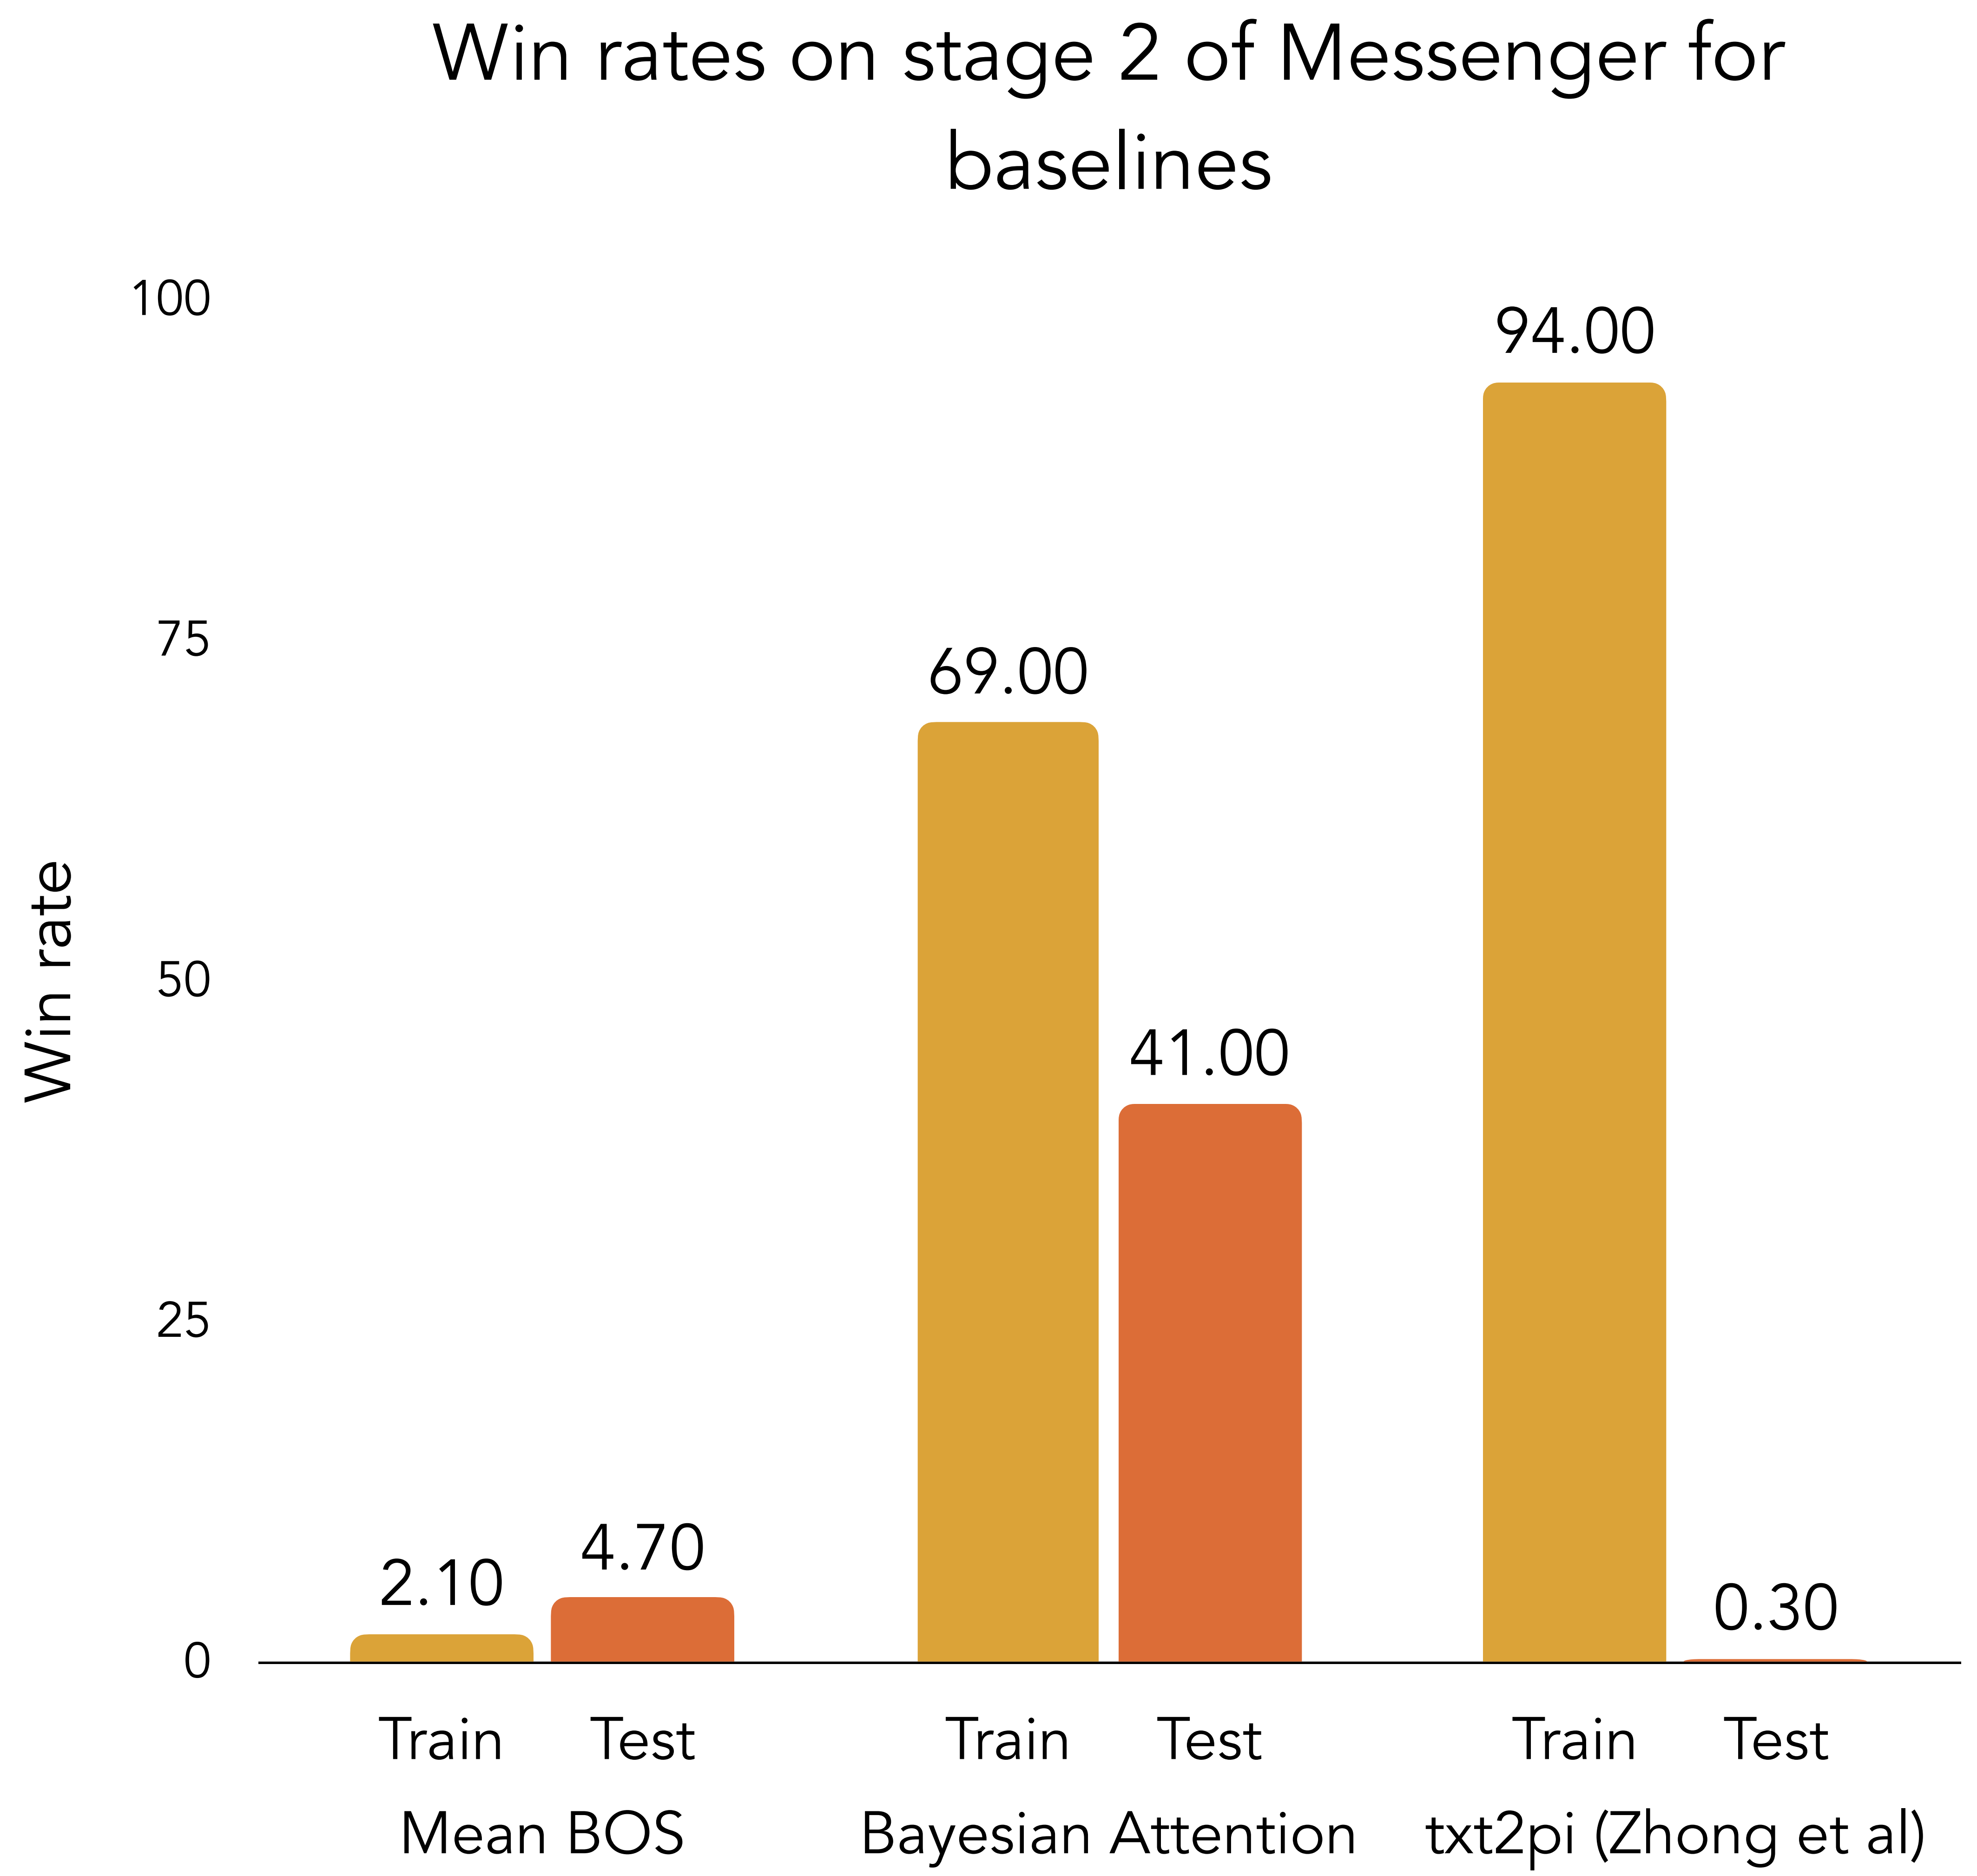

Our model: Entity Mapper with Multimodal Attention (EMMA)

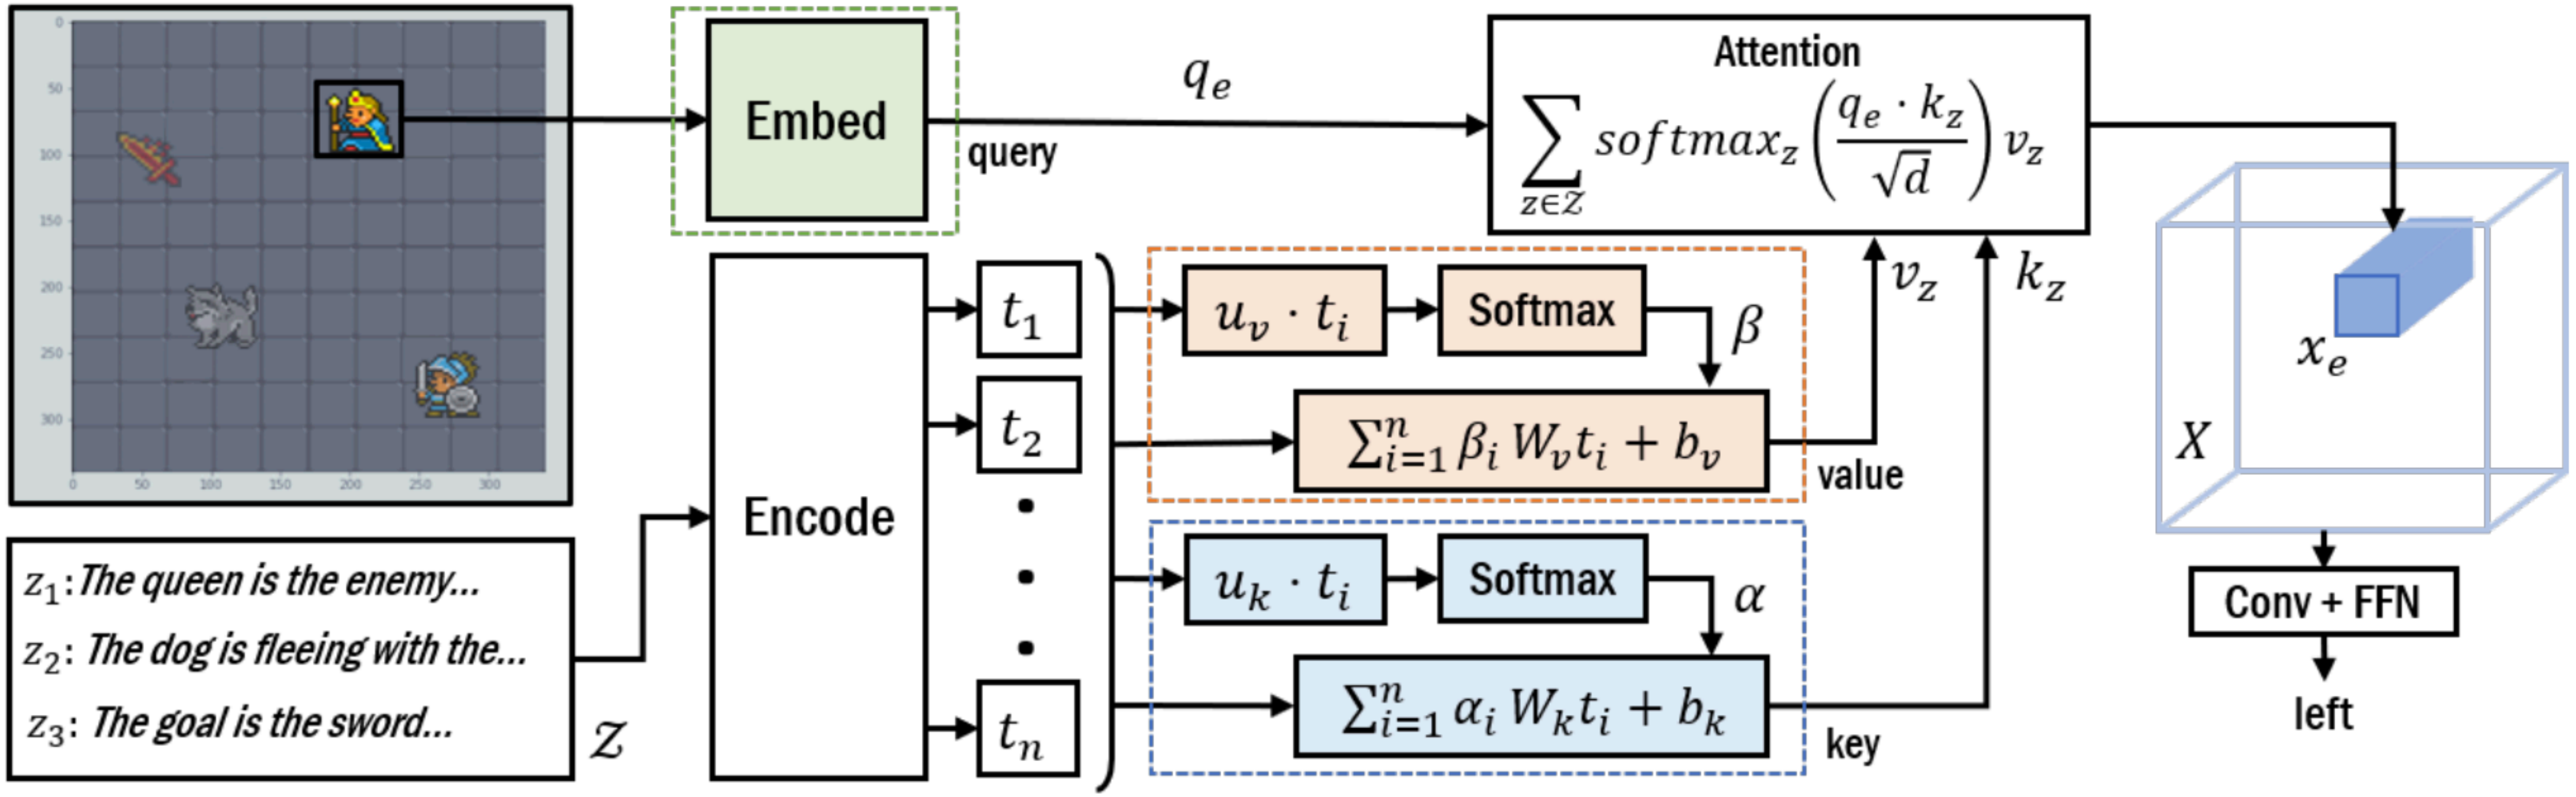

Jointly process observations with text manual for control policy

EMMA does better on Messenger…

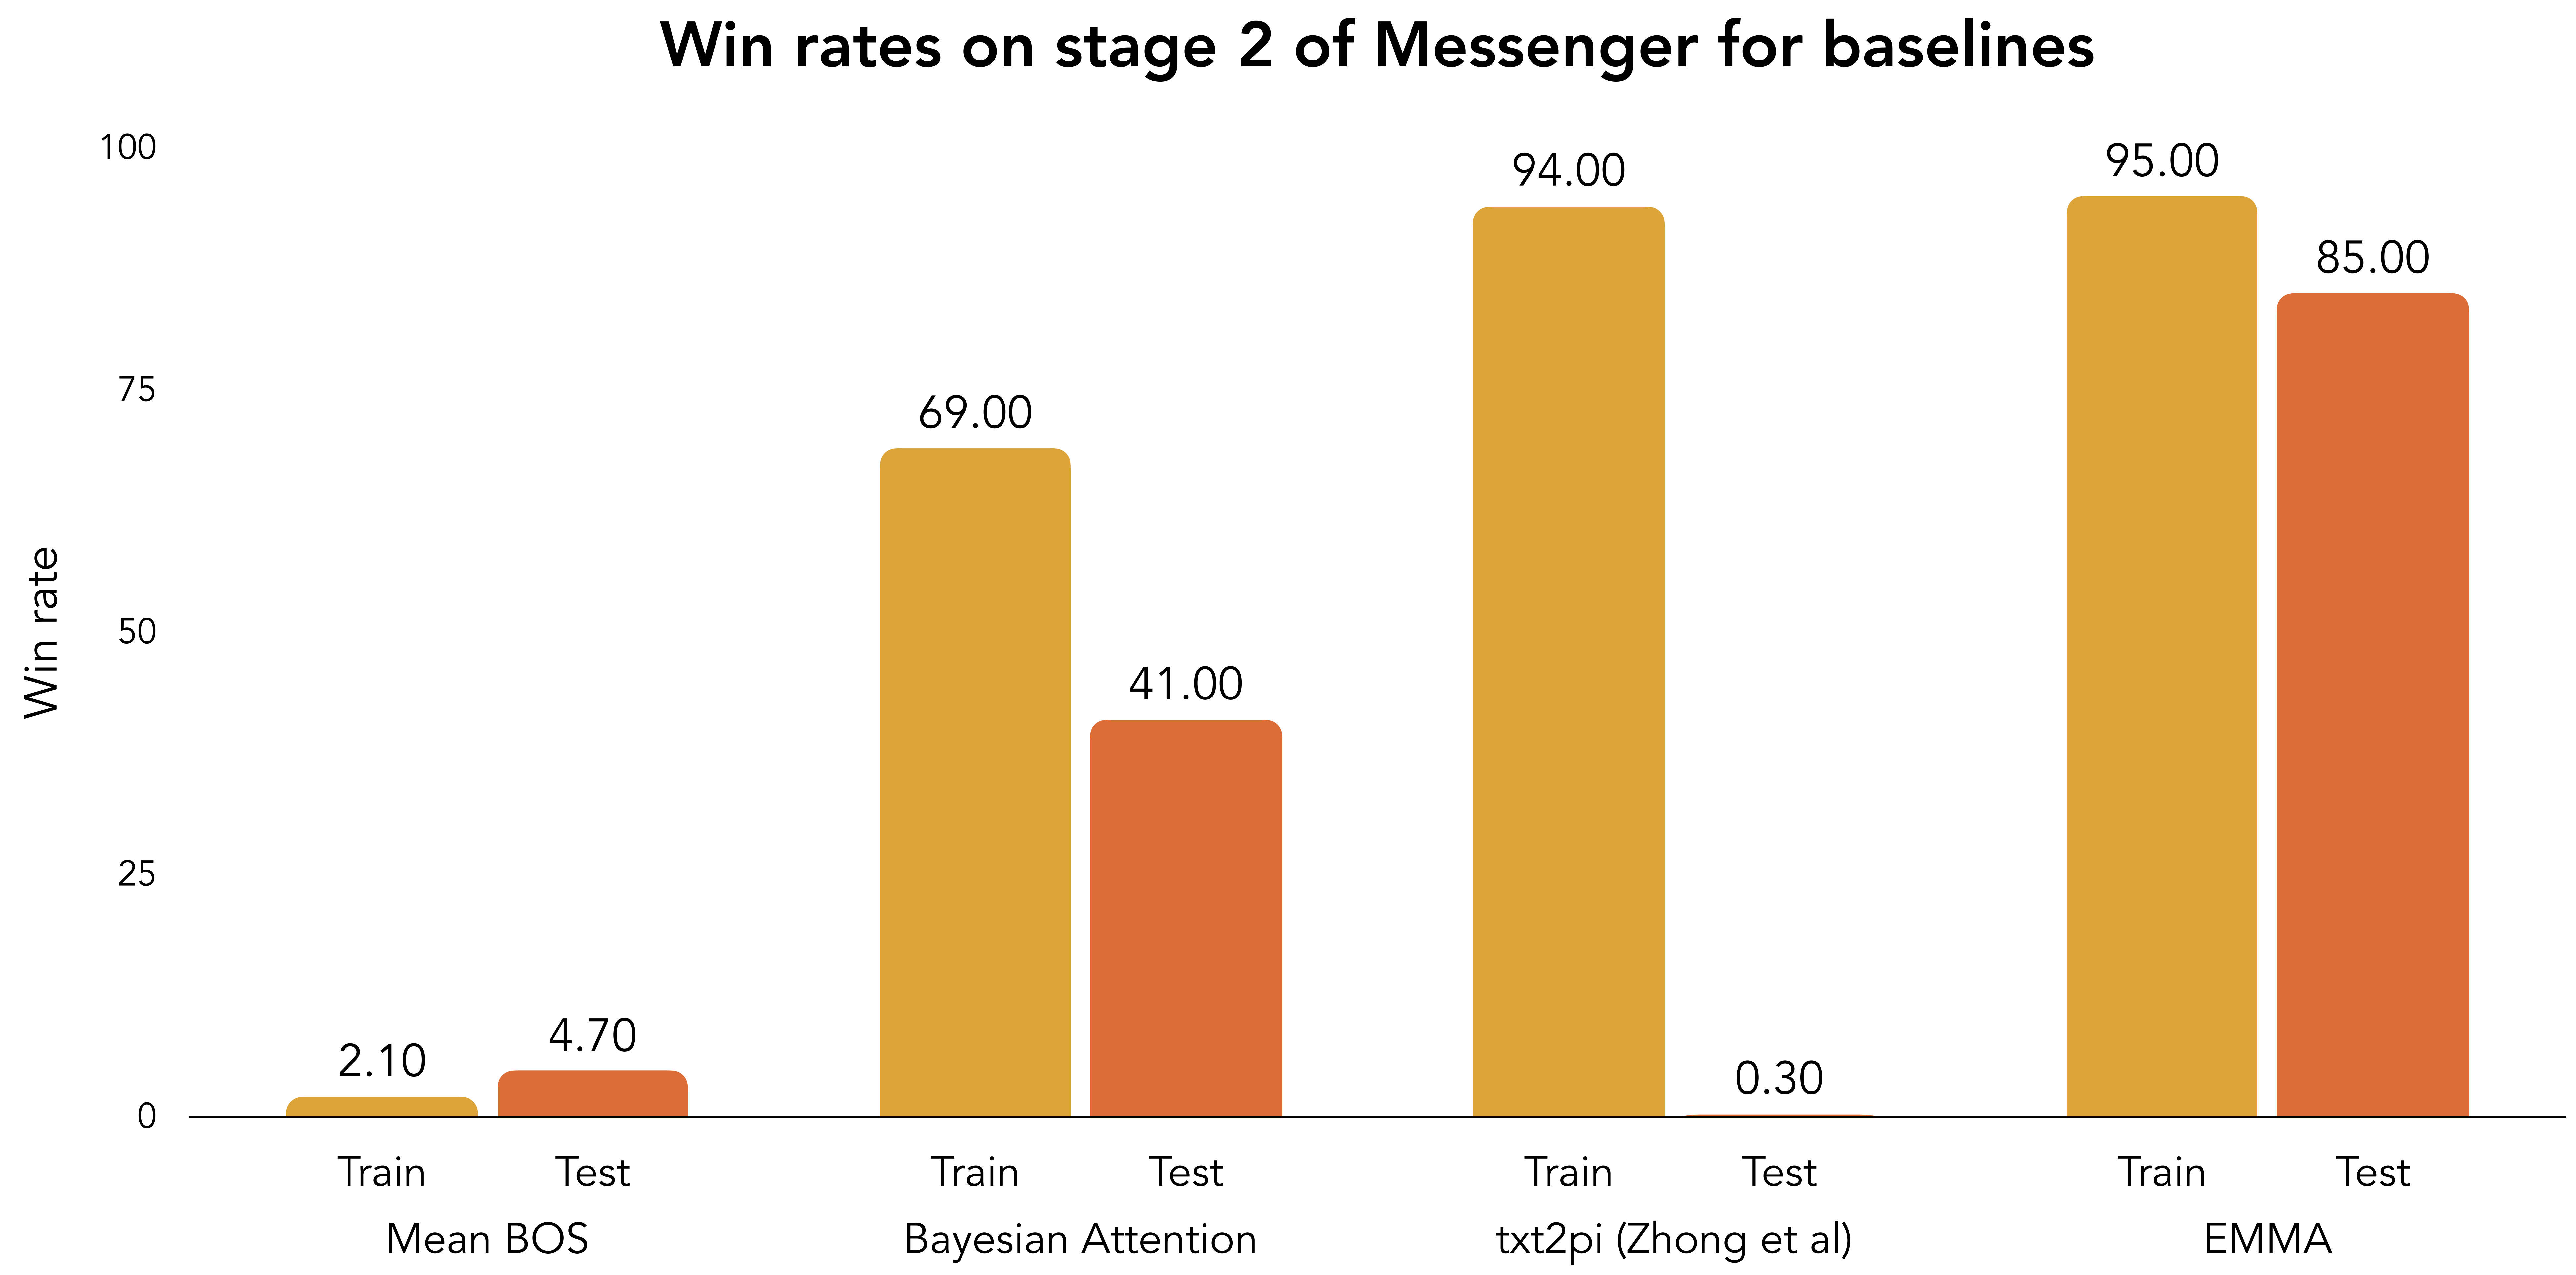

… but some stages continue to prove challenging

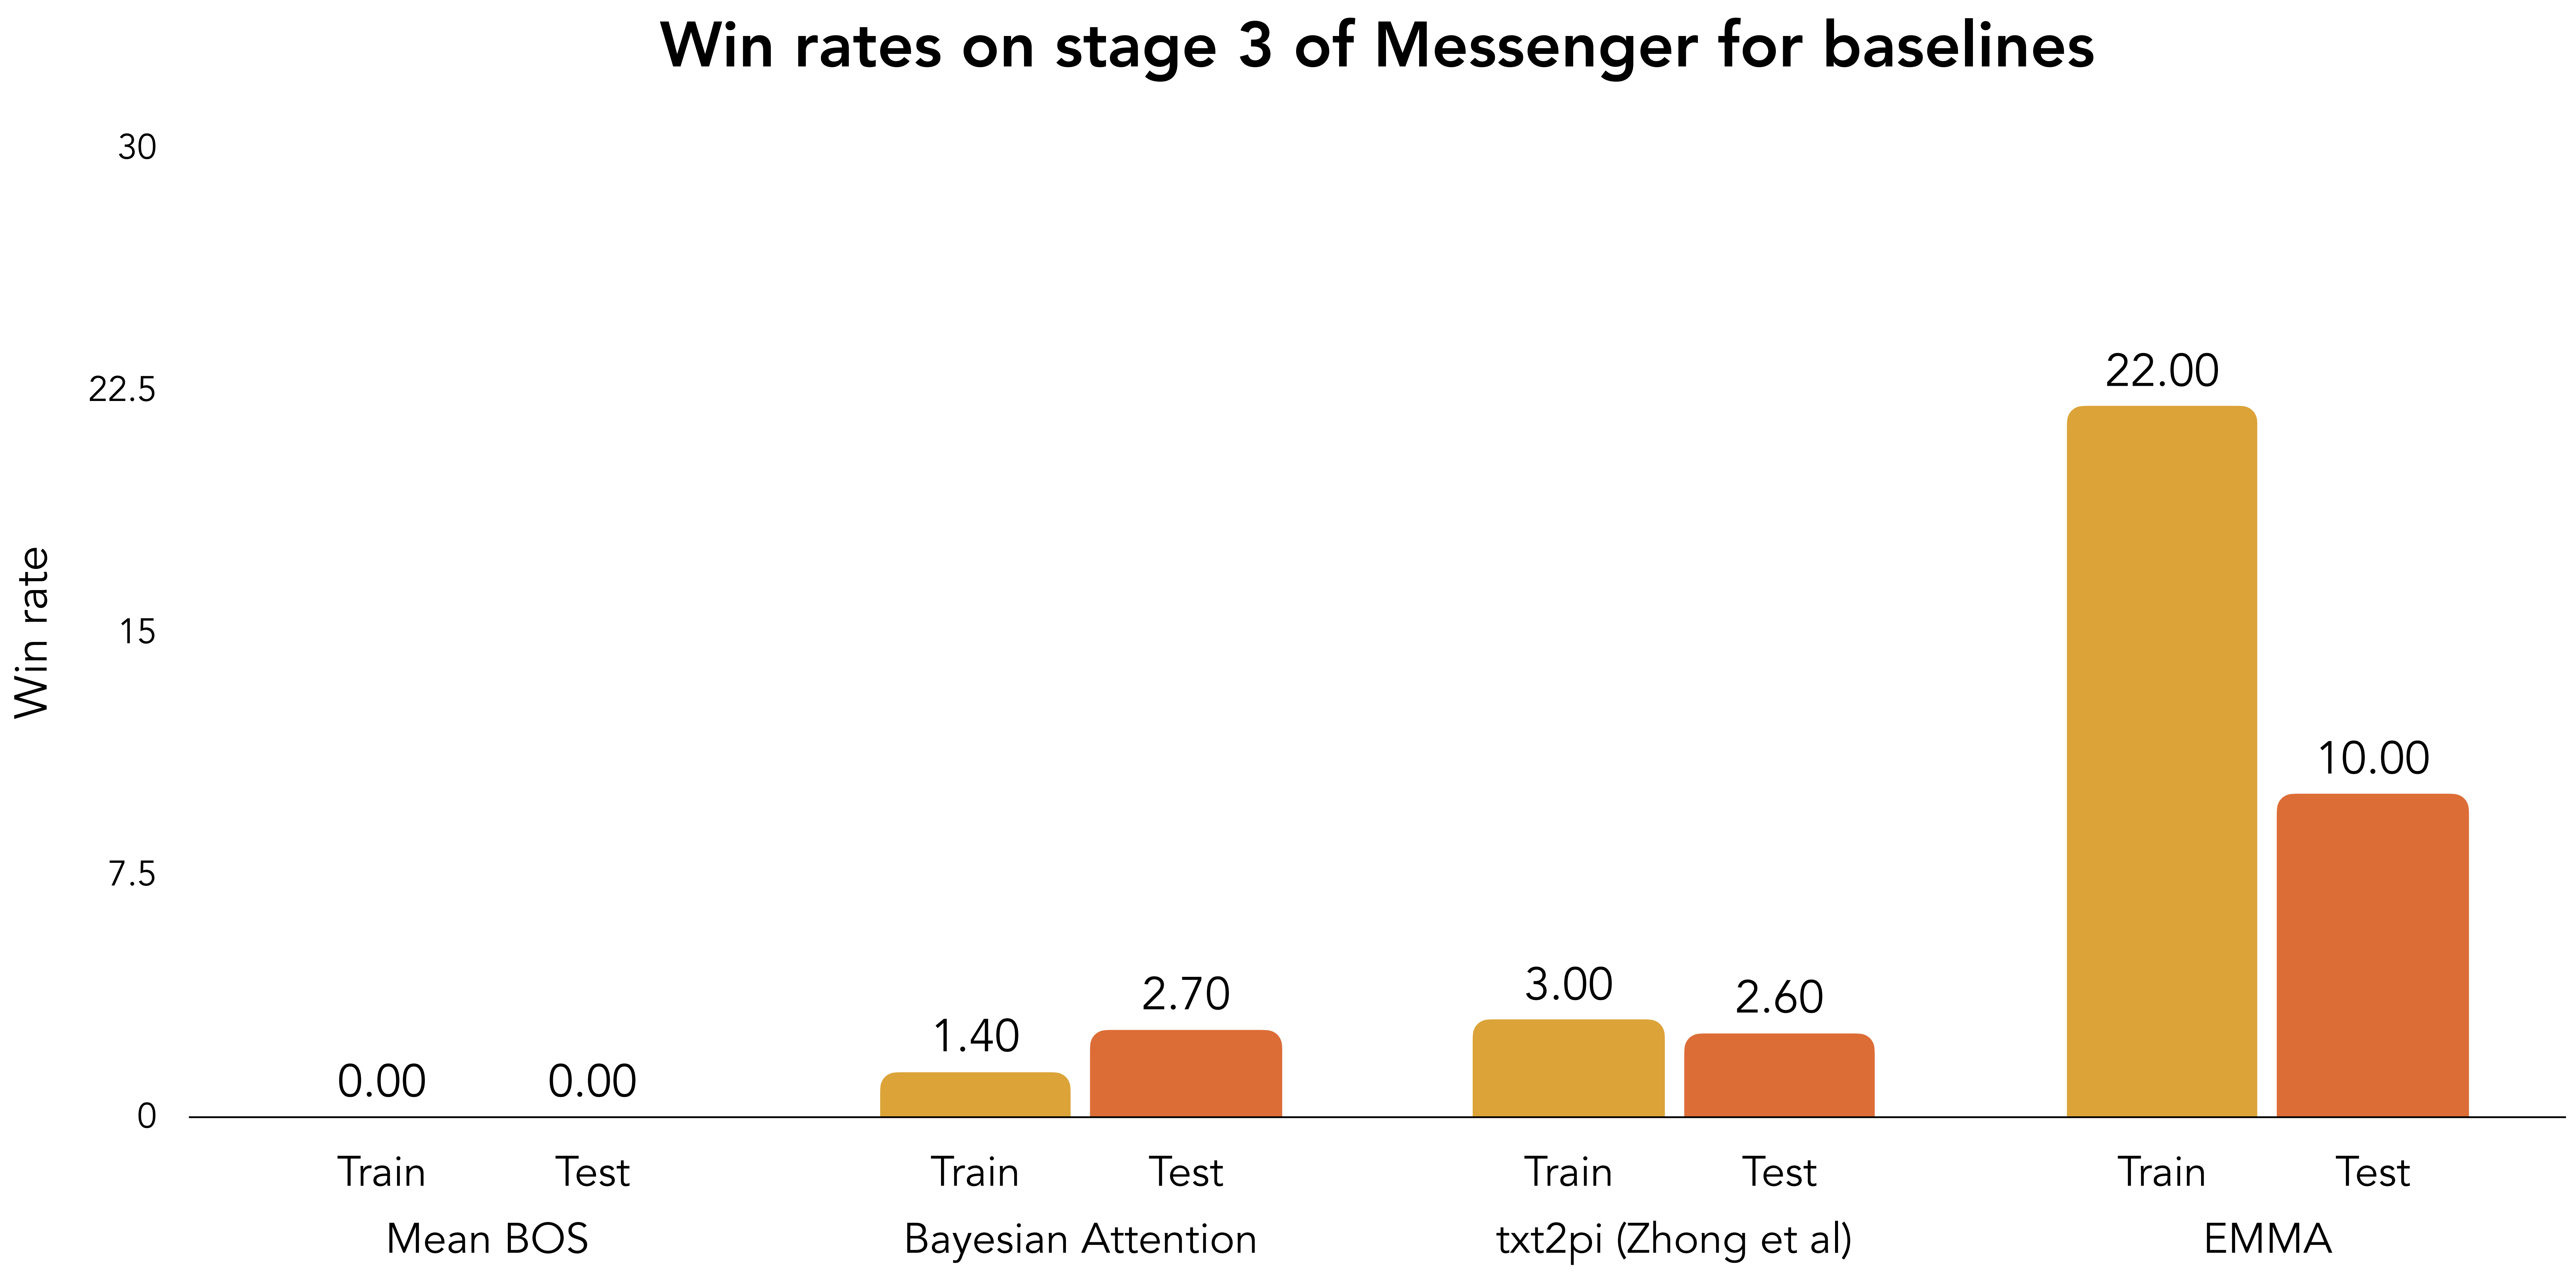

Analysis

Attention weights of **EMMA**

EMMA learns to map each description to the entity it describes

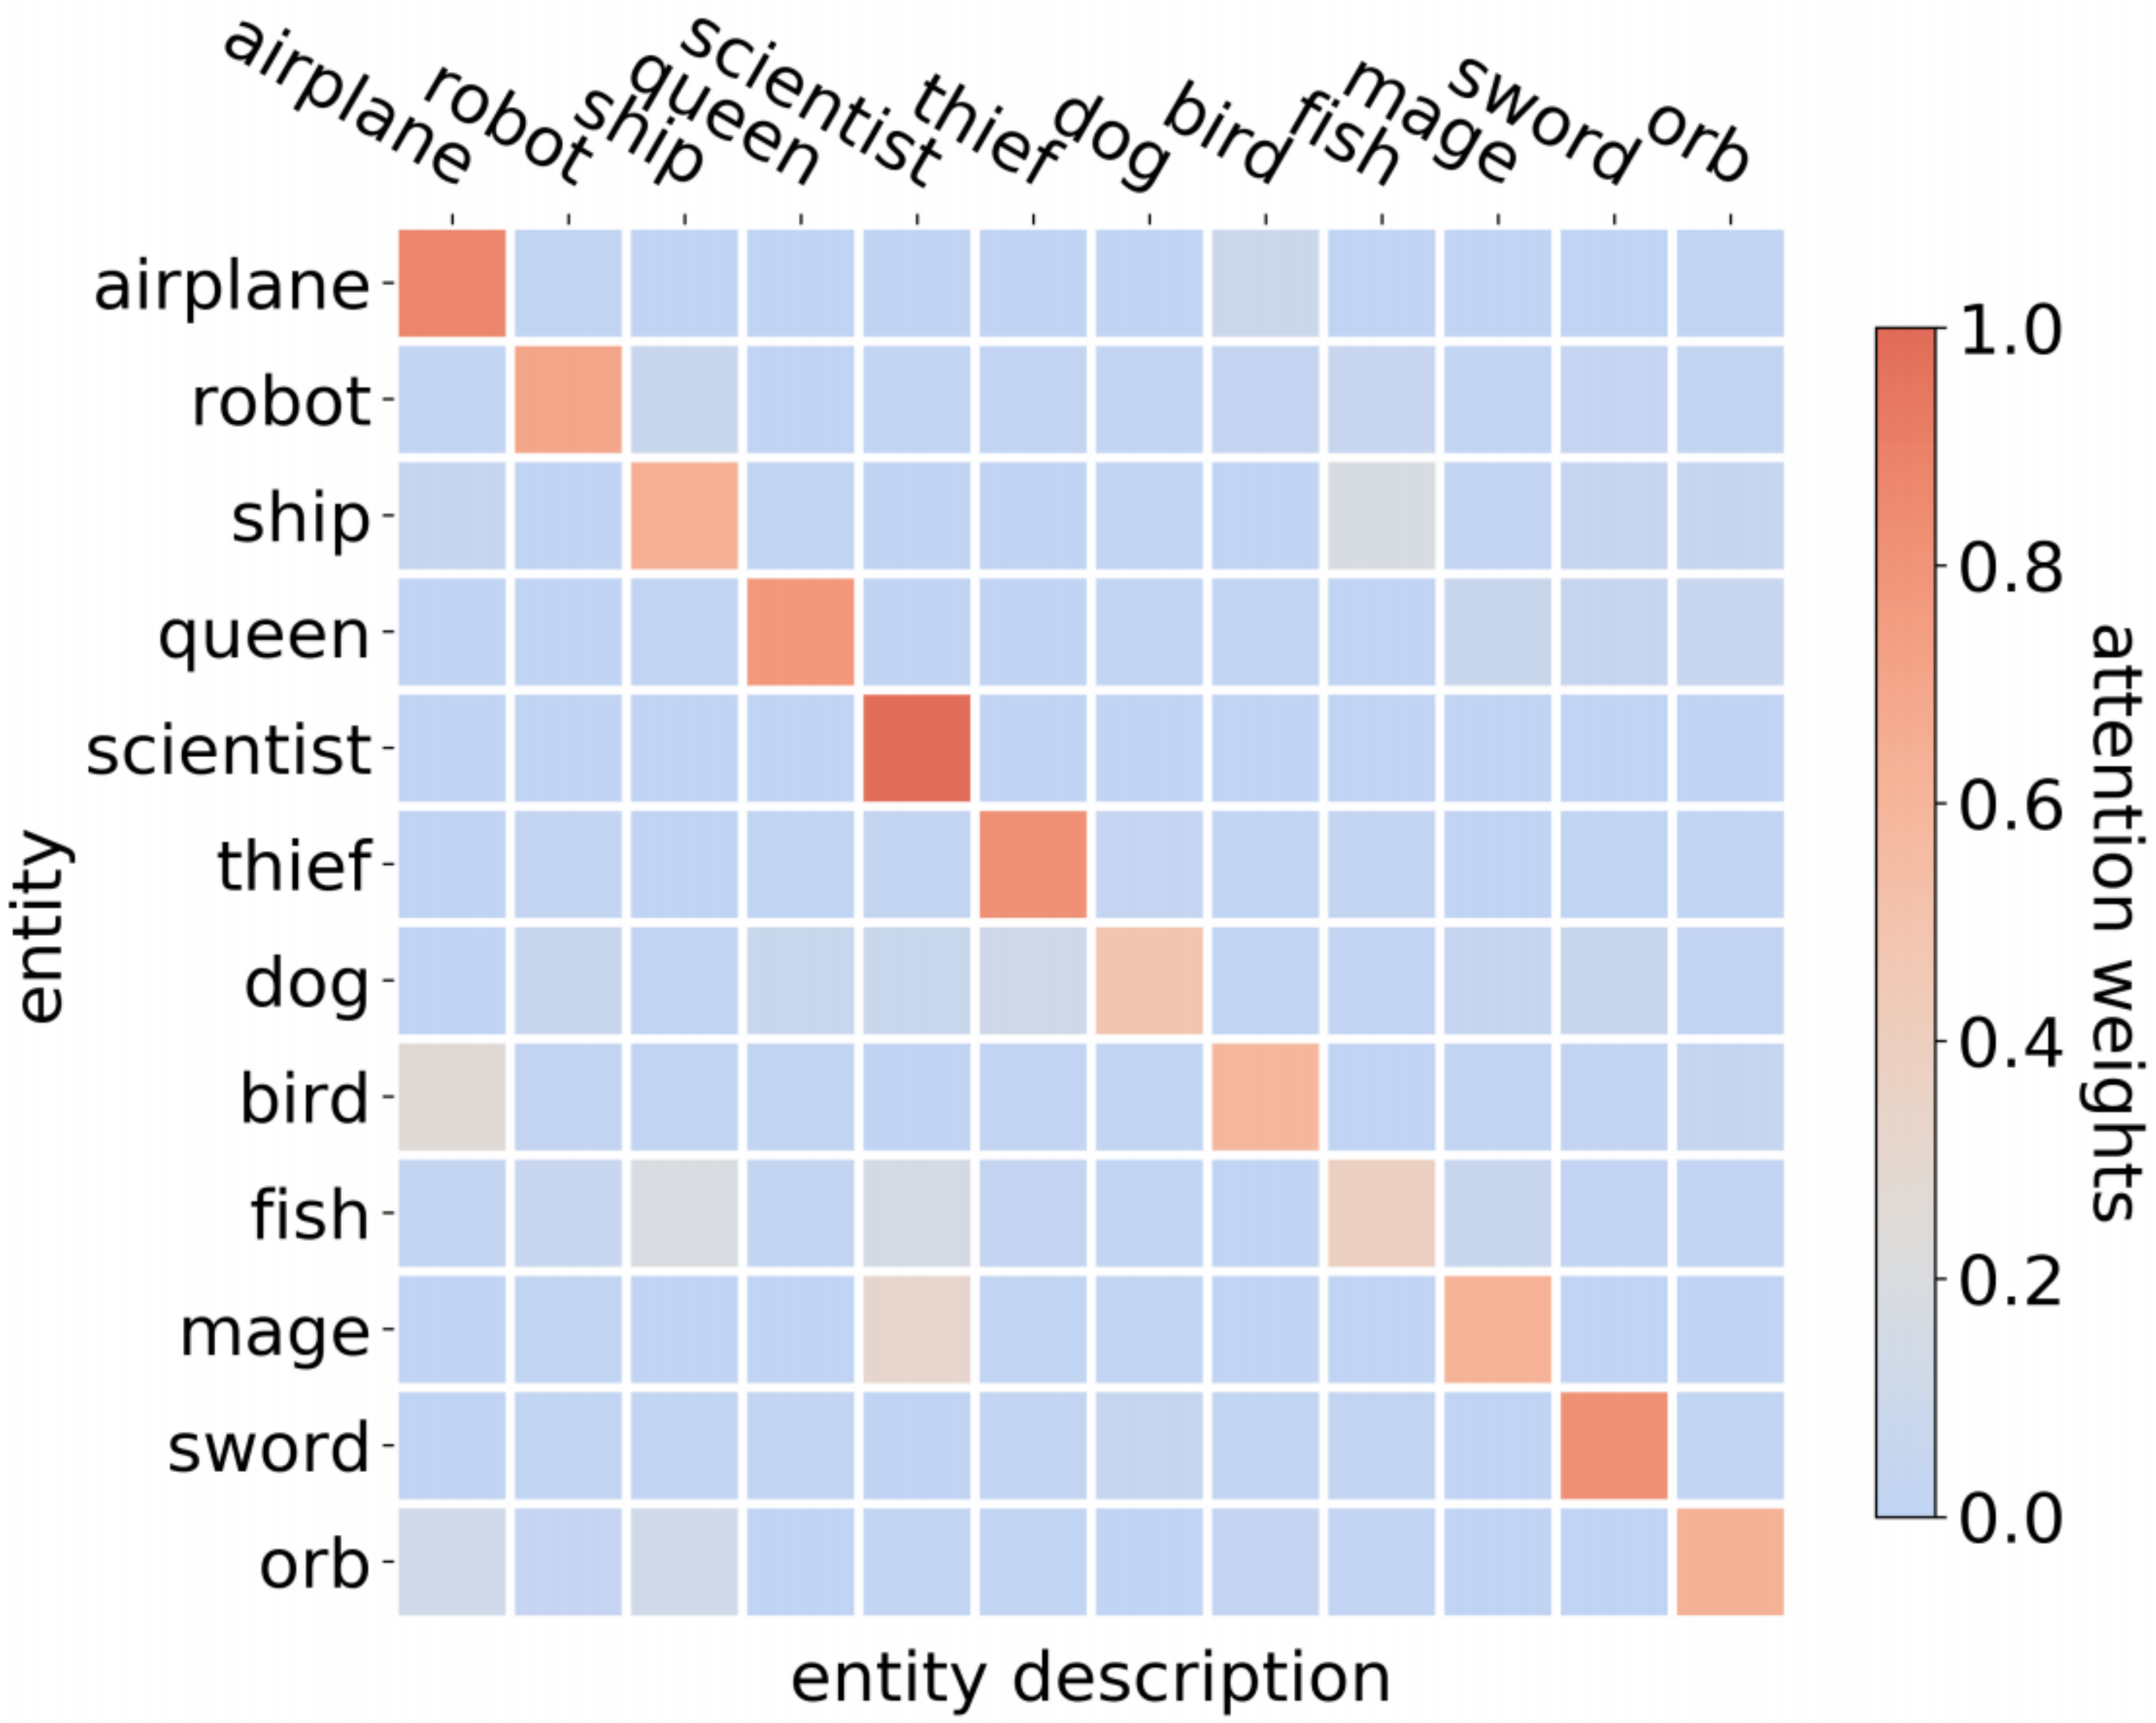

Text-based games

You are standing in an open field west of a white house, with a boarded front door. There is a small mailbox here.

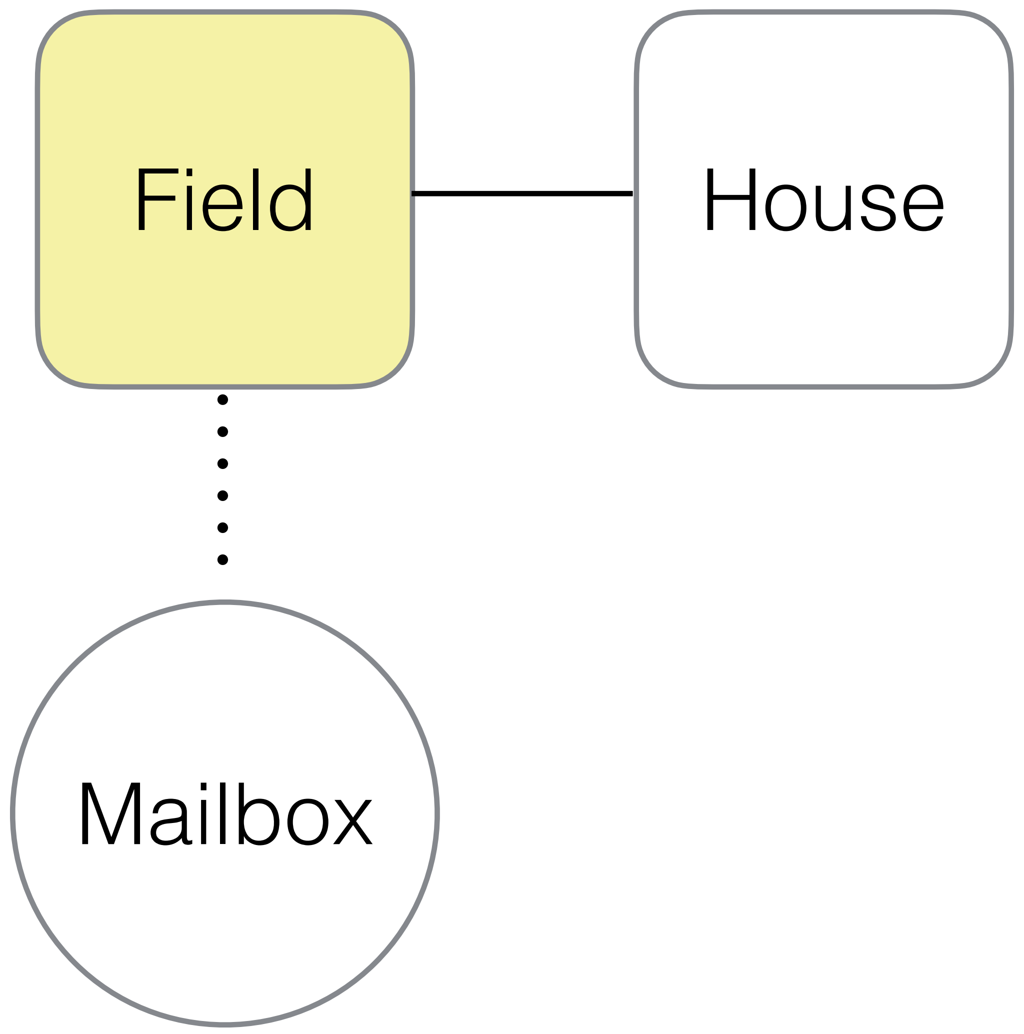

- open mailbox
- go east
- search field

Underlying game state (h1)

(Narasimhan et al., 2015)

Text-based games

You are standing in an open field west of a white house, with a boarded front door. There is a small mailbox here.

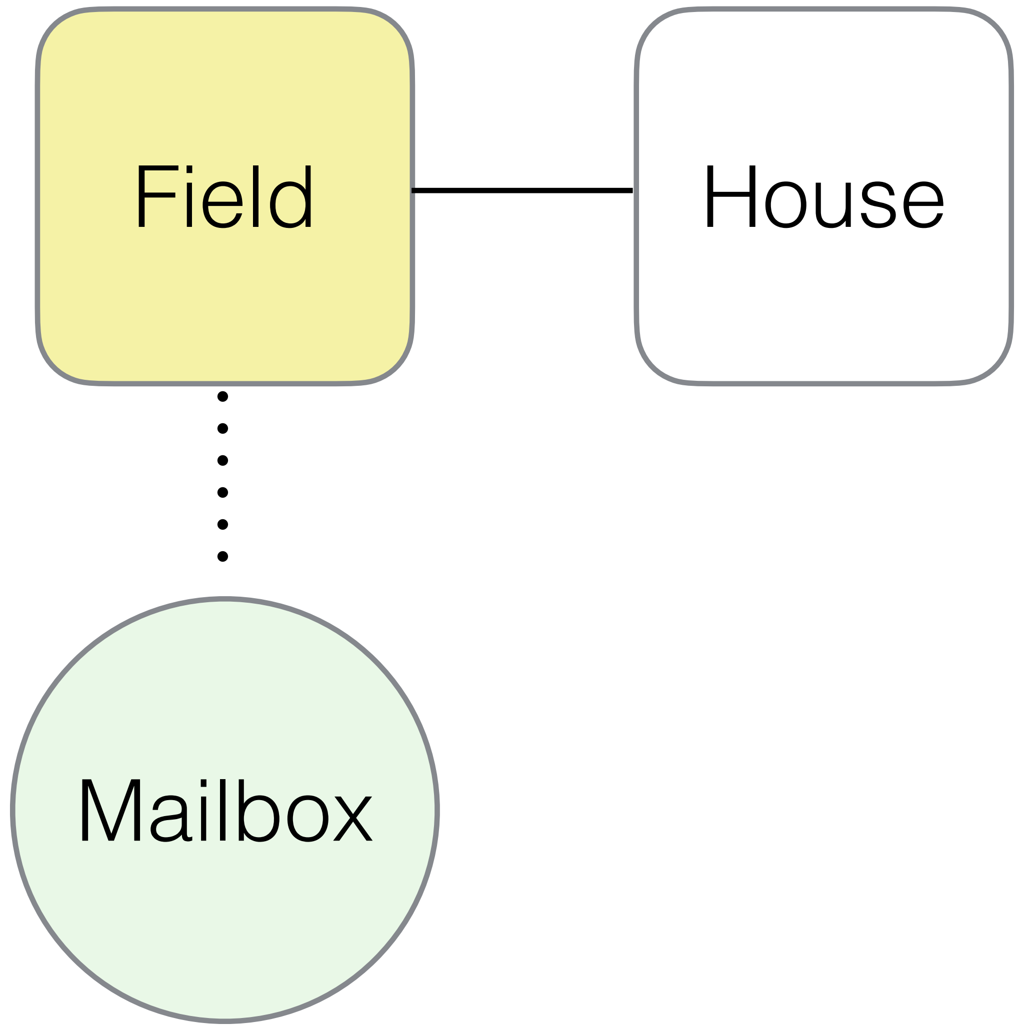

✦ open mailbox

Text-based games

Opening the mailbox reveals a leaflet.

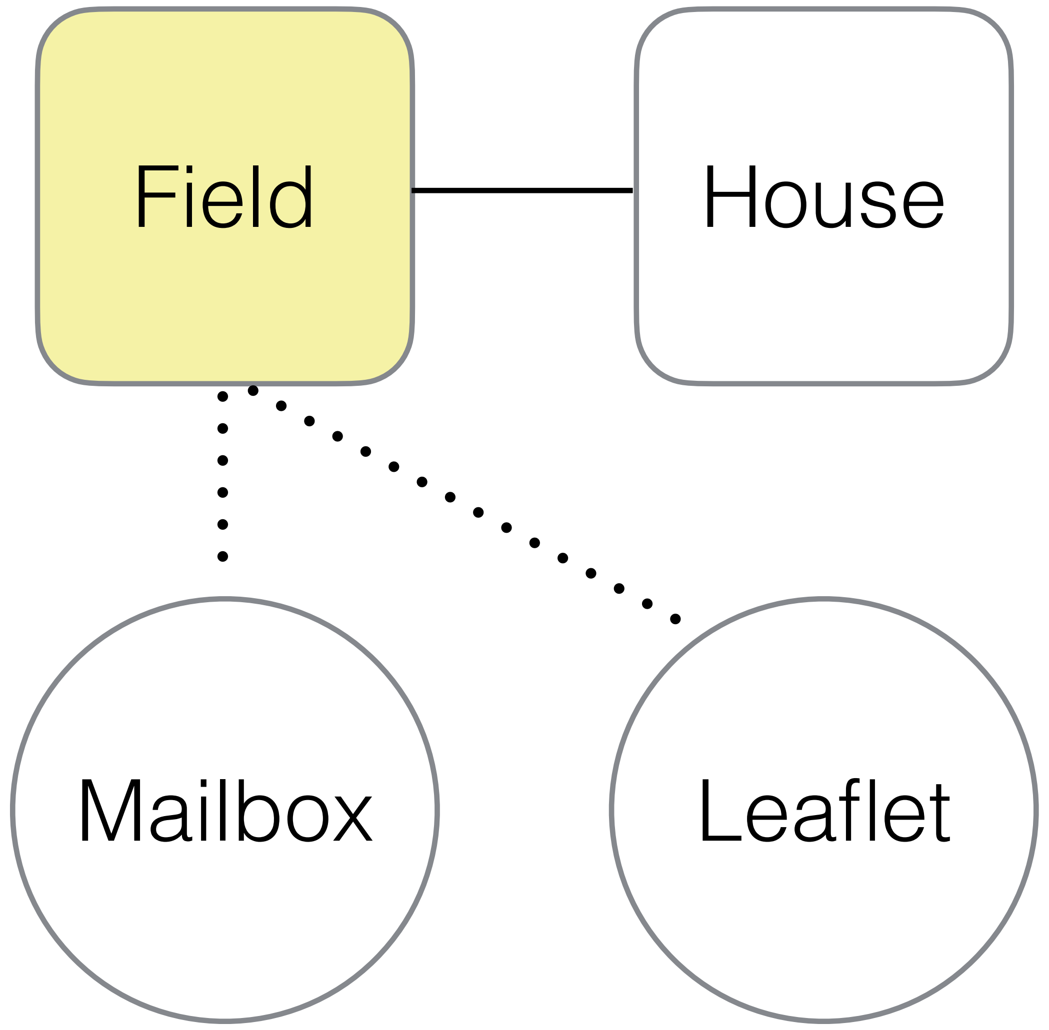

Underlying game state (h2)

No symbolic representation available

You are standing in an open field west of a white house, with a boarded front door. There is a small mailbox here.

~~*Location:* Field~~

~~*Wind level:* 3~~

~~*Time:* 12pm~~

Varying text descriptions

You are standing in an open field west of a white house, with a boarded front door. There is a small mailbox here.

You are in an open field next to a white house. The house's front door is boarded shut. You see a small mailbox here.

Opportunity

Grounded language learning

You are standing in an open field west of a white house, with a boarded front door. There is a small mailbox here.

✦ open mailbox

Opportunity

Grounded language learning

In-game rewards provide unstructured feedback

+10 gold

+5 health

Opportunity

Learn langu… …rming tasks

In-game rewards prov… …tructured feedback to learn

+10 gold

+5 health

Recurrent Neural Network

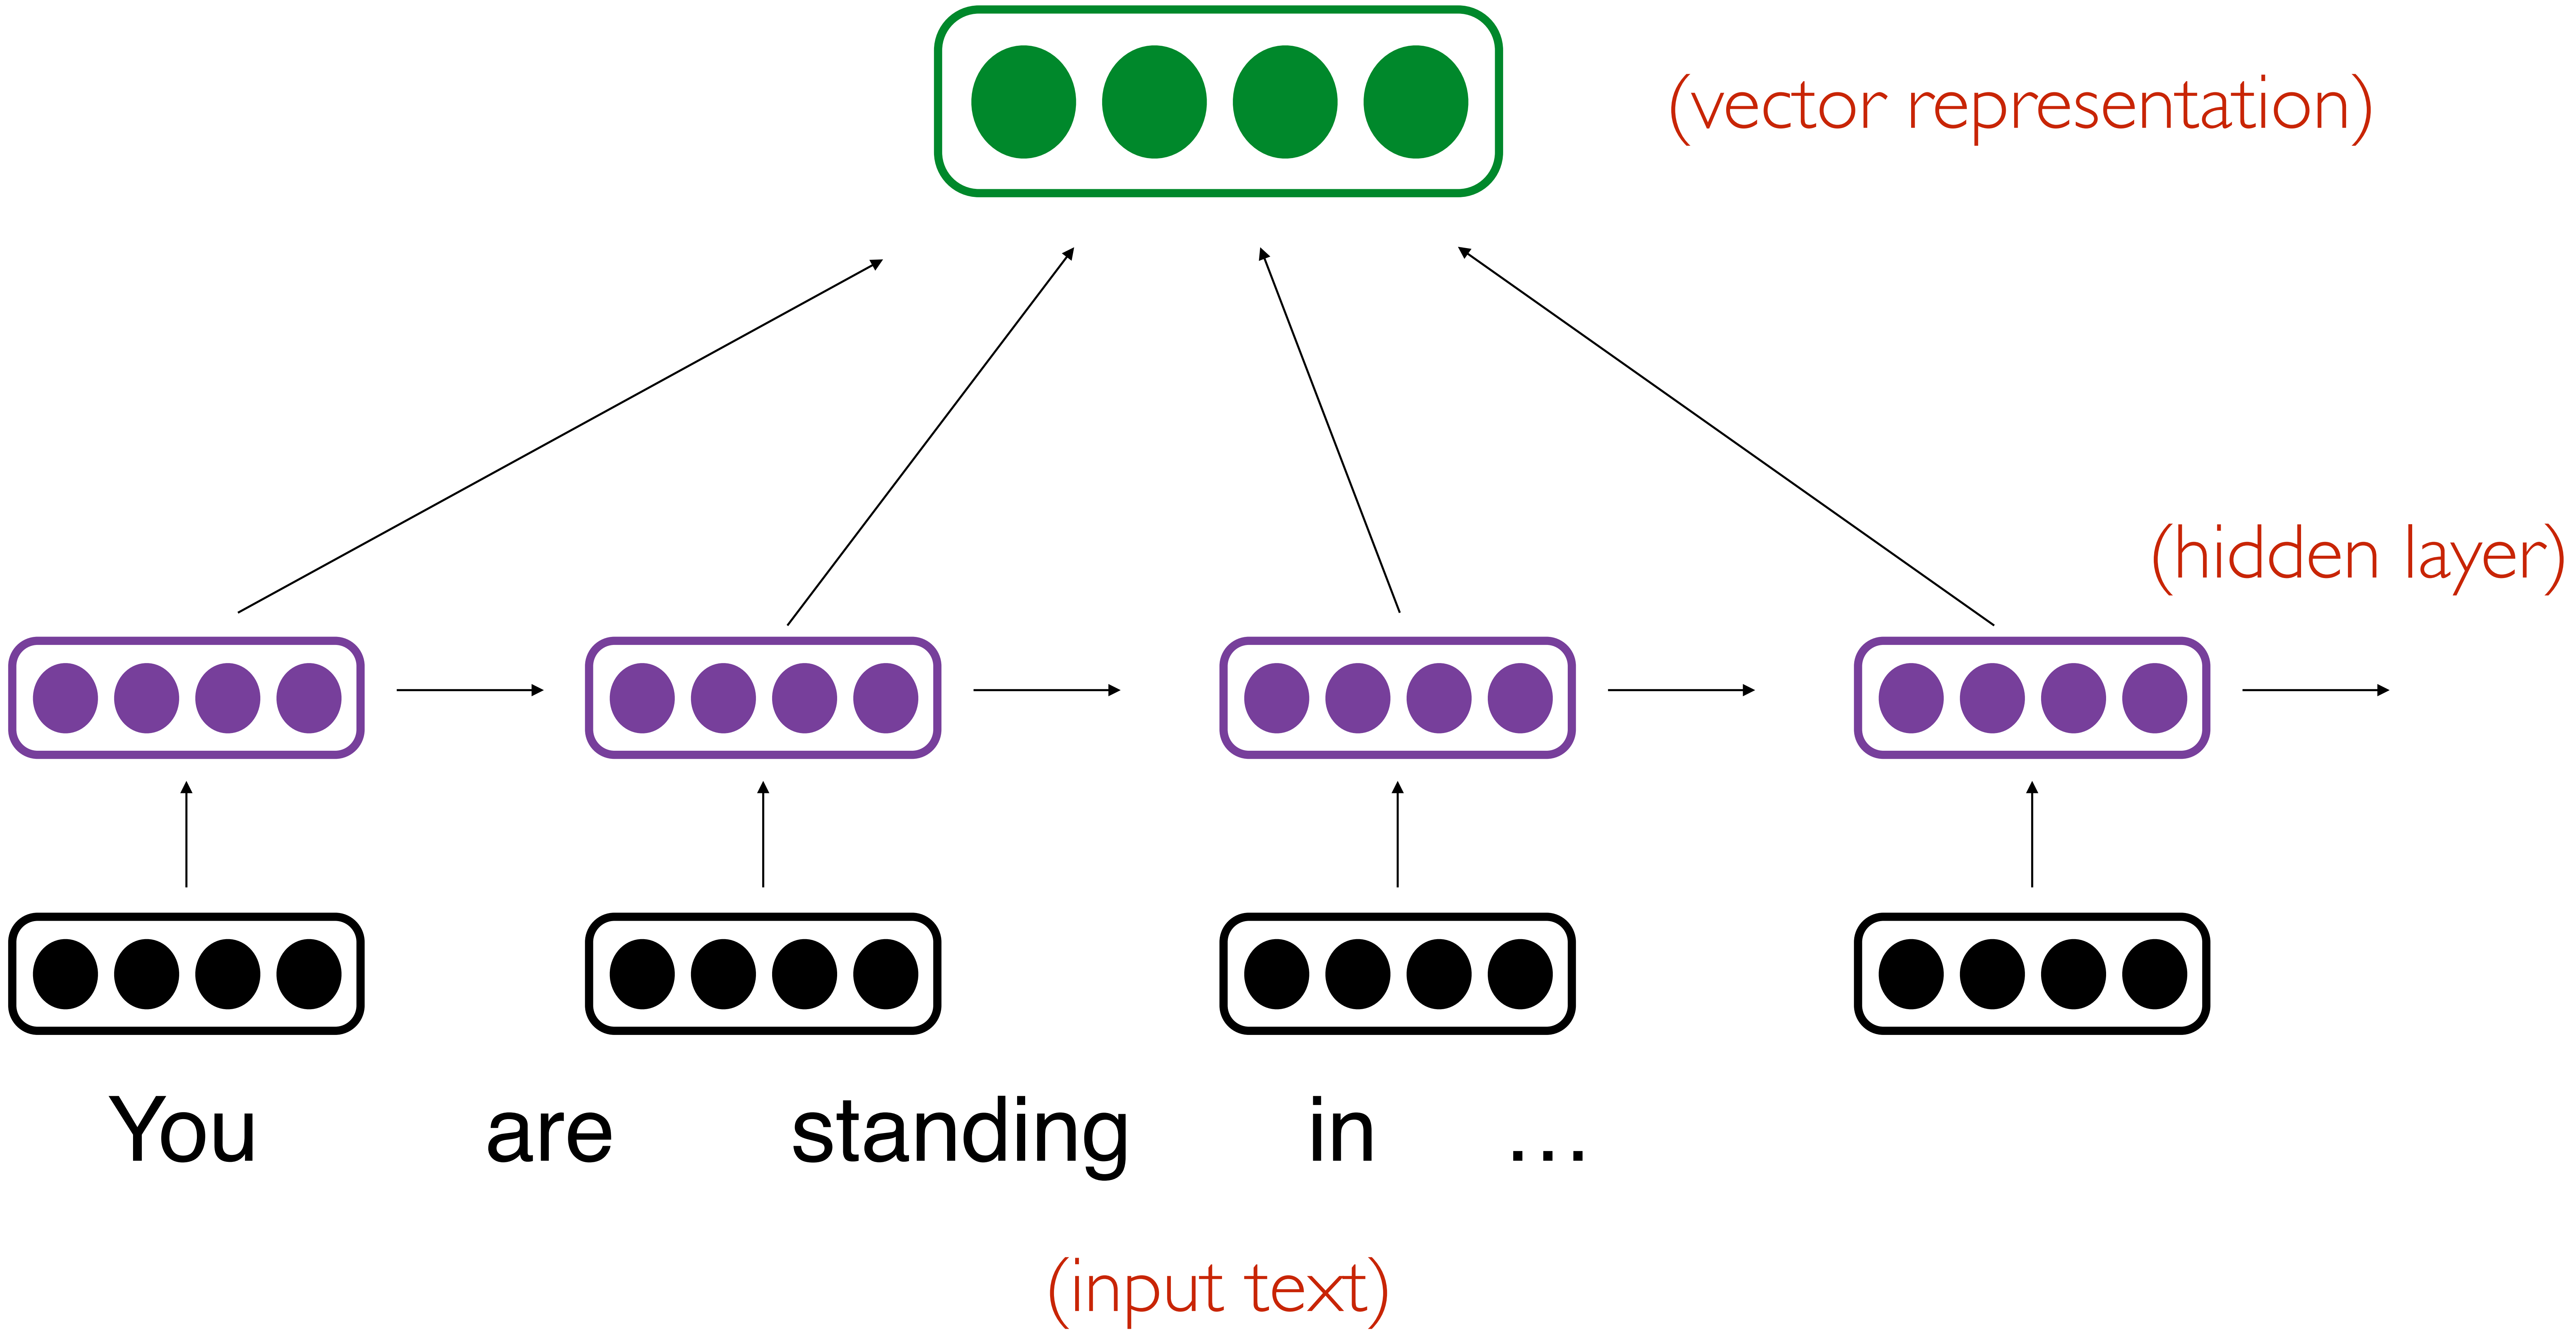

(Narasimhan et al., 2015)

LSTM-DQN: Action Scorer

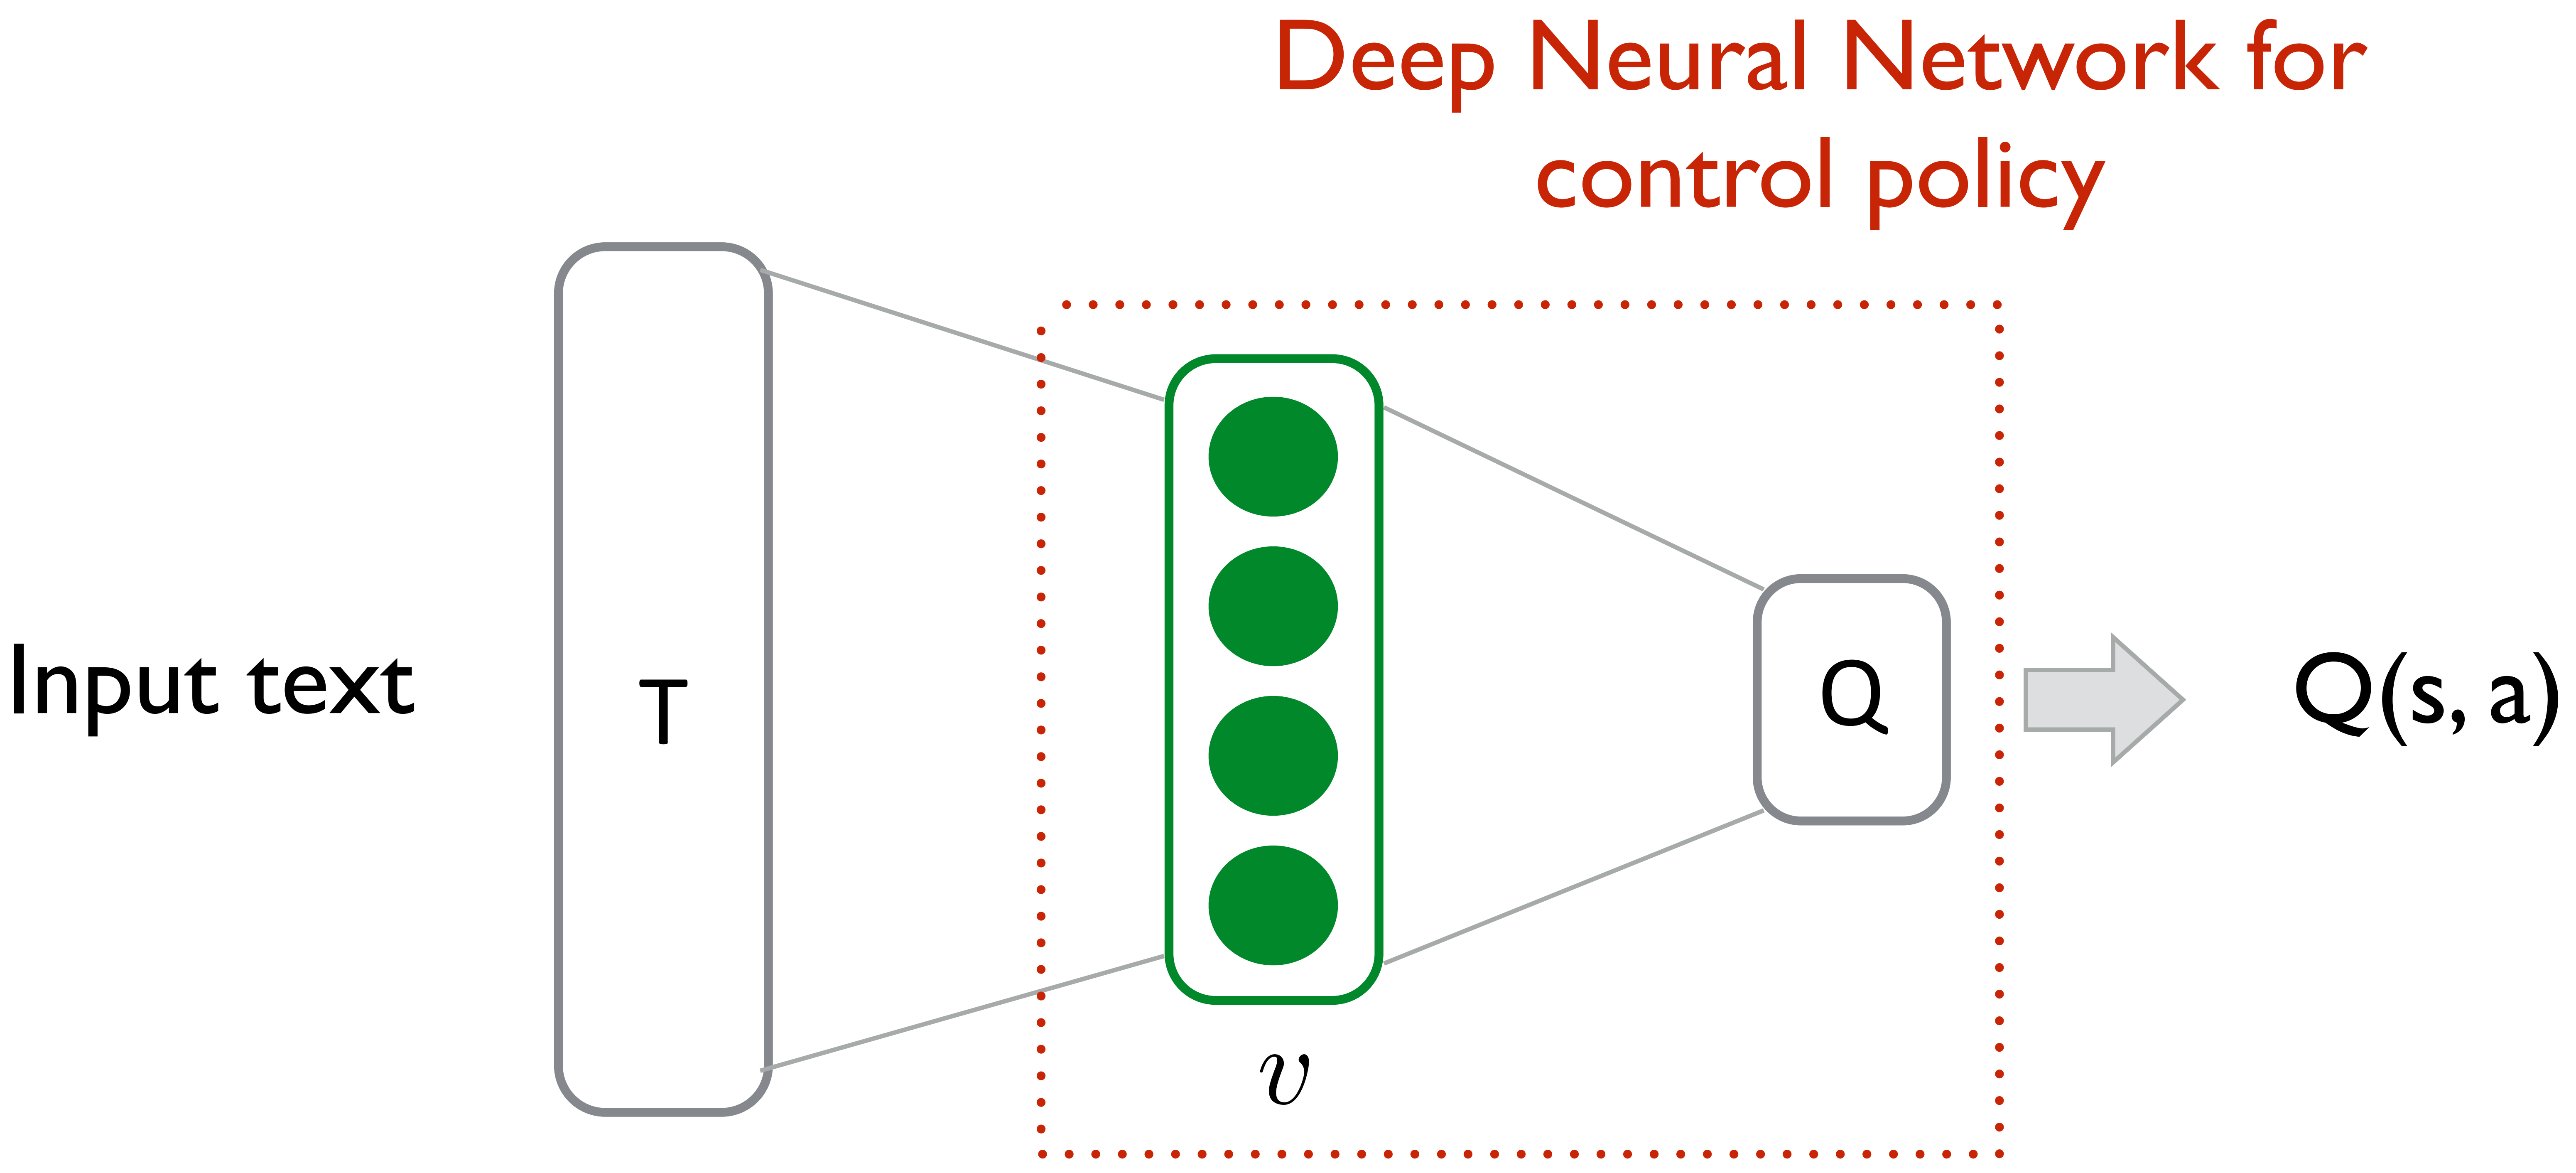

Recurrent Neural Network
to map text to vector
representation

Learn parameters using
Q-learning

Results

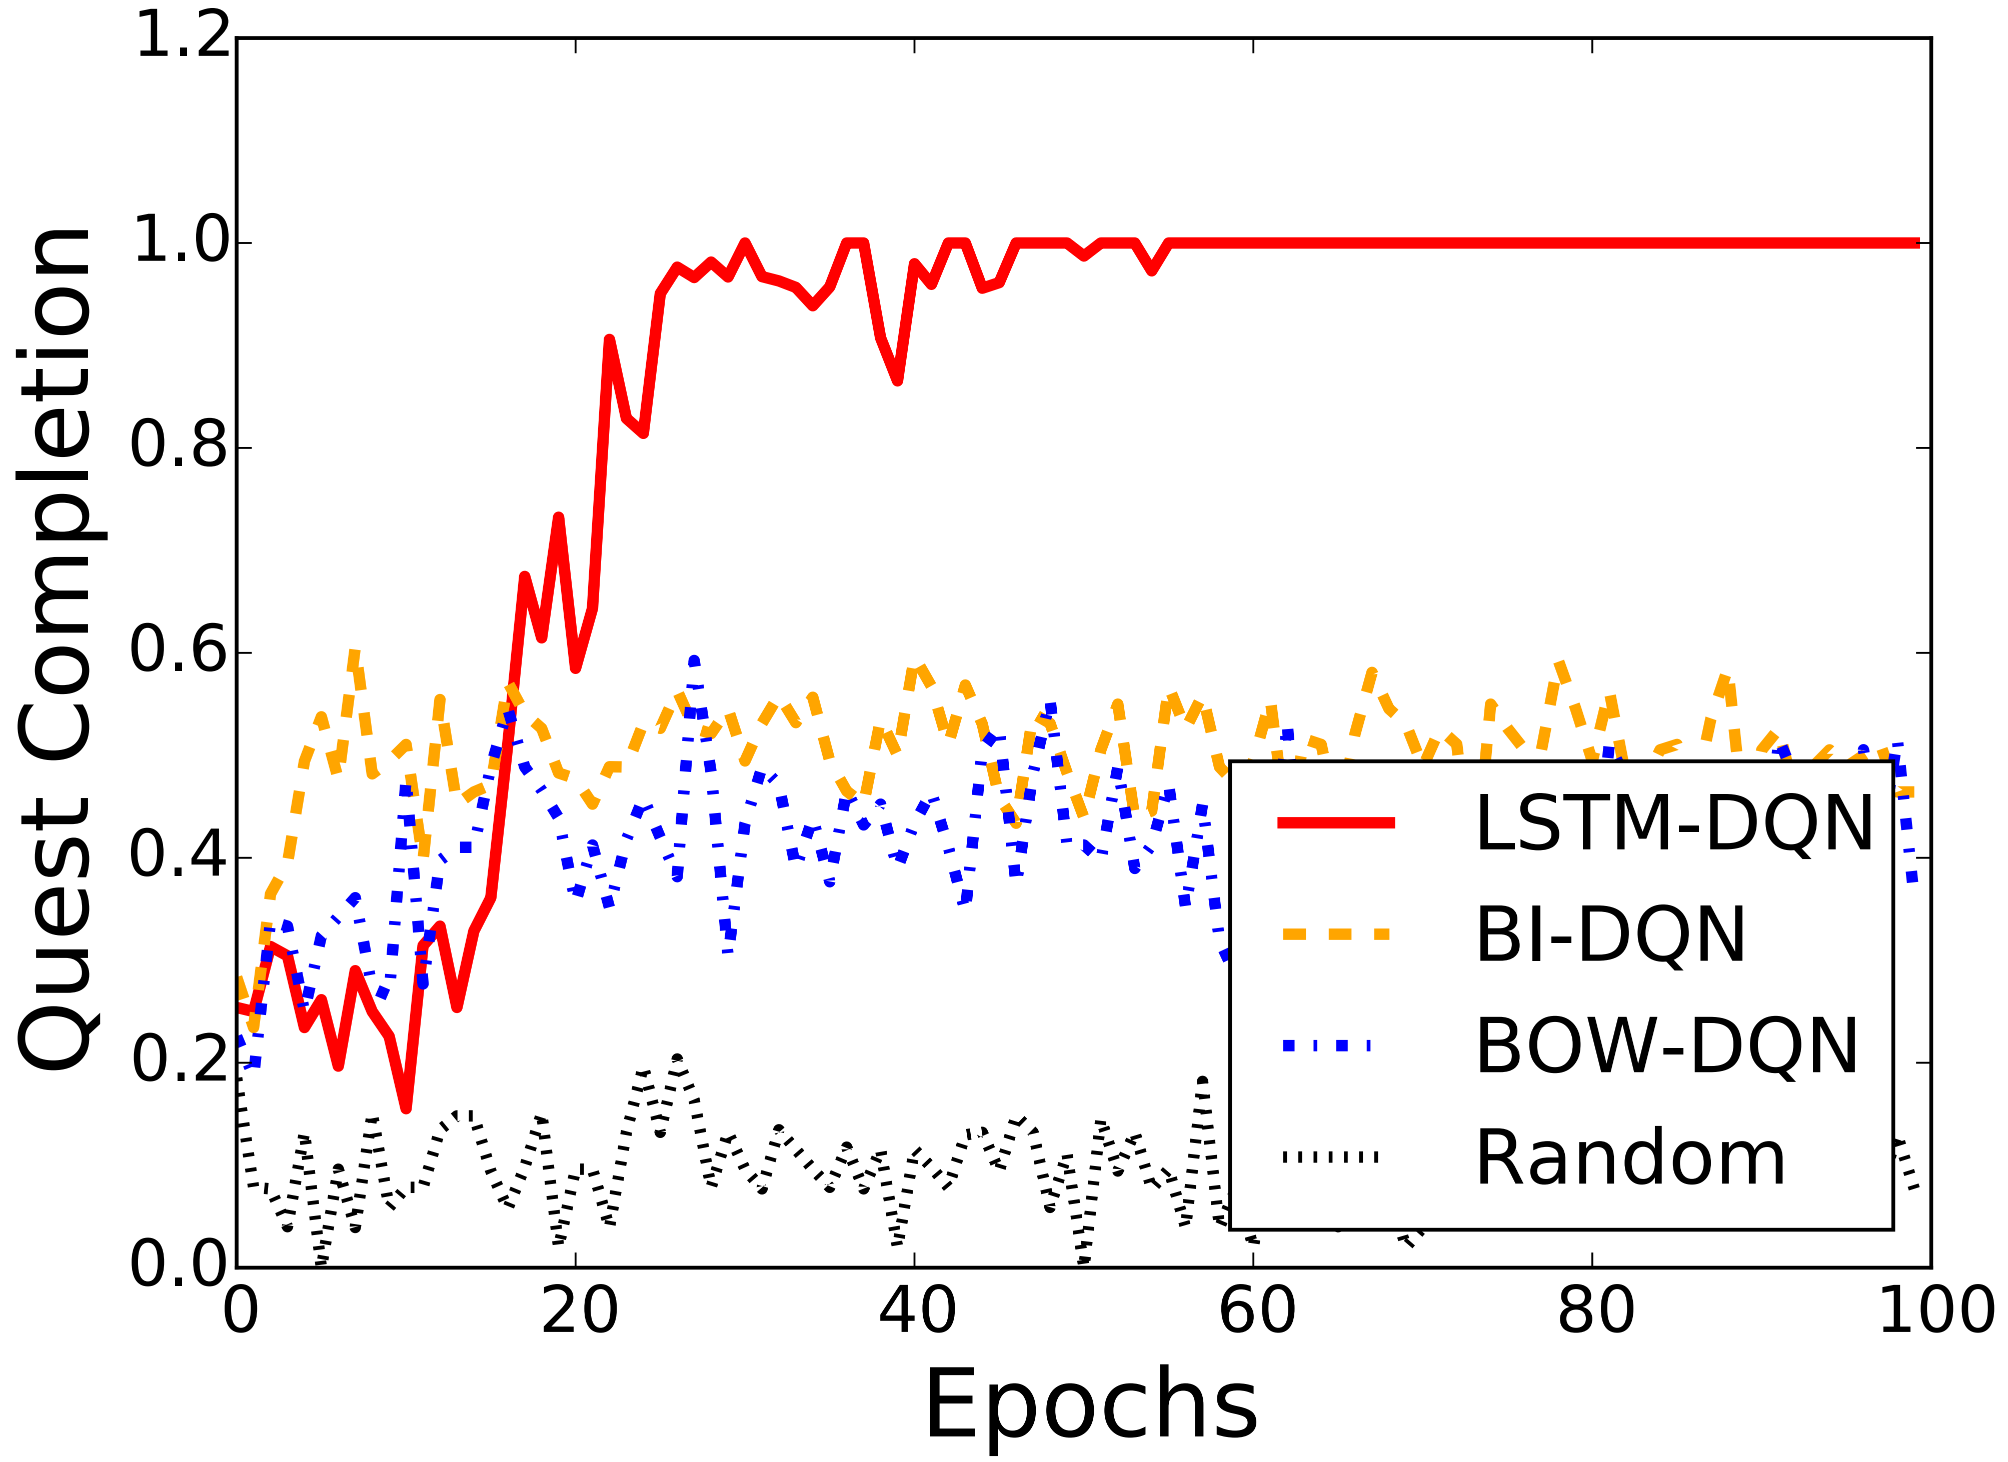

Quest Completion
Epochs
LSTM-DQN
BI-DQN
BOW-DQN
Random

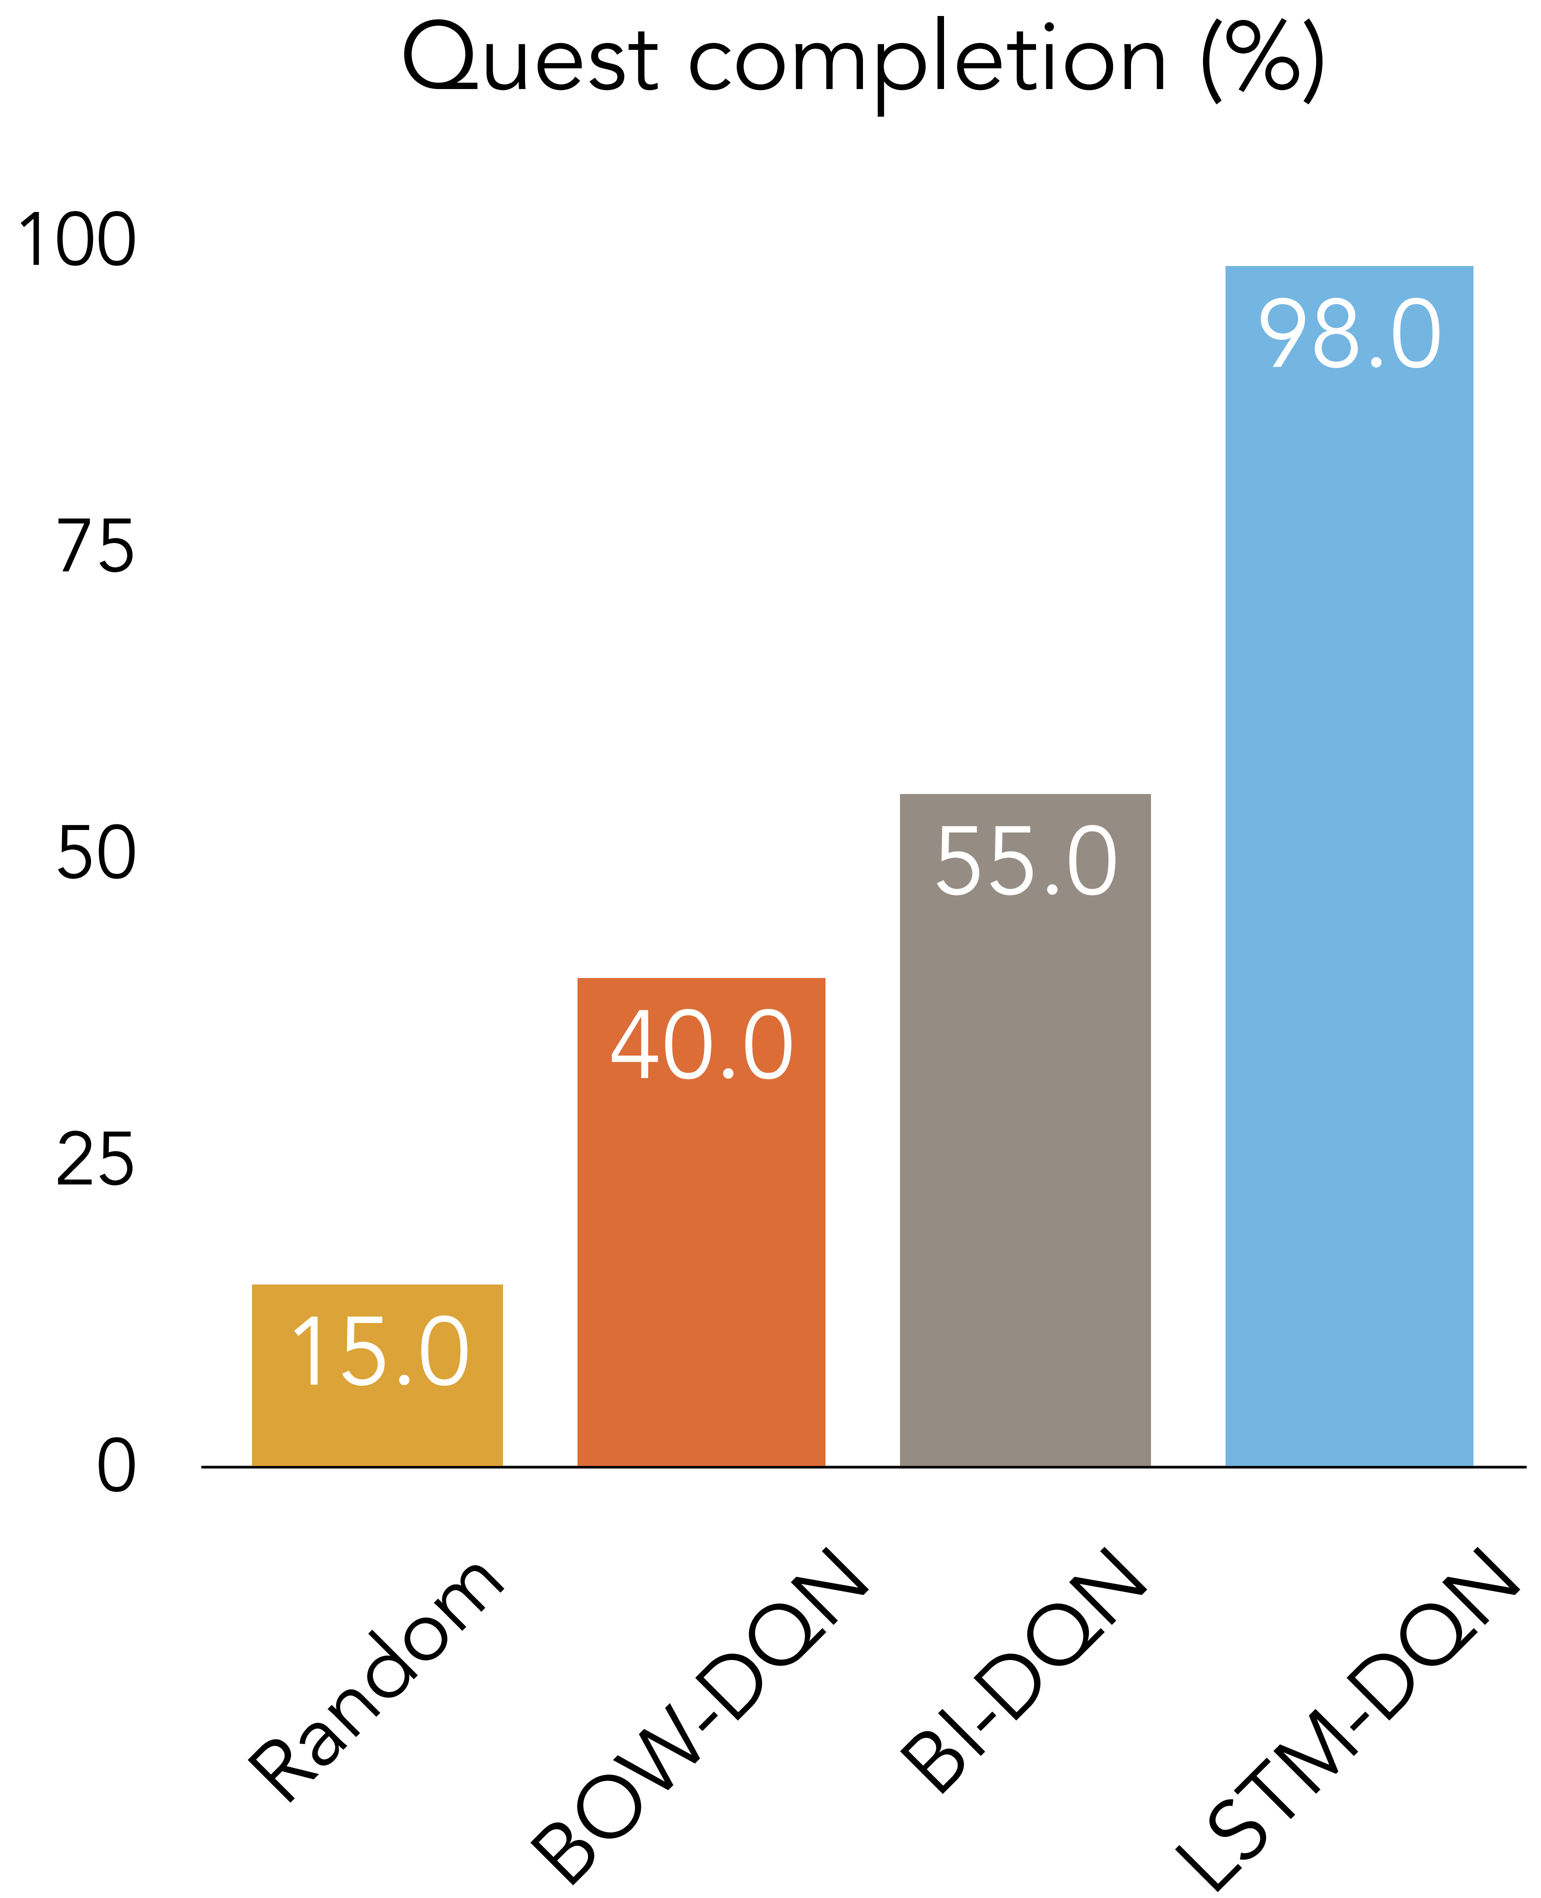

Quest completion (%)
15.0
40.0
55.0
98.0
Random
BOW-DQN
BI-DQN
LSTM-DQN

Visualizing Learnt Representations

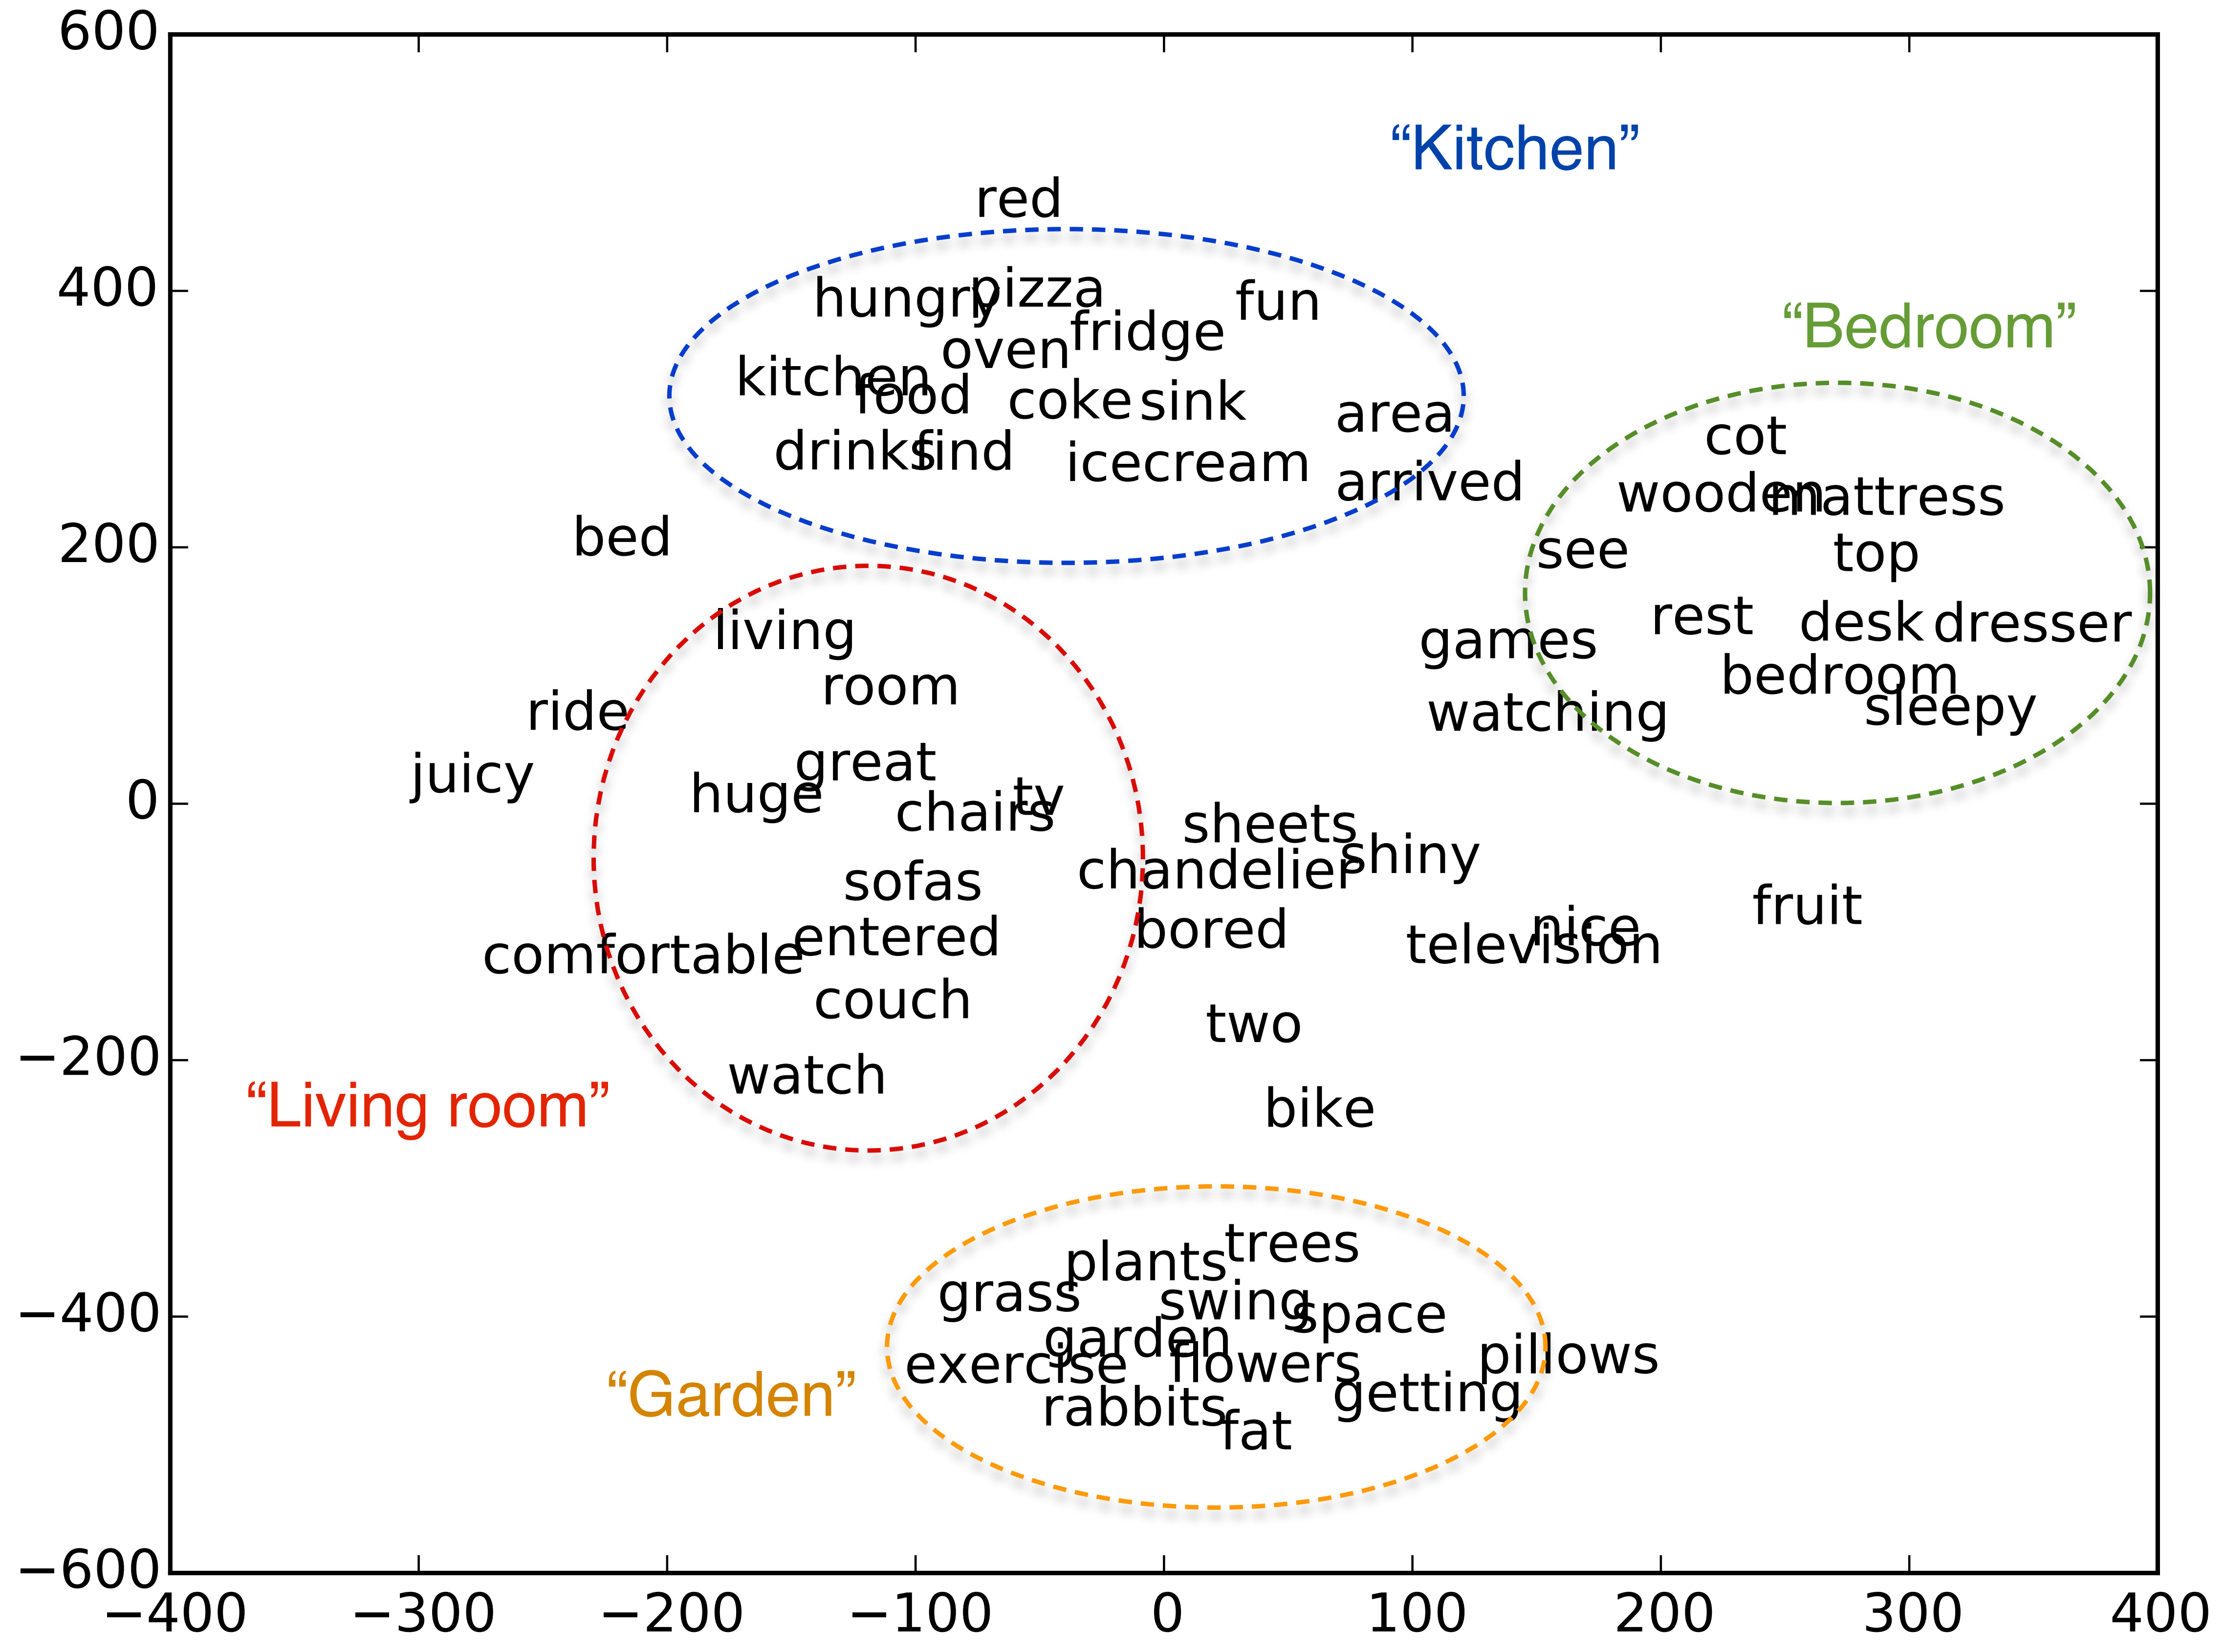

t-SNE visualization of vectors learnt by agent

Contextual Action Language Model (CALM)

- **Want:** Generate sensible action commands
- **Idea:** Train a *single* language model to generate action candidates for any game
  - Actions are subsequently reranked by an RL agent using game-specific rewards

Observation: You are in the living room. There is a doorway to the east, a wooden door with strange gothic lettering to the west, which appears to be nailed shut, a trophy case, and a large oriental rug in the center of the room. You are carrying: A brass lantern . . .

Random Actions:
close door, north a, eat troll with egg, . . .
CALM (n-gram) Actions:
enter room, leave room, lock room, open door, close door, knock on door, . . .
CALM (GPT-2) Actions:
east, open case, get rug, turn on lantern, move rug, unlock case with key, . . .

Next Observation: With a great effort, the rug is moved to one side of the room, revealing the dusty cover of a closed trap door...

(Yao et al., 2020)

Semantics does not exist in isolation

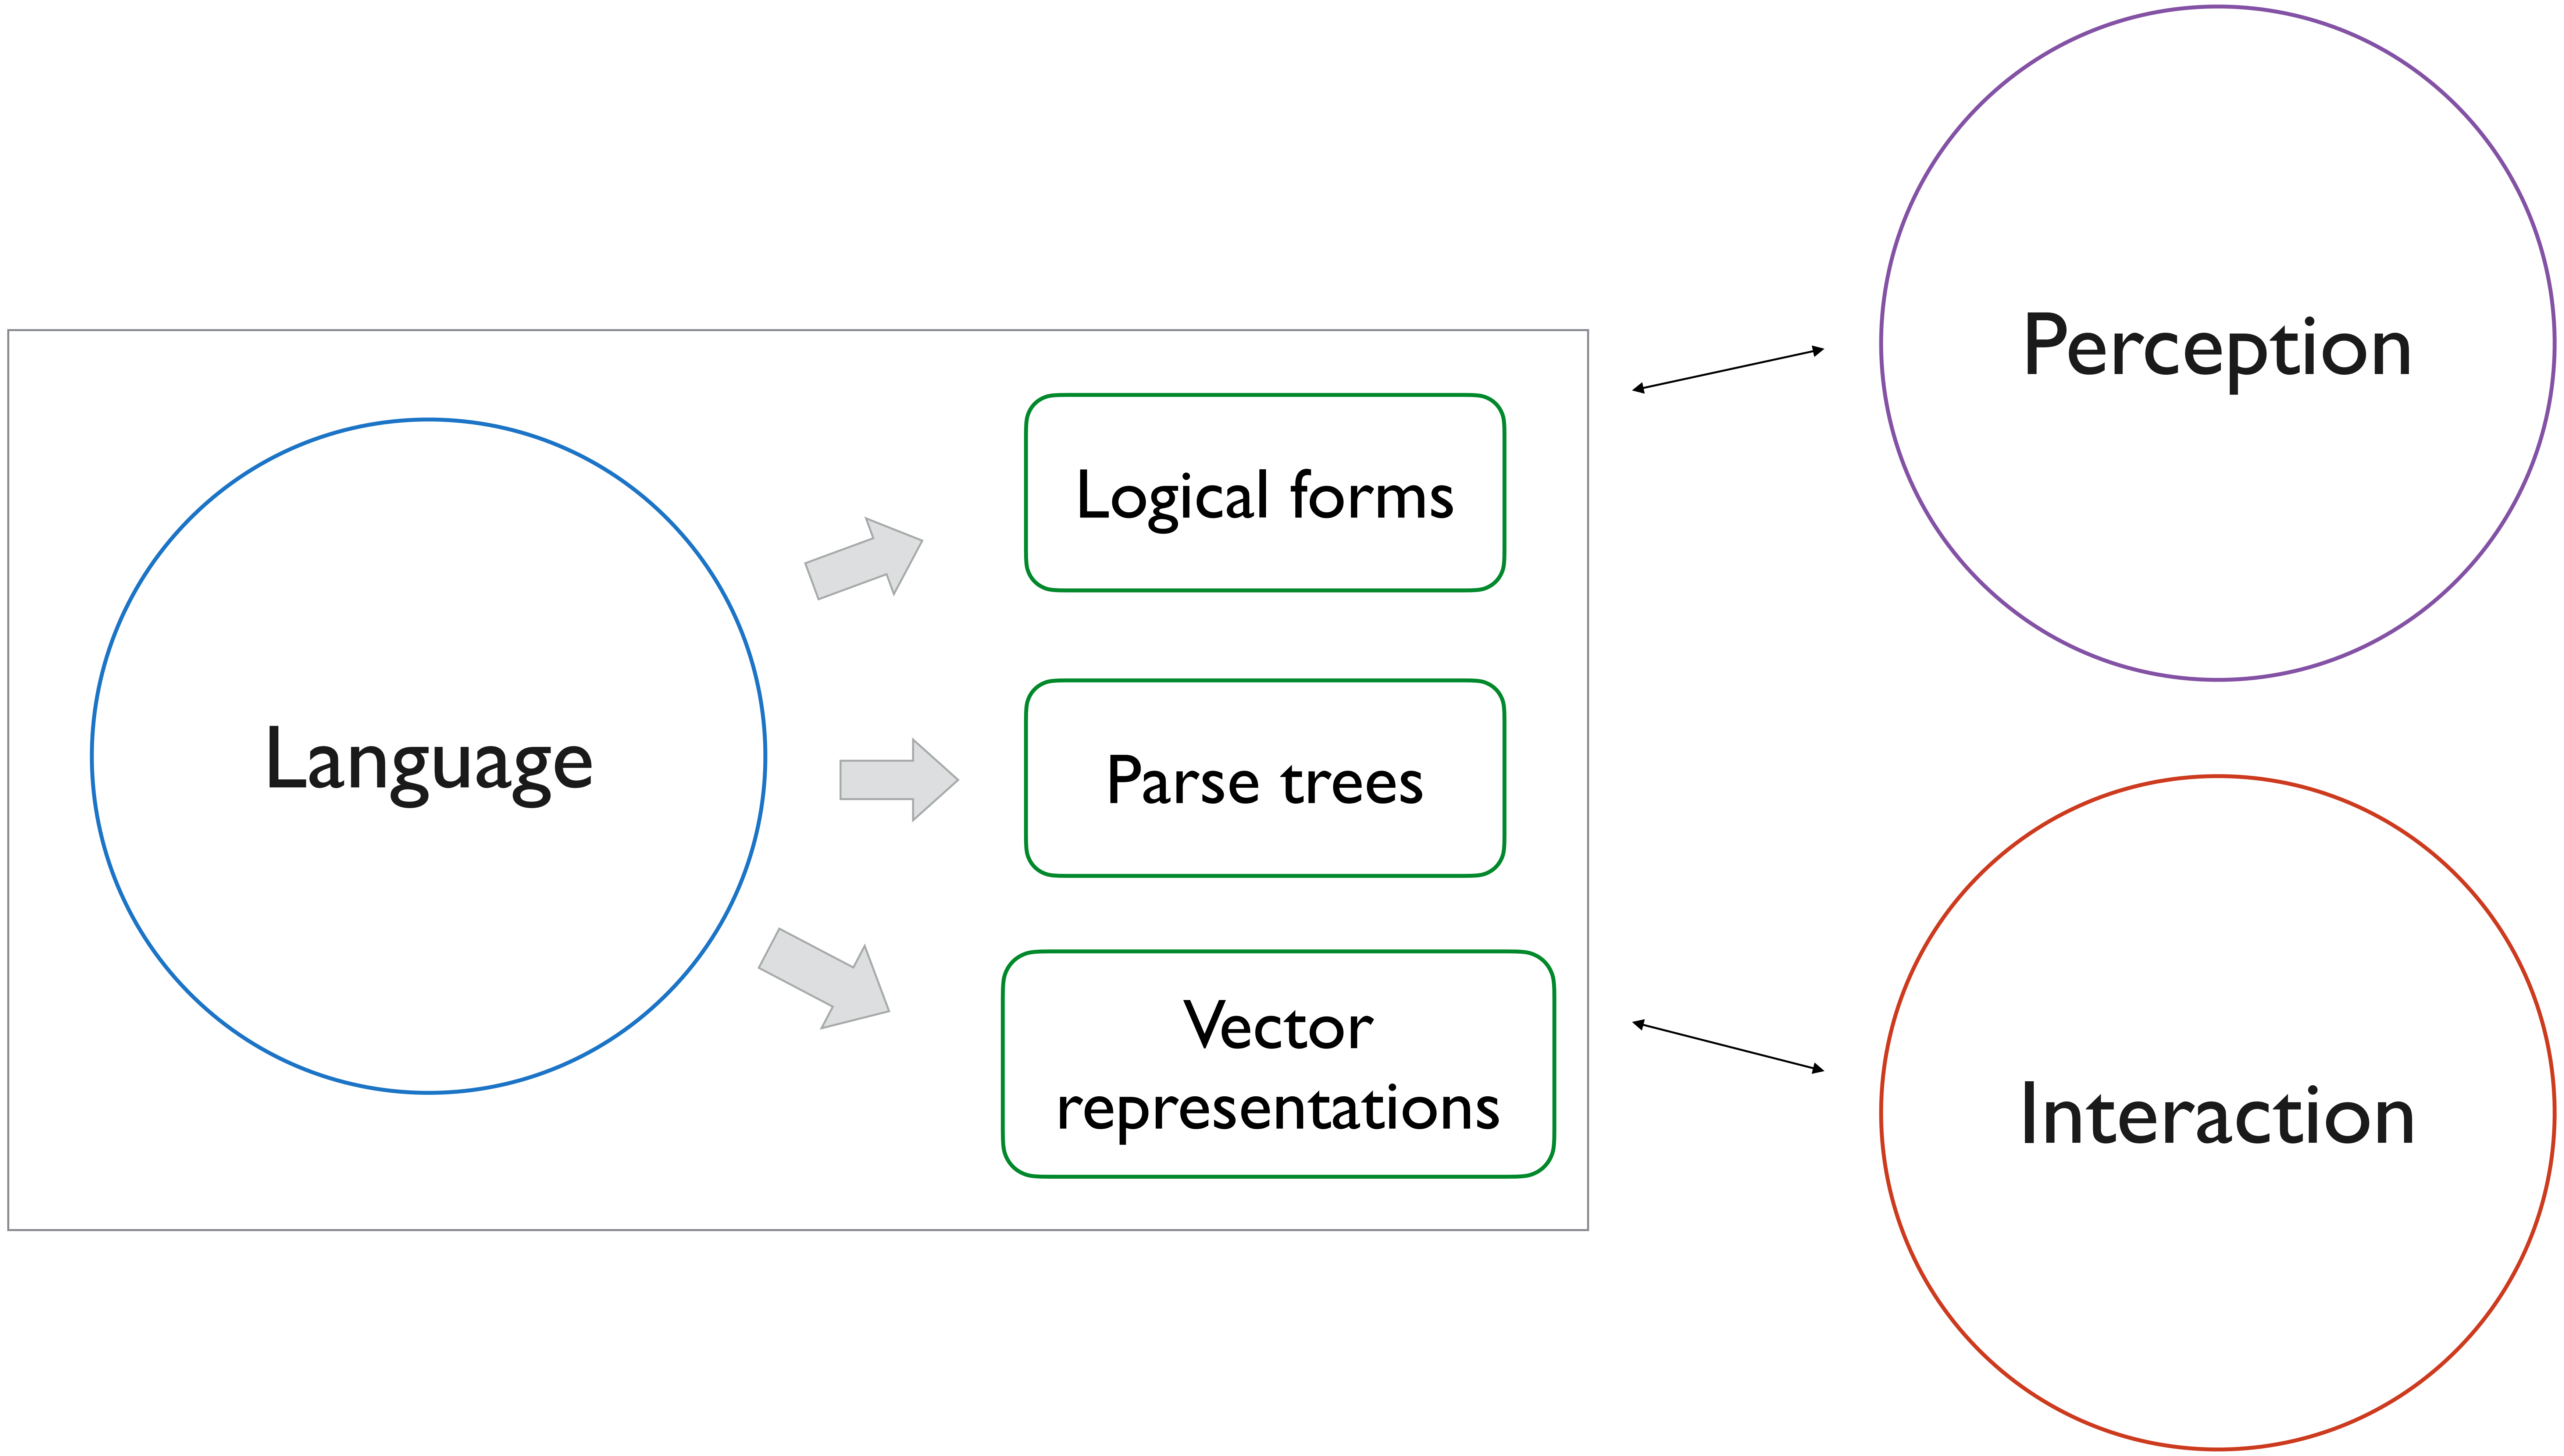

Is coffee a carcinogen?

Coffee significantly reduced ER and cyclin D1 abundance in ER(+) cells
…
Coffee reduced the pAkt levels in both ER(+) and ER(-) cells.

Hard to understand!

Information Extraction: State of the Art

Dependence on large training sets

ACE: 300K words

Freebase: 24M relations

Not available for many domains (ex. medicine, crime)

Even large corpora do not guarantee high performance

~ 75% F1 on relation extraction (ACE)

~ 58% F1 on event extraction (ACE)

IE: A hard reading task for machines

CBS Chicago

A 2 year old girl and <u>four</u> other people were wounded in a shooting in West Englewood Thursday night, police said

Extraction (NumWounded)

four ❌

IE: A hard reading task (not always!)

		Extraction (NumWounded)
CBS Chicago	A 2 year old girl and four other people were wounded in a shooting in West Englewood Thursday night, police said	four ✗
abc 7	The last shooting left five people wounded.	five ✓

Incorporate external evidence

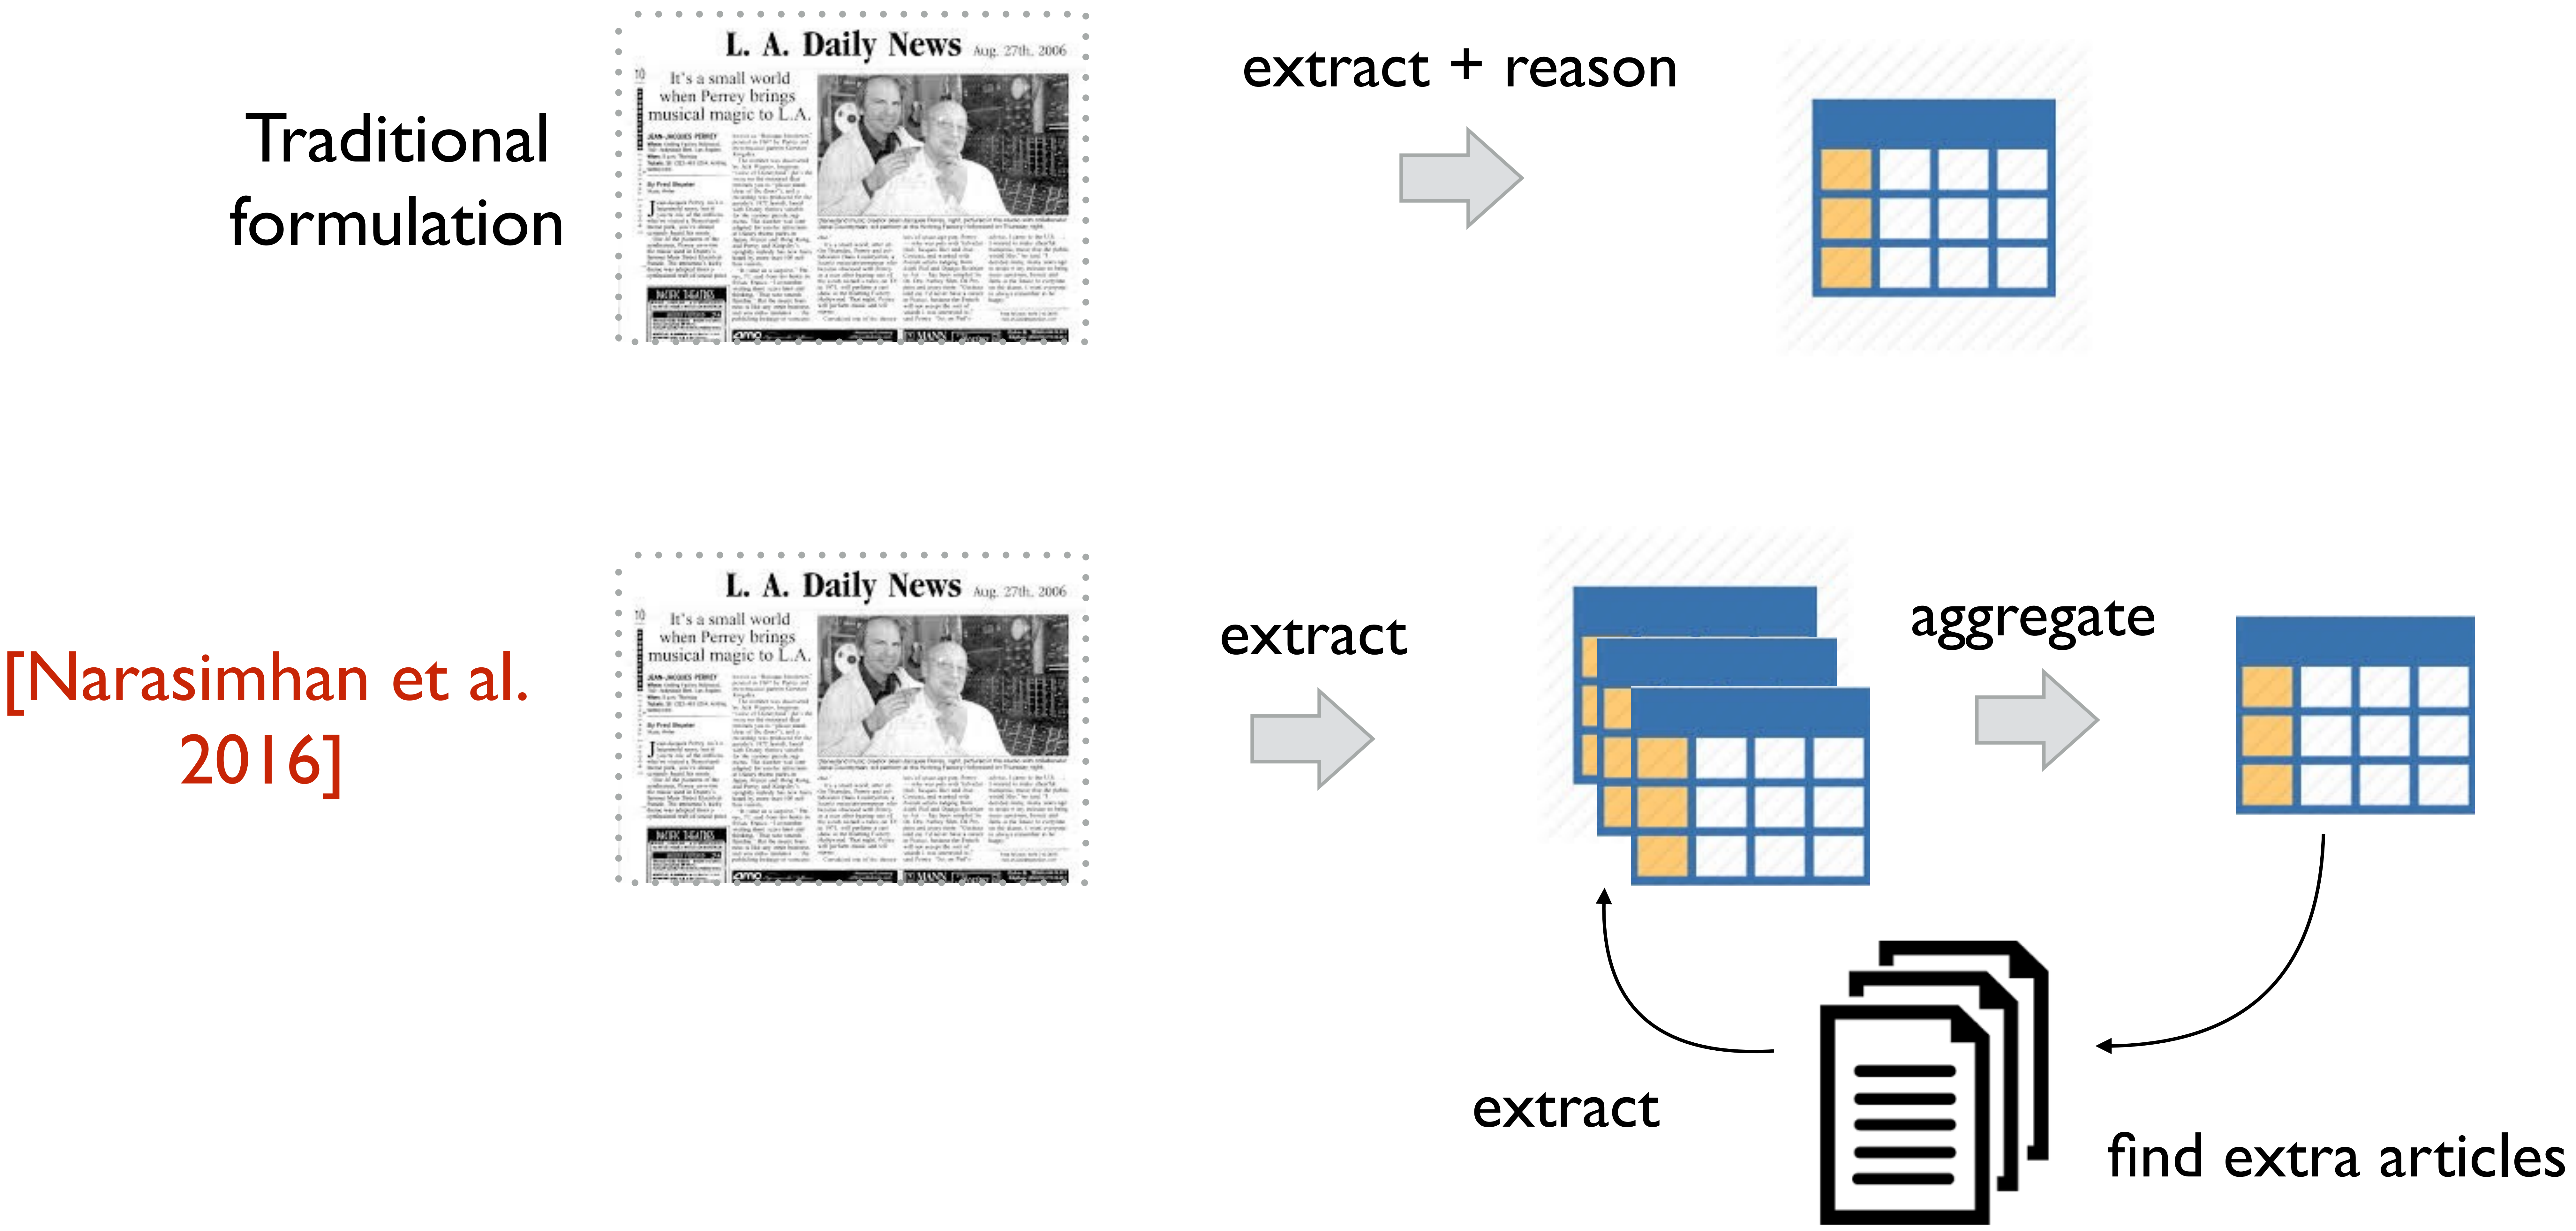

Challenges

1. Event Coreference

4 adults, 1 teenager shot in west Baltimore

All News Shopping Images Videos More ▾ Search tools

About 16,200,000 results (0.63 seconds)

4 adults, 1 teenager shot in west Baltimore | Maryland News ...
www.wbaltv.com/news/...**shot-in-west-baltimore**/32156116 ▾ WBAL-TV ▾
Apr 3, 2015 - Five people were **shot** Thursday afternoon in **west Baltimore**.

1 killed, 3 injured in Baltimore shooting, police say ... - WBAL
www.wbaltv.com/news/...**shot-in-west-baltimore**.../36588266 ▾ WBAL-TV ▾
Nov 21, 2015 - **2 teens**, **2 adults shot** on Stricker Street ... man was **killed** and three others were injured in a **shooting** Saturday morning in **west Baltimore**, police said. ... Mom tries to buy baby **for** her 14-year-old daughter; WBALTV.com. Undo.

10-year-old boy shot in West Baltimore - Baltimore Sun
www.**baltimoresun**.com/.../**baltimore**.../bs-md-ci-**shoot**... ▾ The Baltimore Sun ▾
Sep 3, 2015 - A 10-year-old boy was **shot** Thursday night, along with two **adult** ... **Baltimore** police report 6 shootings, including one of a **teenage** boy. ... The homicide occurred about **4**:30 p.m. at Ninth and East Jeffrey streets in Brooklyn, police said. ... At **1**:20 a.m., officers found a 32-year-old **Baltimore** man **shot** in the ...

Several irrelevant articles!

2. Reconciling Predictions

Shooter: Scott Westerhuis

NumKilled: 4

Location: S.D

Shooter: Scott Westerhuis

NumKilled: 6

Location: Platte

Inconsistent extractions

Learning through reinforcement

original

SECTIONS

DAILY NEWS | NEWS

Crime U.S. World Politics

S.D. dad killed wife, four kids with shotgun setting house ablaze and killing self: authori

THE ASSOCIATED PRESS Tuesday, November 3, 2015, 8:59 PM

PLATTE, S.D. — Financial issues appear to have contributed to an educational cooperative business manager's decision to kill his wife and four children with a shotgun before setting the family home ablaze and then shooting himself, South Dakota's attorney general said Tuesday.

Attorney General Marty Jackley released the results of his office's investigation of the September deaths at a news conference in Platte, a few miles north of the burned ruins of the home where the bodies of Scott and

extract

Shooter: Scott Westerhuis

NumKilled: 4

Location: S.D

Start with traditional extraction system

Learning through reinforcement

original

DAILY NEWS | NEWS

SECTIONS

Crime U.S. World Politics

S.D. dad killed wife, four kids with shotgun before setting house ablaze and killing self: authori

THE ASSOCIATED PRESS Tuesday, November 3, 2015, 8:59 PM

PLATTE, S.D. — Financial issues appear to have contributed to an educational cooperative business manager's decision to kill his wife and four children with a shotgun before setting the family home ablaze and then shooting himself, South Dakota's attorney general said Tuesday.

Attorney General Marty Jackley released the results of his office's investigation of the September deaths at a news conference in Platte, a few miles north of the burned ruins of the home where the bodies of Scott and

extract

Shooter: Scott Westerhuis

NumKilled: 4

Location: S.D

query

Platte Fire: Westerhuis Family Apparent Murder-Suicide, Offi

by HENRY AUSTIN

SHARE

FROM SEPT. 19TH: Westerhuis Neighbor Reacts to Home Fire 0:28

A couple and four children found dead in their burning South Dakota home had been shot in an apparent murder-suicide, officials said Monday.

extract

Shooter: Scott Westerhuis

NumKilled: 6

Location: Platte

Perform a query and extract from a new article

Learning through reinforcement

original

DAILY NEWS | NEWS

Crime U.S. World Politics

S.D. dad killed wife, four kids with shotgun
setting house ablaze and killing self: authori

THE ASSOCIATED PRESS Tuesday, November 3, 2015, 8:59 PM

PLATTE, S.D. — Financial issues appear to have contributed to an educational cooperative business manager's decision to kill his wife and four children with a shotgun before setting the family home ablaze and then shooting himself, South Dakota's attorney general said Tuesday.

Attorney General Marty Jackley released the results of his office's investigation of the September deaths at a news conference in Platte, a few miles north of the burned ruins of the home where the bodies of Scott and

extract

query

Platte Fire: Westerhuis Family
Apparent Murder-Suicide, Offi

by HENRY AUSTIN

SHARE

FROM SEPT. 19TH: Westerhuis Neighbor Reacts to Home Fire 0:28

A couple and four children found dead in their burning South Dakota home had been shot in an apparent murder-suicide, officials said Monday.

extract

Current

Shooter: Scott Westerhuis

NumKilled: 4

Location: S.D

Shooter: Scott Westerhuis

NumKilled: 6

Location: Platte

State

New

RL: Actions

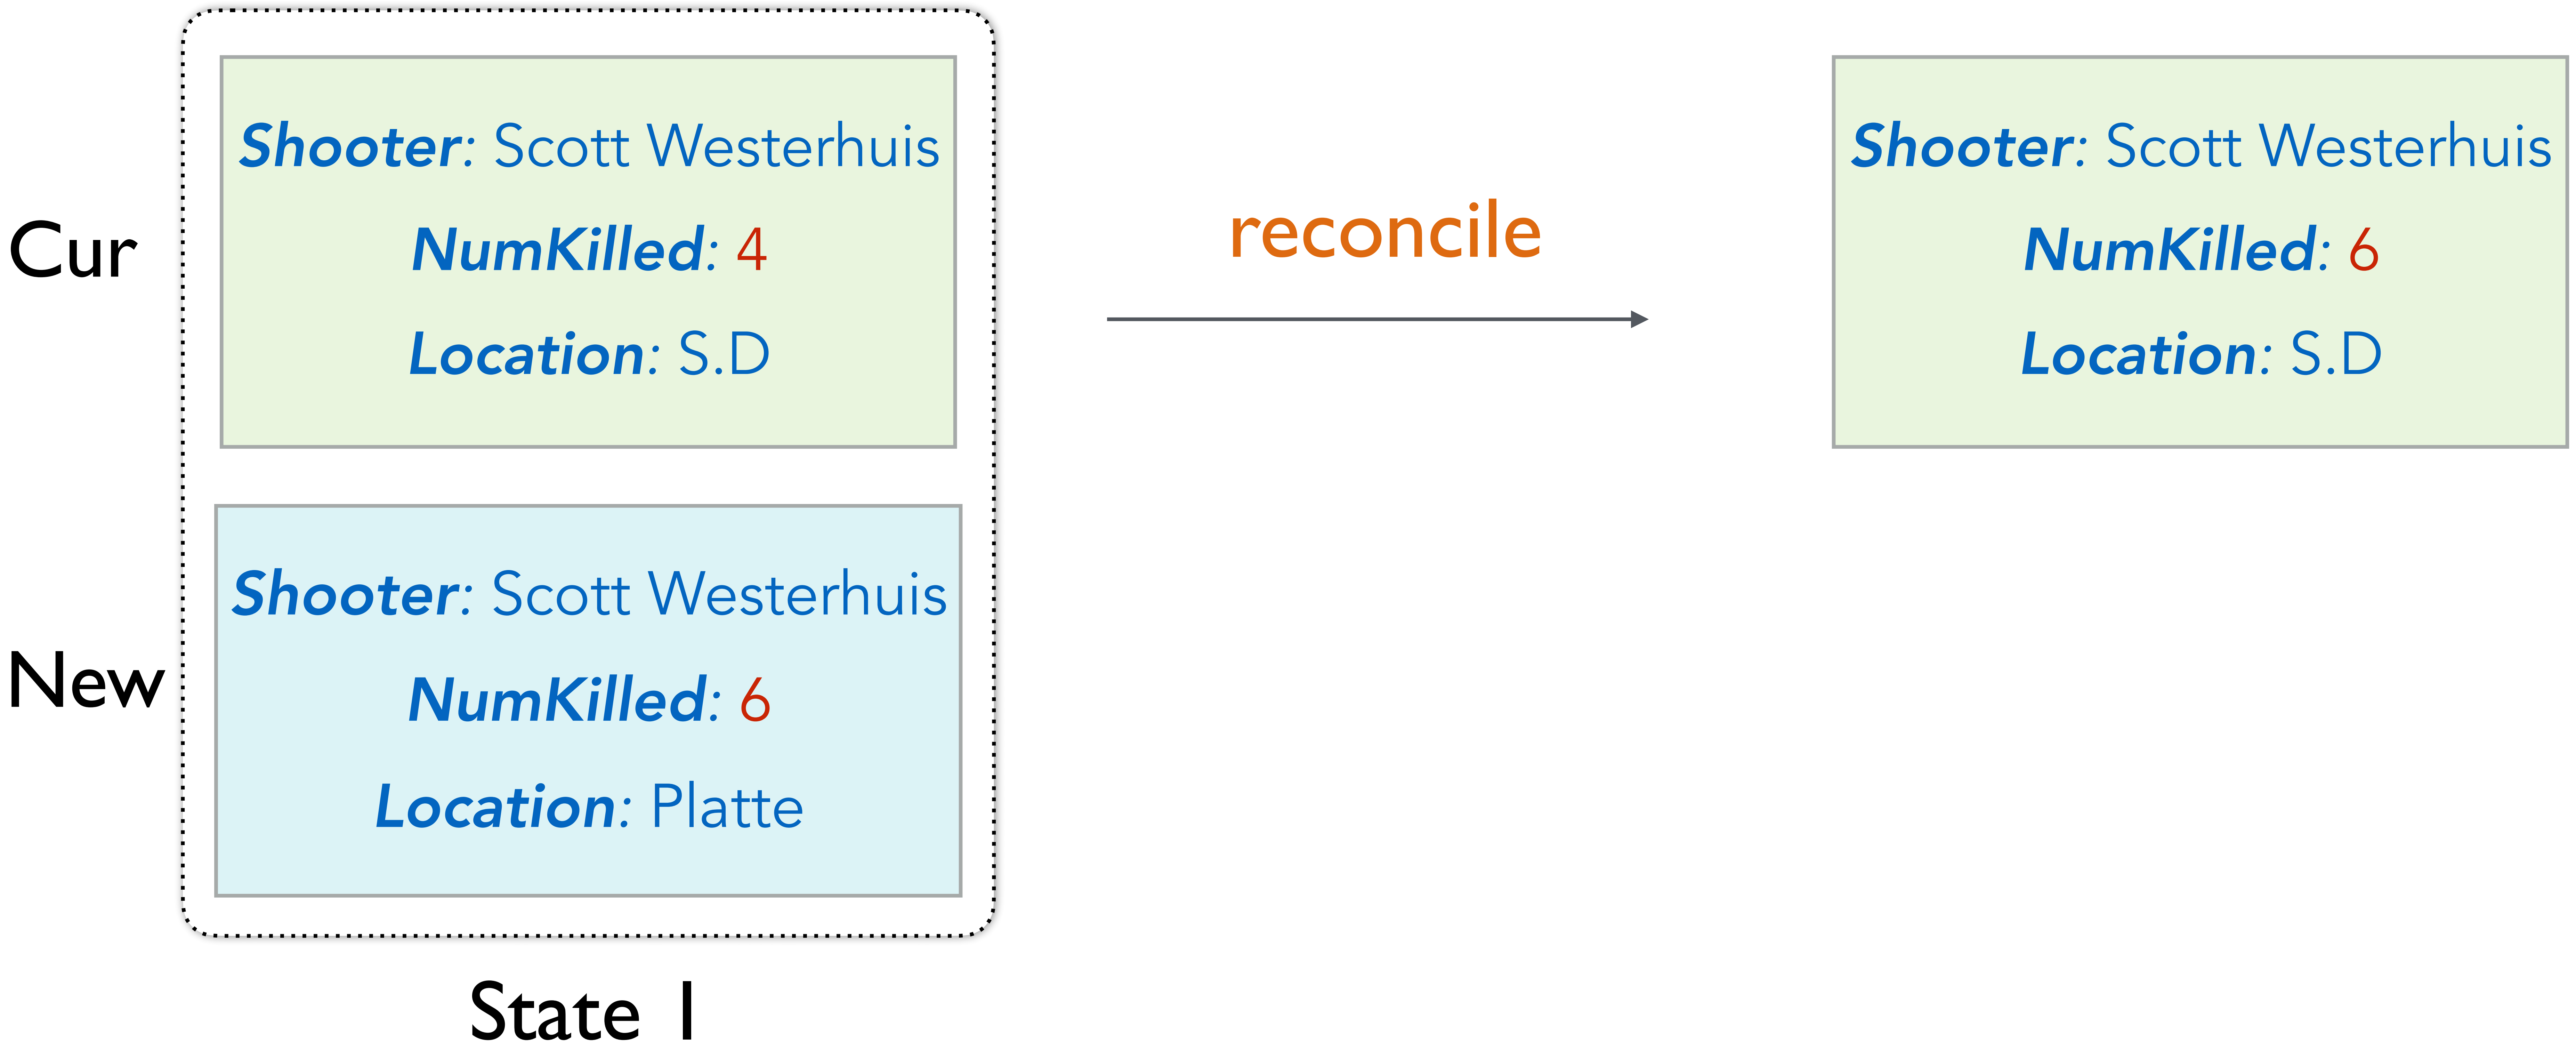

1. Reconcile (d) old values and new values.
   - Pick a single value, all values or no value from new set

RL: Actions

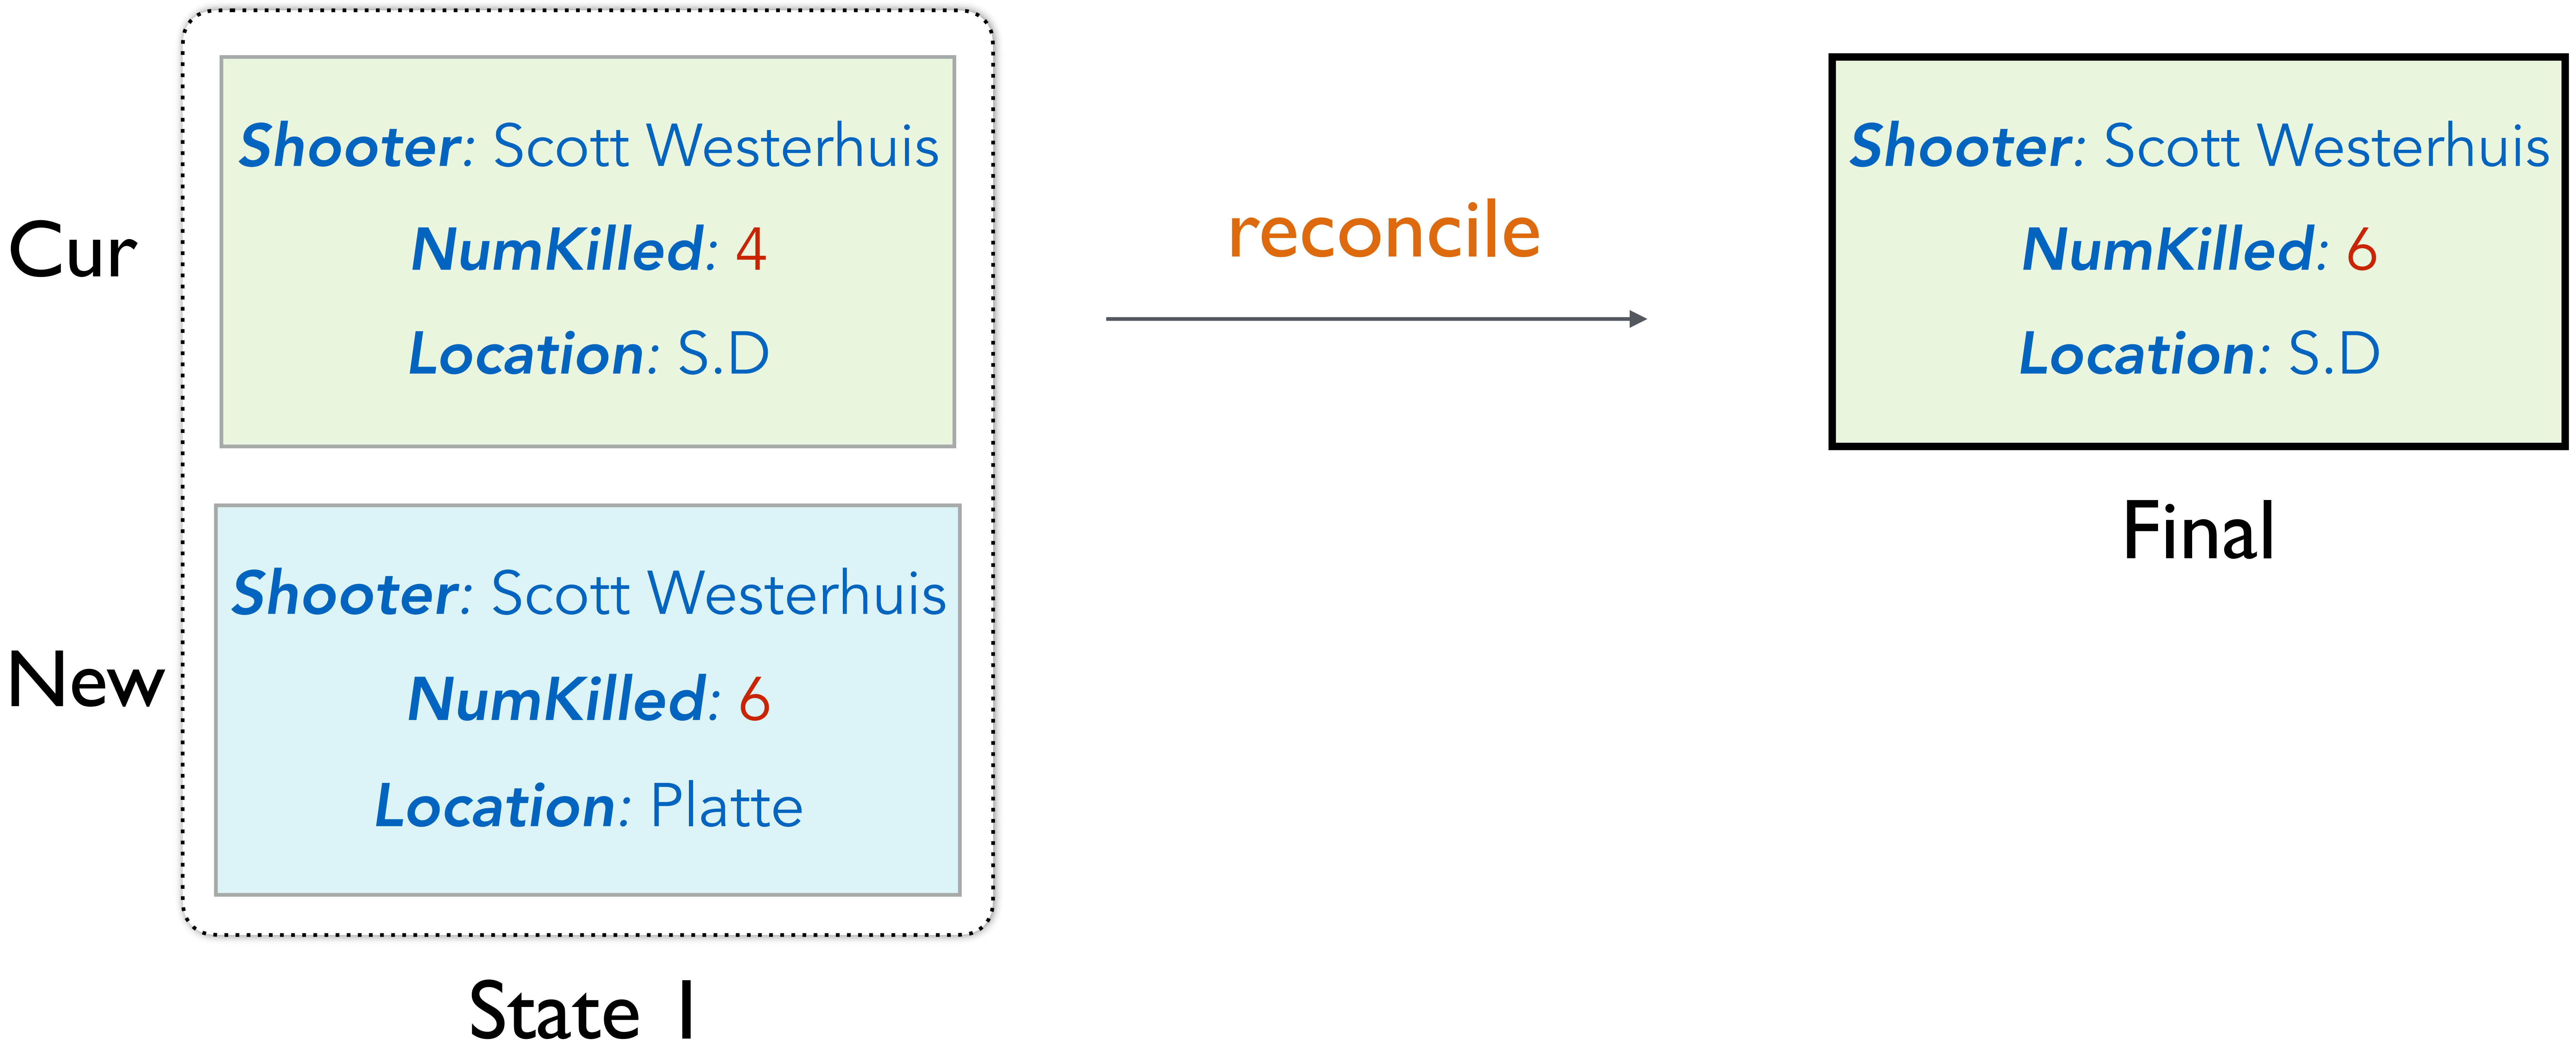

2. Decide how to proceed:
   - Stop

RL: Actions

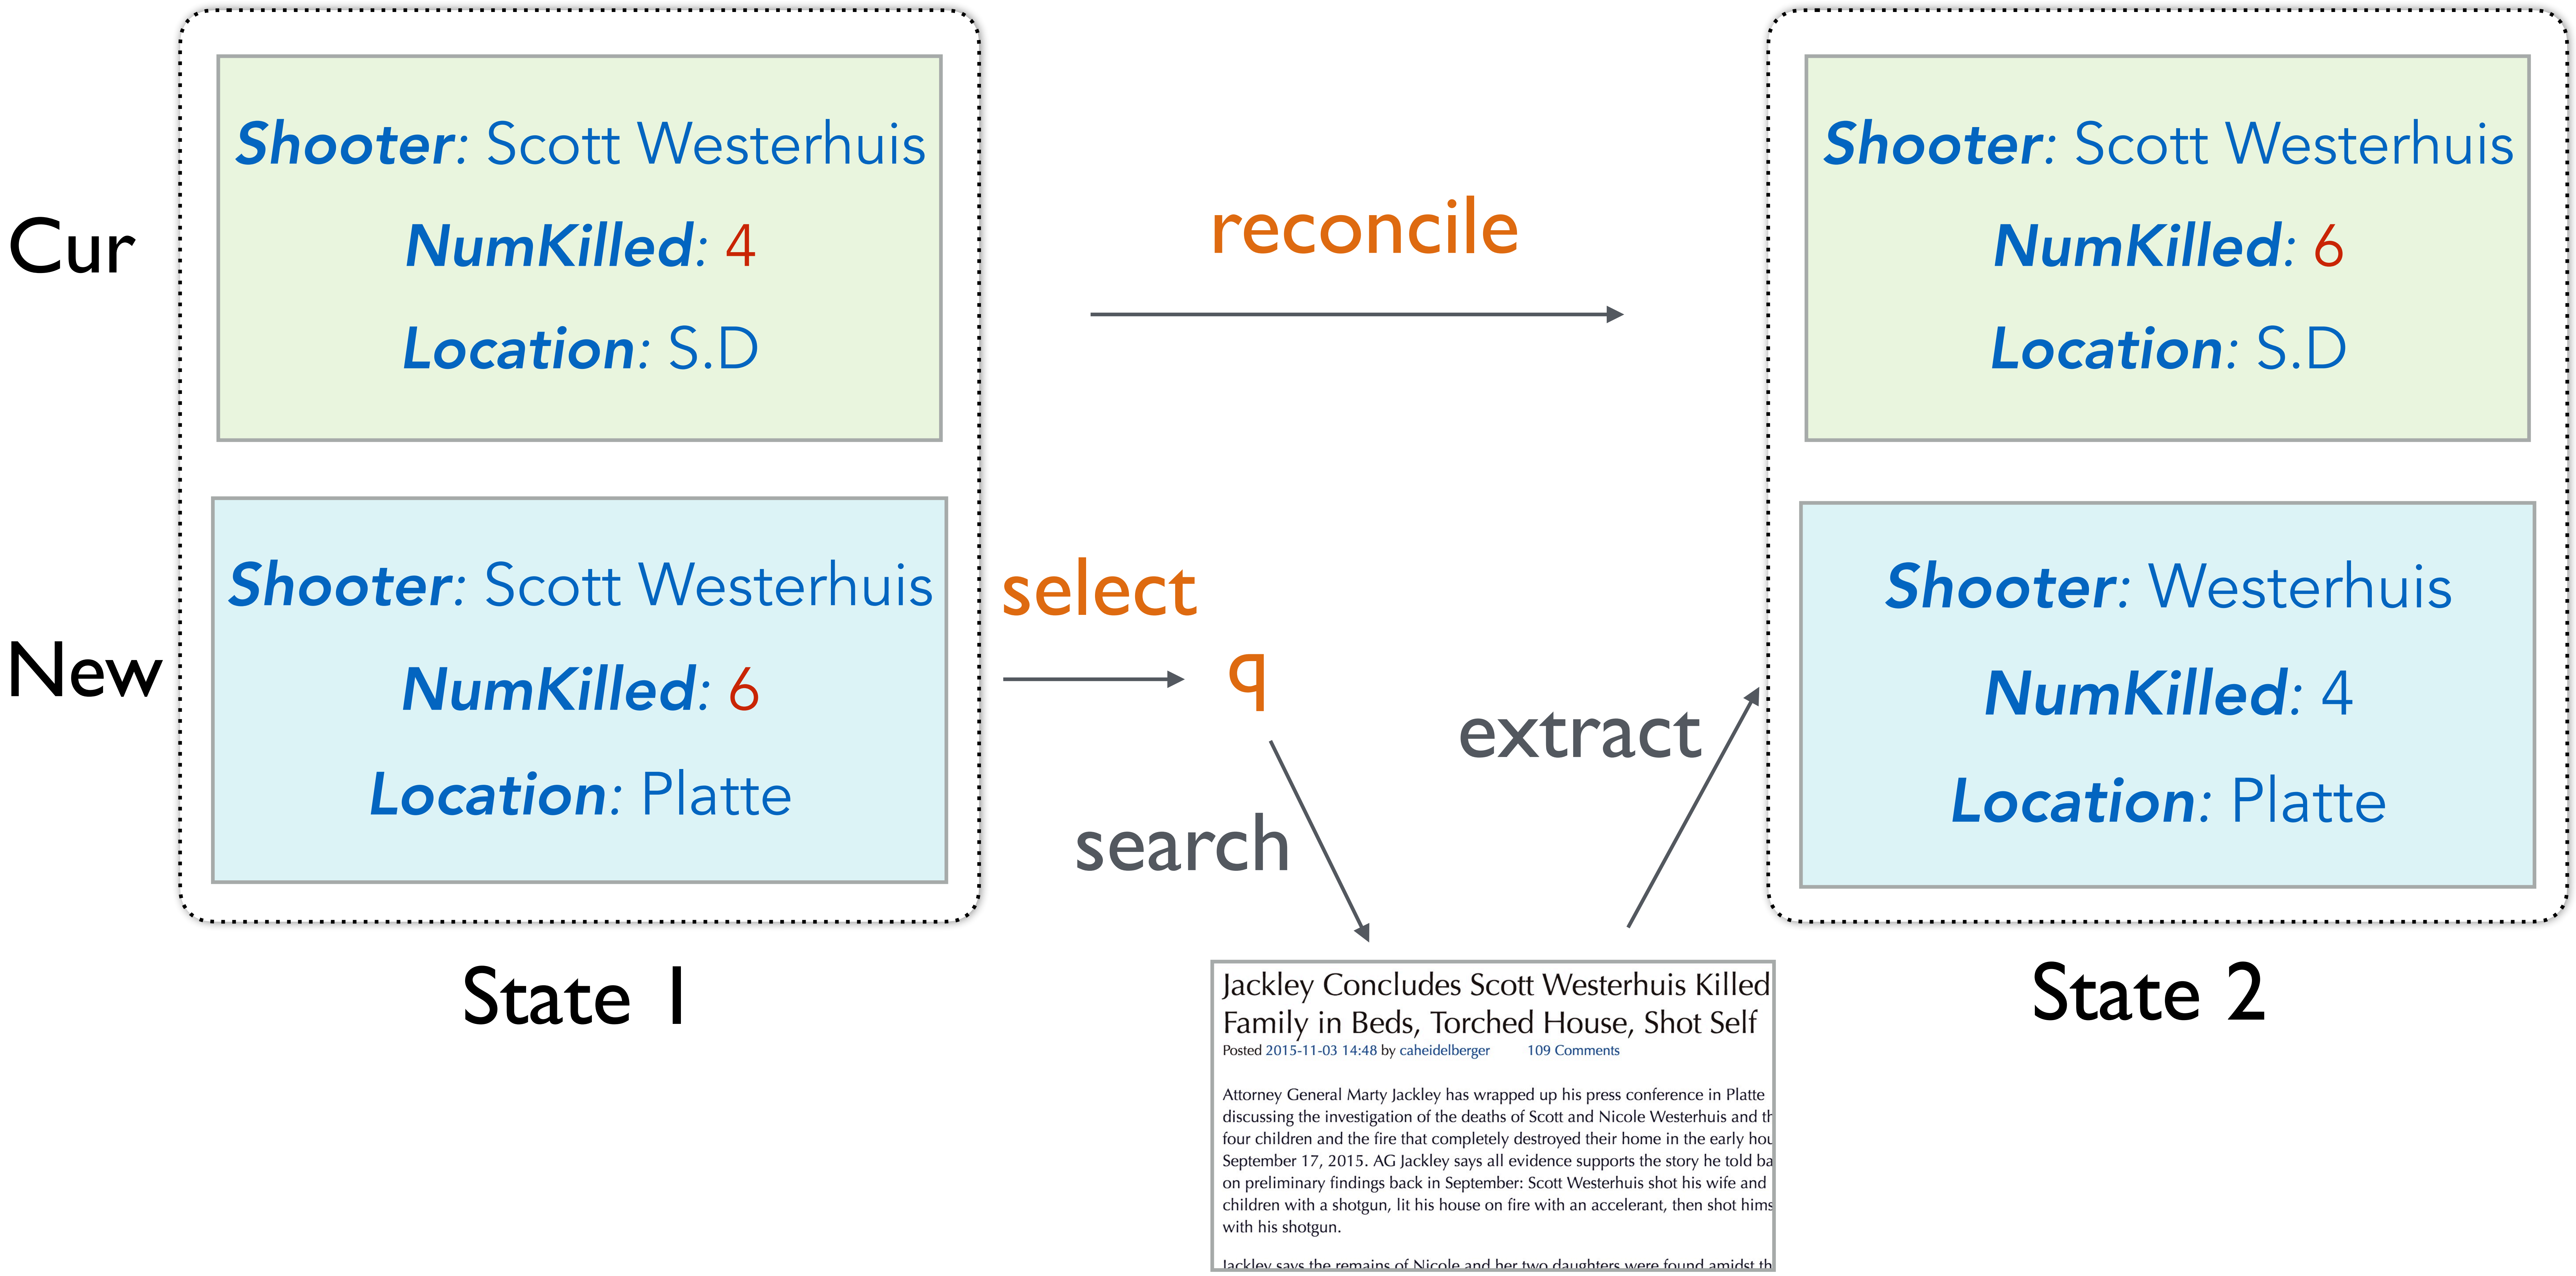

2. Decide how to proceed:

- ✦ Select next query (q)

Acquiring external evidence

1. Select a query to search for articles on the same event

2. Use base extractor to obtain values for entities of interest


Platte Fire: Westerhuis Family Died in Apparent Murder-Suicide, Officials Say

A couple and four children found dead in their burning South Dakota home had been shot in an apparent murder-suicide, officials said Monday.

extract

Shooter: Scott Westerhuis

NumKilled: 6

Location: Platte

3. Reconcile old and new extractions

Shooter: Scott Westerhuis

NumKilled: 4

Location: S.D

Shooter: Scott Westerhuis

NumKilled: 6

Location: Platte

Learning from rewards

- Change in accuracy

Previous Values

Shooter: Scott Westerhuis
NumKilled: 6
NumWounded: 1
Location: Platte

Current Values

Shooter: Scott Westerhuis
NumKilled: 6
NumWounded: 0
Location: Platte

$$R(s,a) = \sum_{\text{entity}\, j} \mathrm{Acc}(e^j_{cur}) - \mathrm{Acc}(e^j_{prev}) \quad = 1$$

- Small penalty for each transition

Mass shootings in the United States

GUN VIOLENCE Archive

HOME | CHARTS & MAPS | LAST 72 HOURS | REPORTS

→

- Shooter Name
- Num Killed
- Num Wounded
- City

~300 training instances

Adulteration incidents from Foodshield EMA

FDA U.S. FOOD & DRUG ADMINISTRATION

→

- Food
- Adulterant
- Location

Accuracy

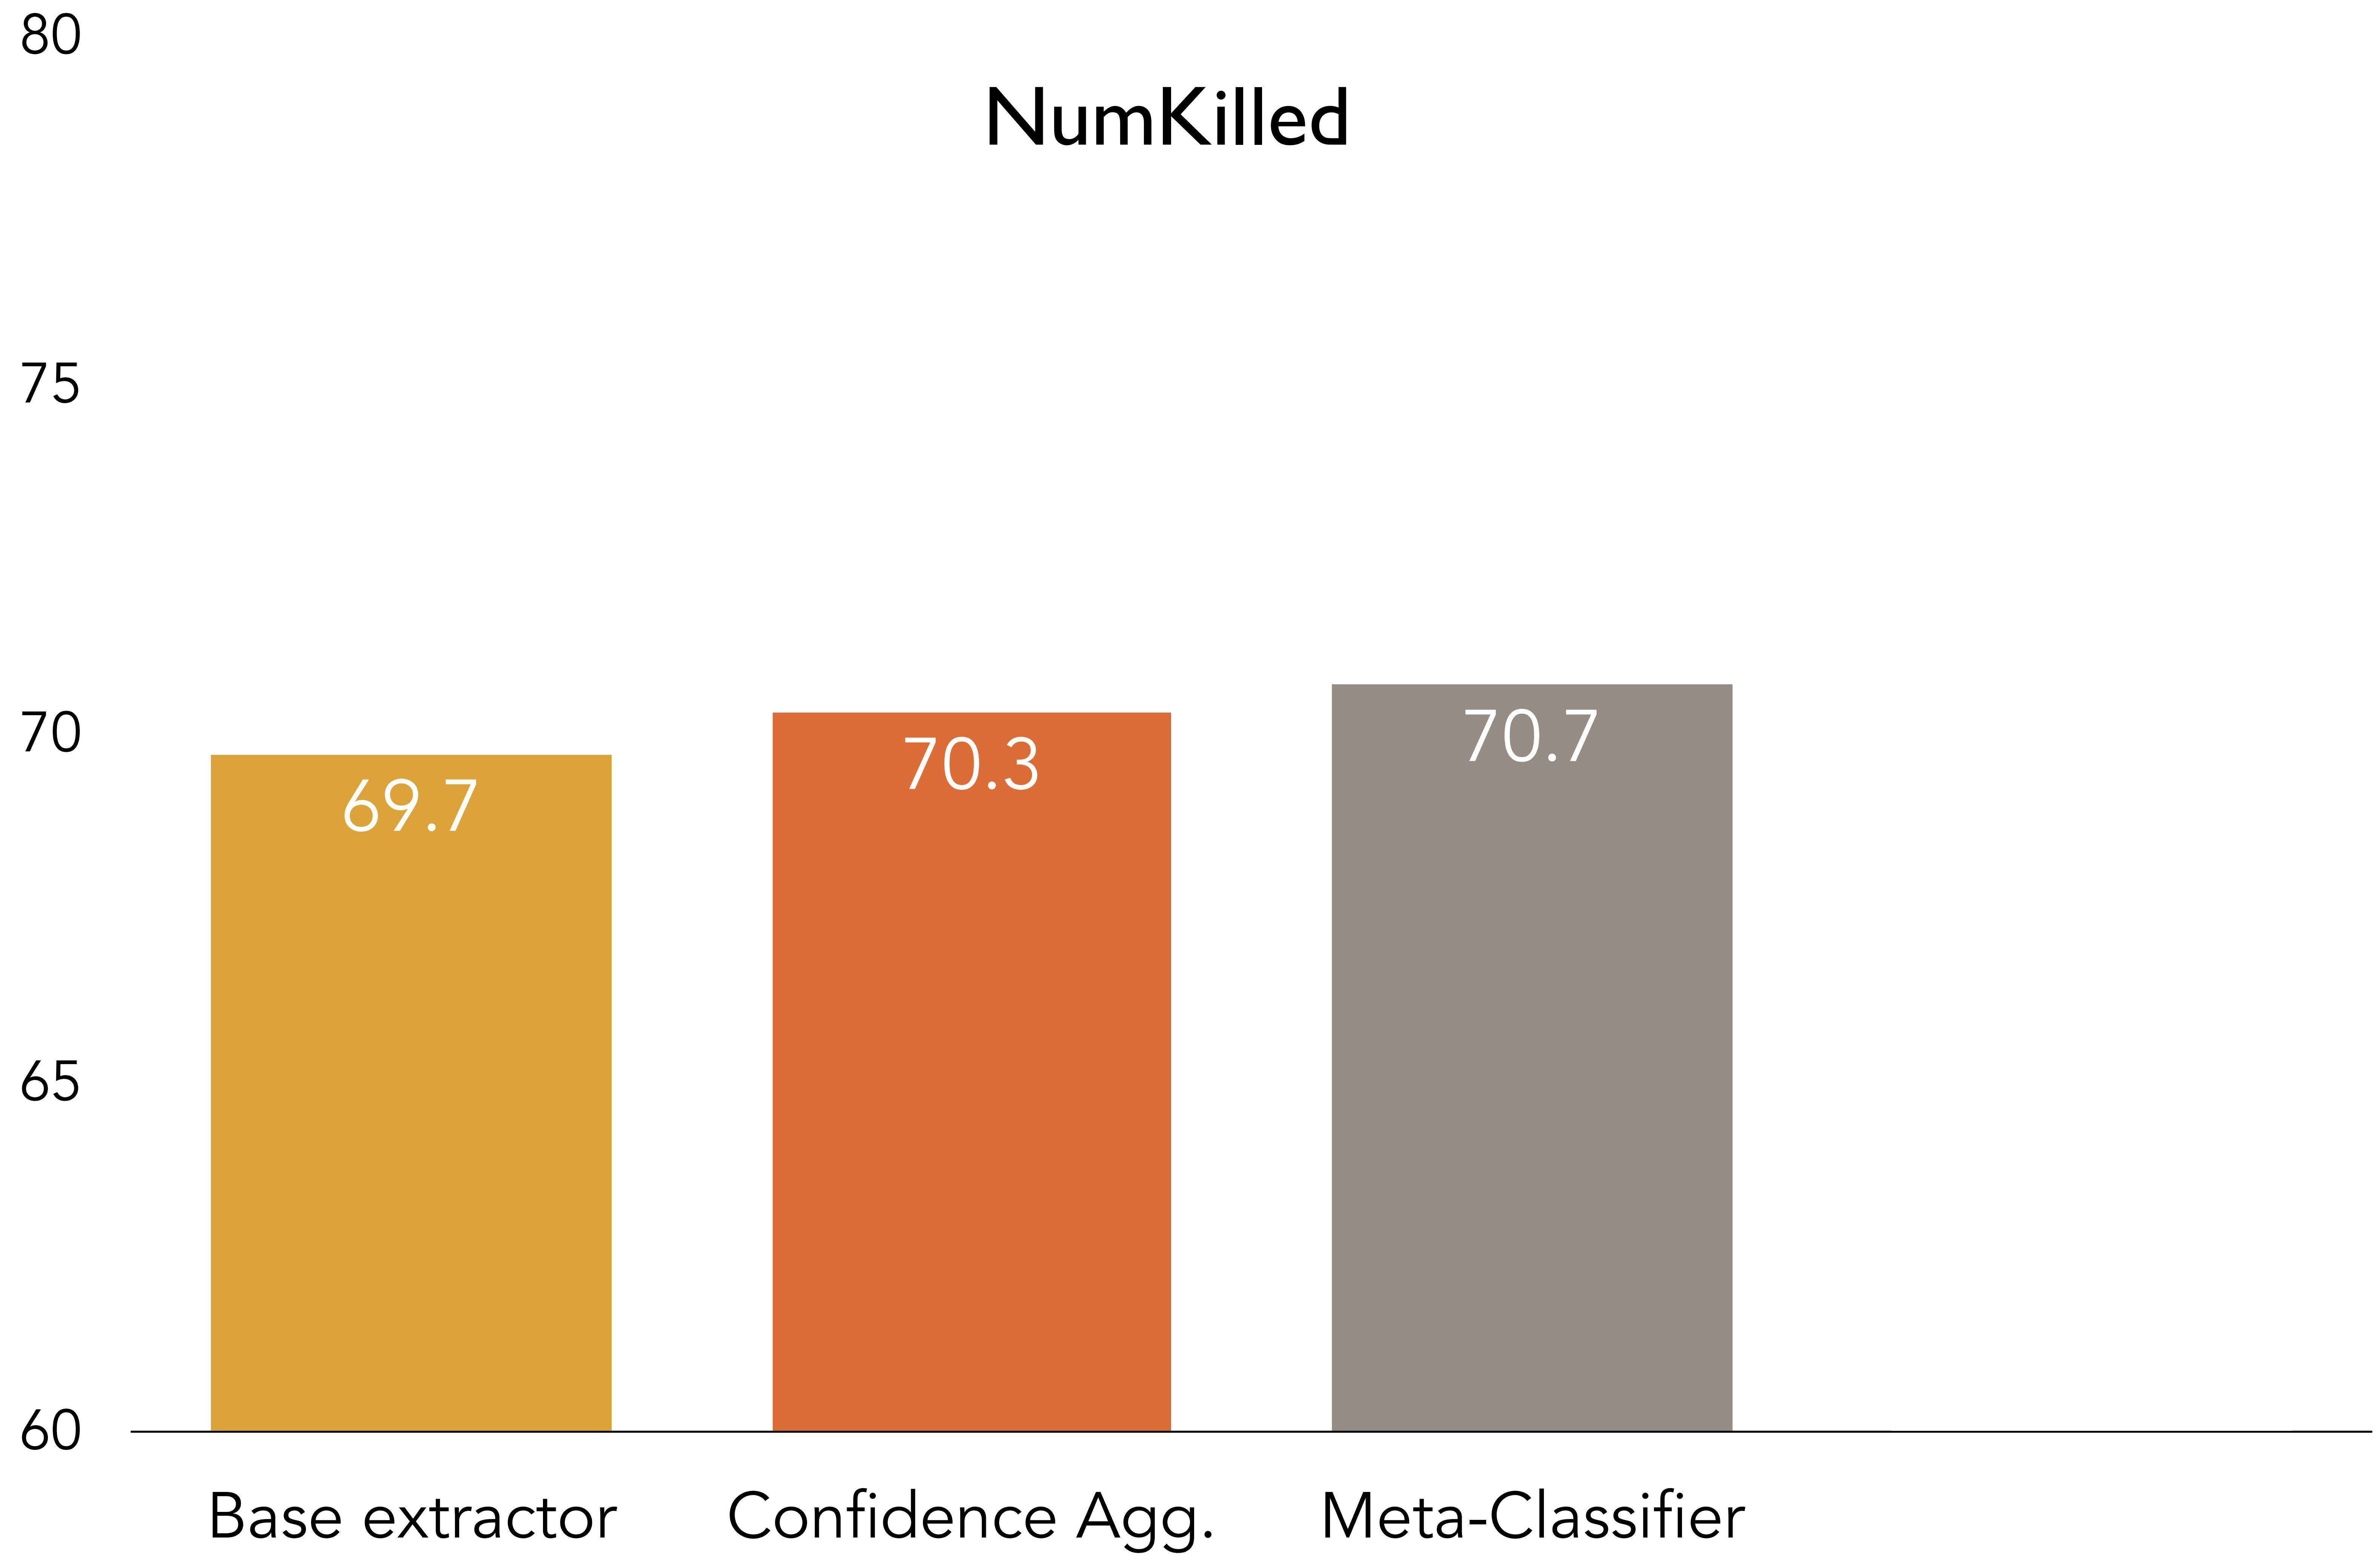
80
75
70
65
60
NumKilled
69.7
70.3
70.7
Base extractor
Confidence Agg.
Meta-Classifier

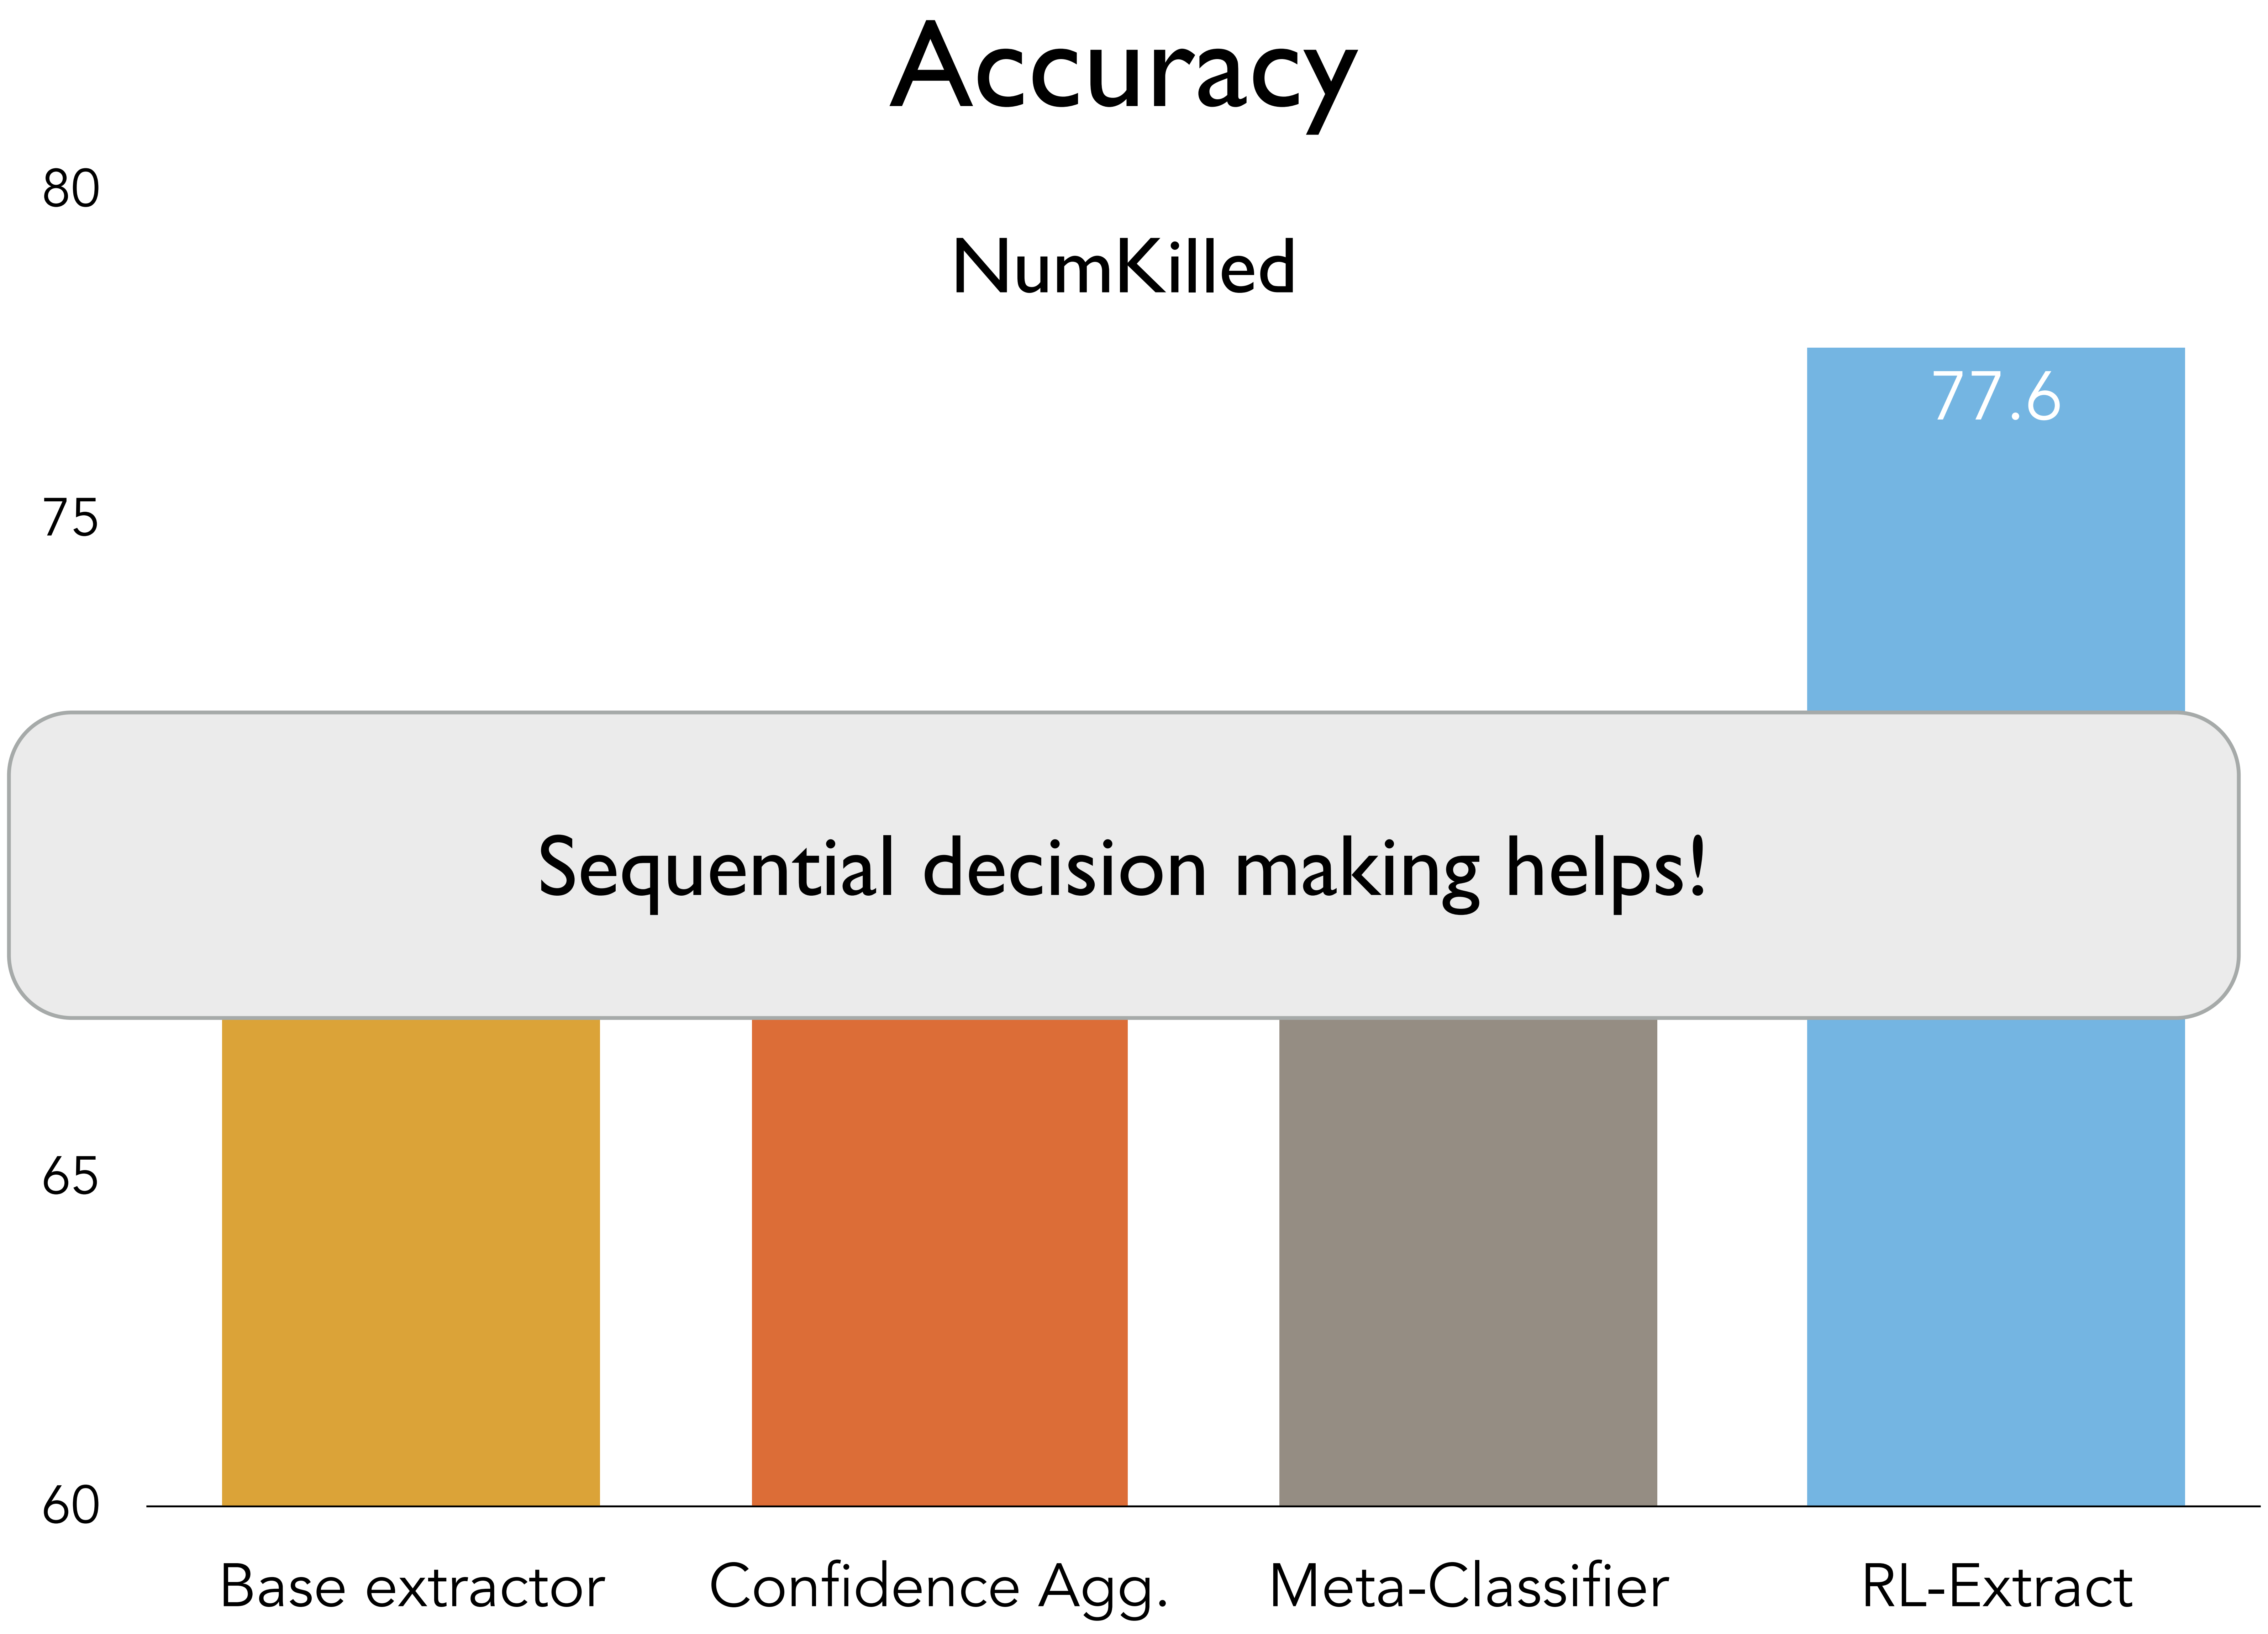

Accuracy
NumKilled
80
75
65
60
77.6
Sequential decision making helps!
Base extractor
Confidence Agg.
Meta-Classifier
RL-Extract