Multi-Hop Knowledge Graph Reasoning with Deep Reinforcement Learning

Richard Socher
Work by Victoria Lin et al.

Real-world Knowledge Graphs

Real-world Knowledge Graphs

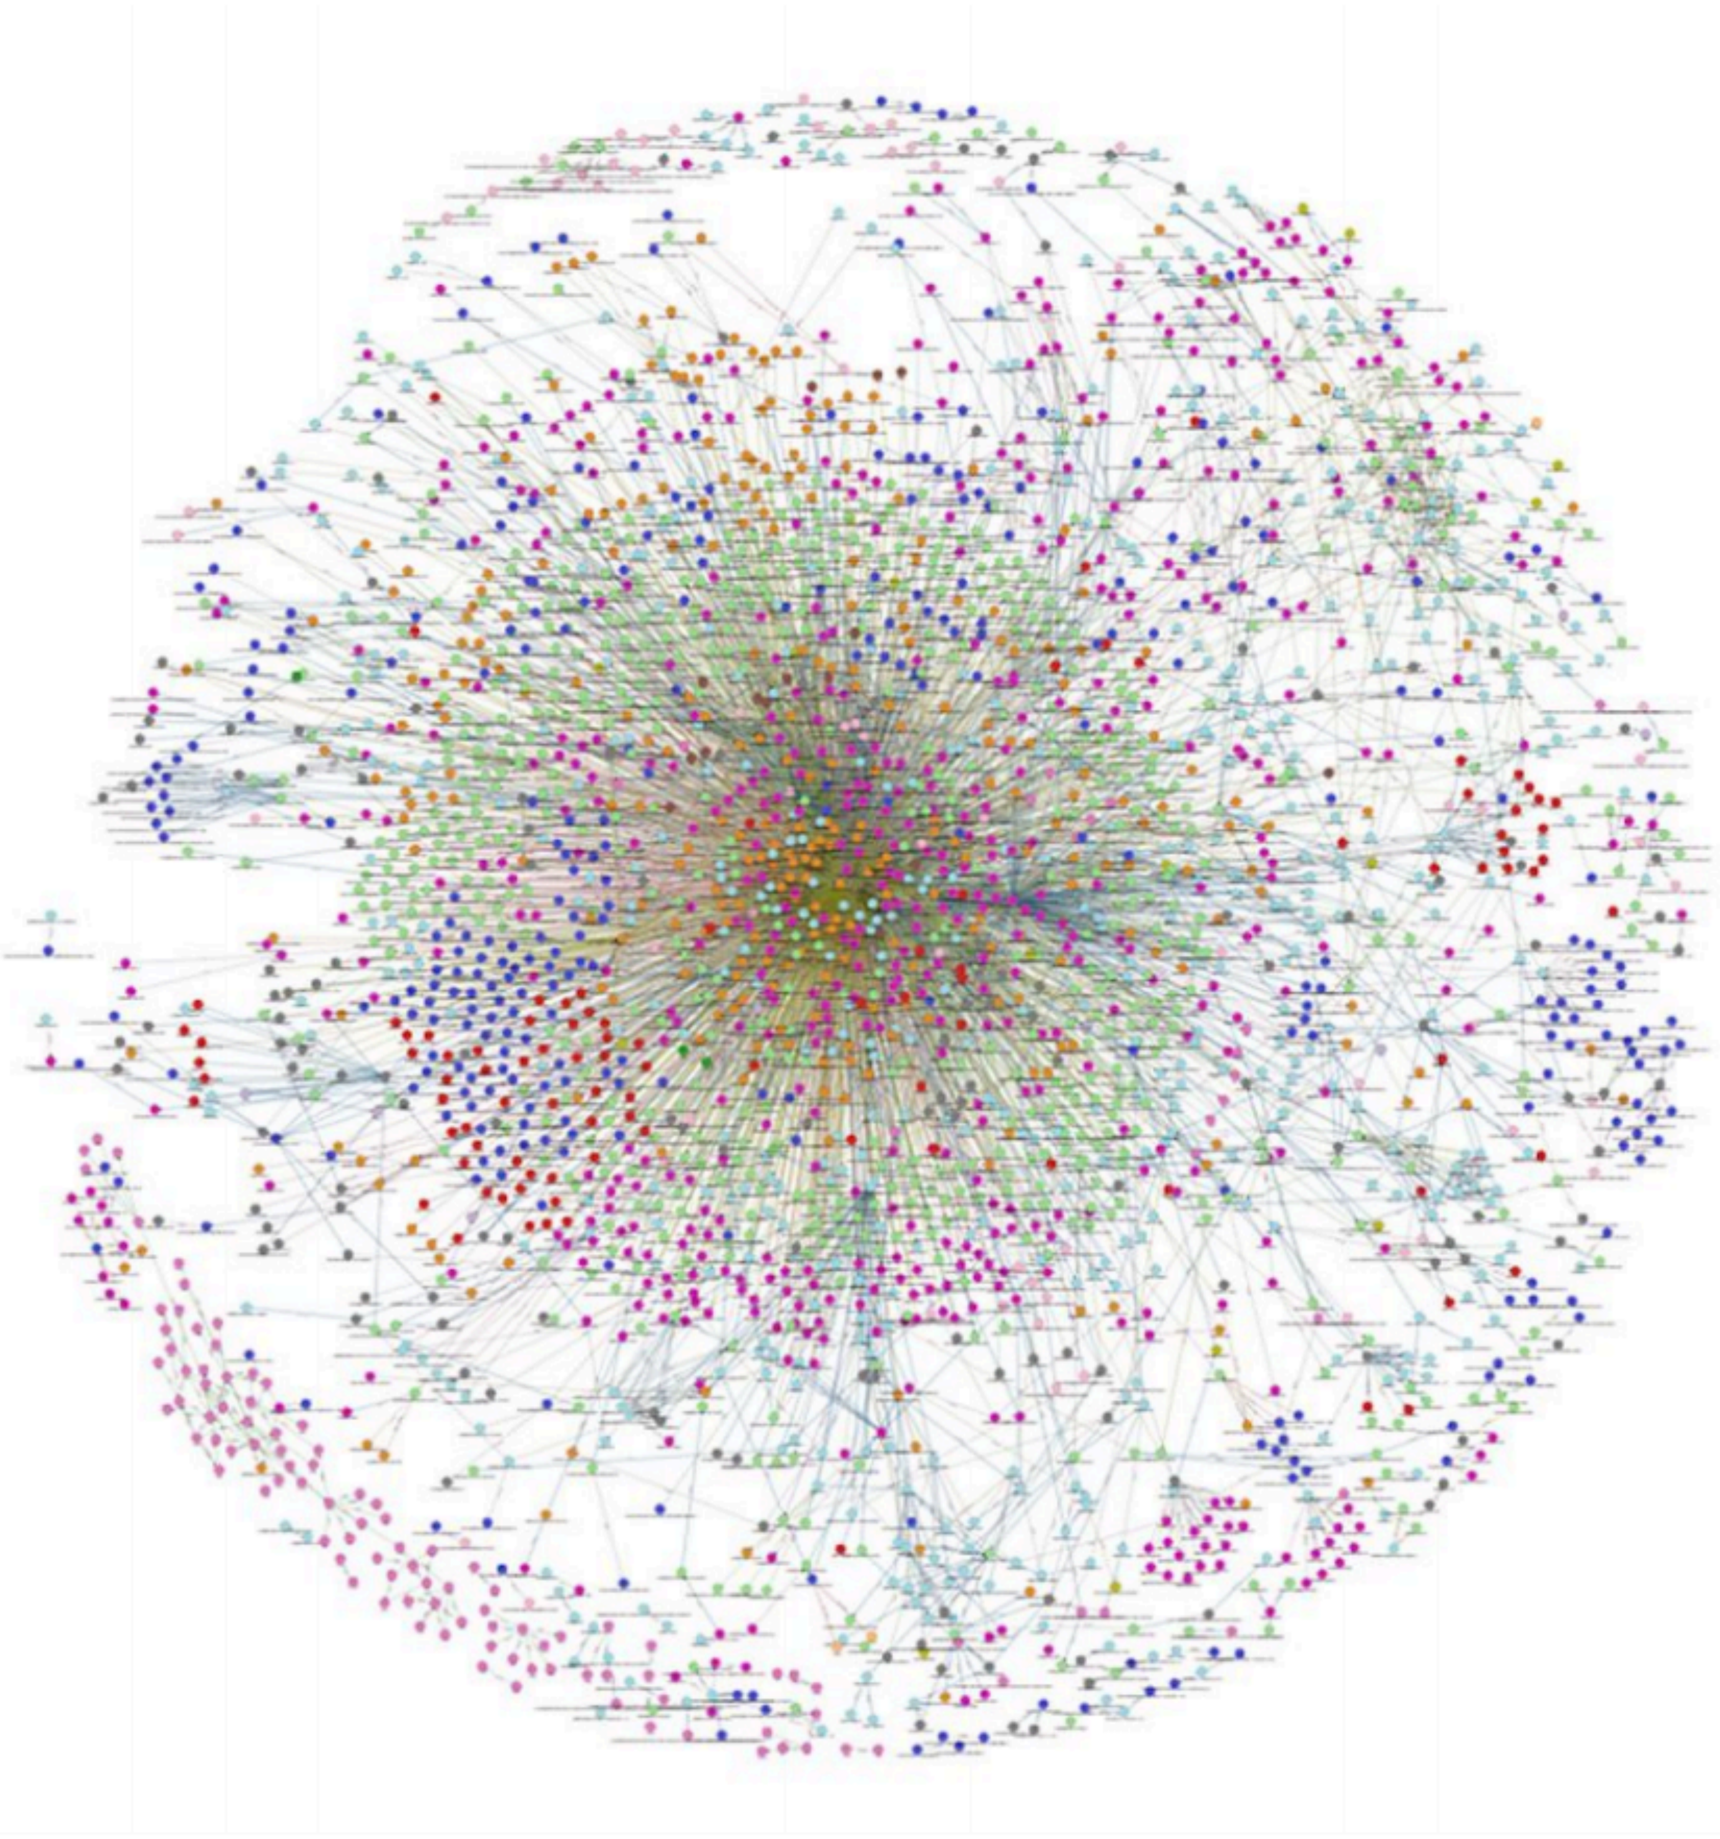

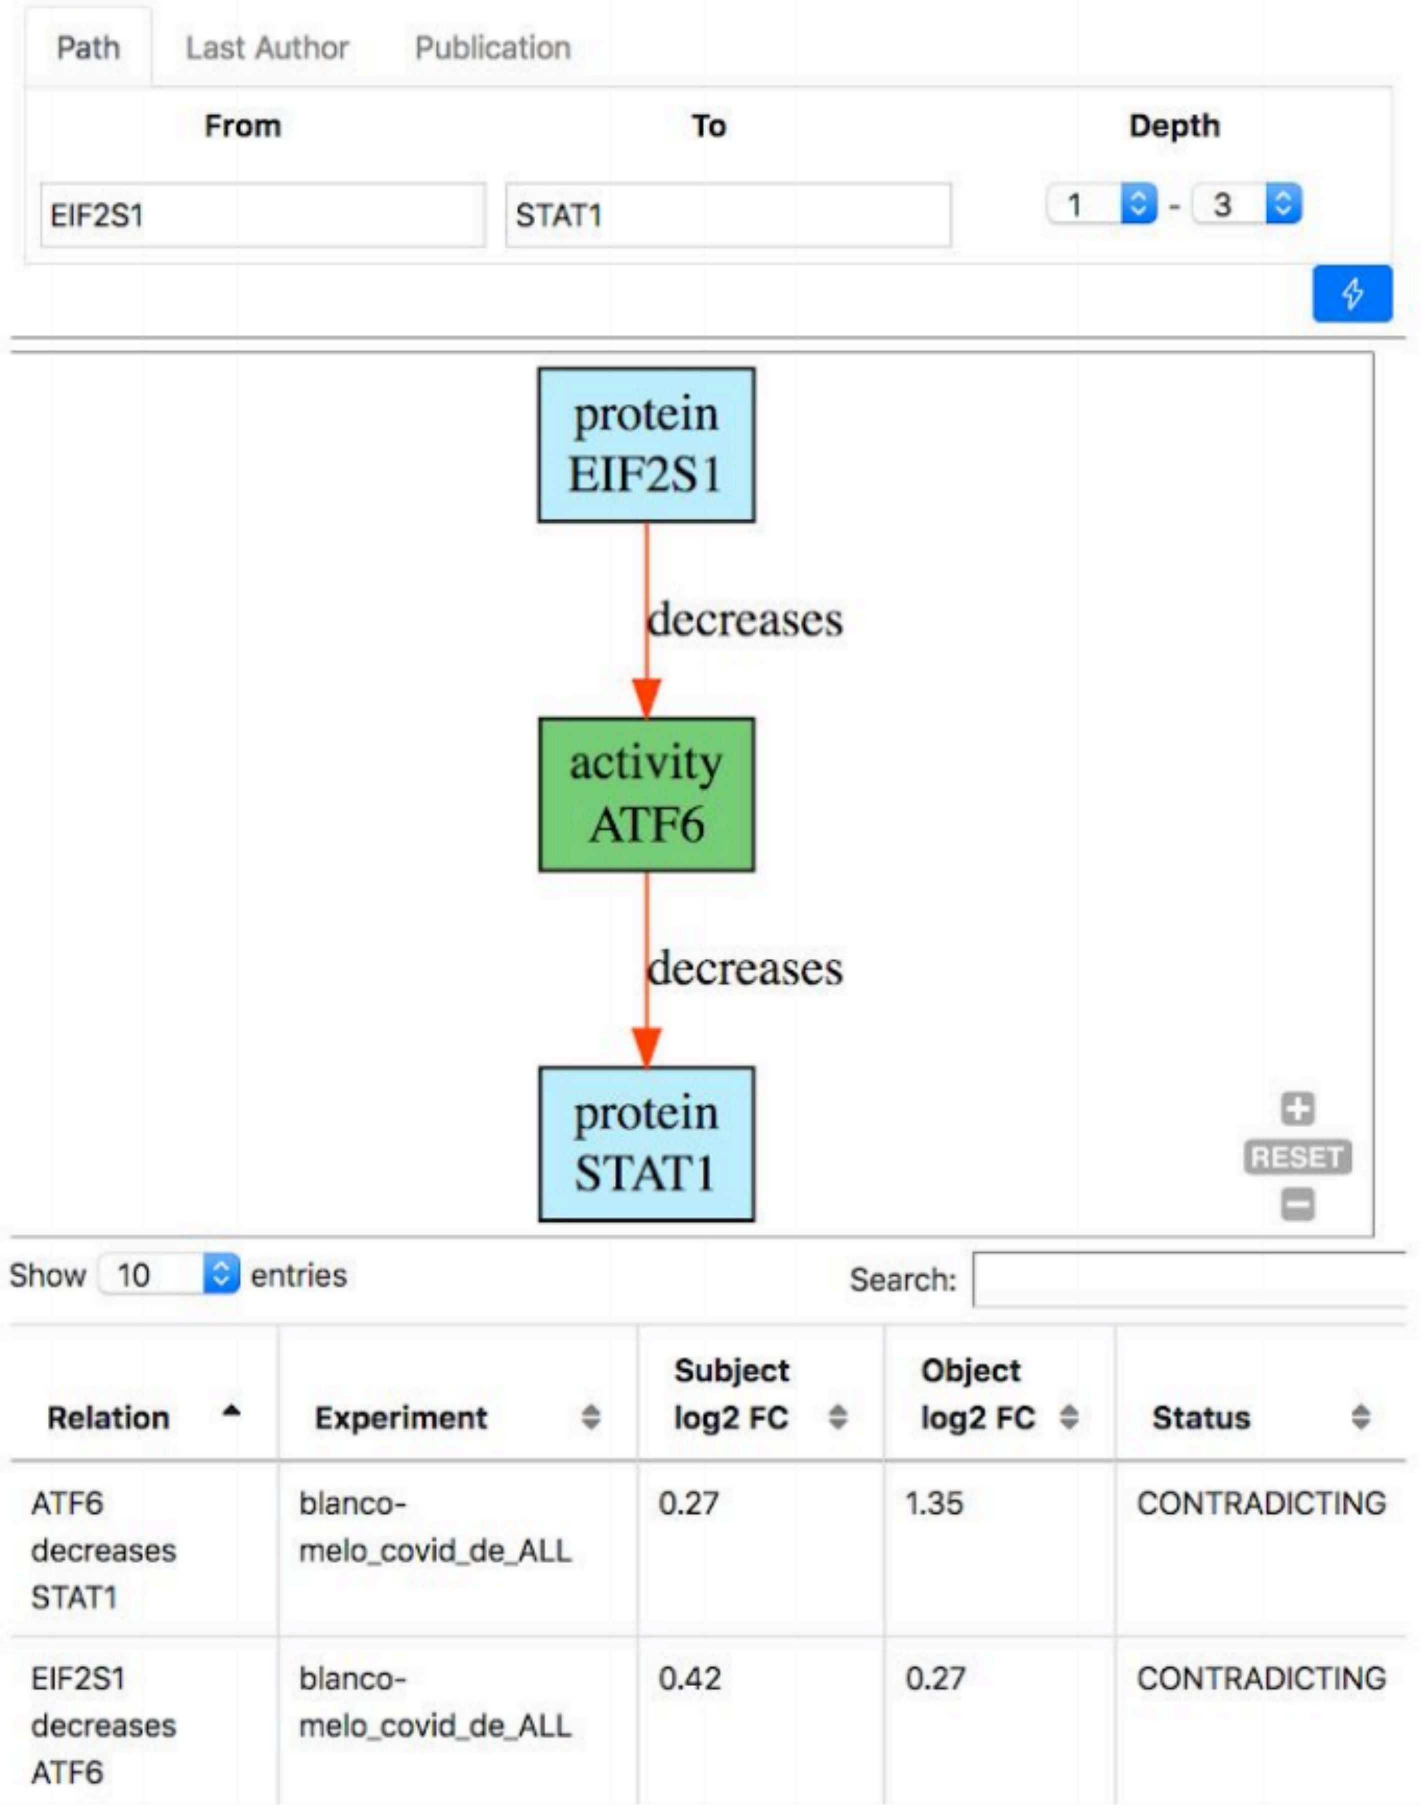

COVID-19 Knowledge Graph: a computable, multi-modal, cause-and-effect knowledge model of COVID-19 pathophysiology. (Domingo-Fernandez et. al. 2020)

Chatbot

Text

Images

Knowledge Graph

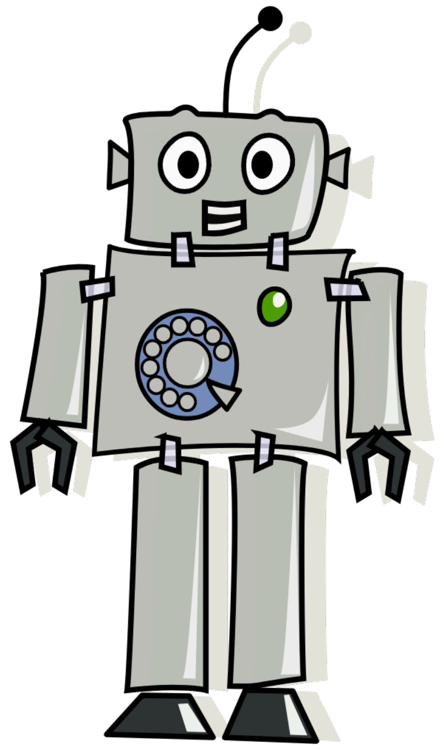

User

Which directors has Tom Hanks collaborated with?

Querying a Knowledge Graph

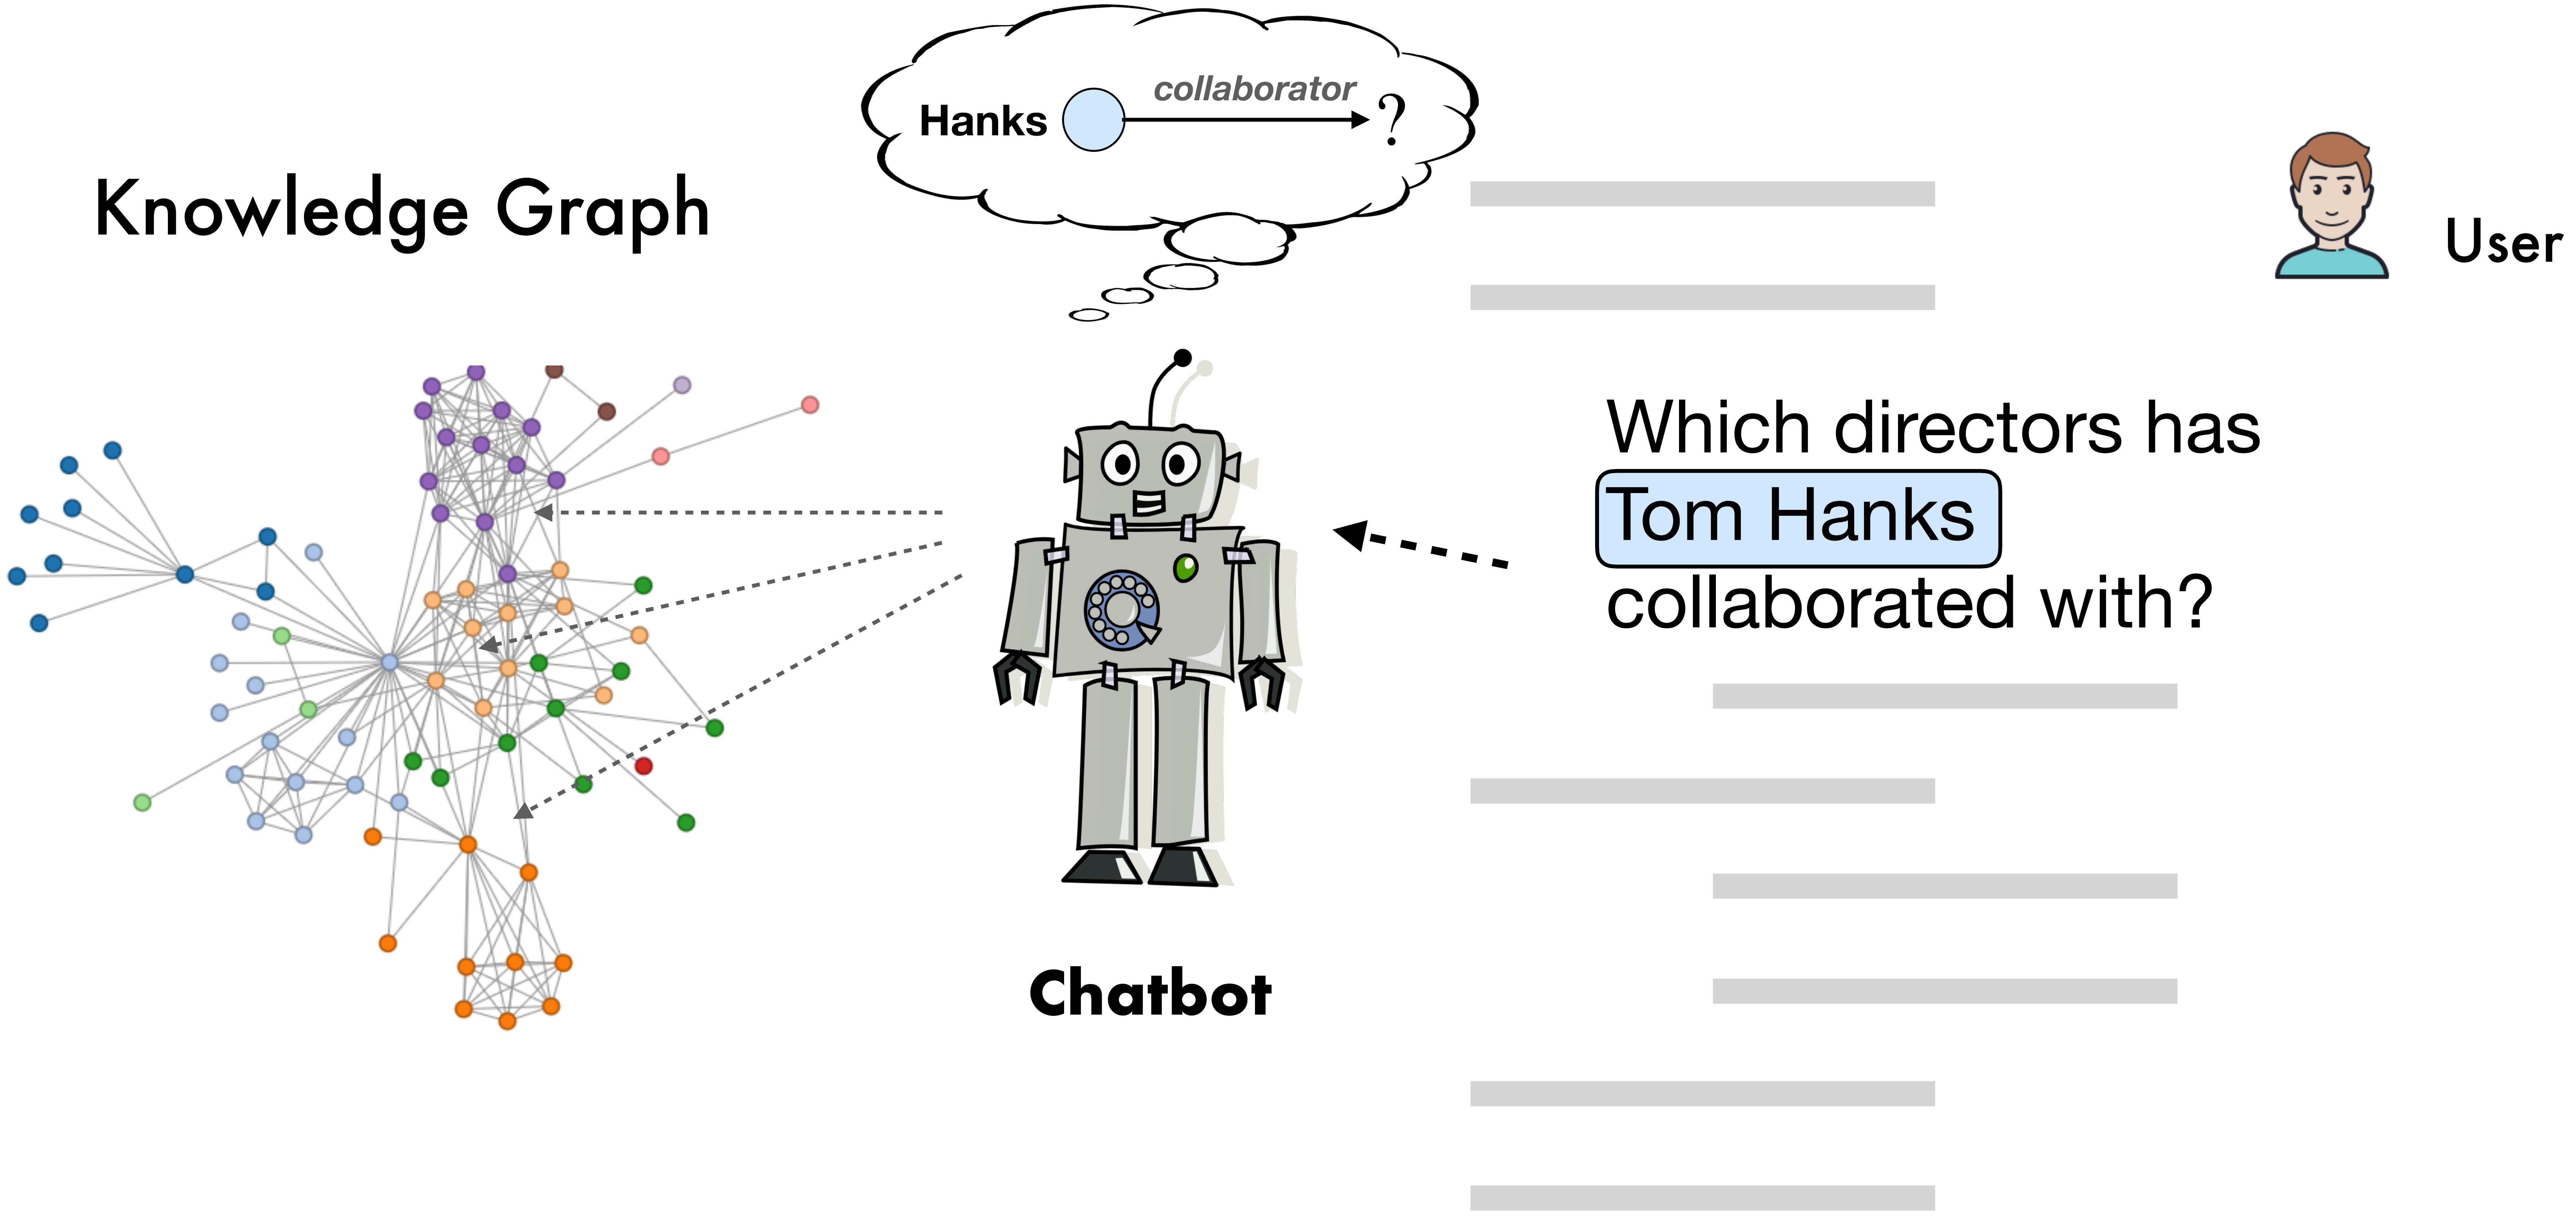

Reason over Discrete Entities

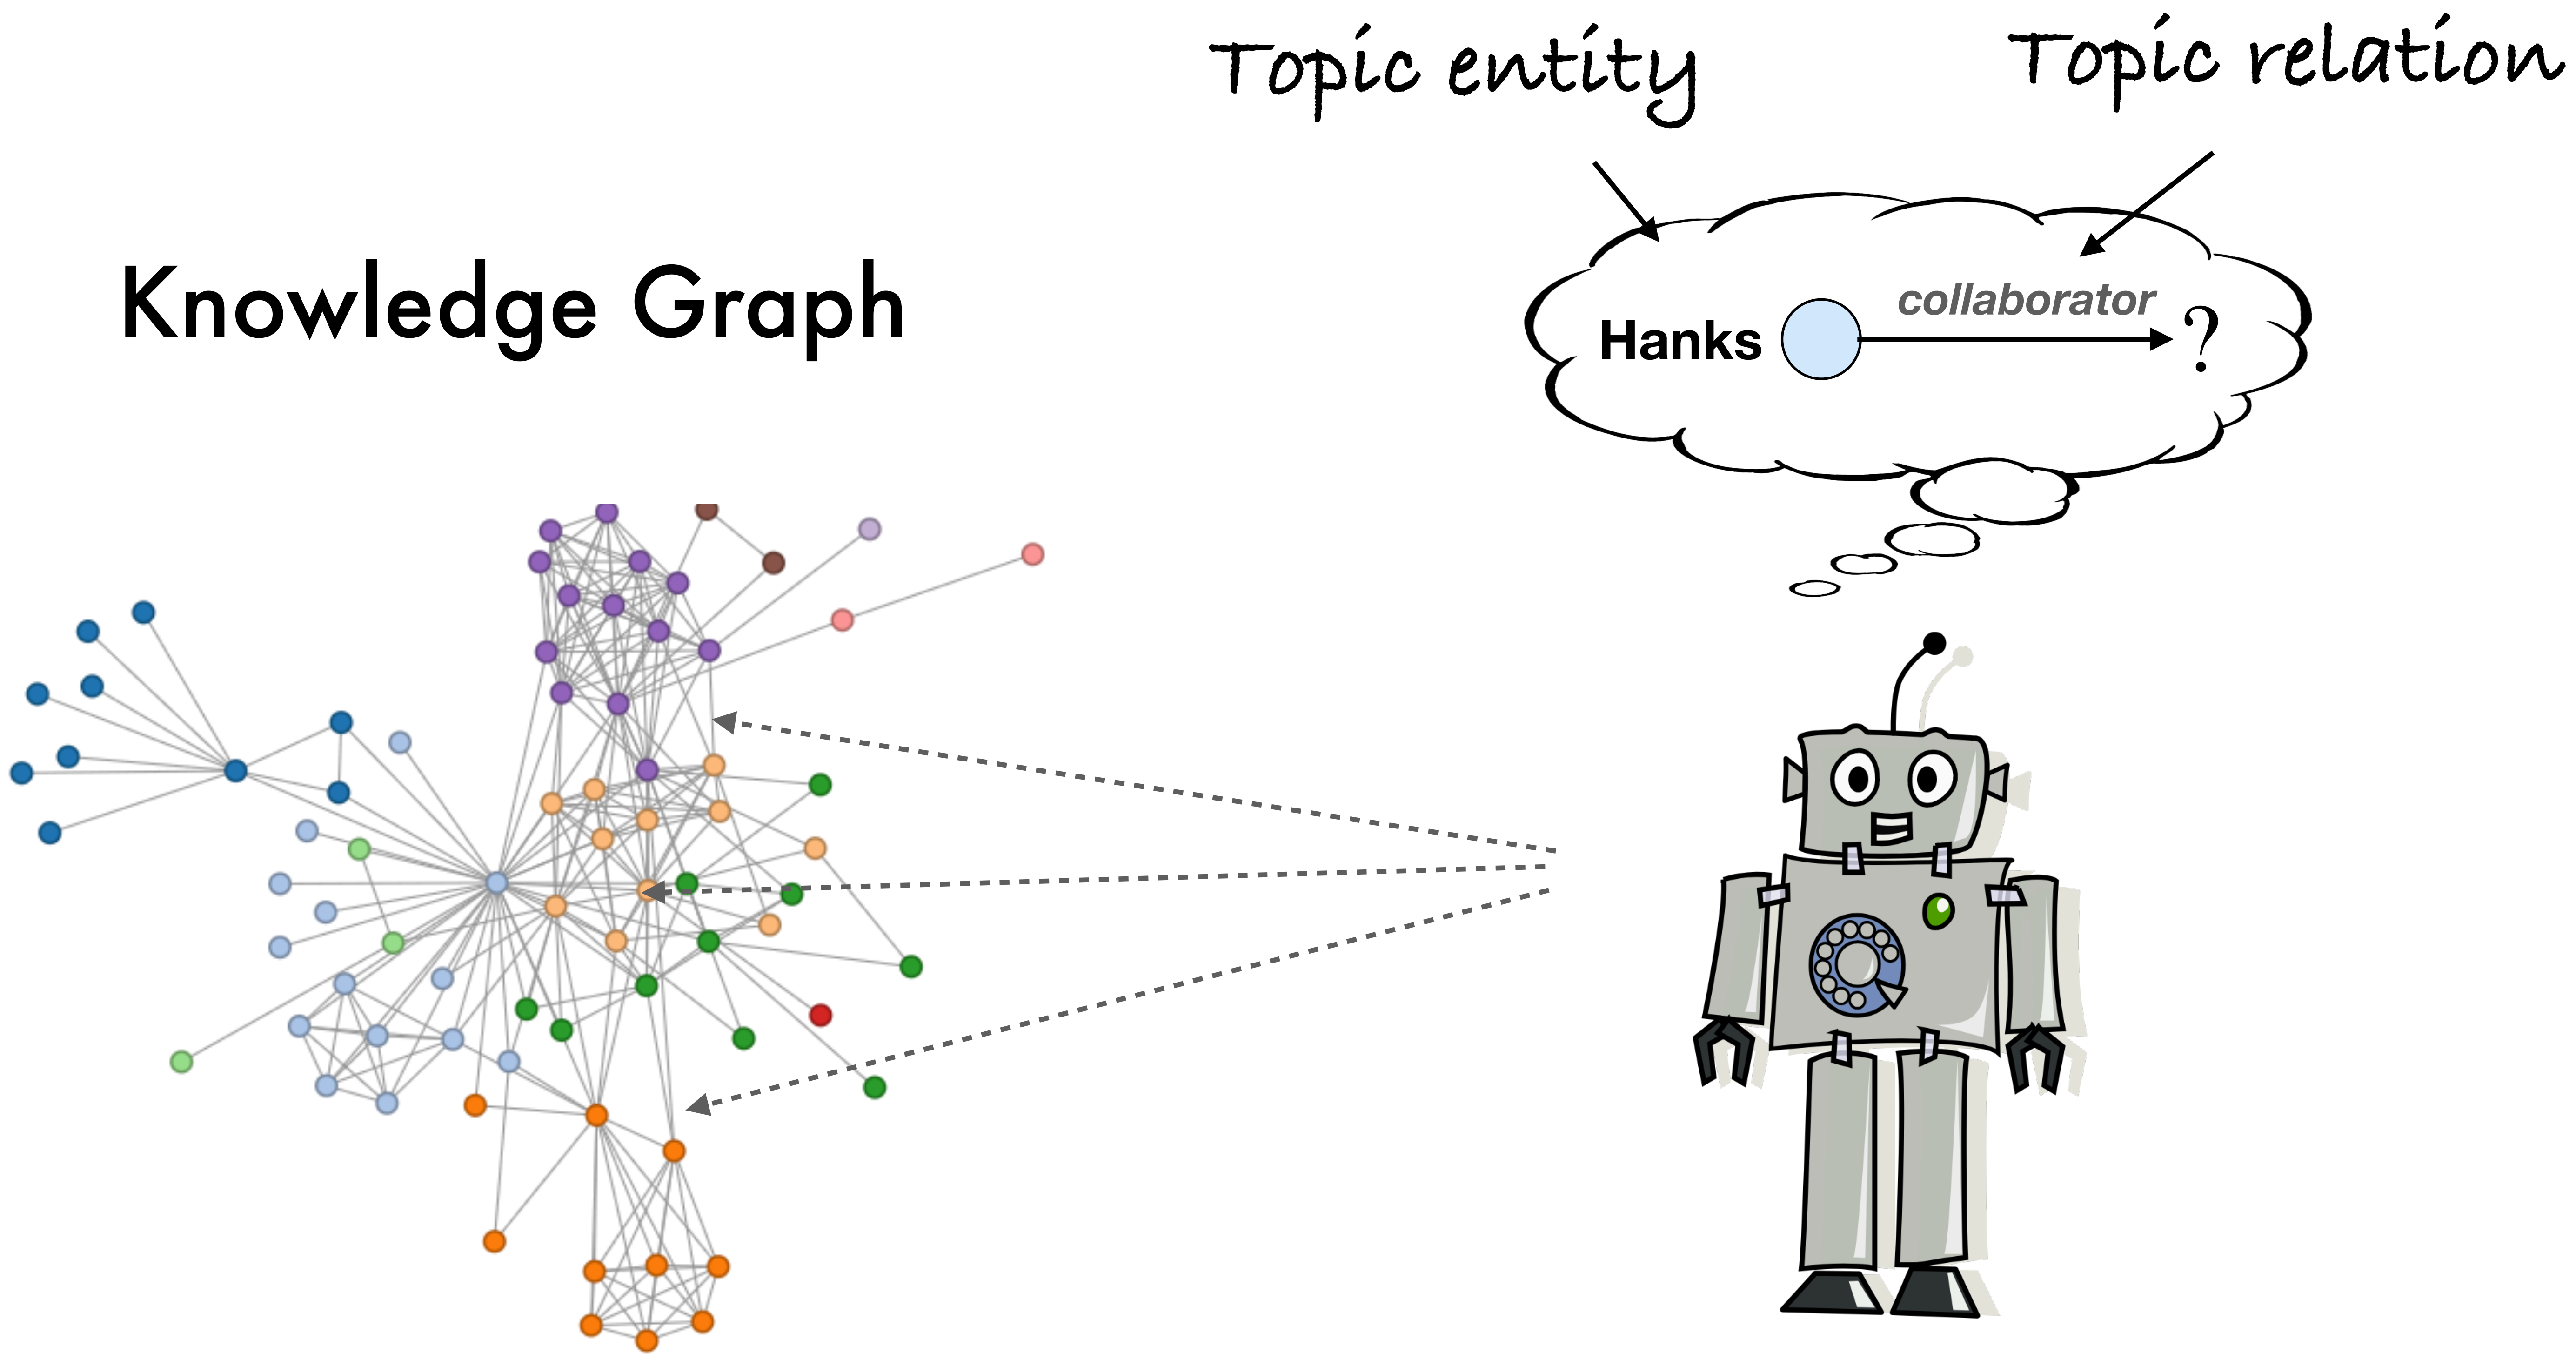

Reason over Discrete Entities

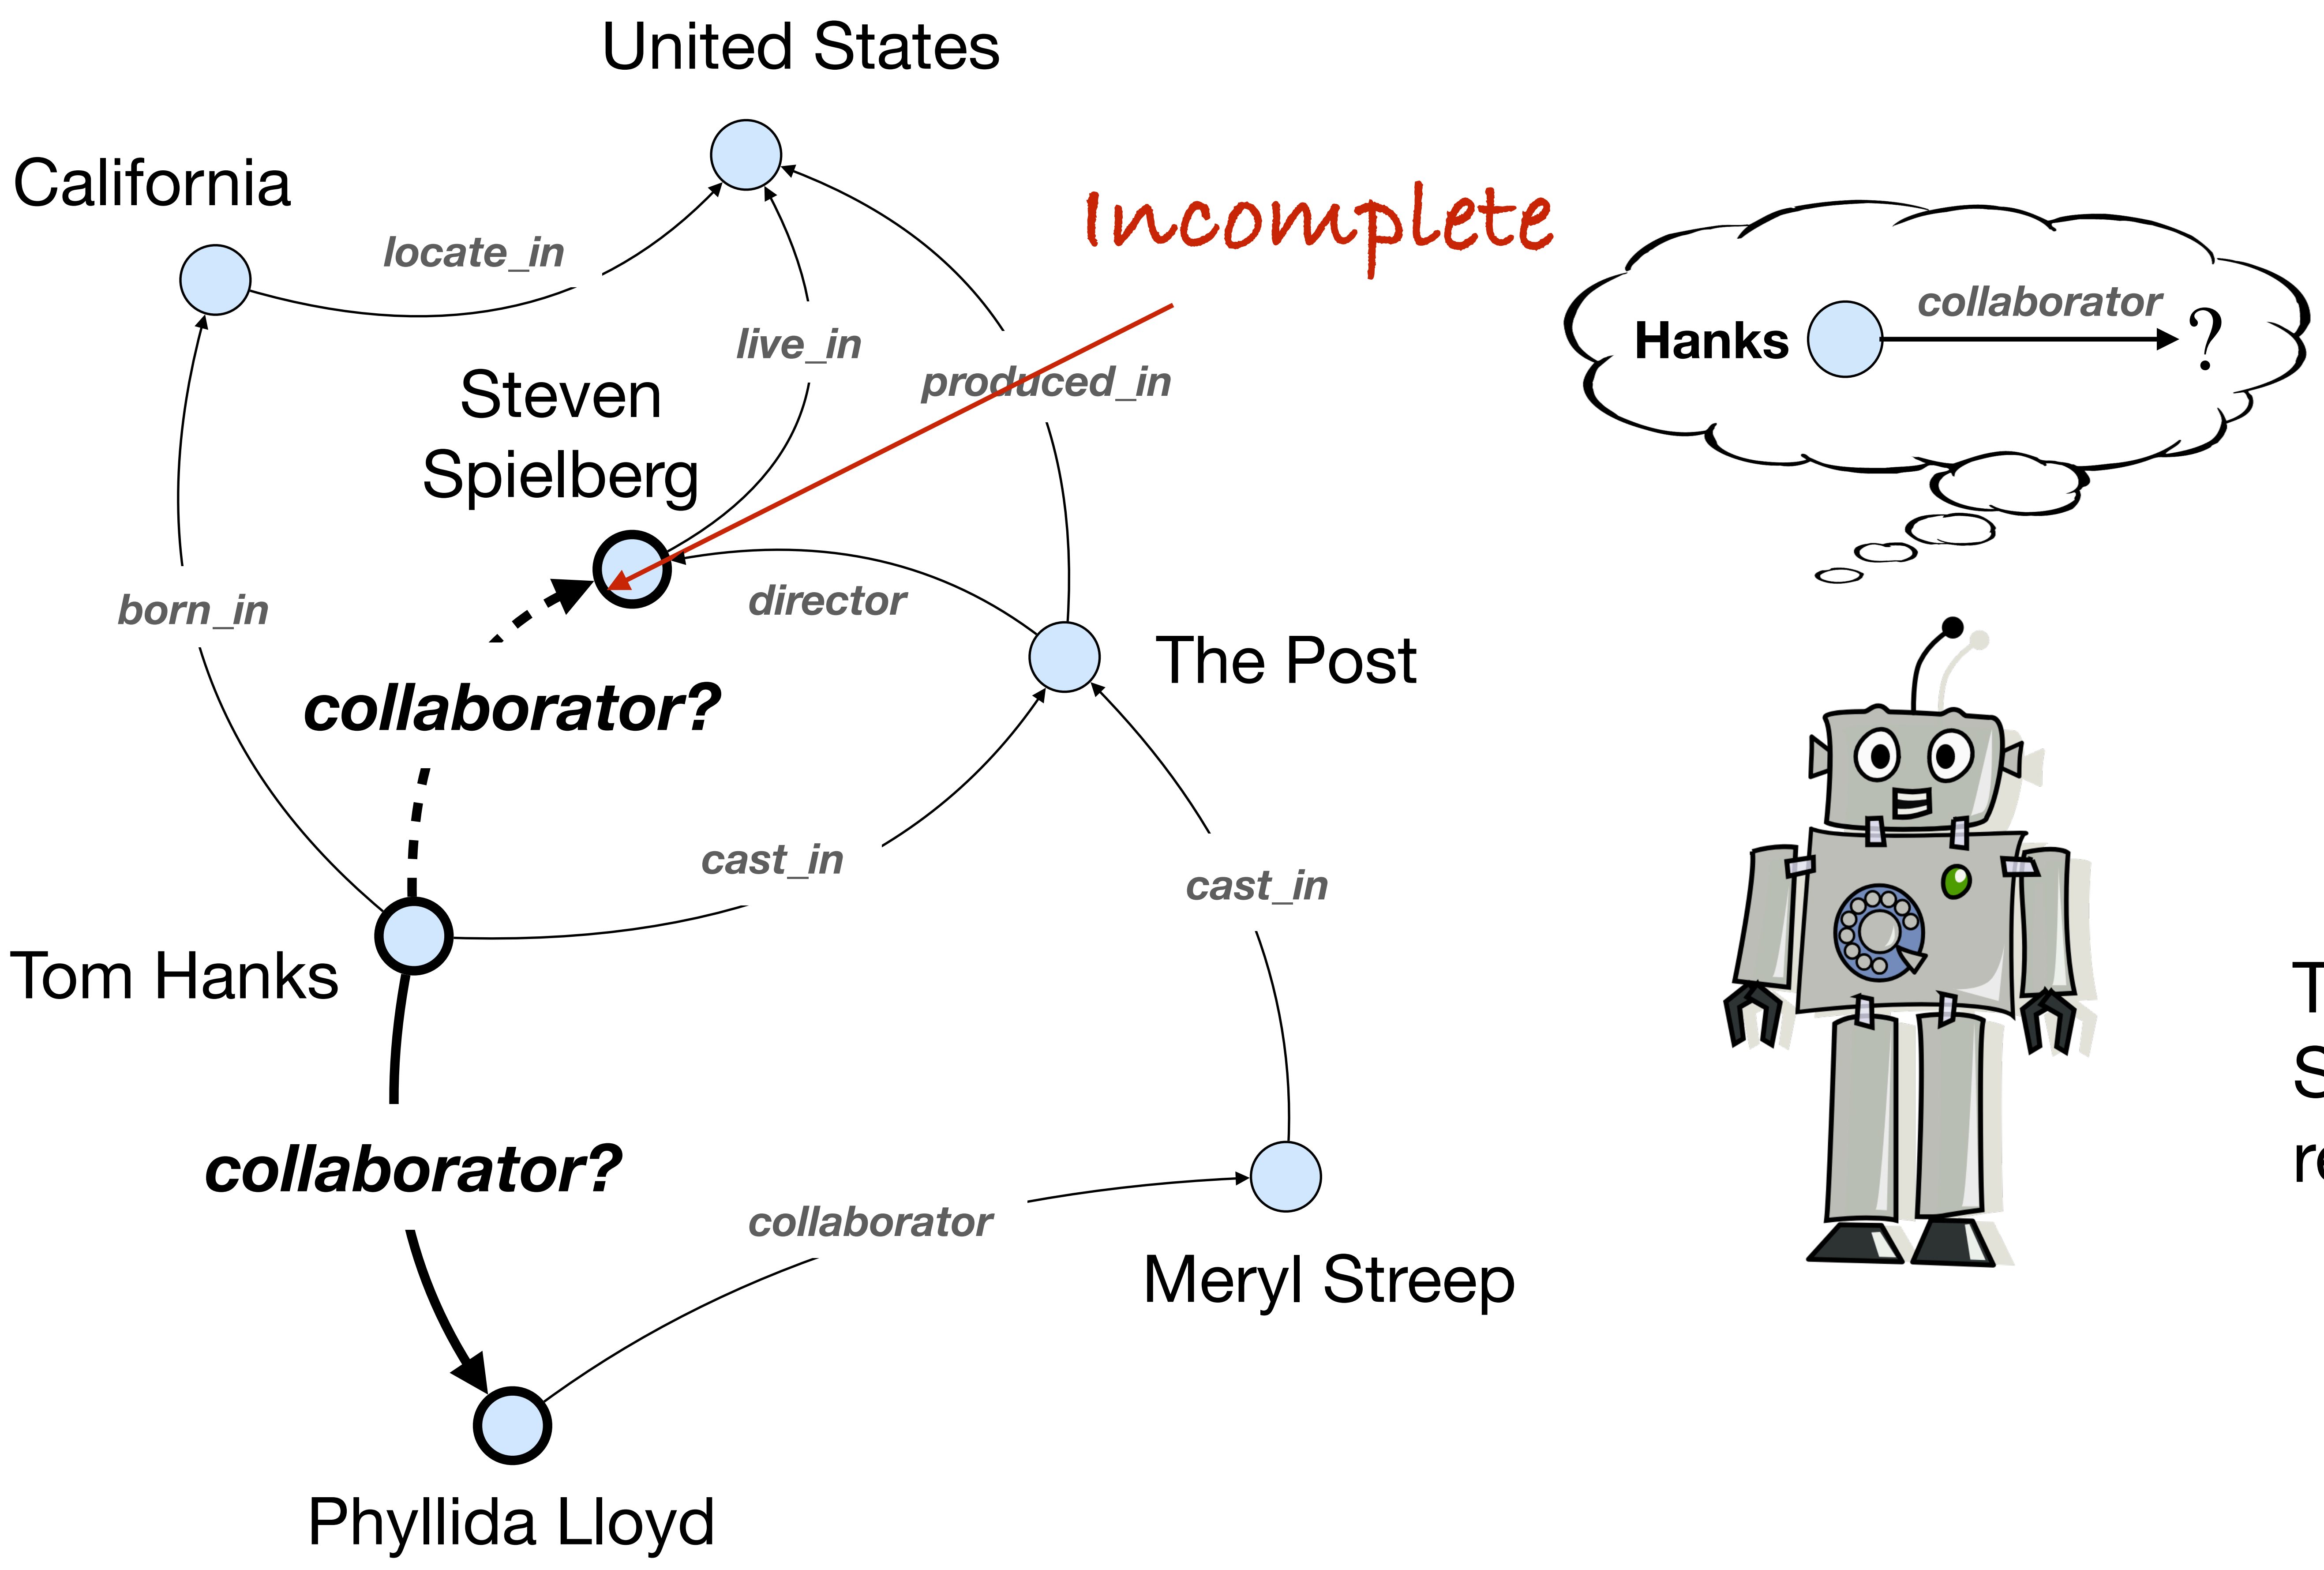

The answer “Steven Spielberg” cannot be directly retrieved from the KG

Knowledge Graph Embeddings

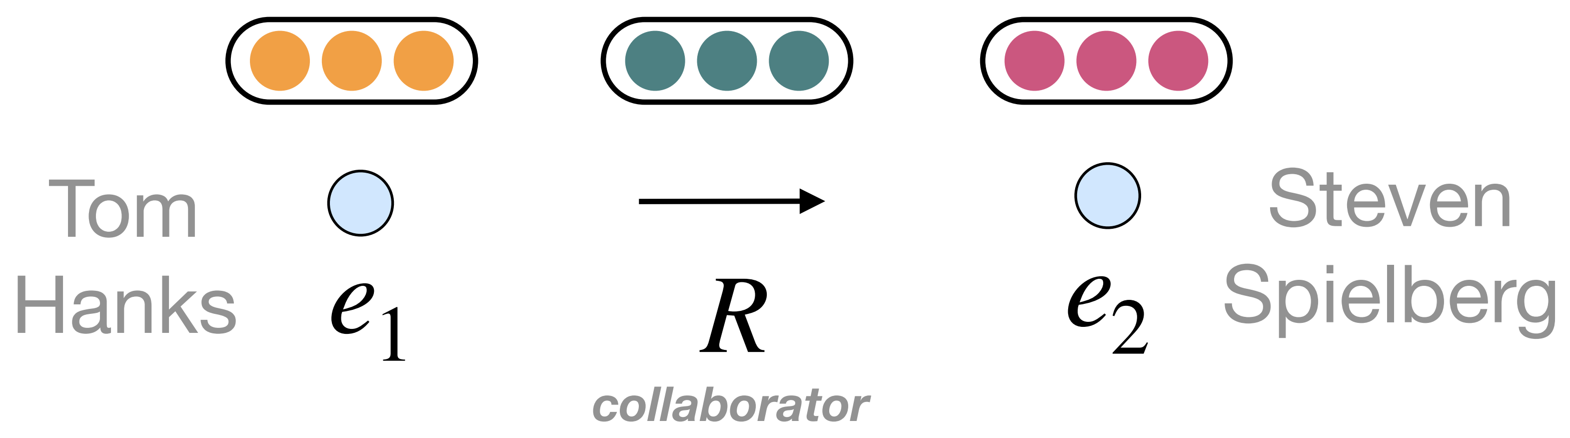

Highly accurate &
Efficient

	Acc
WordNet	86.2
Freebase	90.0

Tab 1. NTN KB fact inference performance on the WordNet and Freebase benchmarks (Socher et. al. 2013)

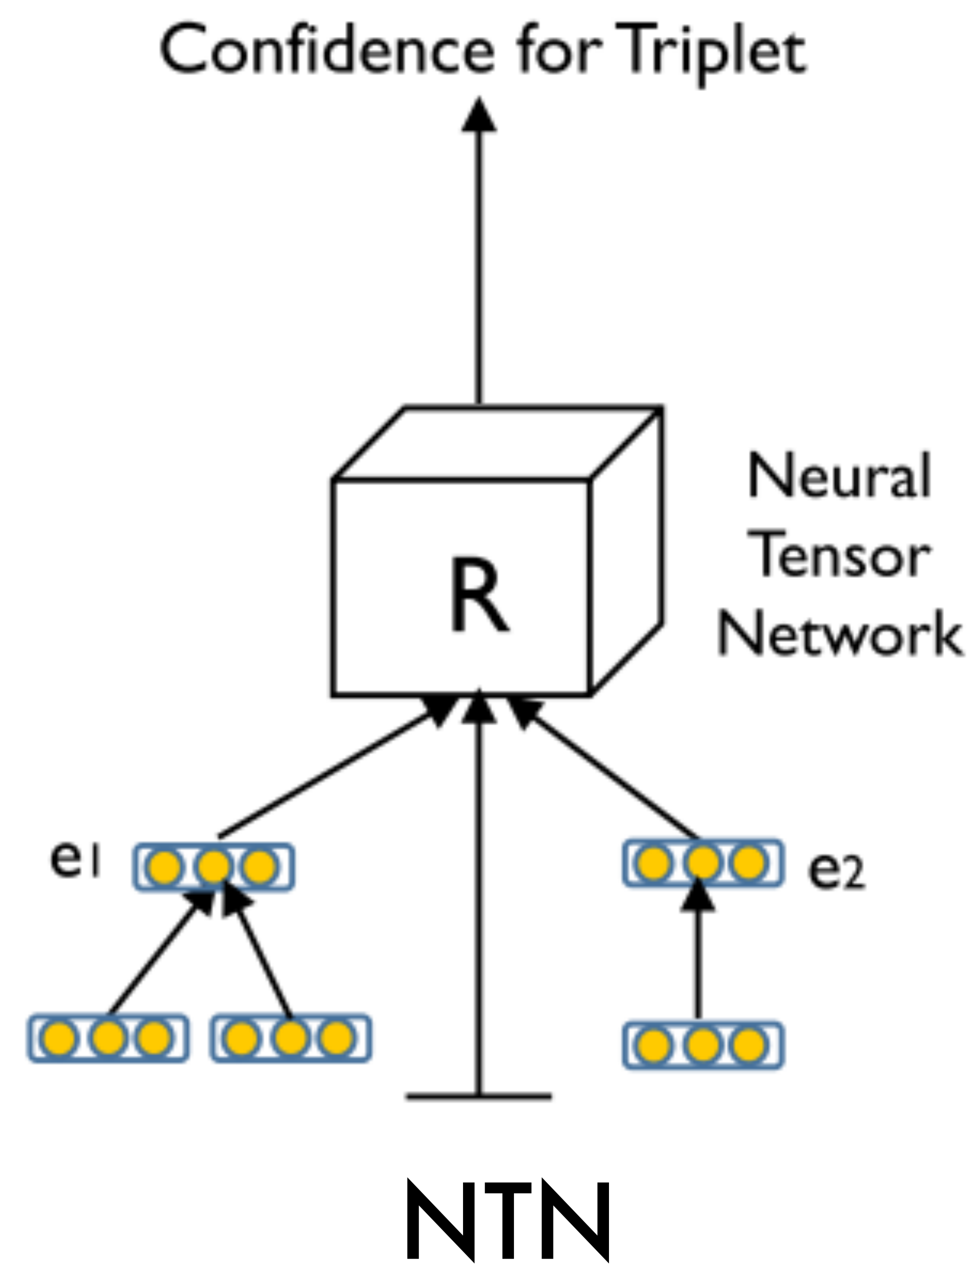

NTN (Socher et. al. 2013), DistMult (Yang et al. 2015), ComplEx (Trouillon et al. 2016), ConvE (Dettmers et al. 2018).

Neural Tensor Network

Contribution 1: Neural Tensor Networks (NTNs) for relation classification

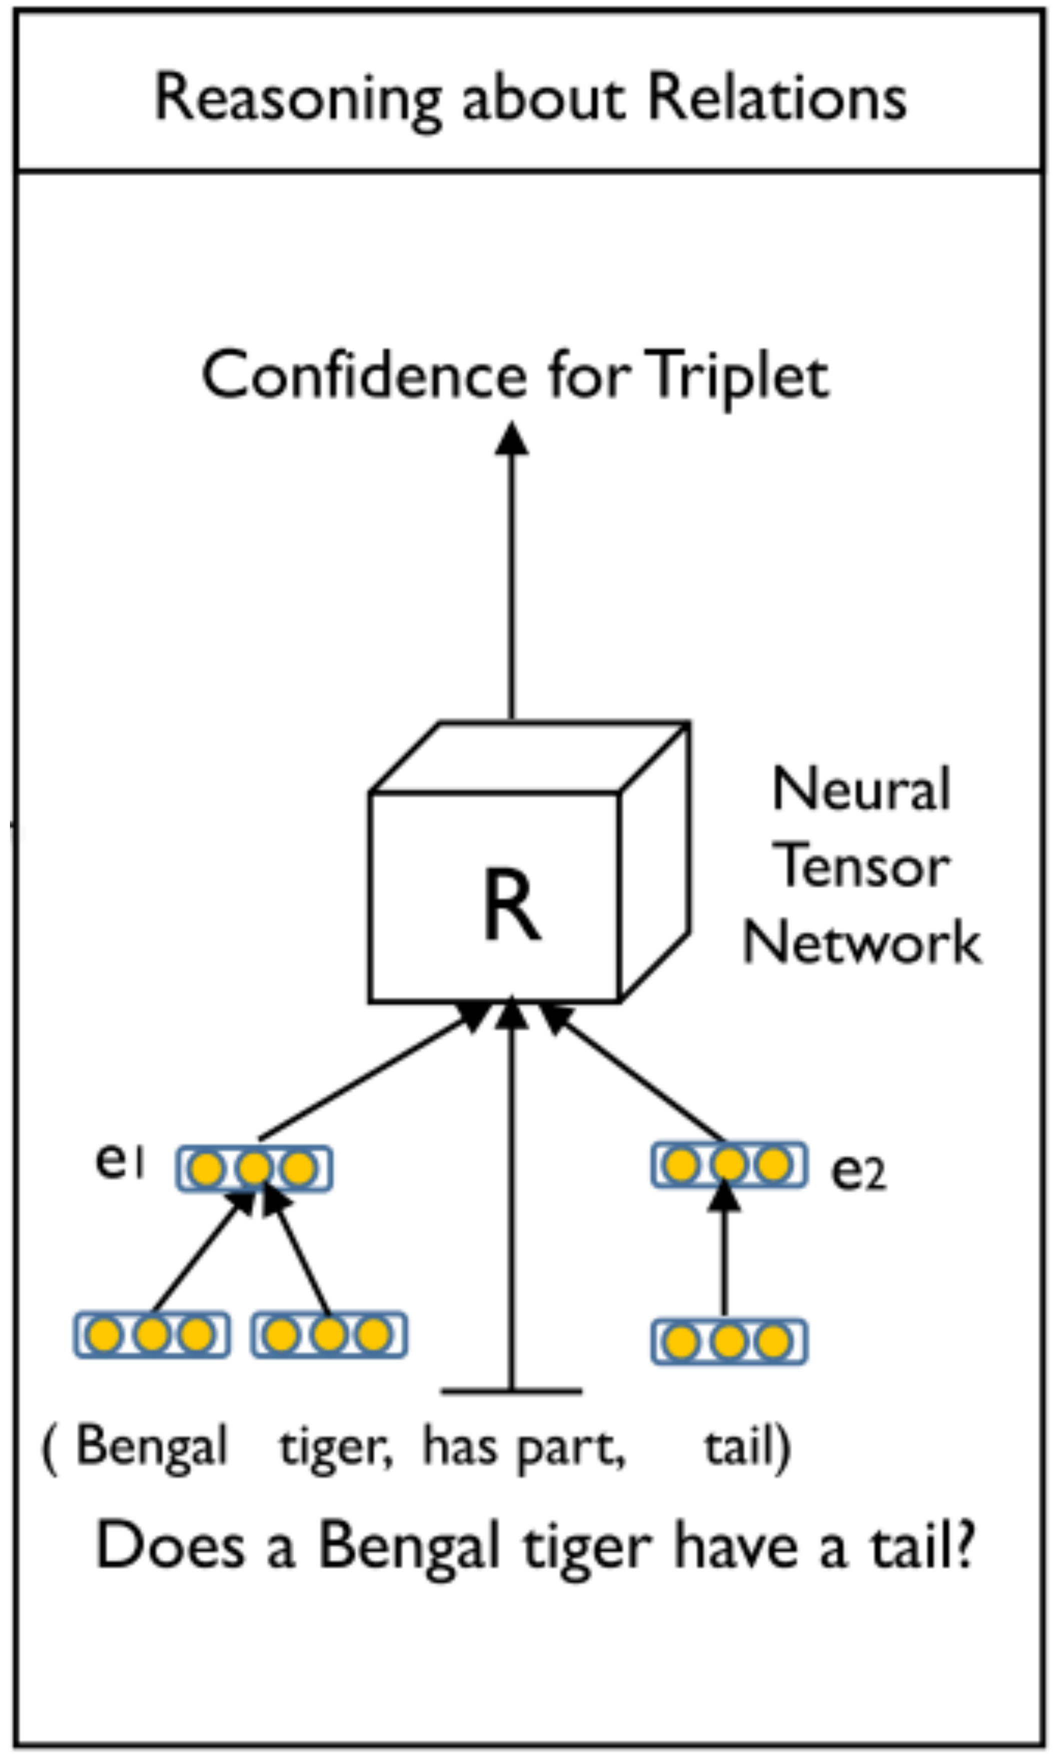

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Neural Tensor Network

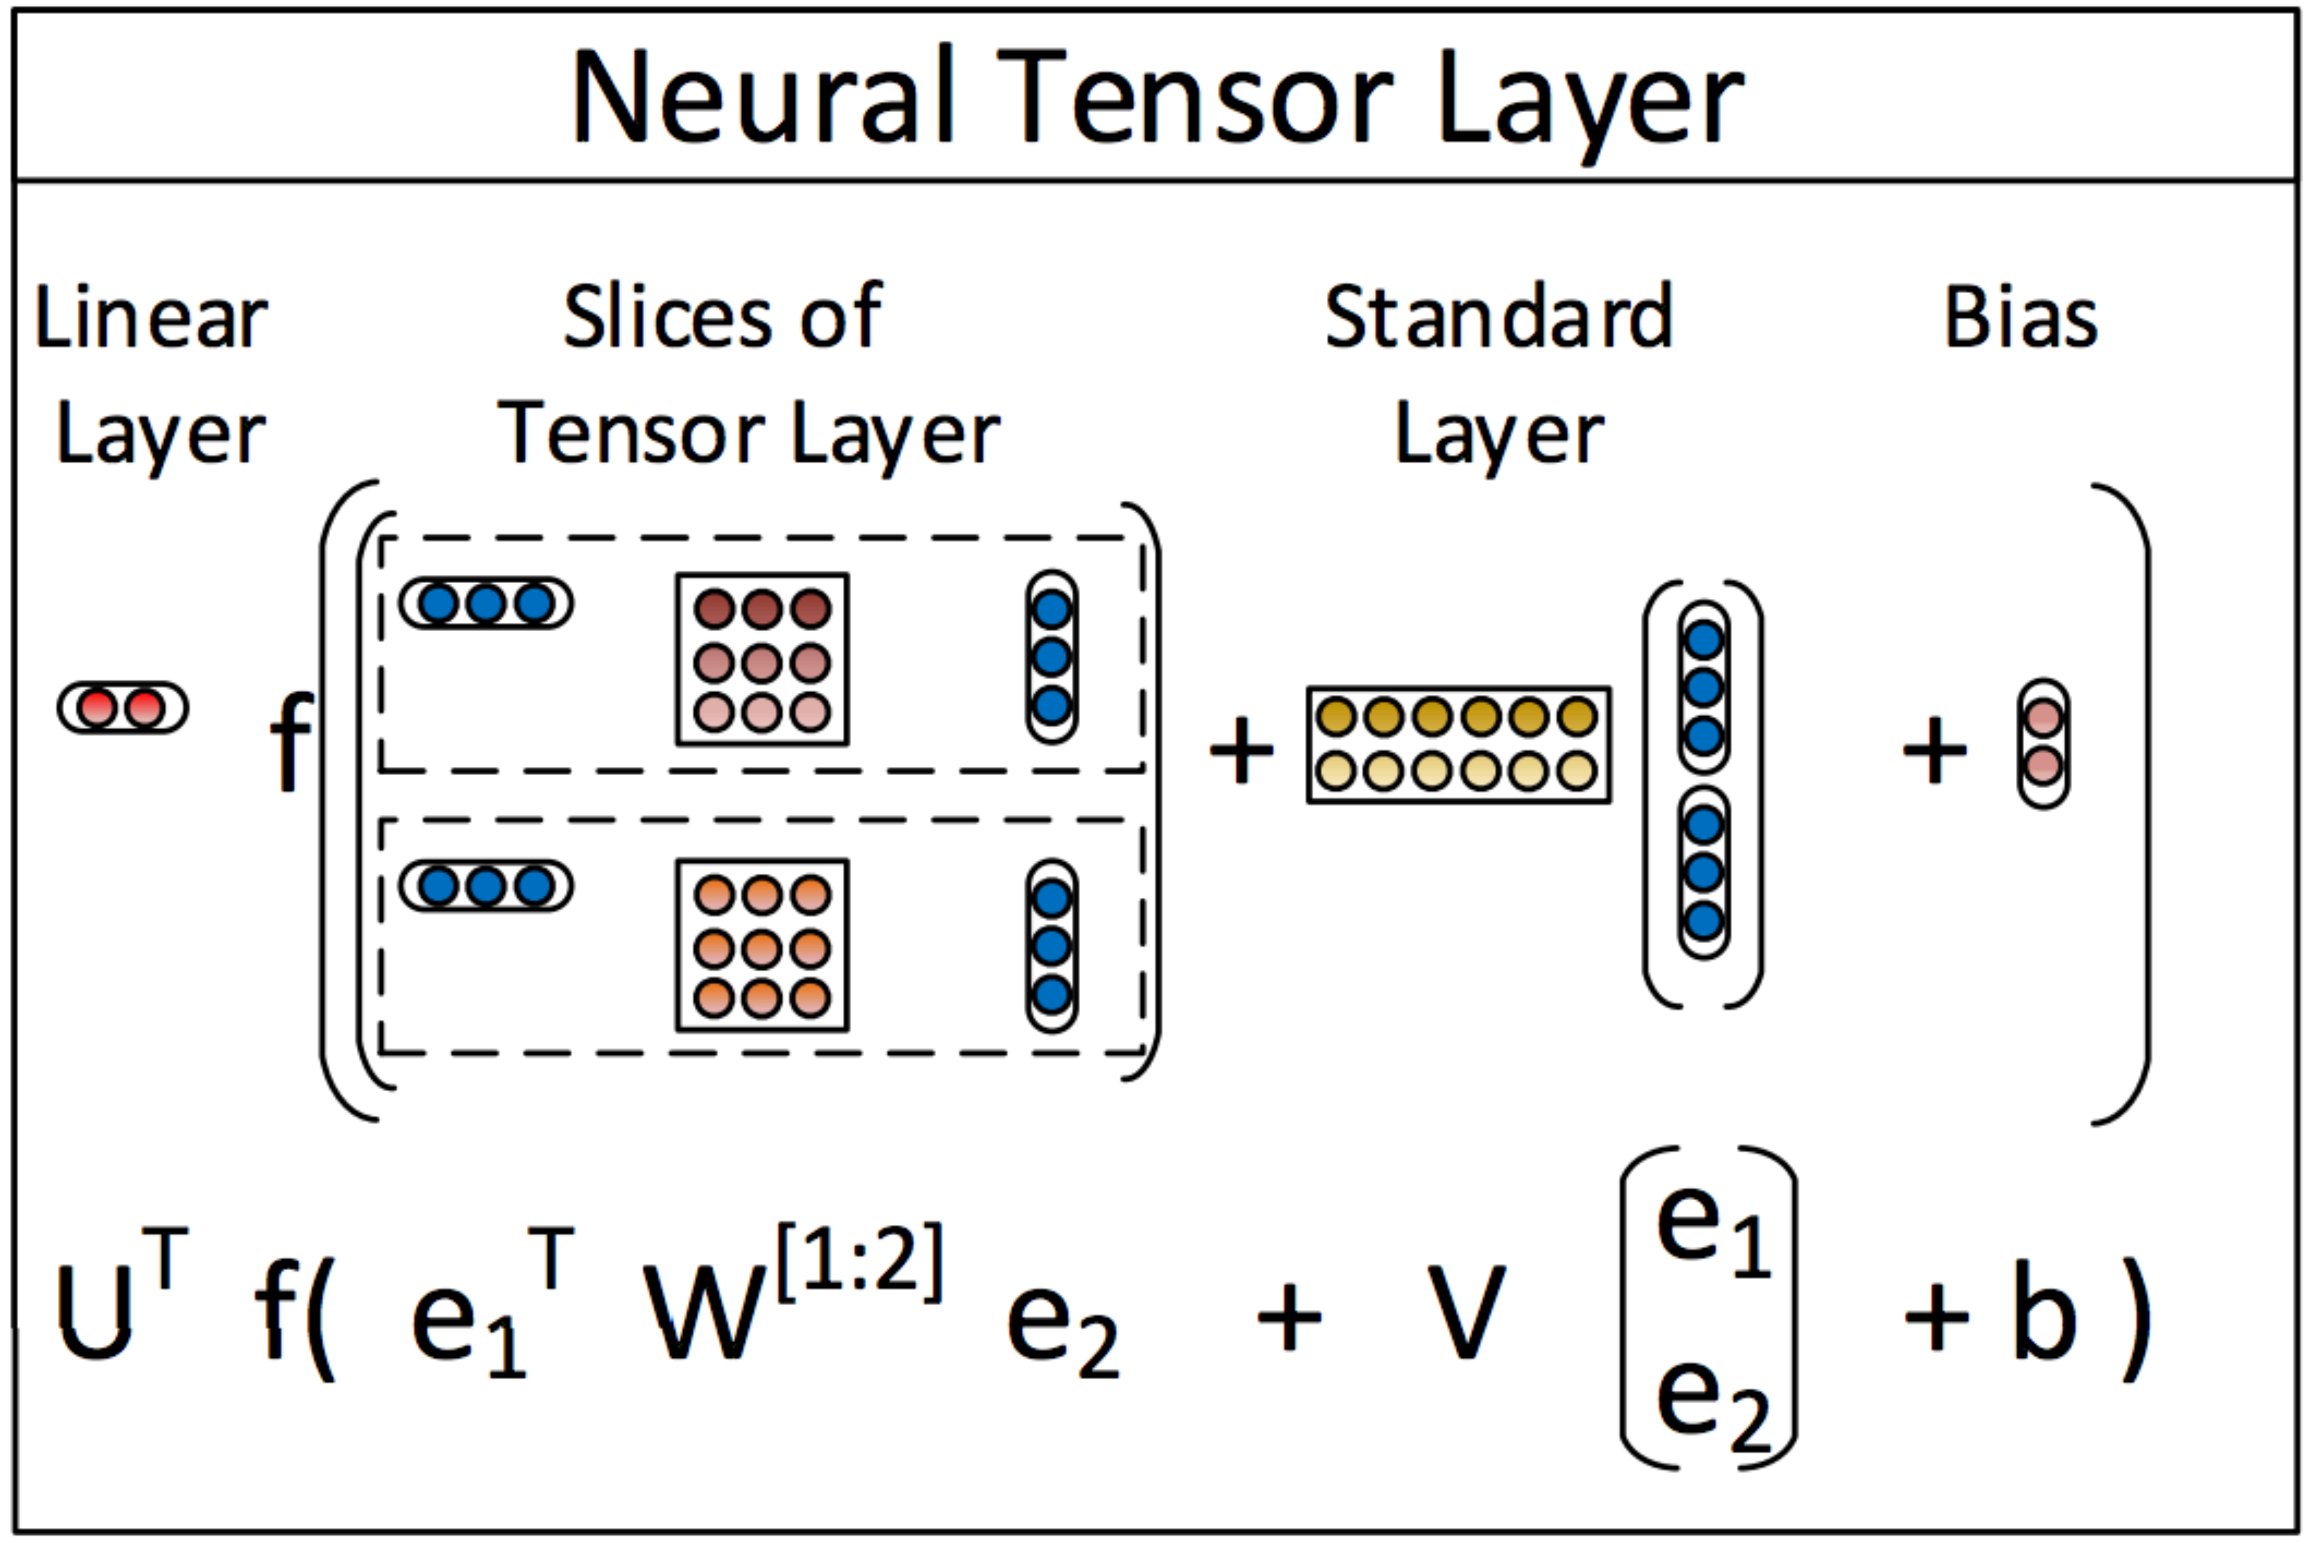

$$g(e_1, R, e_2) = u_R^T f\left(e_1^T W_R^{[1:k]} e_2 + V_R \begin{bmatrix} e_1 \\ e_2 \end{bmatrix} + b_R\right)$$

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Neural Tensor Network

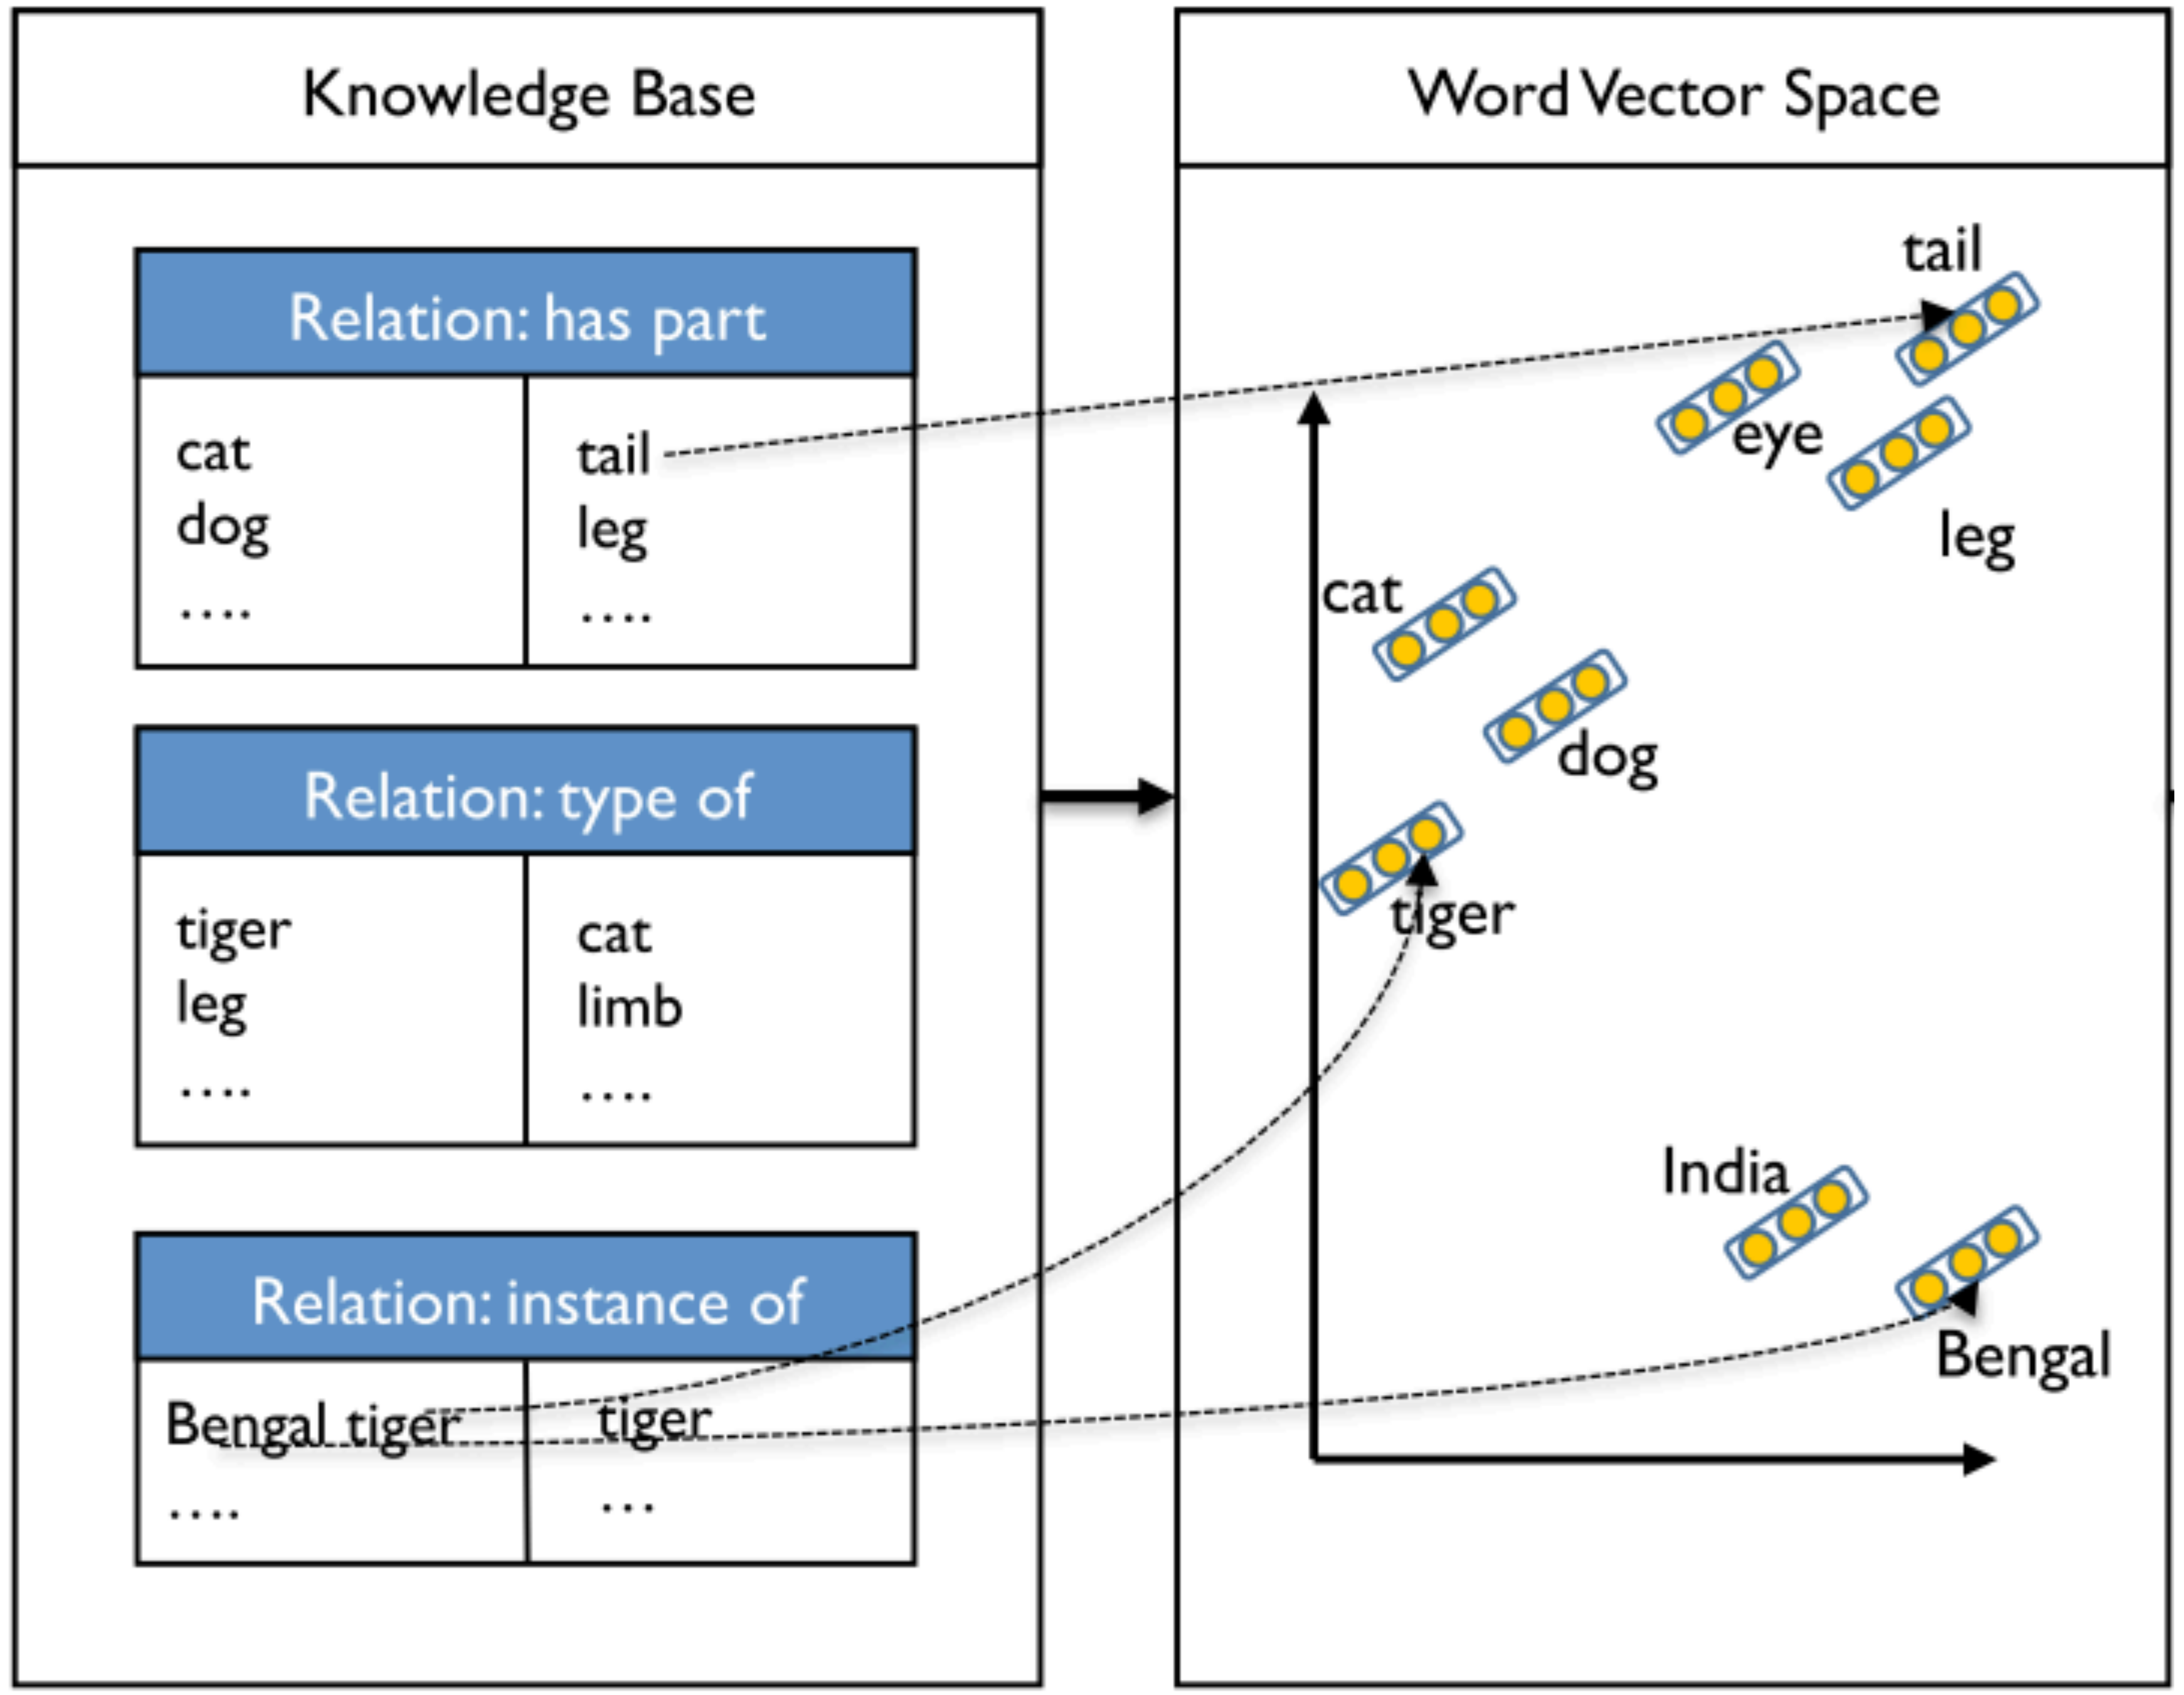

Contribution 2: Compositional Representation of Knowledge Base Entities

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Neural Tensor Network

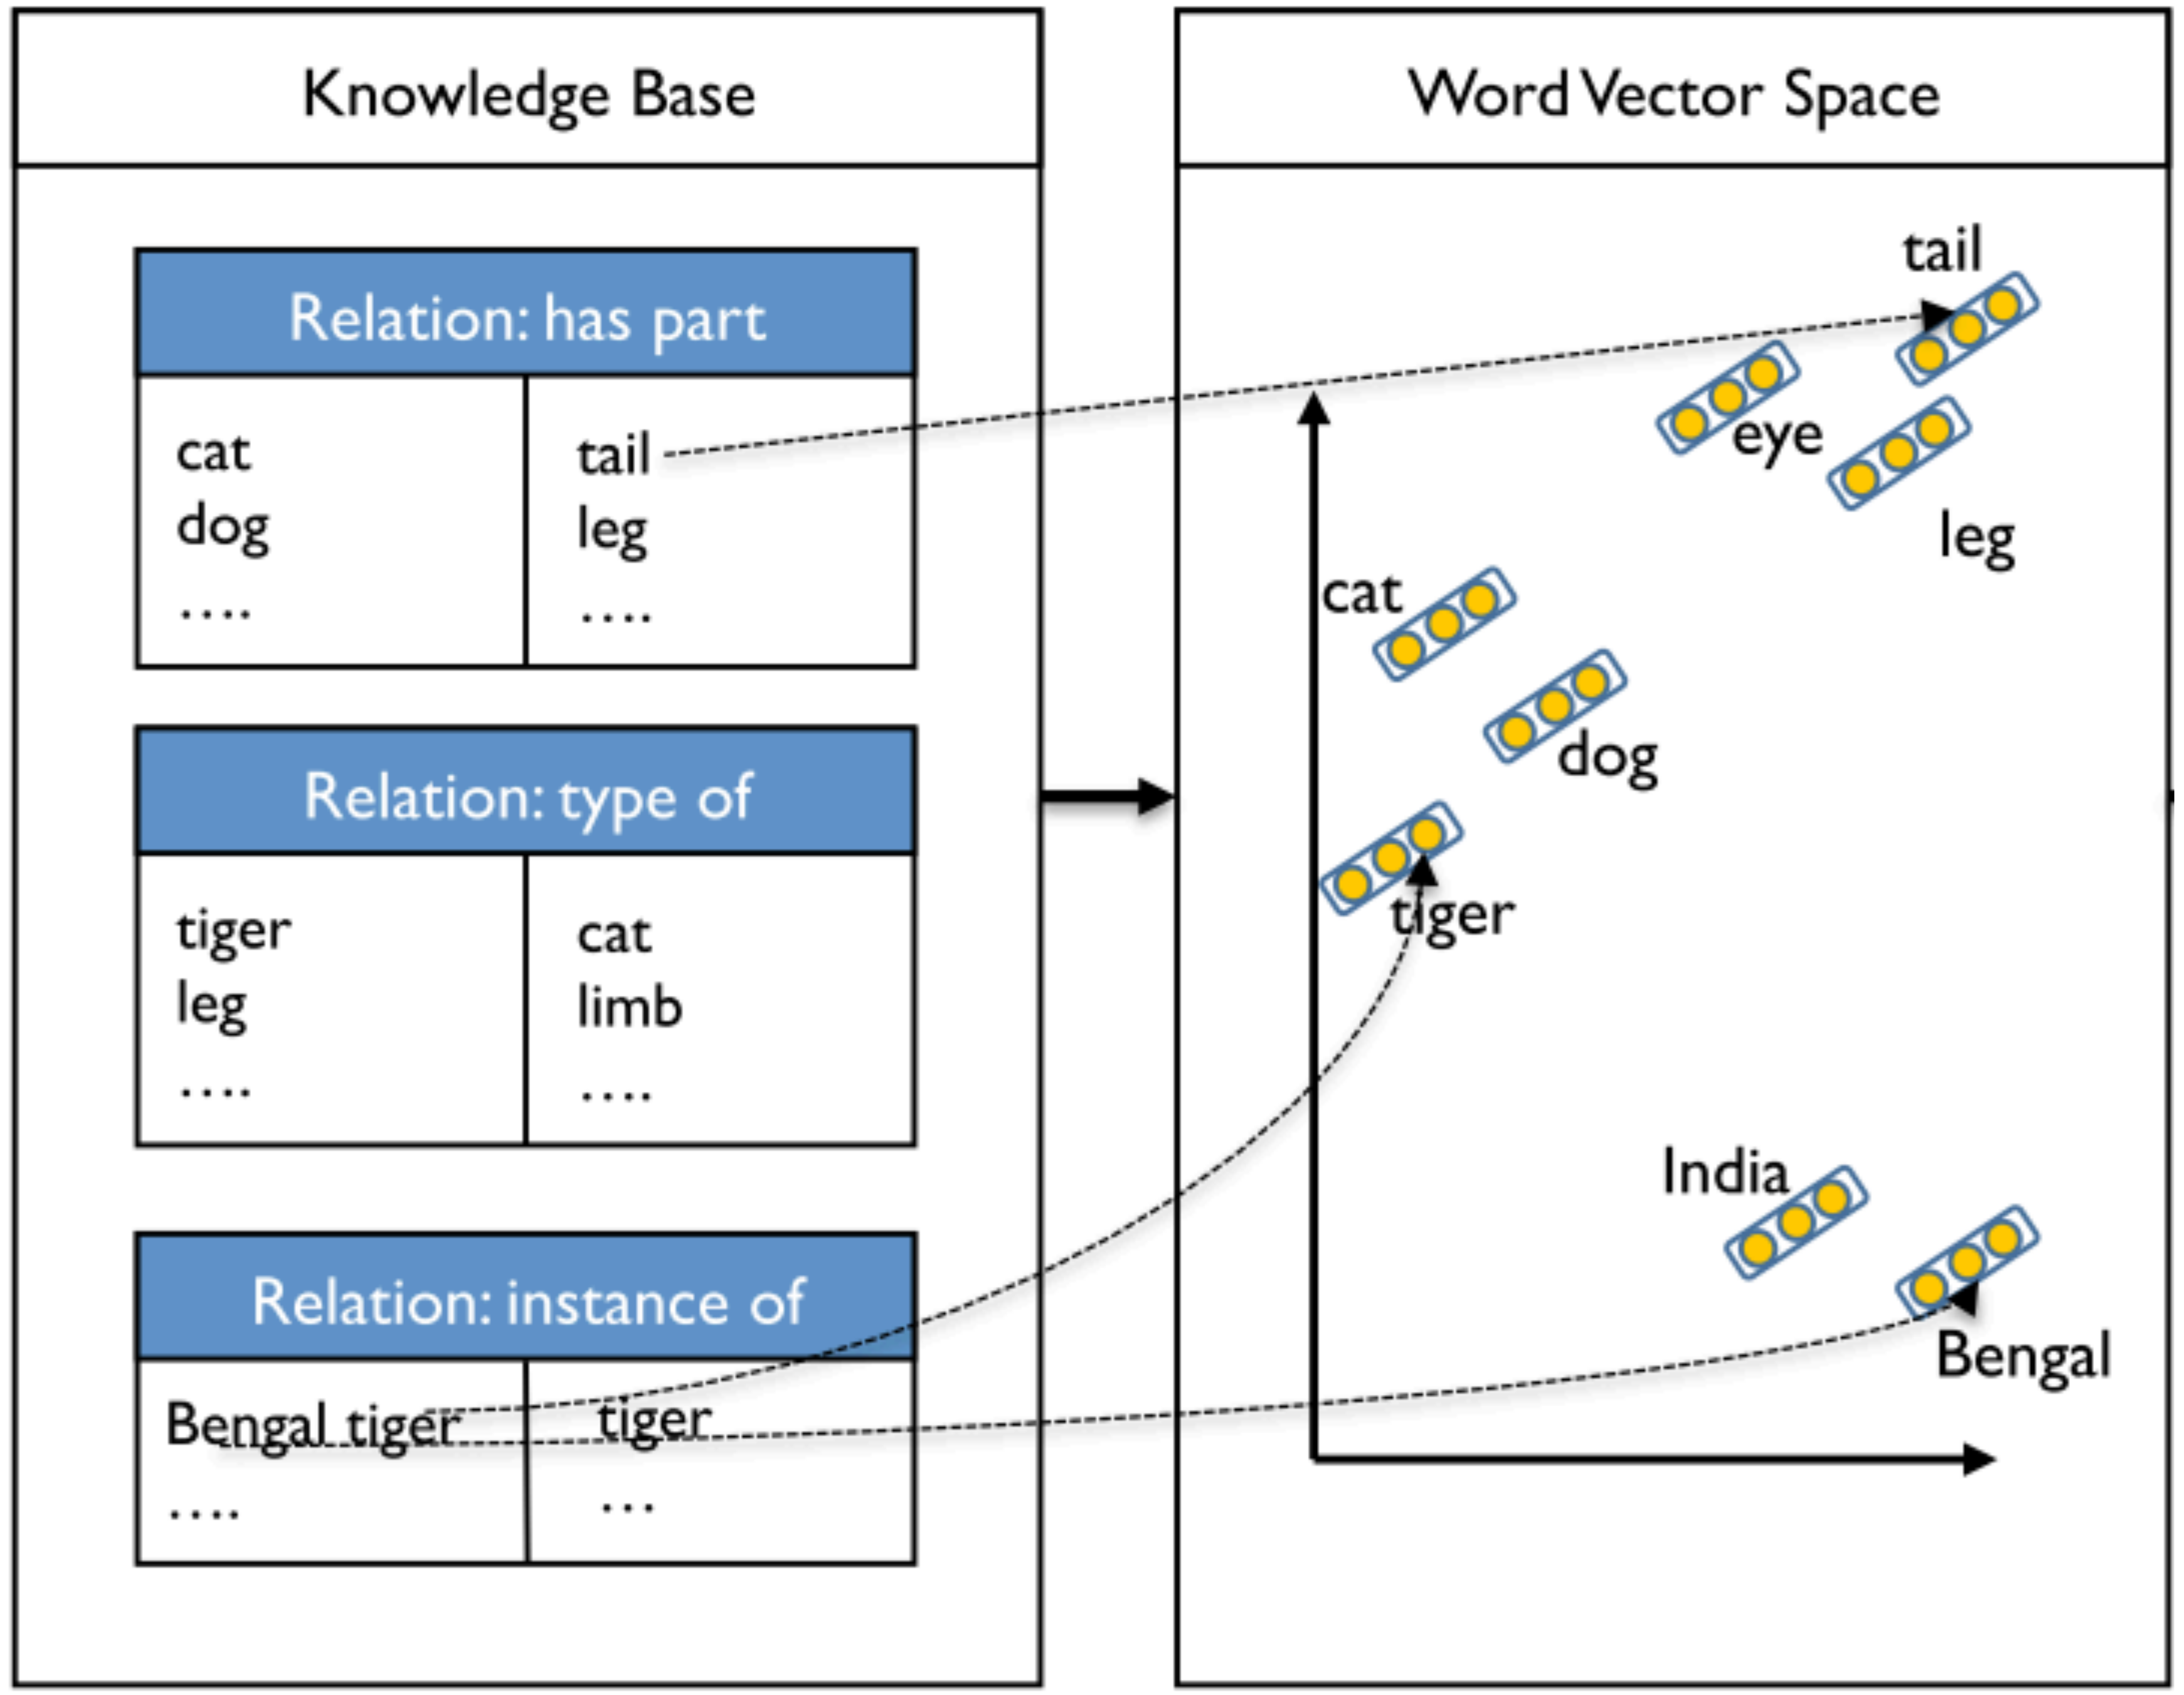

$$v_{\text{Bengal tiger}} = 0.5(v_{\text{Bengal}} + v_{\text{tiger}})$$

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Neural Tensor Network

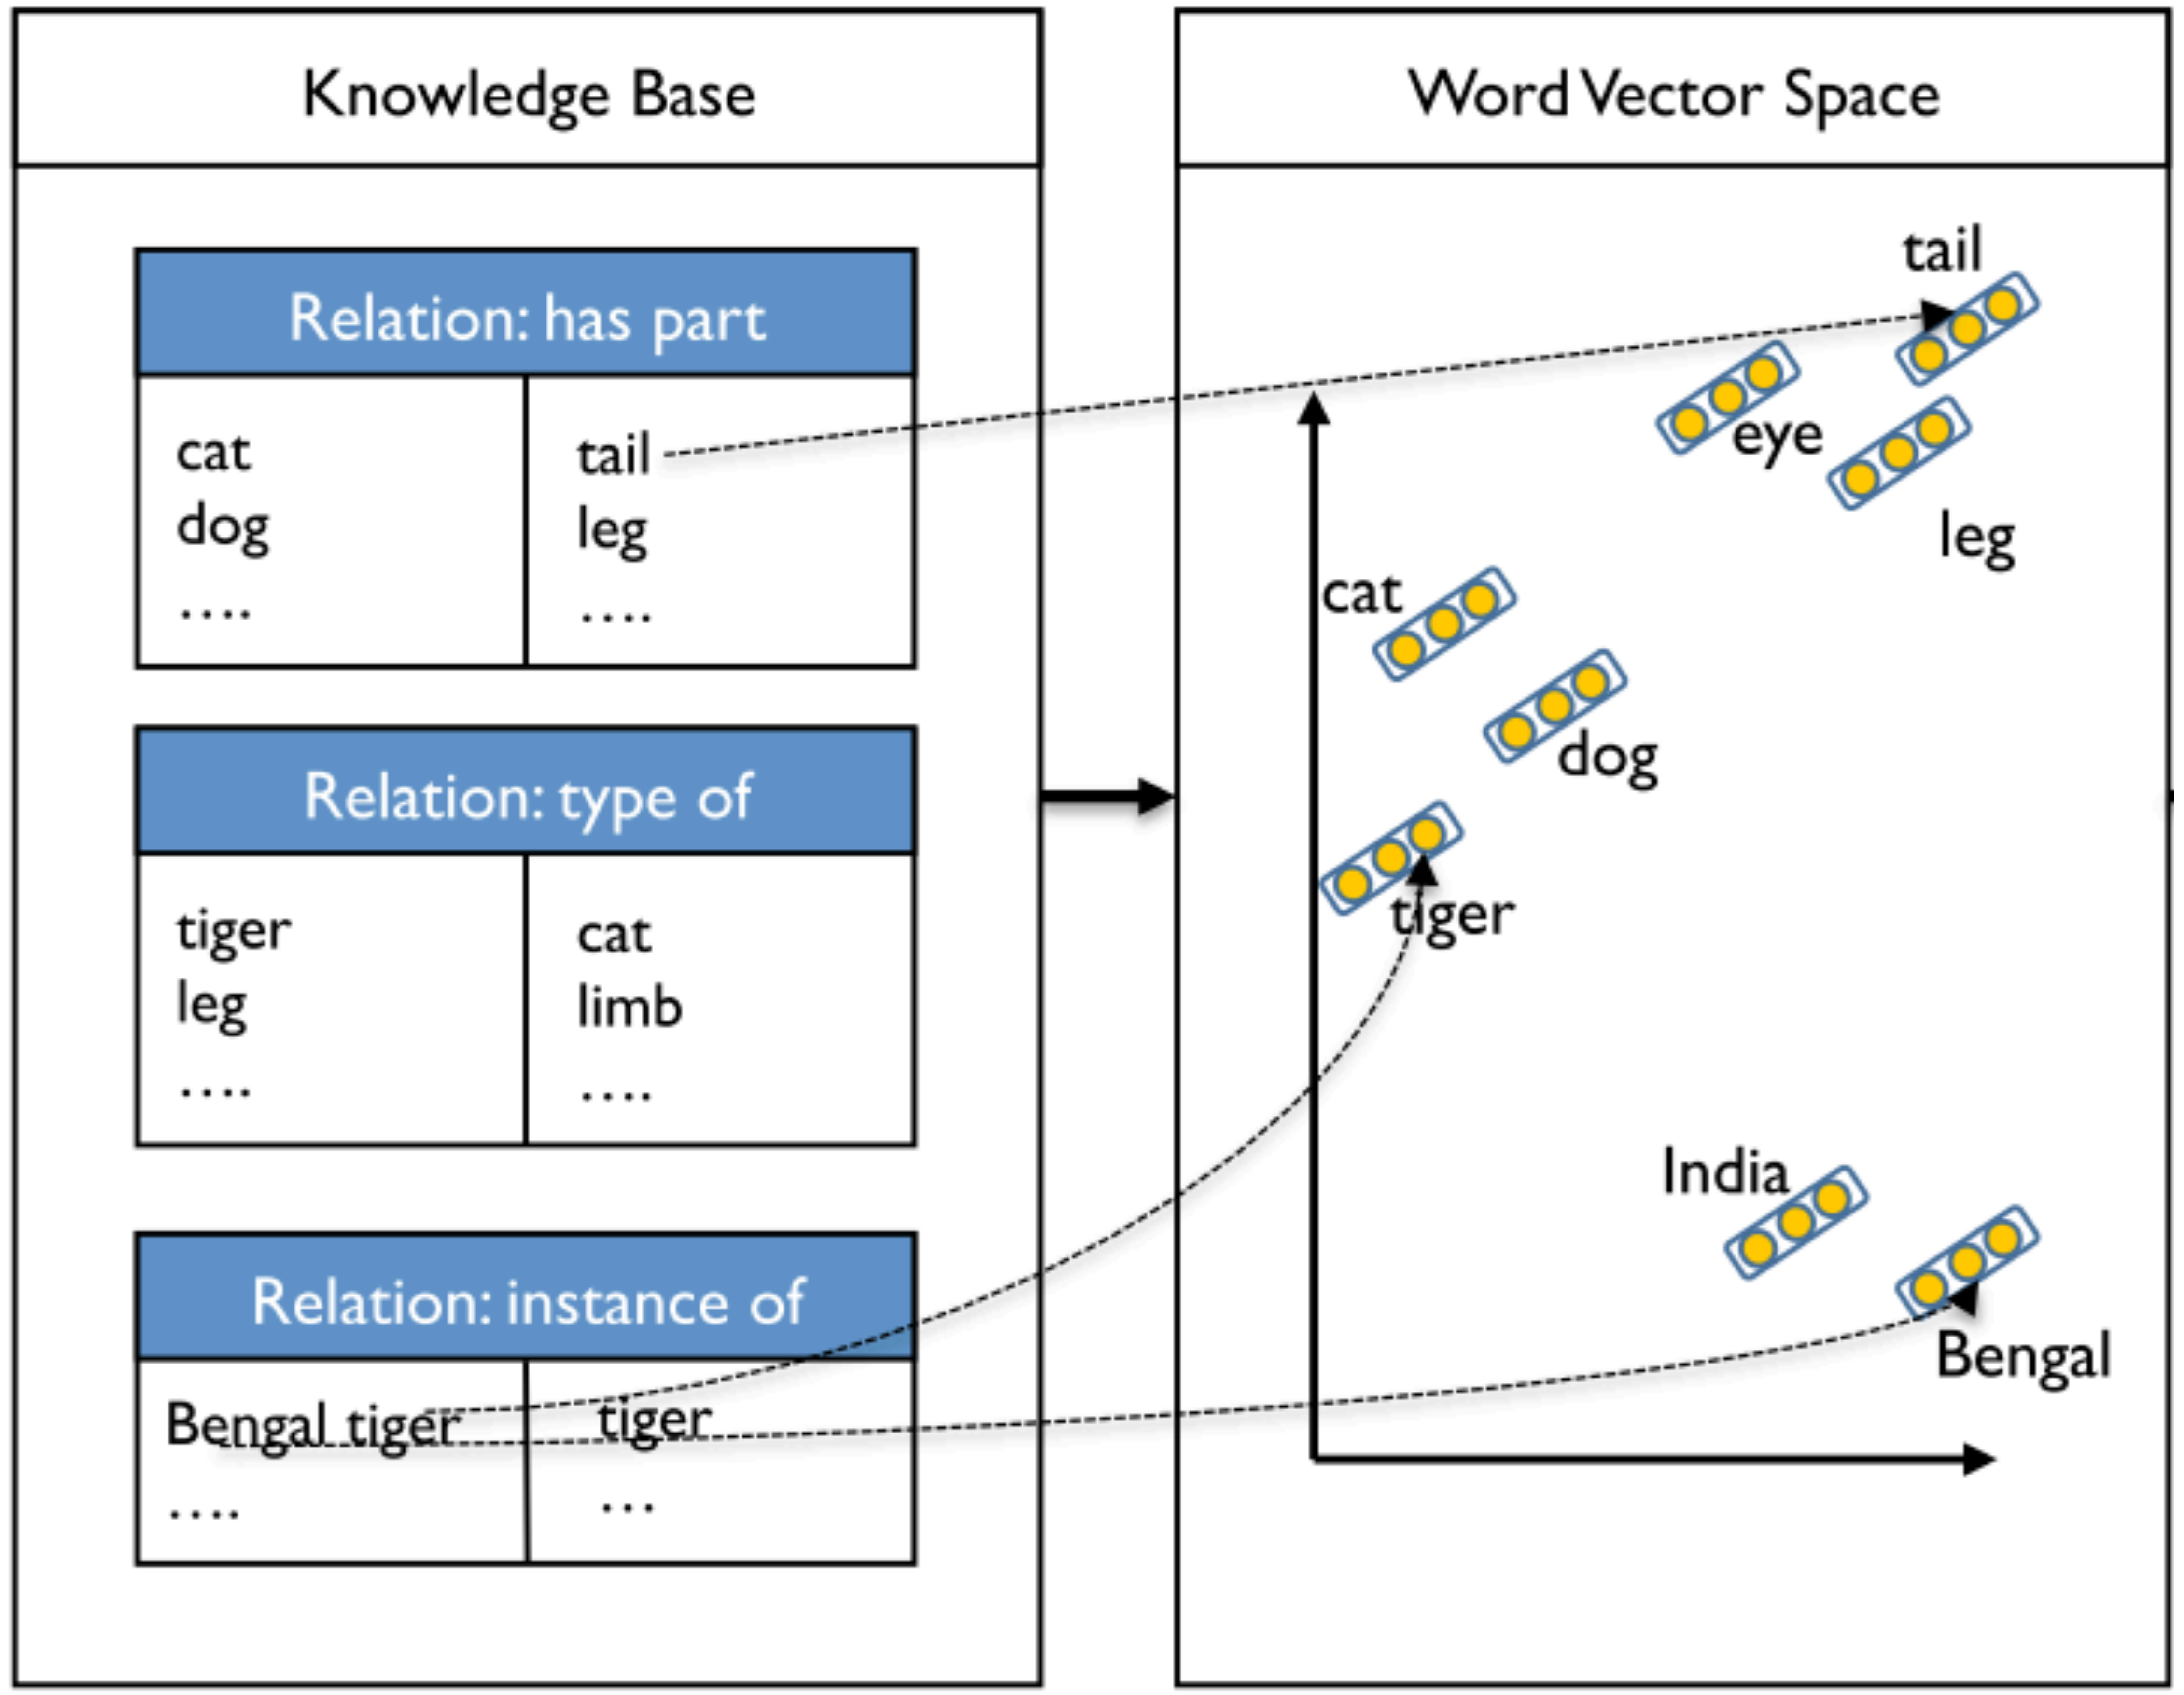

Contribution 3: Initialize entity word vectors using those pre-trained on large unlabeled text

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Neural Tensor Network

Training Objective

\# training triples

\# randomly sampled corrupted triples

$$\mathbf{\Omega} = \mathbf{u}, \mathbf{W}, \mathbf{V}, \mathbf{b}, \mathbf{E}$$

$$J(\mathbf{\Omega}) = \sum_{i=1}^{N} \sum_{c=1}^{C} \max\left(0, 1 - g\left(T^{(i)}\right) + g\left(T_c^{(i)}\right)\right) + \lambda \|\mathbf{\Omega}\|_2^2$$

Correct triple

Corrupted triple

$$T_c^{(i)} = (e_1^{(i)}, R^{(i)}, e_c).$$

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Neural Tensor Network

Experiments

Dataset	#R.	# Ent.	# Train	# Dev	# Test
Wordnet	11	38,696	112,581	2,609	10,544
Freebase	13	75,043	316,232	5,908	23,733

Model	WordNet	Freebase	Avg.
Distance Model	68.3	61.0	64.7
Hadamard Model	80.0	68.8	74.4
Single Layer Model	76.0	85.3	80.7
Bilinear Model	84.1	87.7	85.9
Neural Tensor Network	**86.2**	**90.0**	**88.1**

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Neural Tensor Network

Experiments

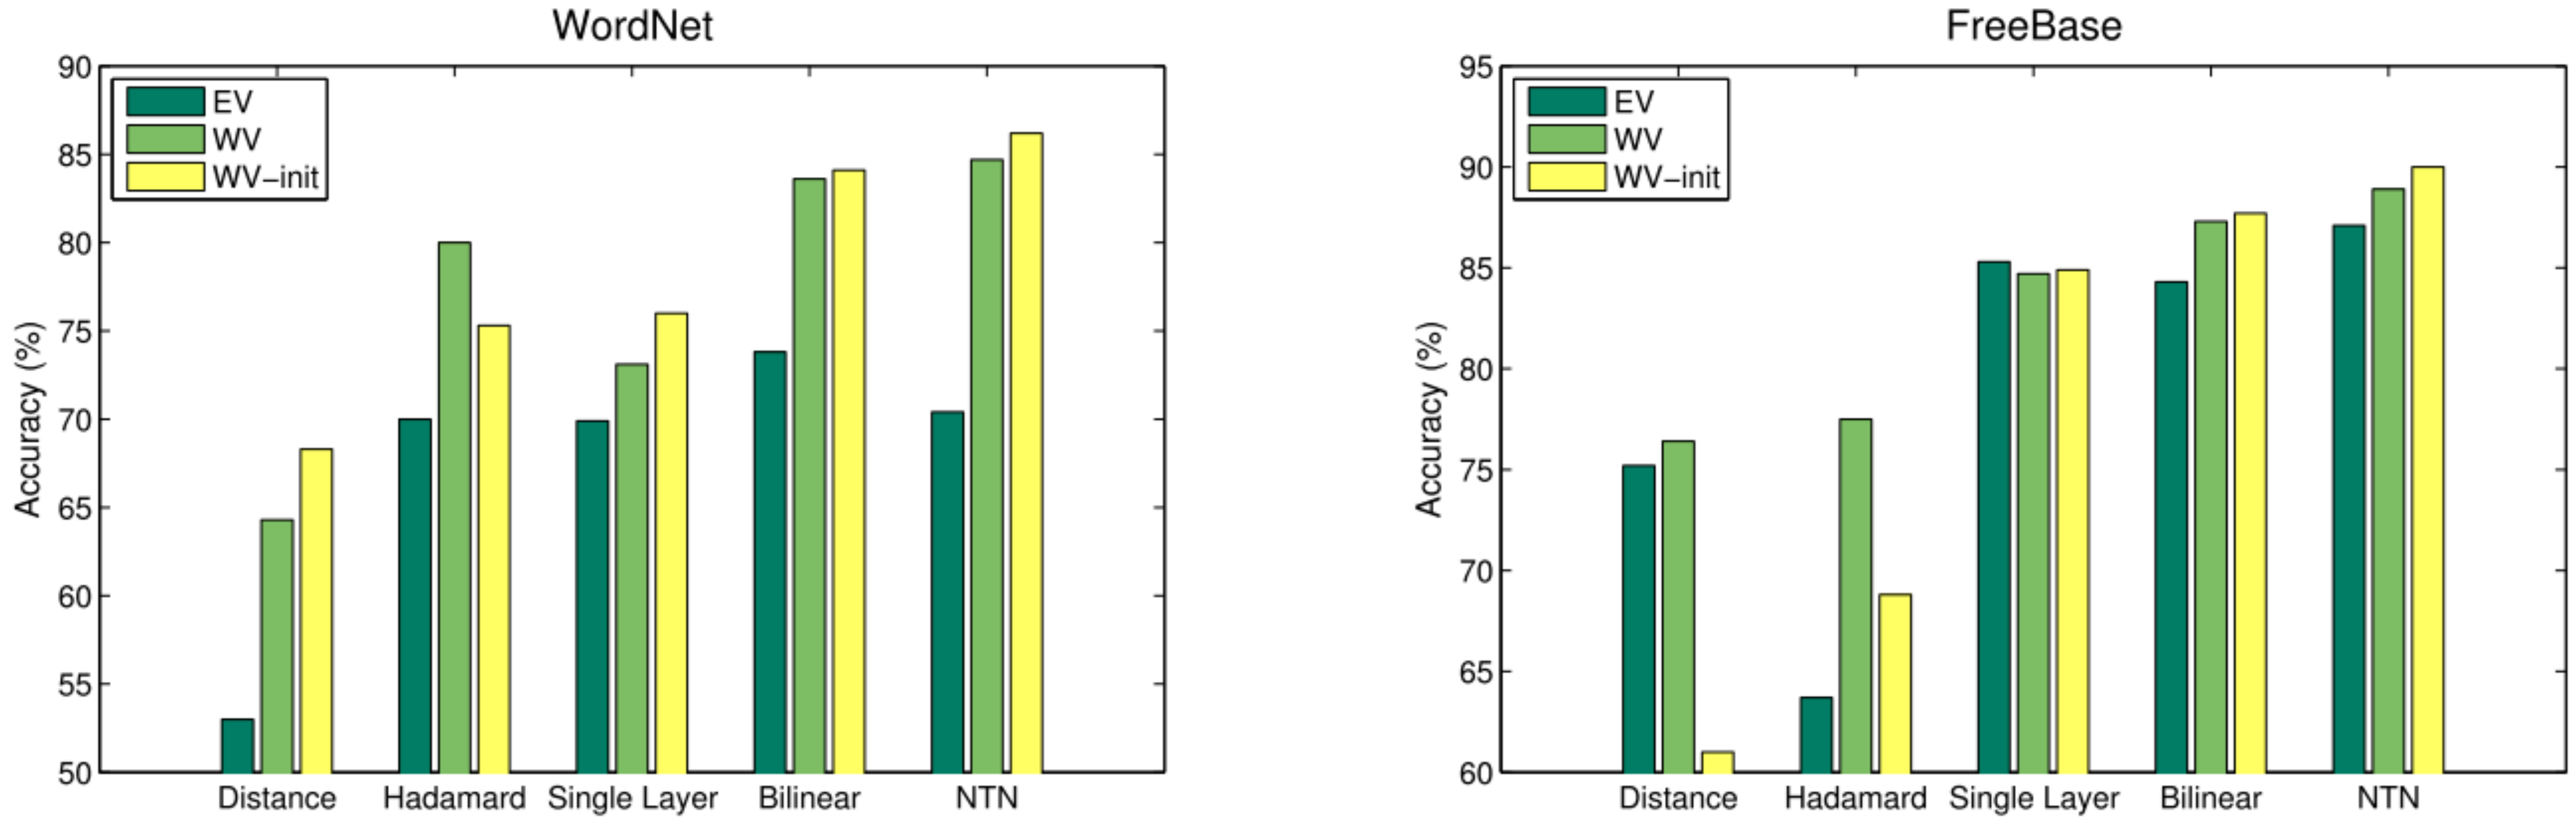

Figure 4: Influence of entity representations. **EV**: entity vectors. **WV**: randomly initialized word vectors. **WV-init**: word vectors initialized with unsupervised semantic word vectors.

Reasoning with Neural Tensor Networks for Knowledge Base Completion (Socher et. al. 2013)

Multi-Hop Reasoning Models

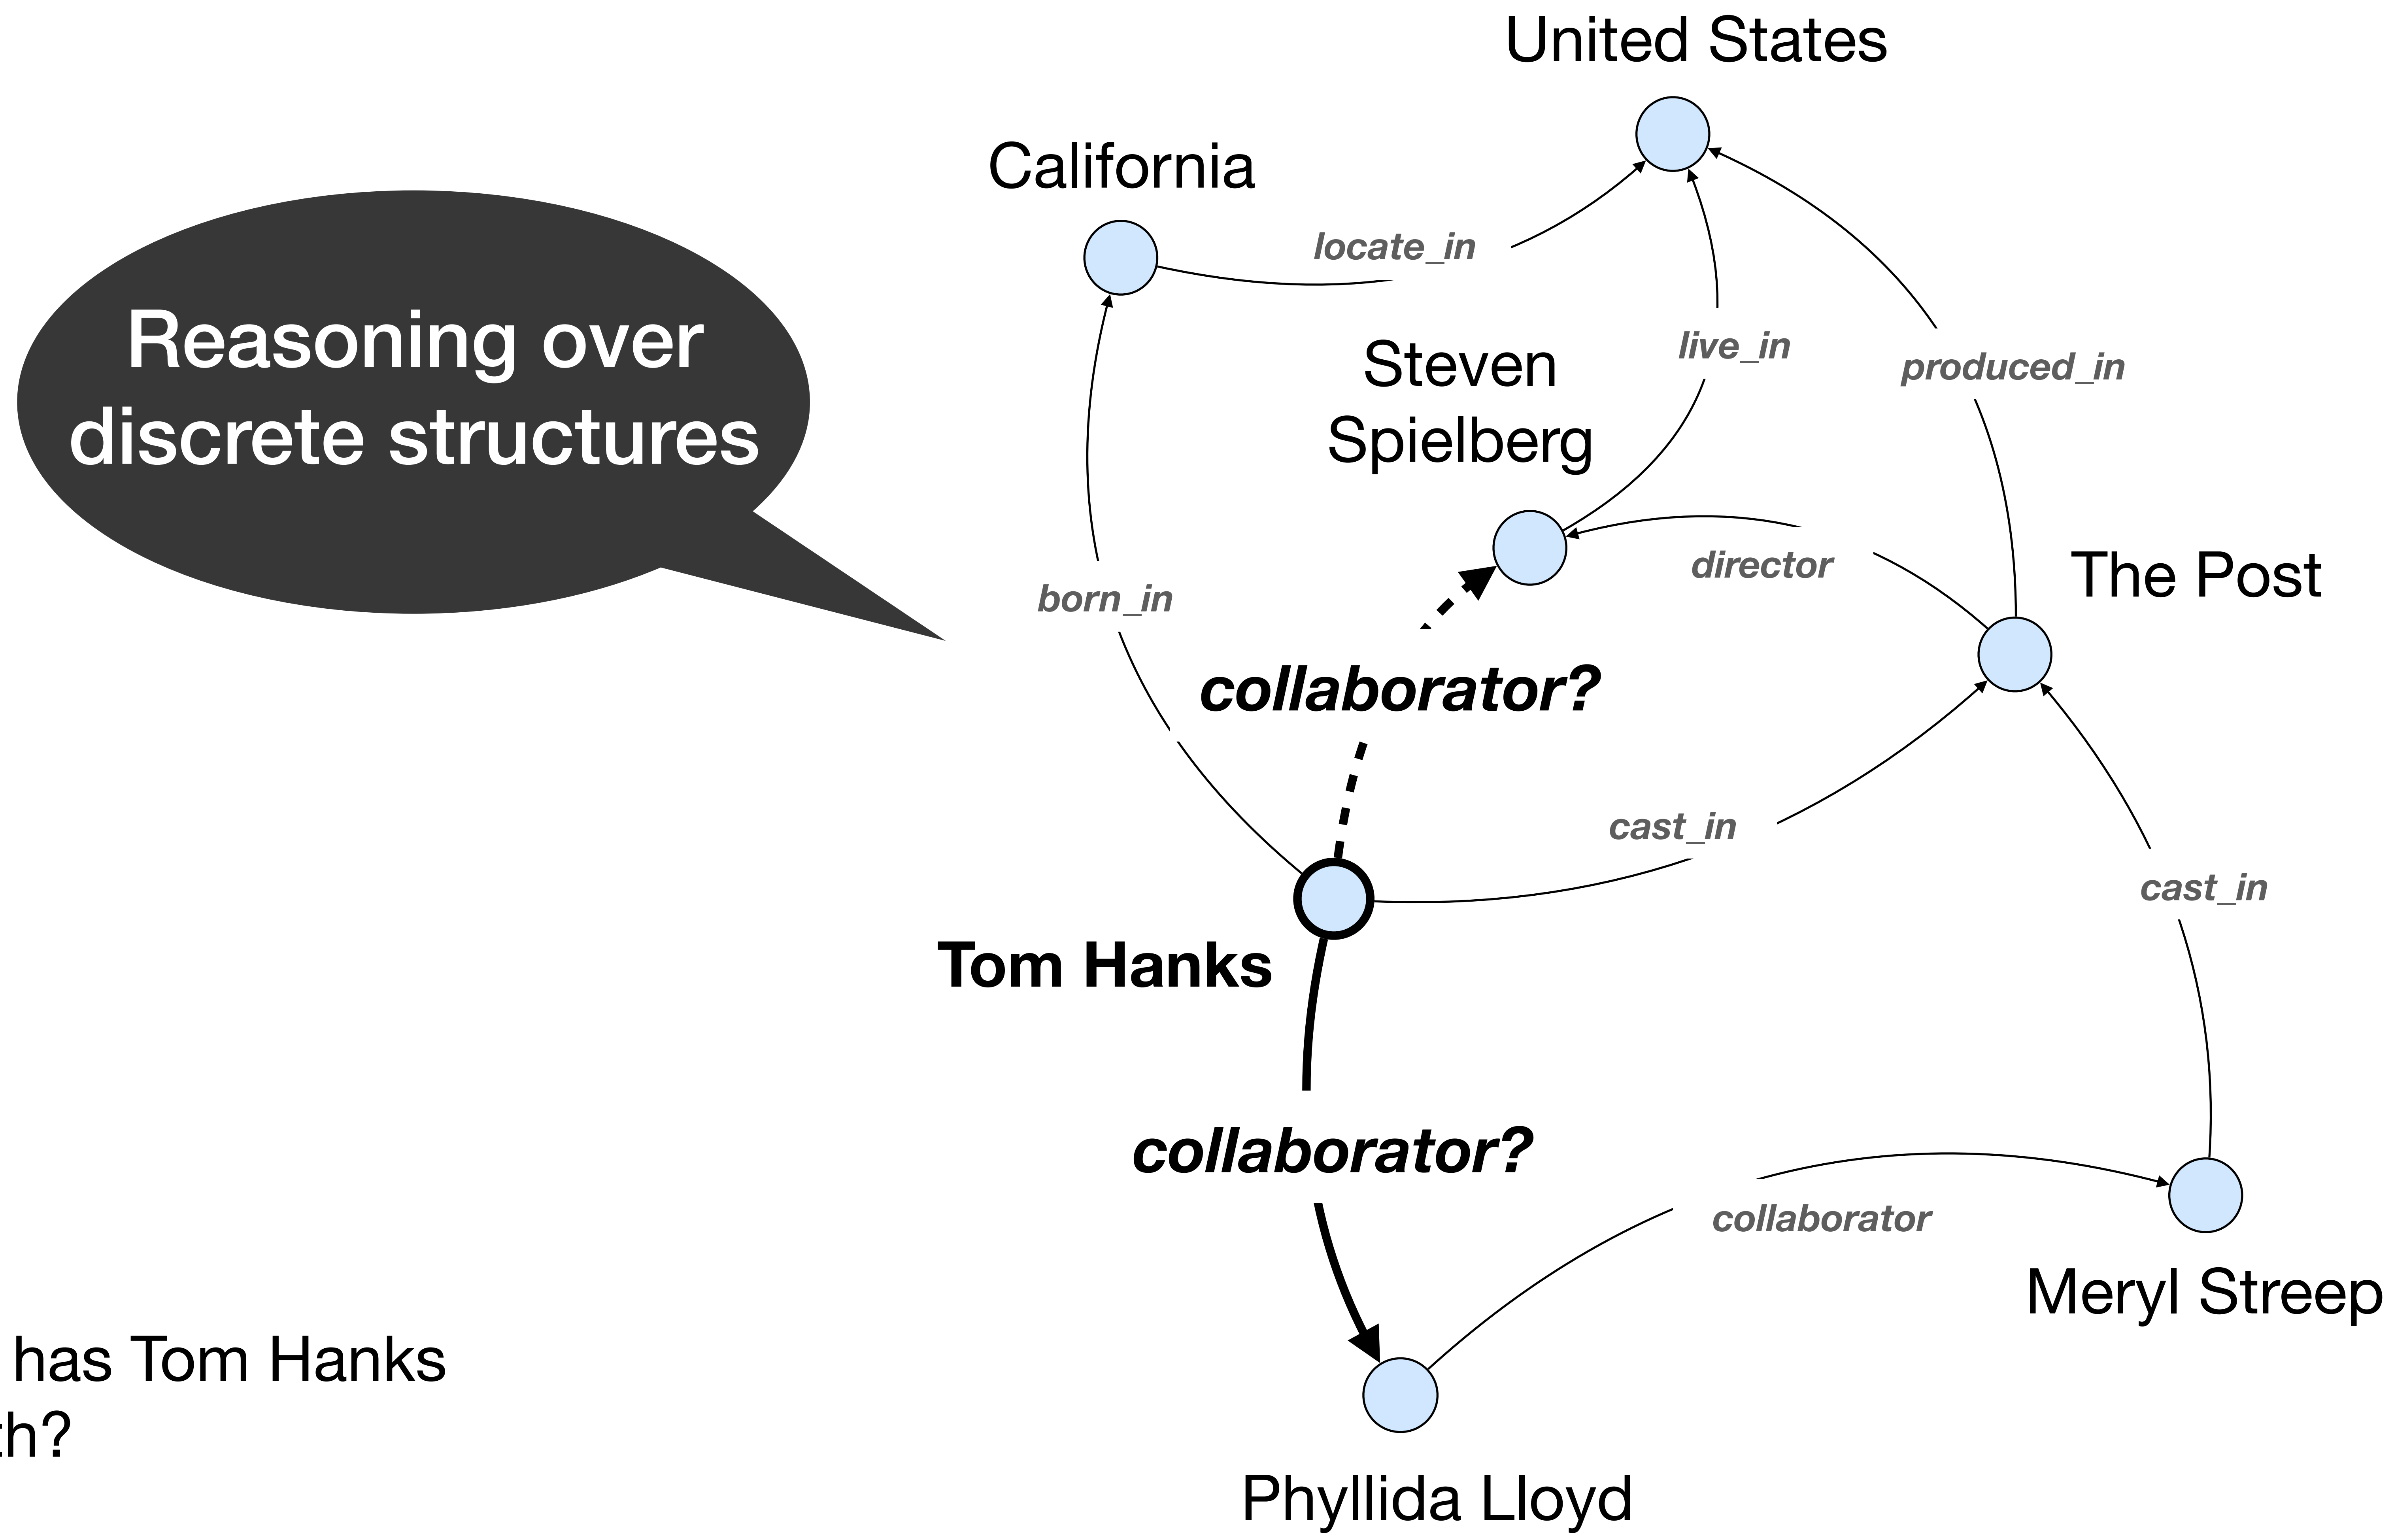

Which directors has Tom Hanks collaborated with?

Multi-Hop Reasoning Models

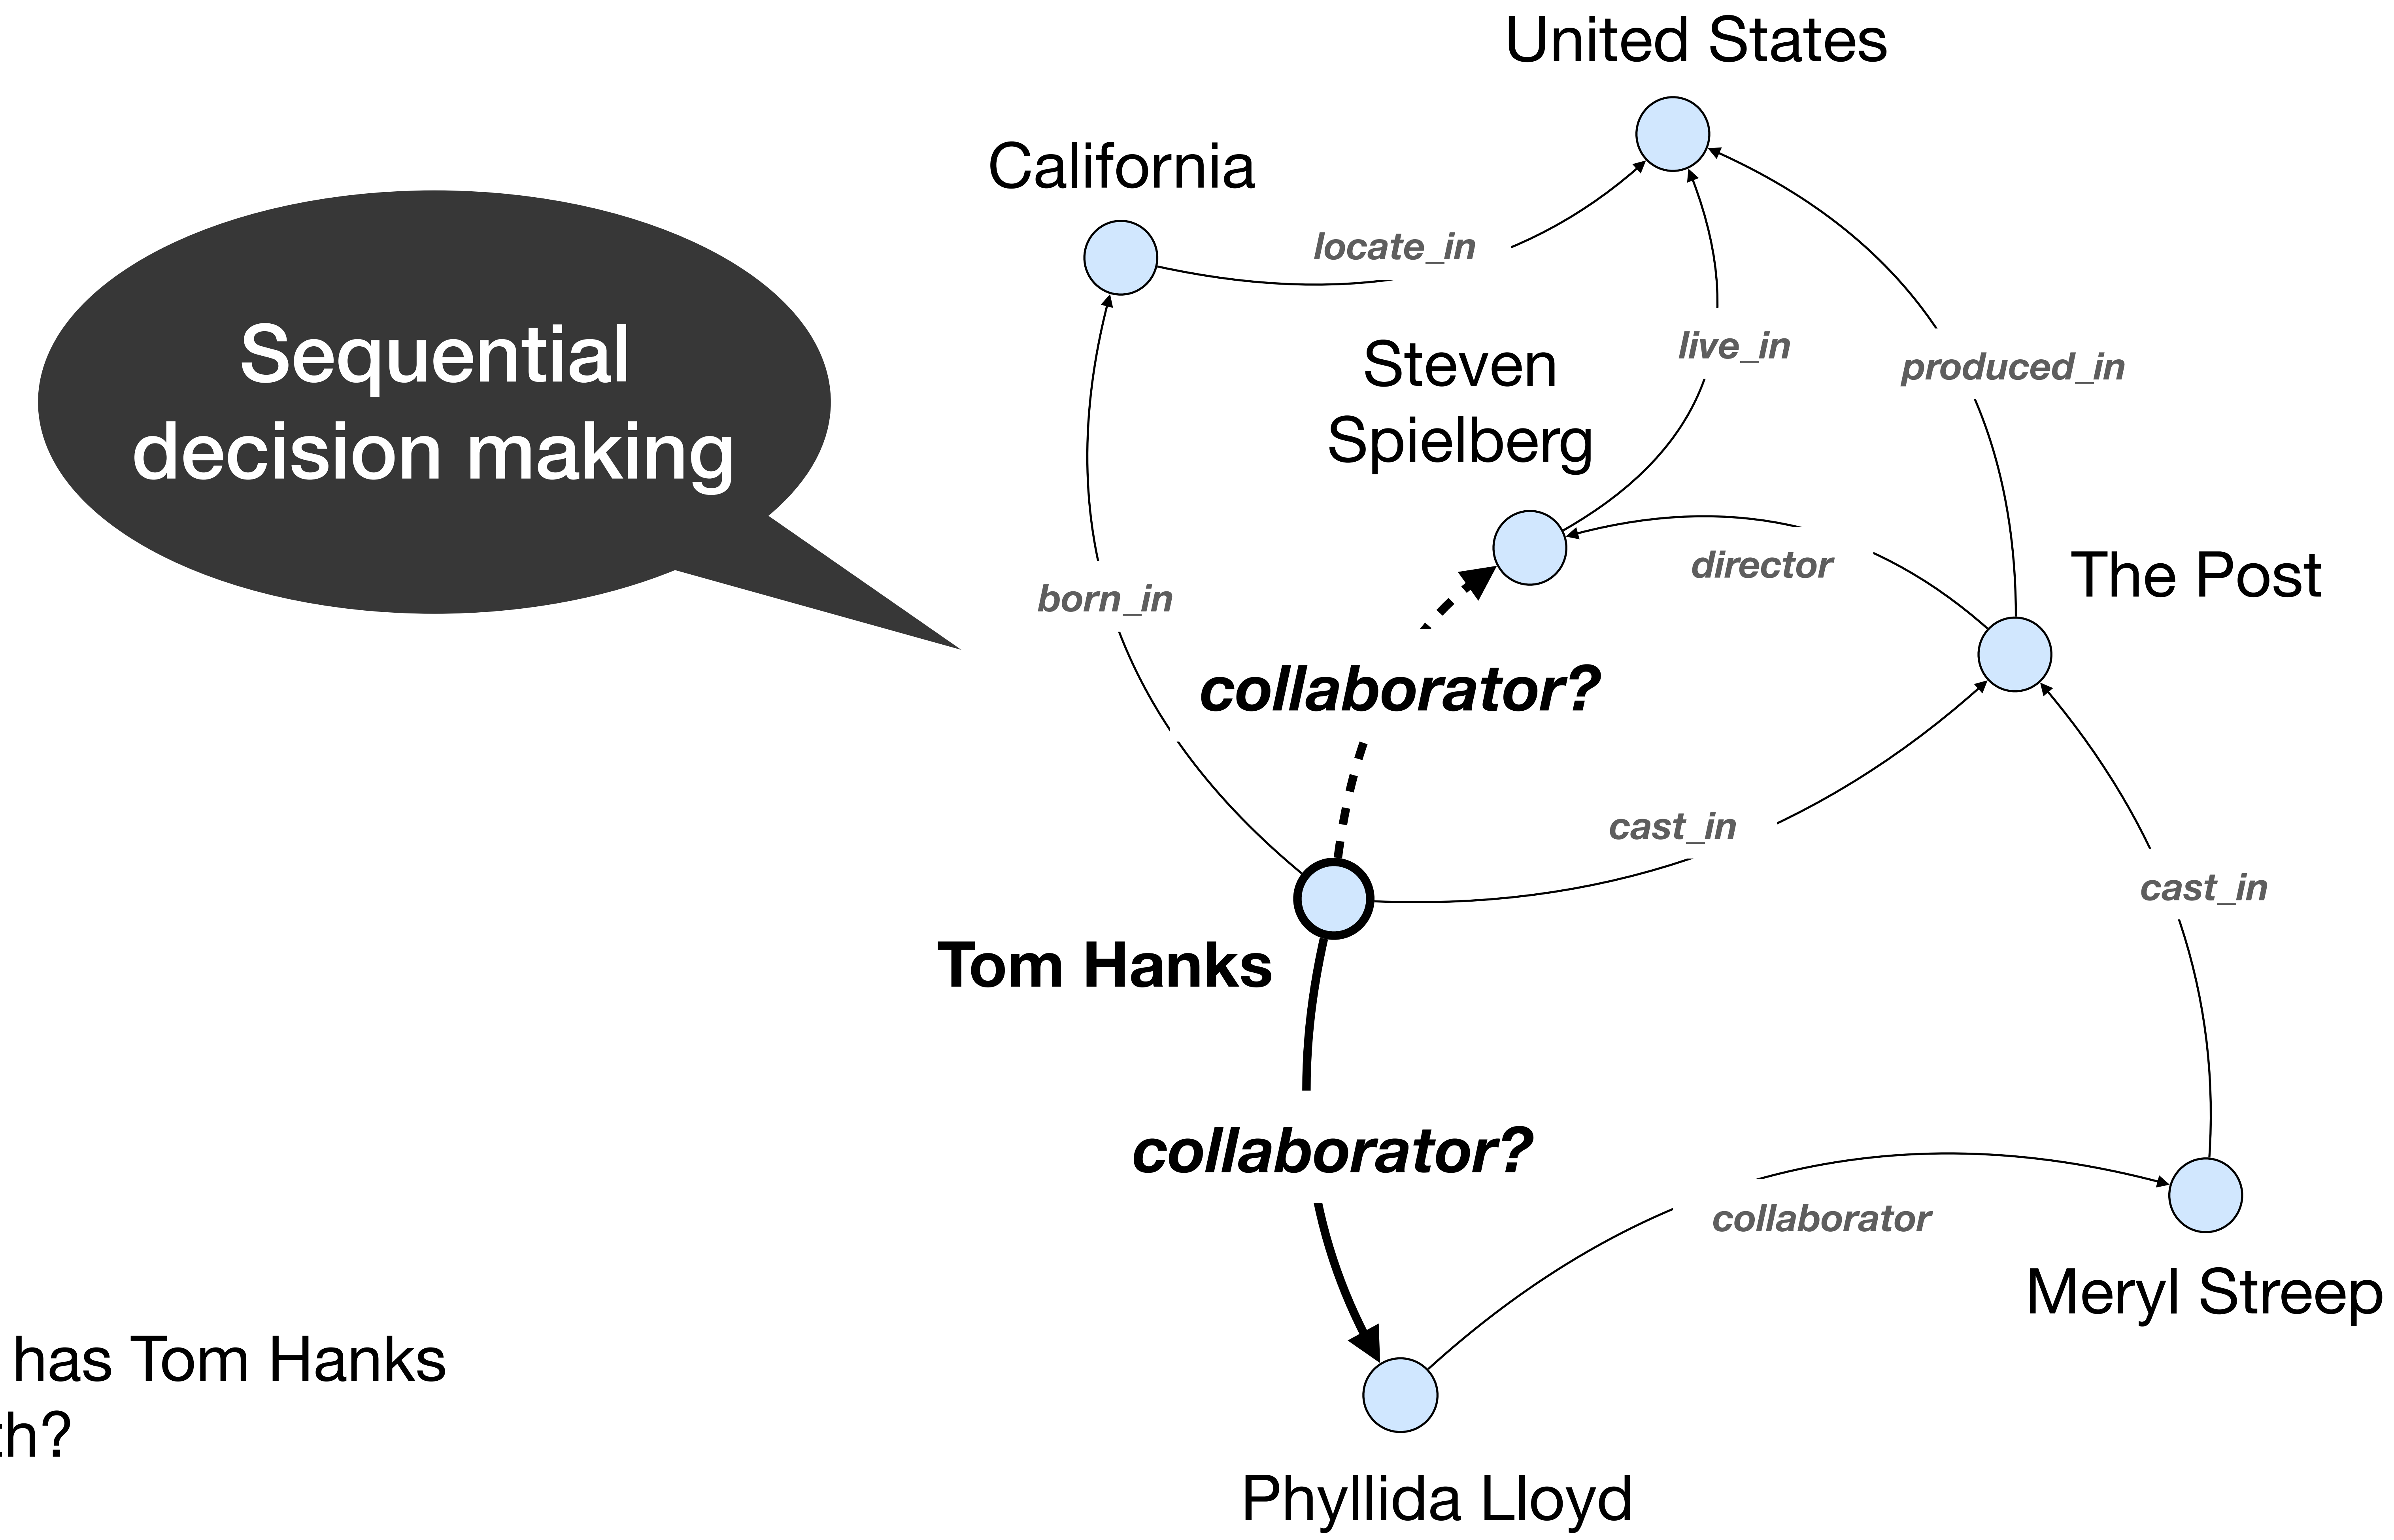

Which directors has Tom Hanks collaborated with?

Multi-Hop Reasoning Models

Topic entity

Tom Hanks

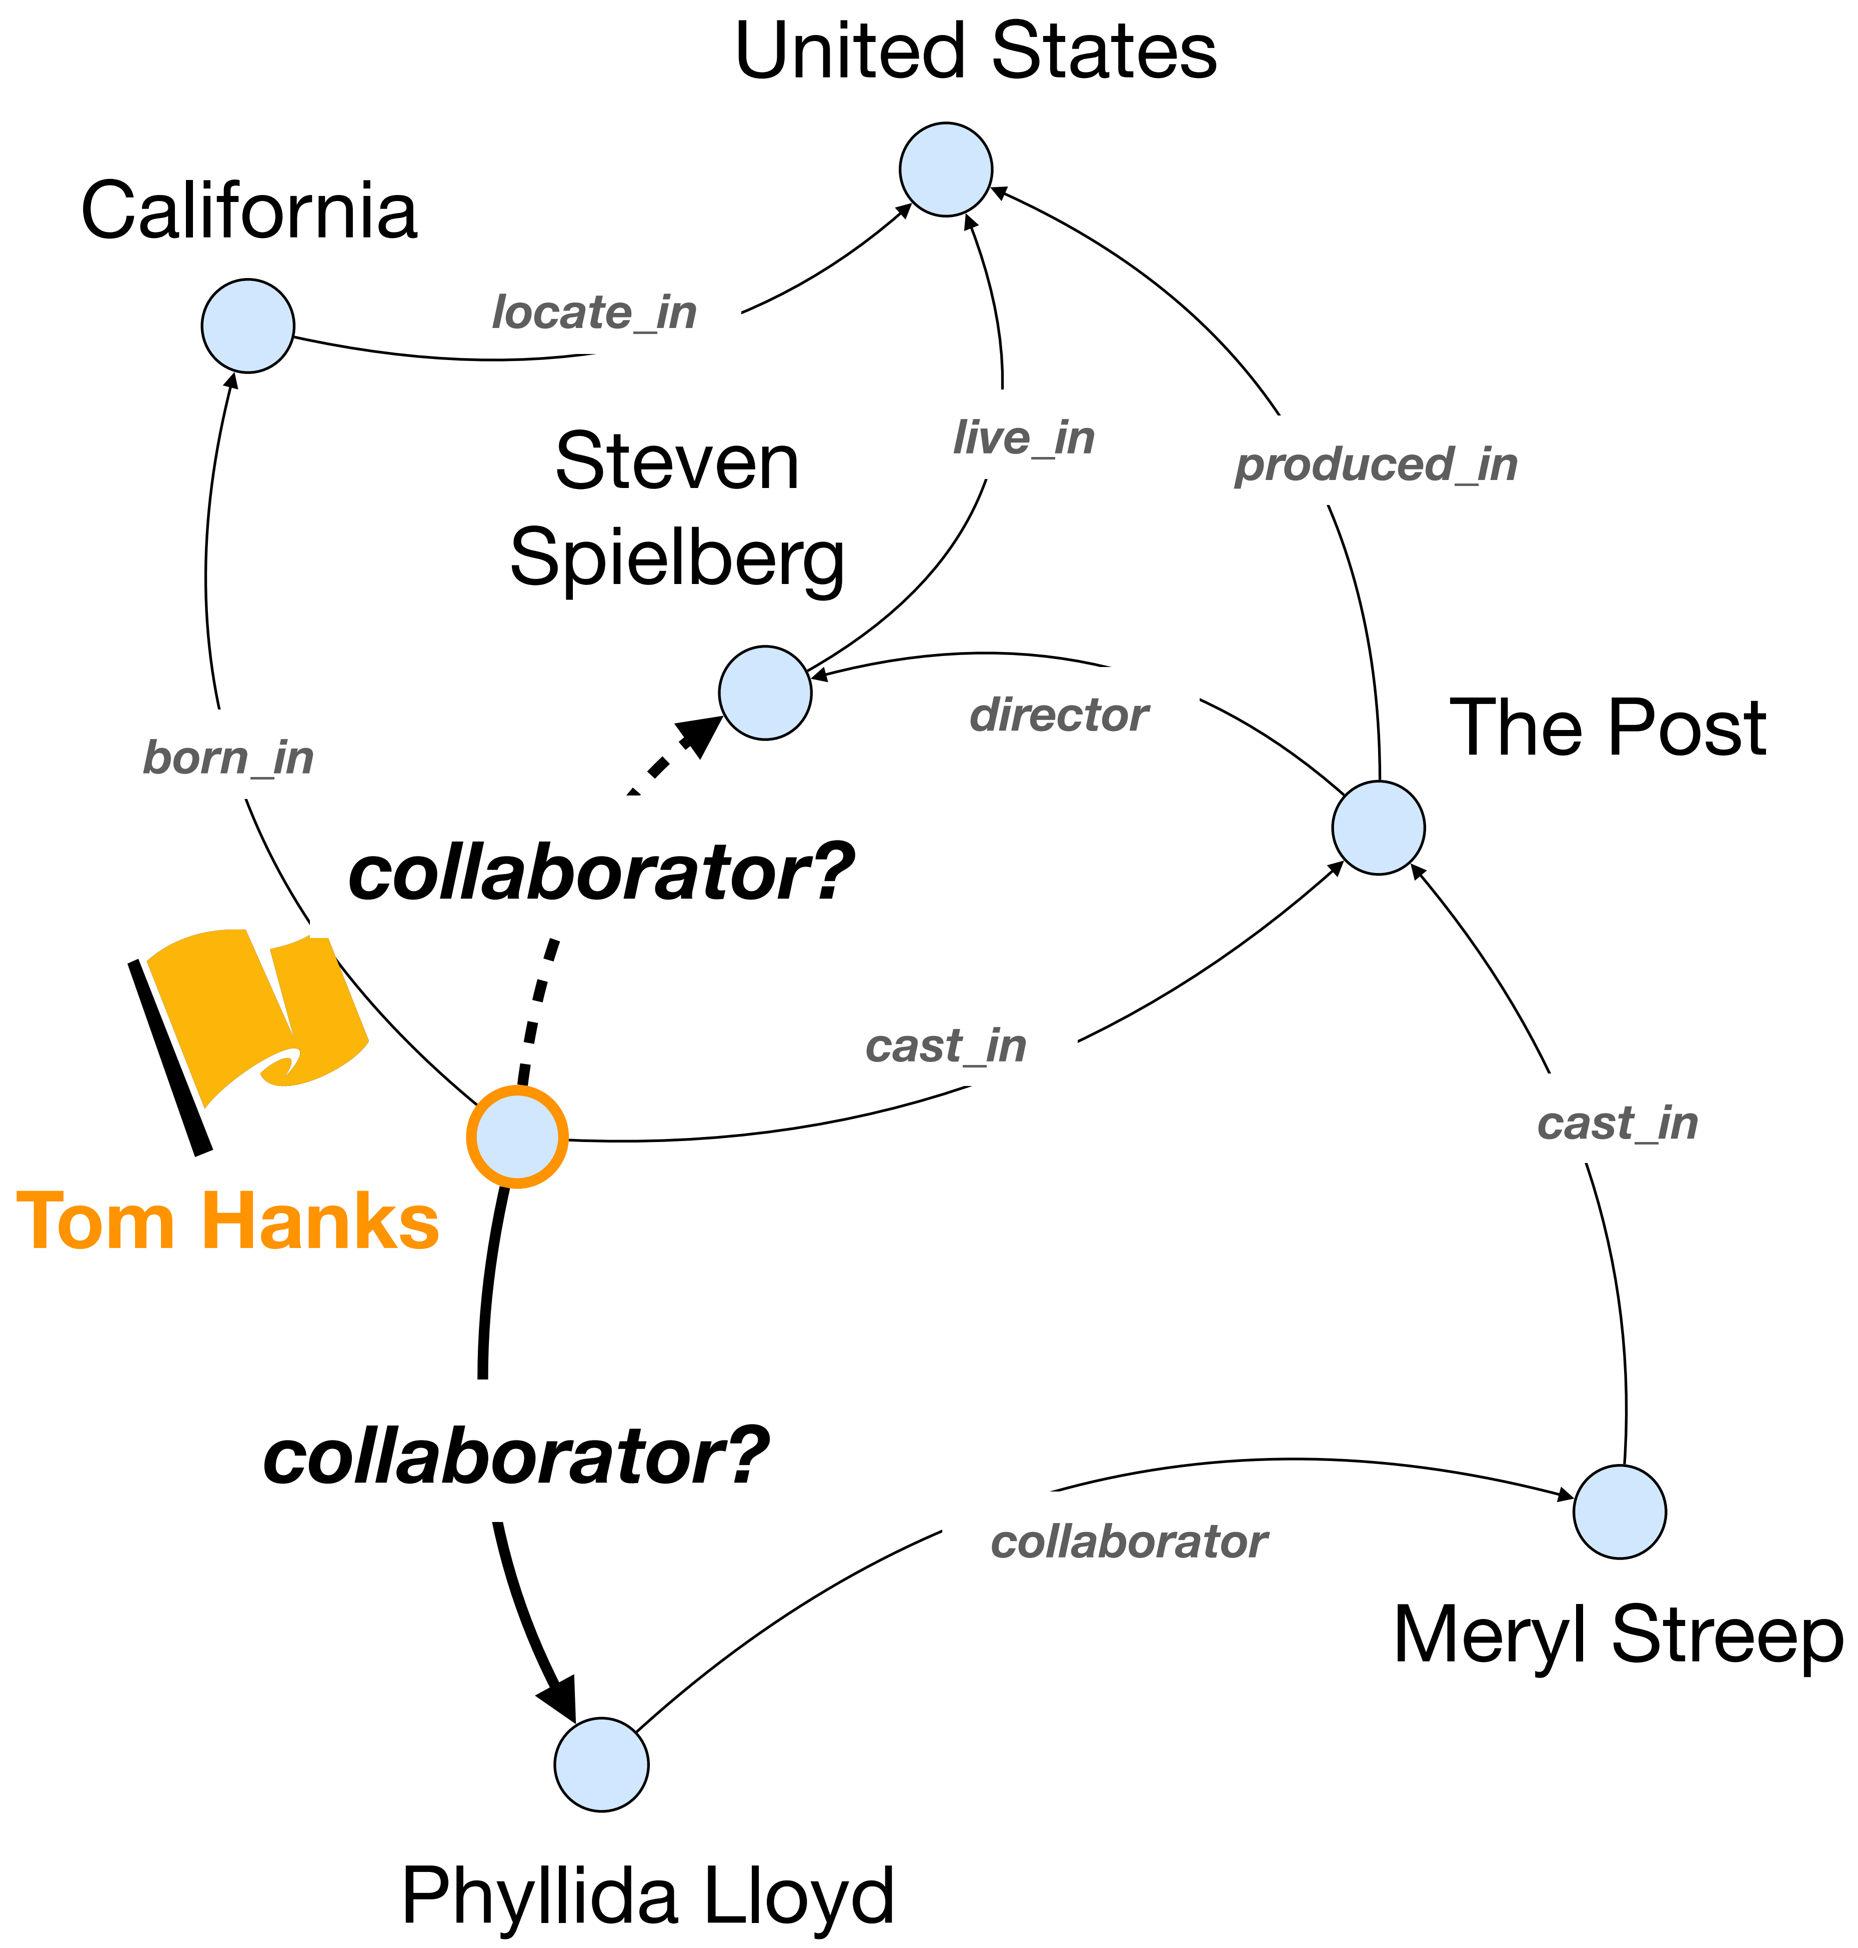

Which directors has Tom Hanks collaborated with?

Multi-Hop Reasoning Models

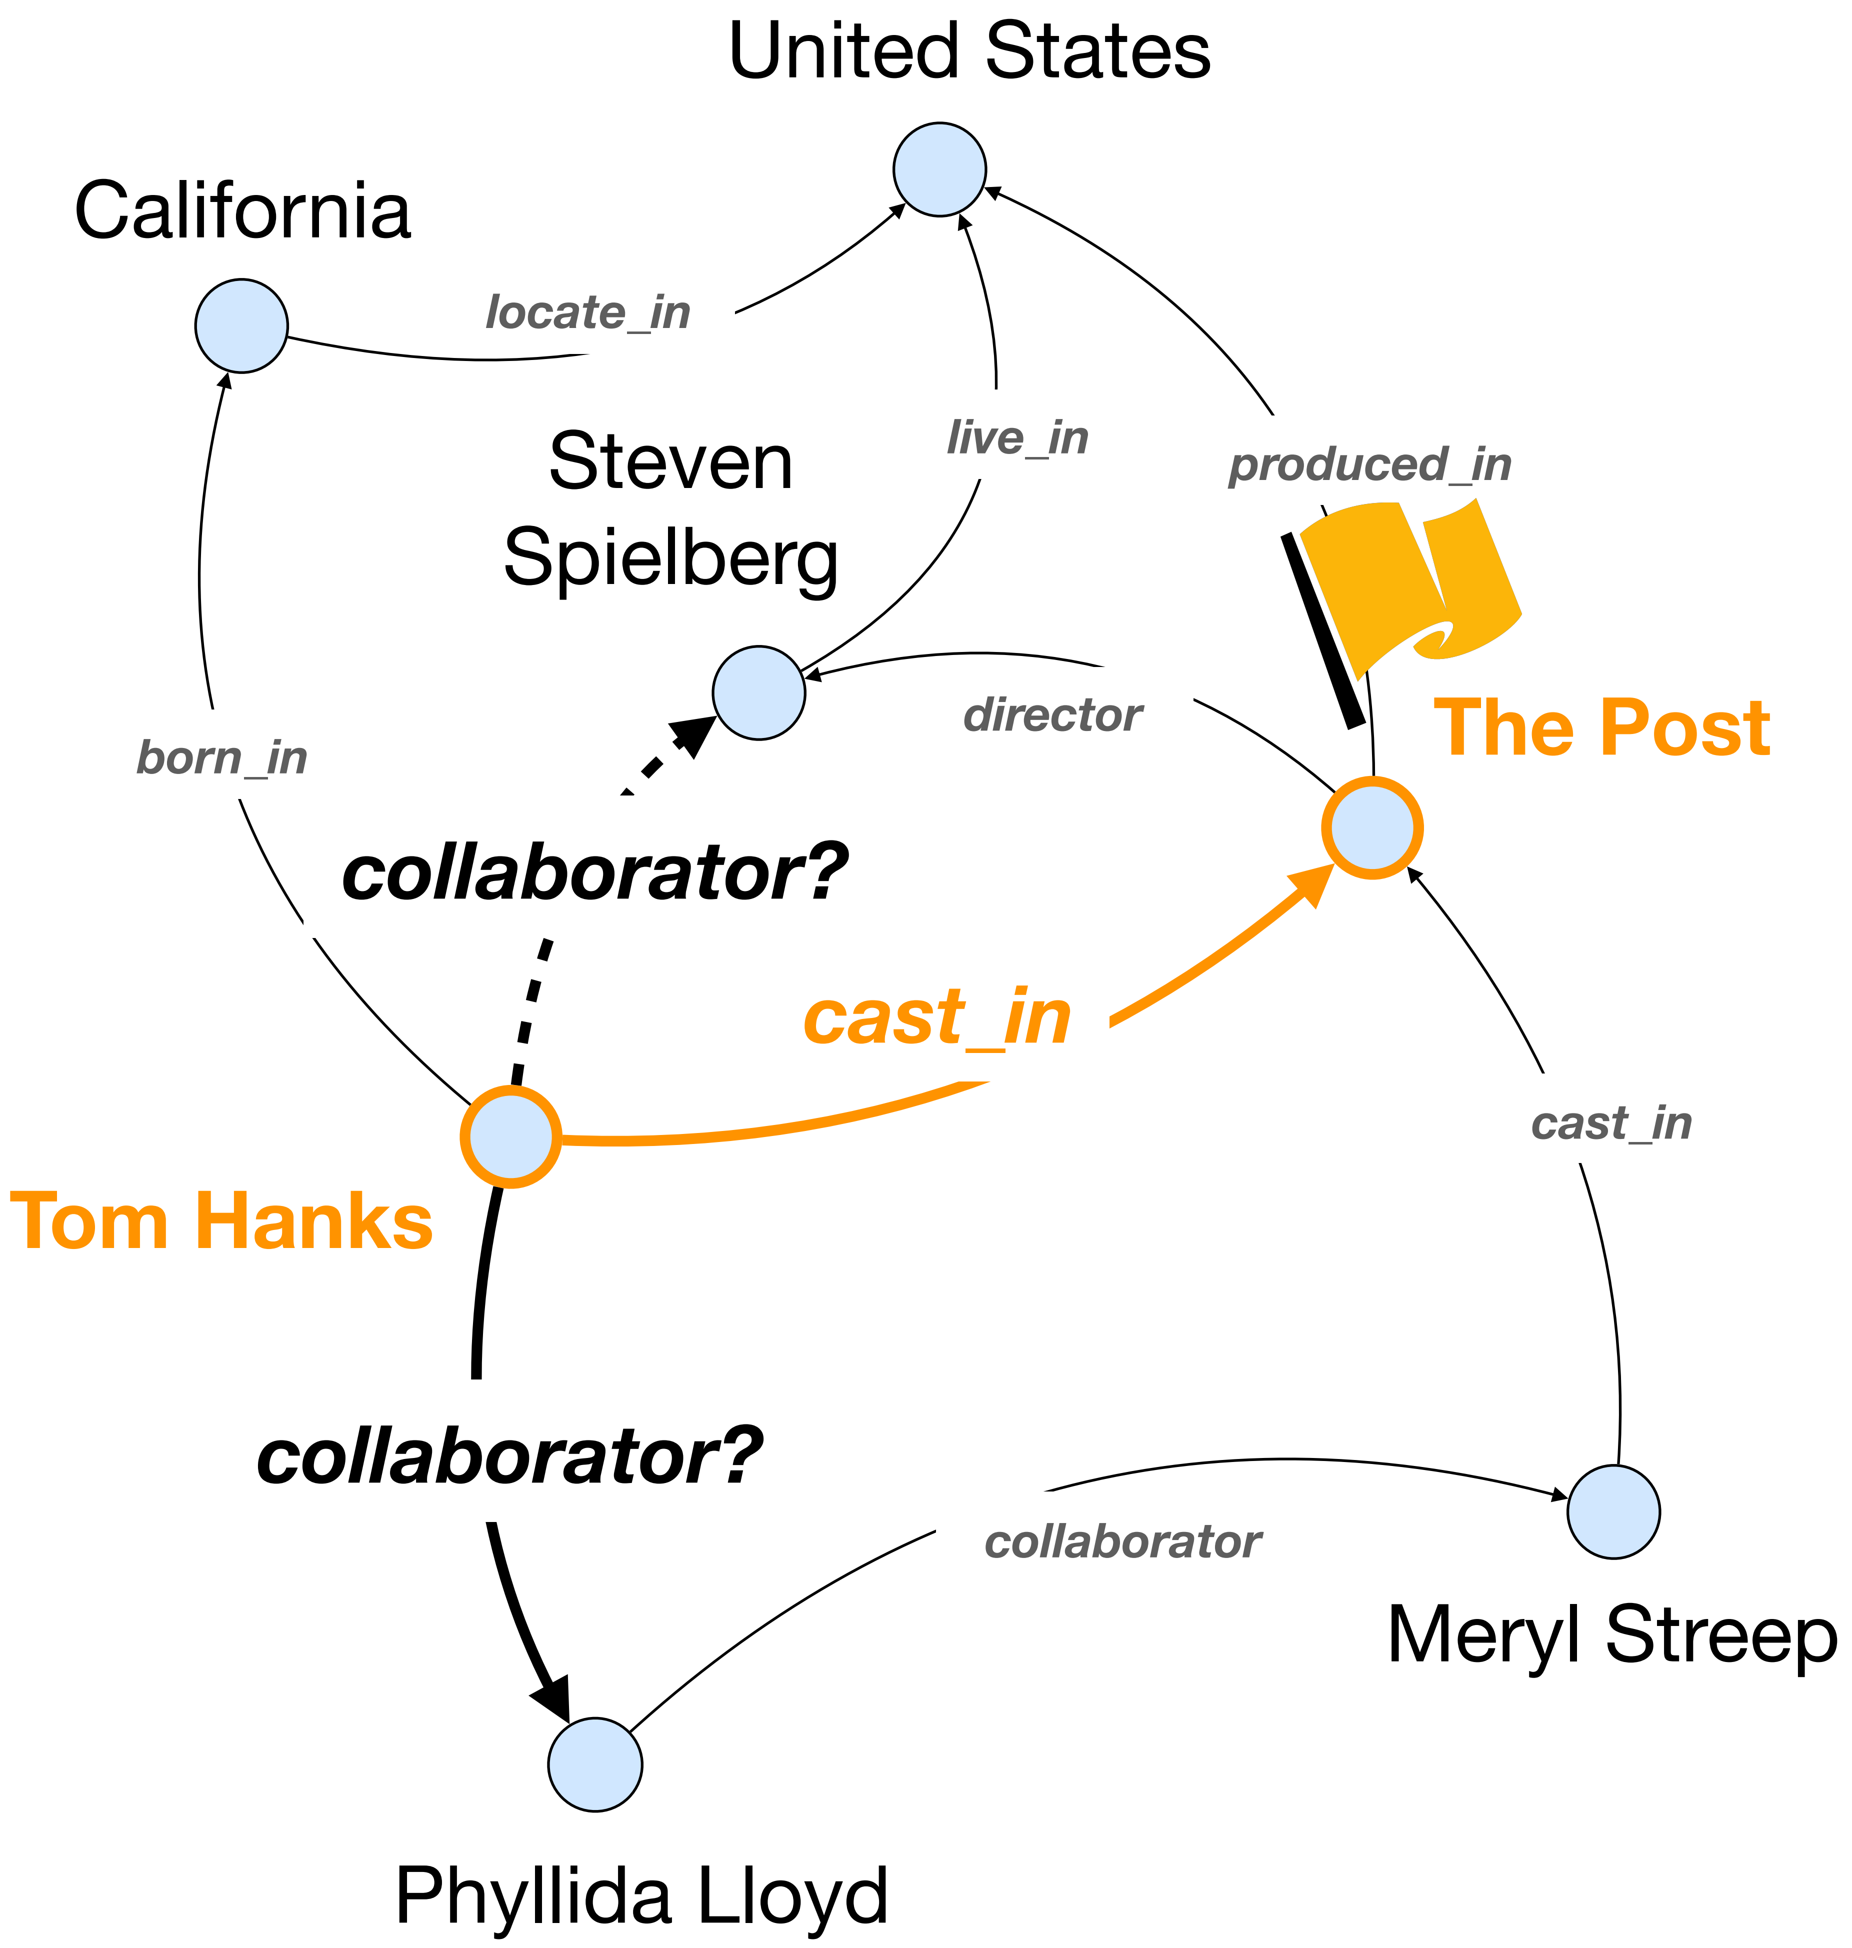

Which directors has Tom Hanks collaborated with?

Multi-Hop Reasoning Models

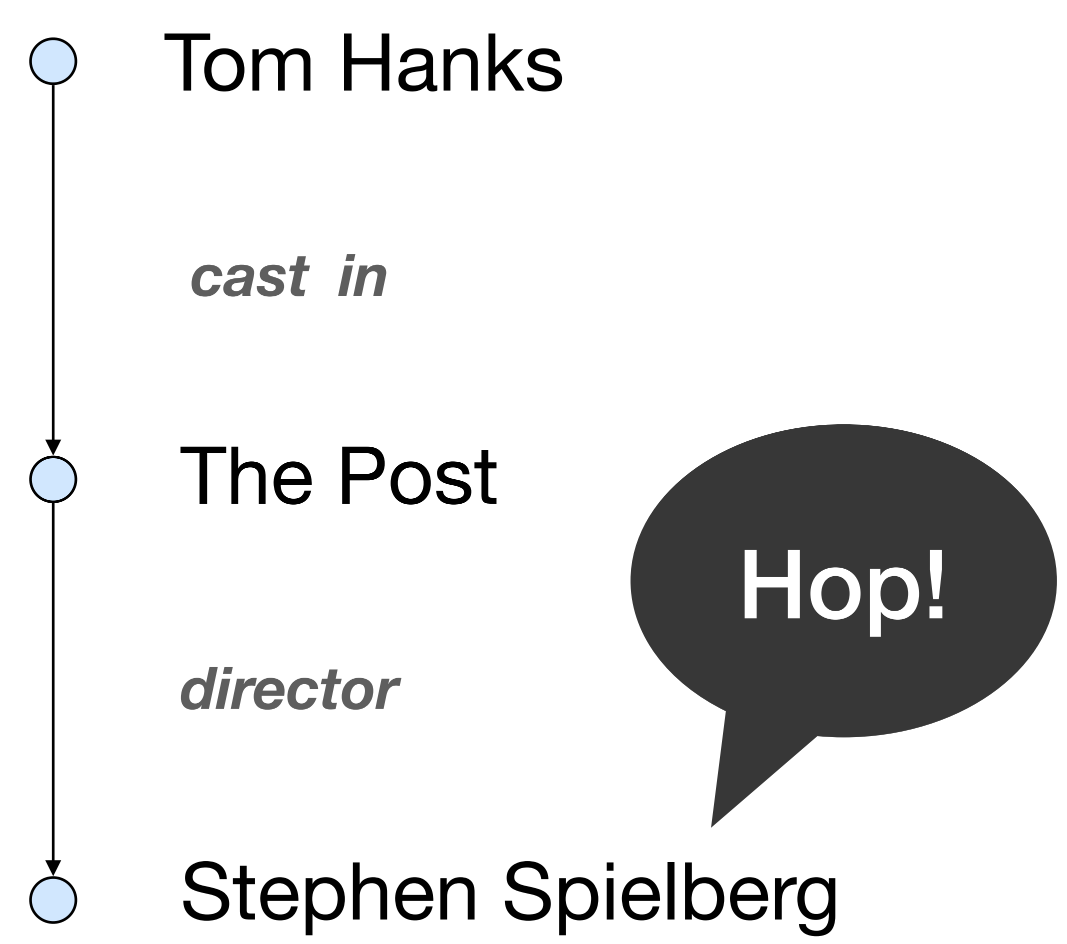

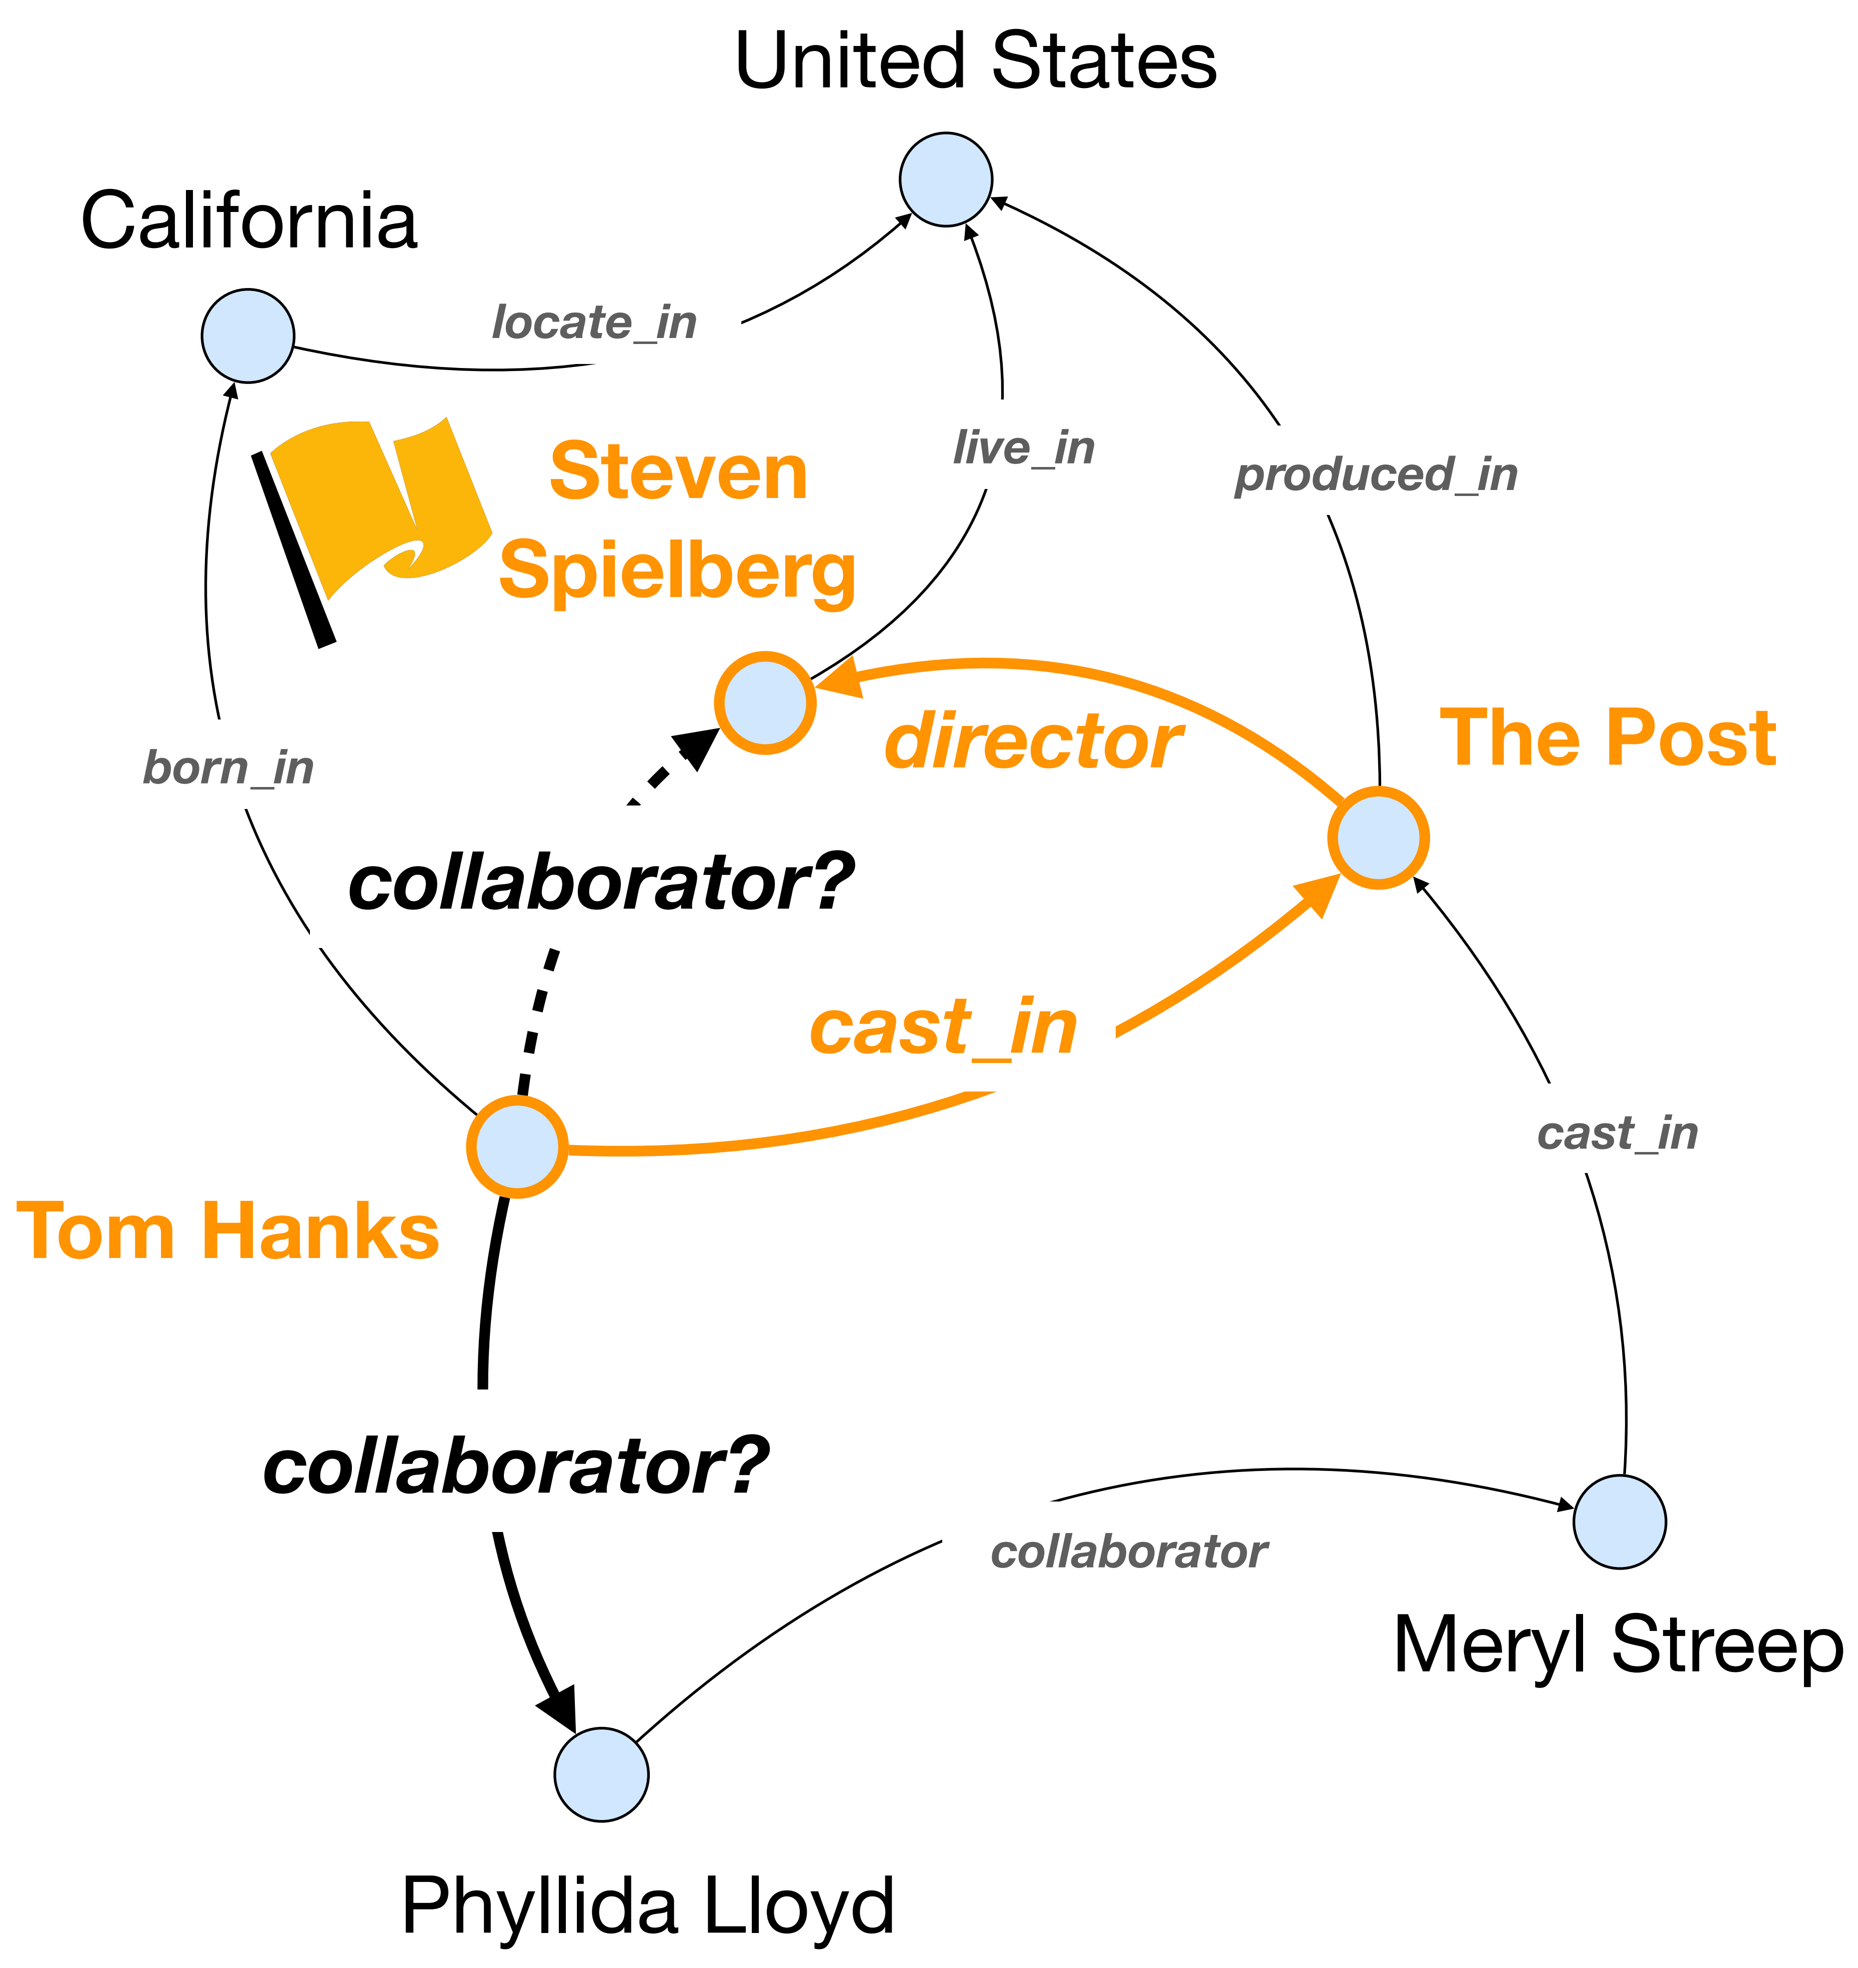

Which directors has Tom Hanks collaborated with?

Multi-Hop Reasoning Models

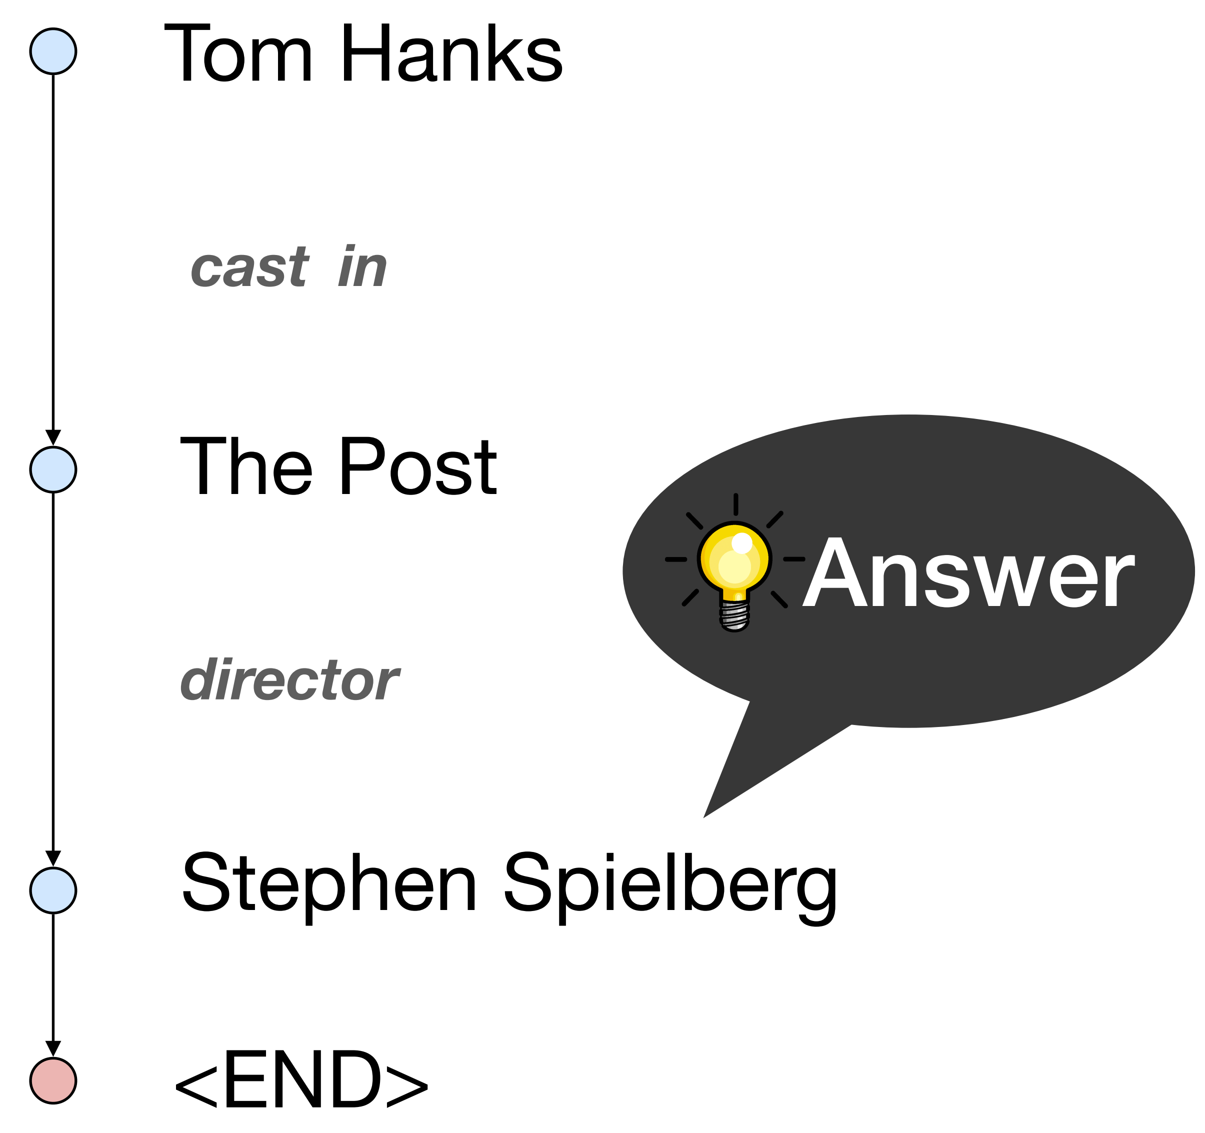

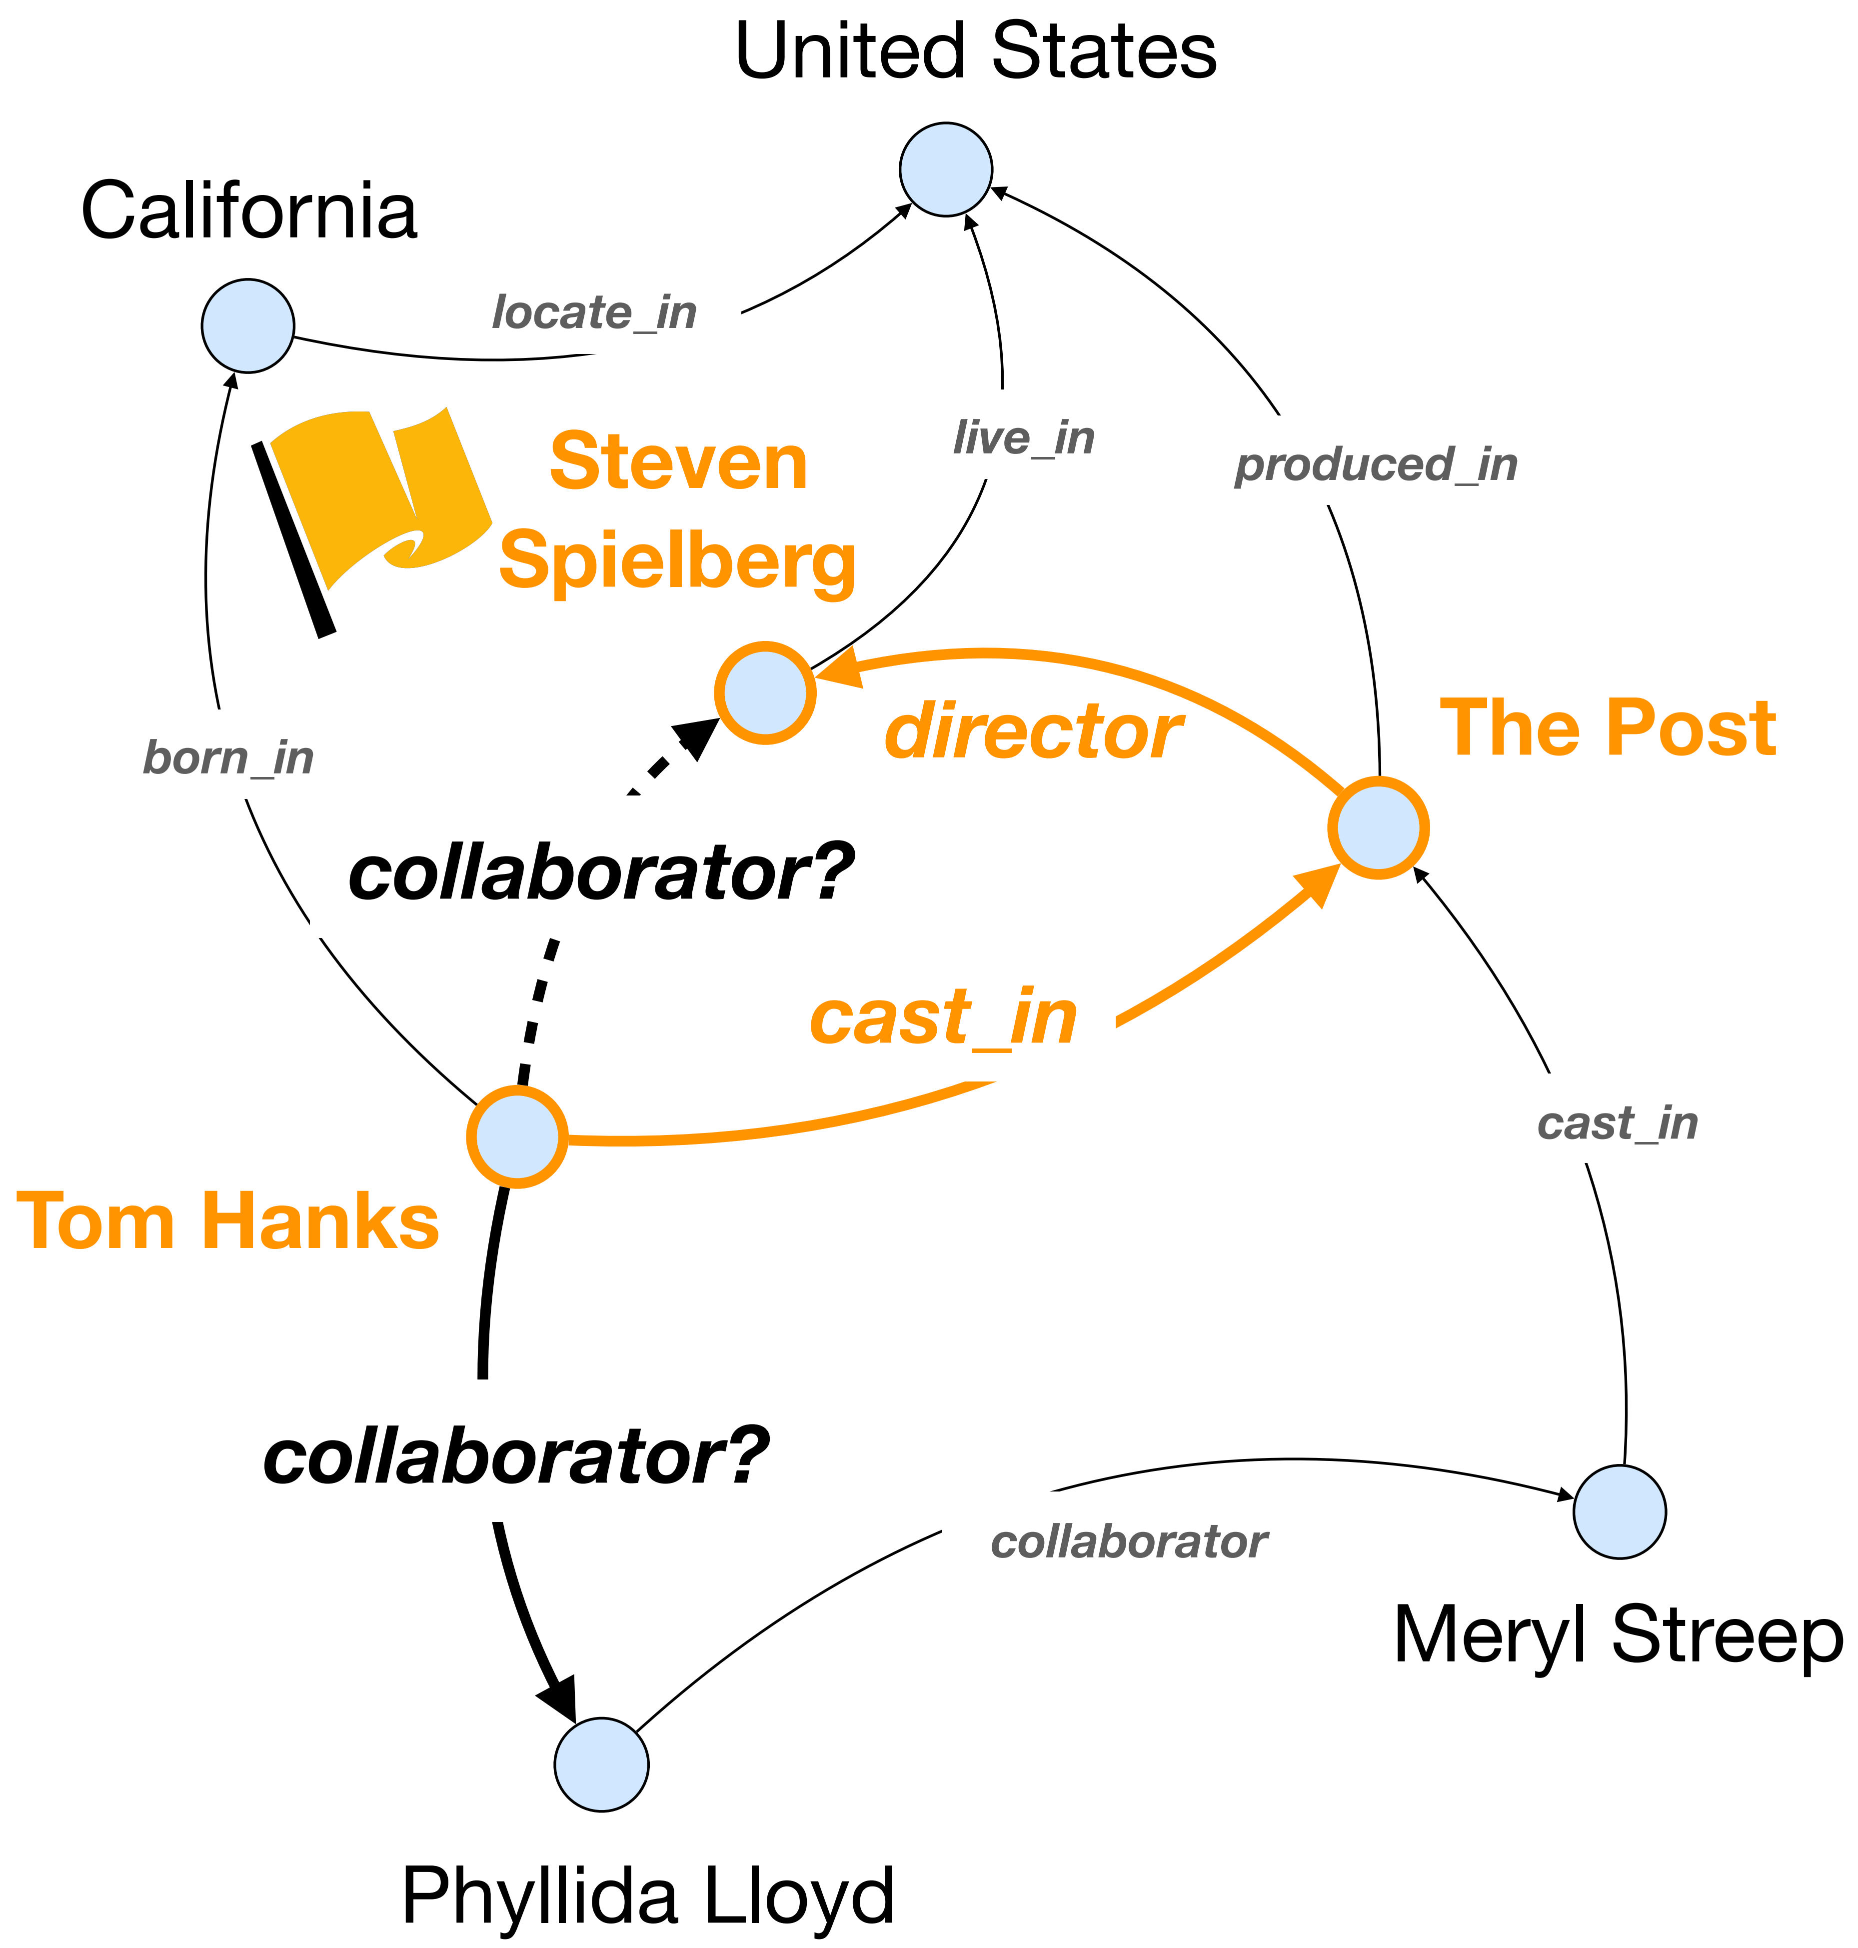

Which directors has Tom Hanks collaborated with?

Multi-Hop Reasoning Models

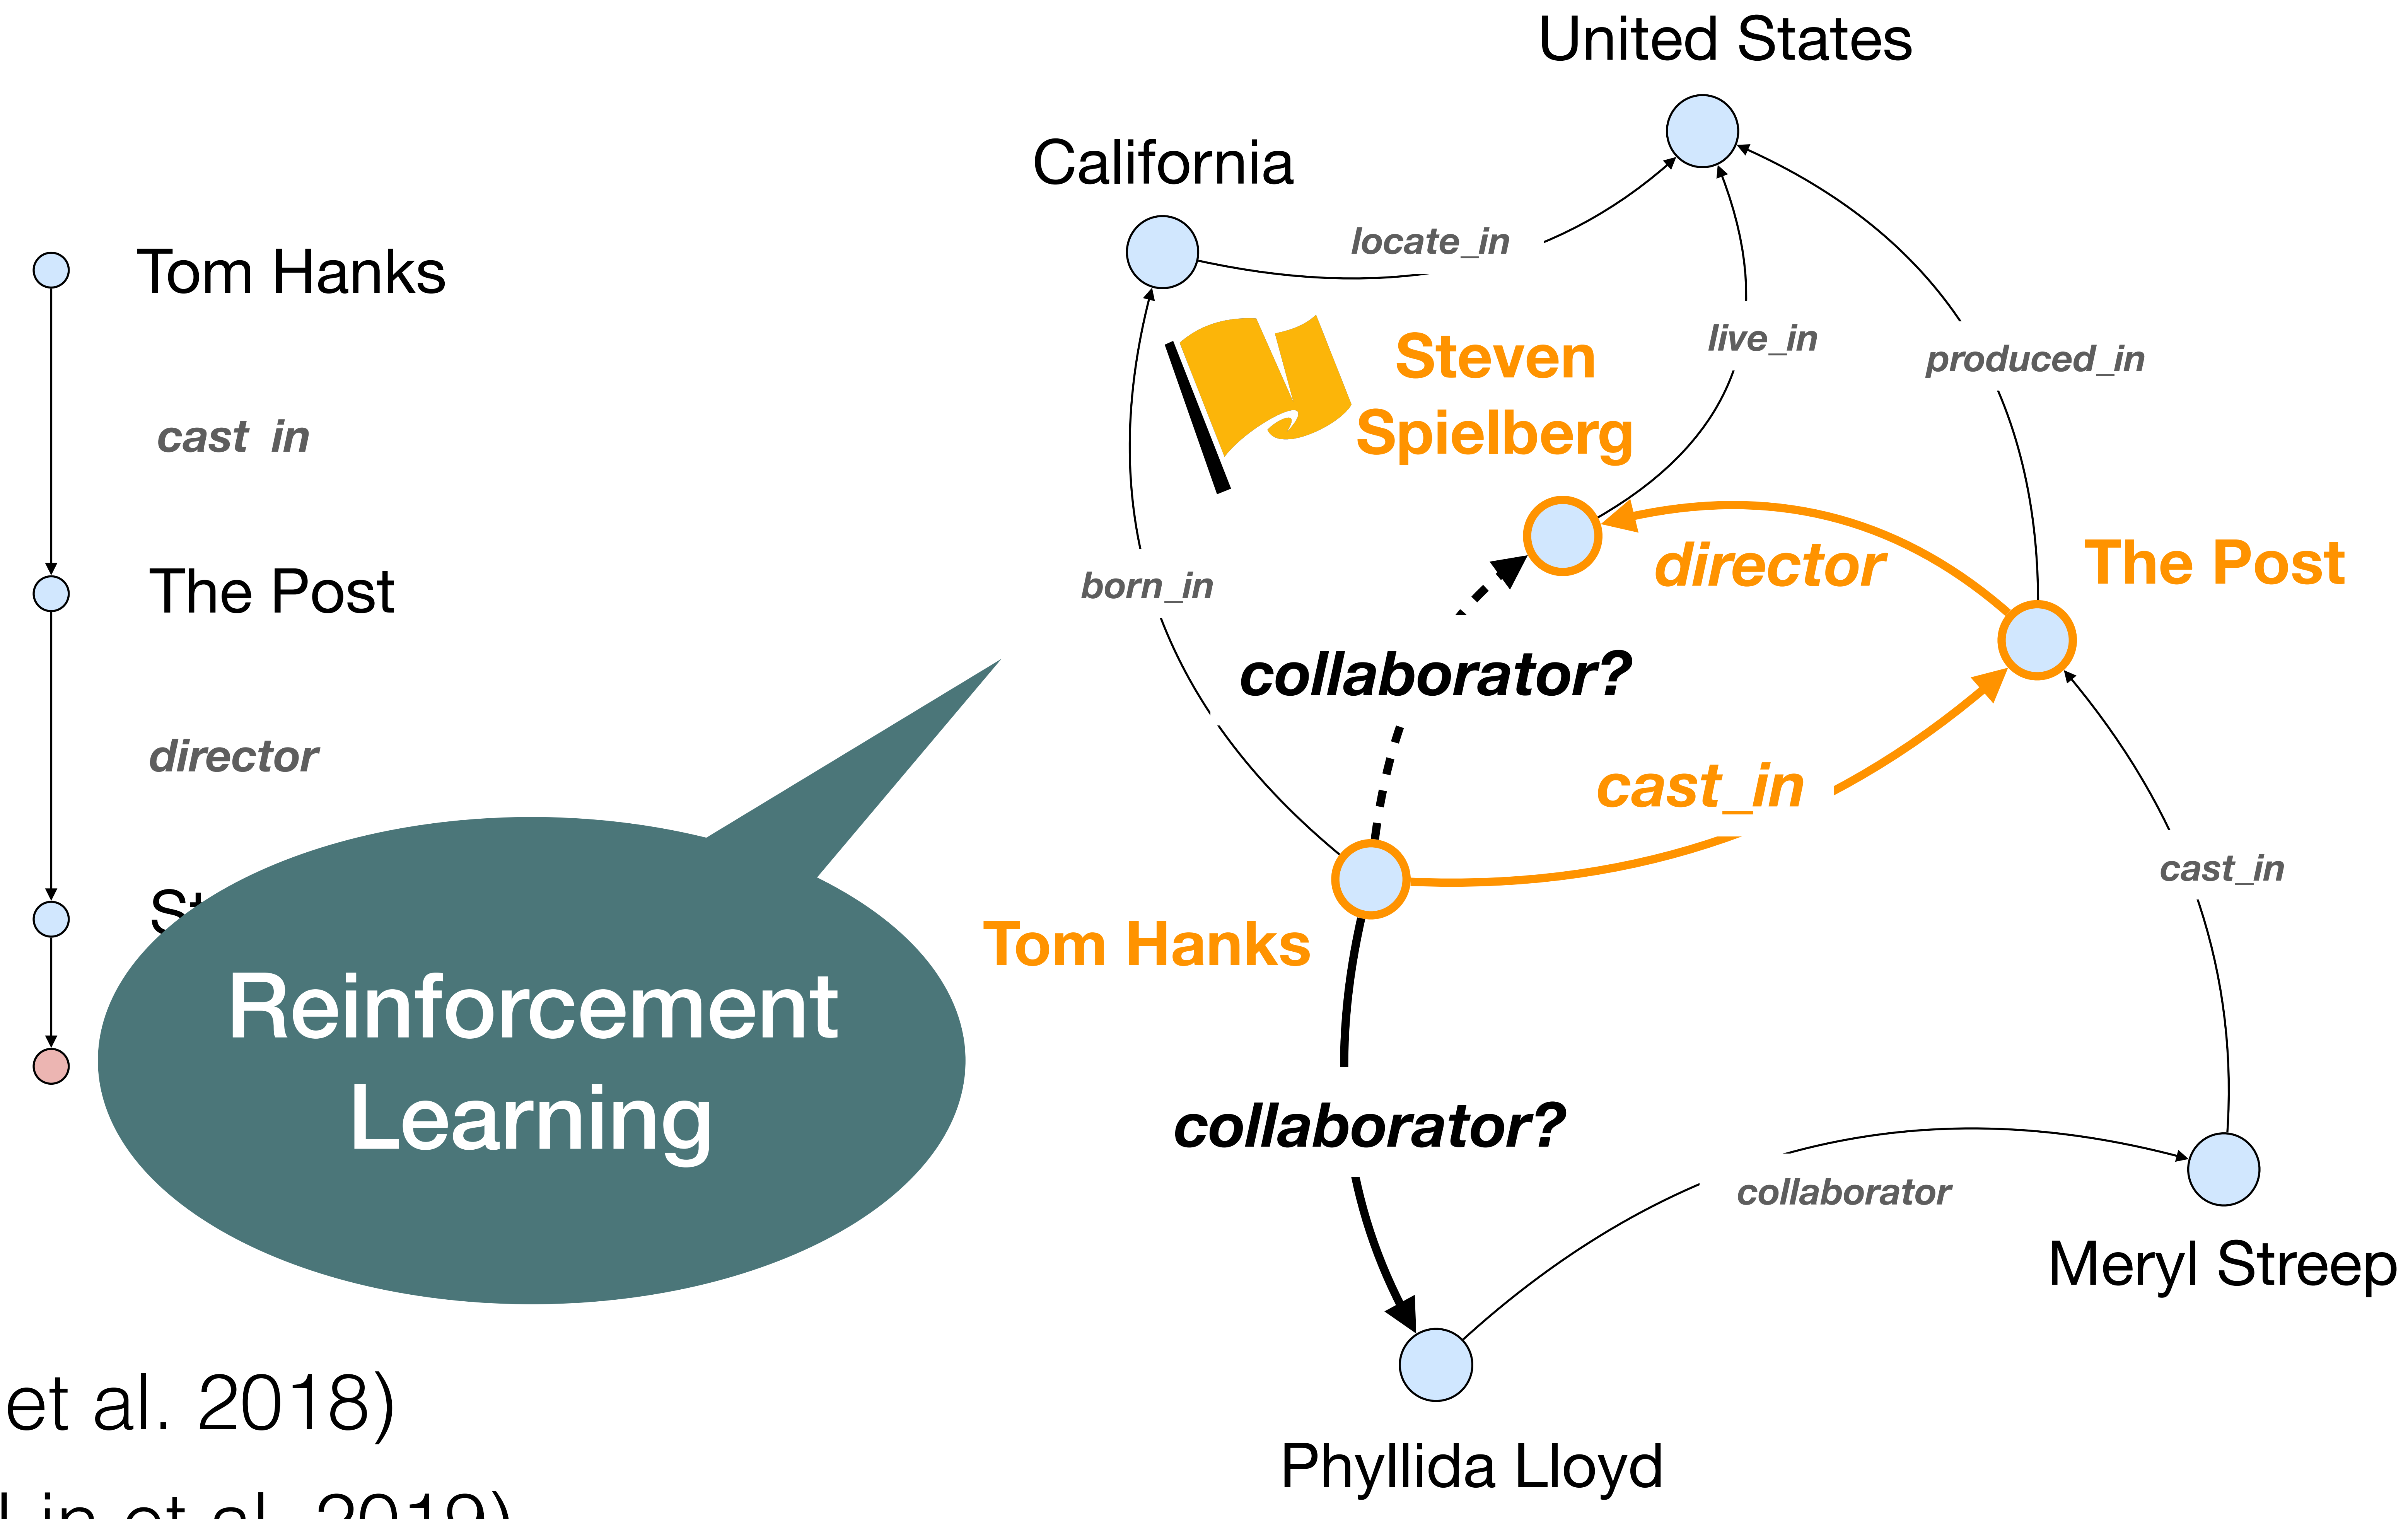

MINERVA (Das et al. 2018)

MINERVA+RS (Lin et al. 2019)

Multi-Hop Reasoning Models

Interpretable

Competitive Accuracy

	MRR
MINERVA	0.825
MINERVA+RS	0.94

Tab 2. MINERVA and MINERVA + Reward Shaping fact inference performance on the UMLS benchmark dataset (Das et. al. 2018; Lin et. al 2018)

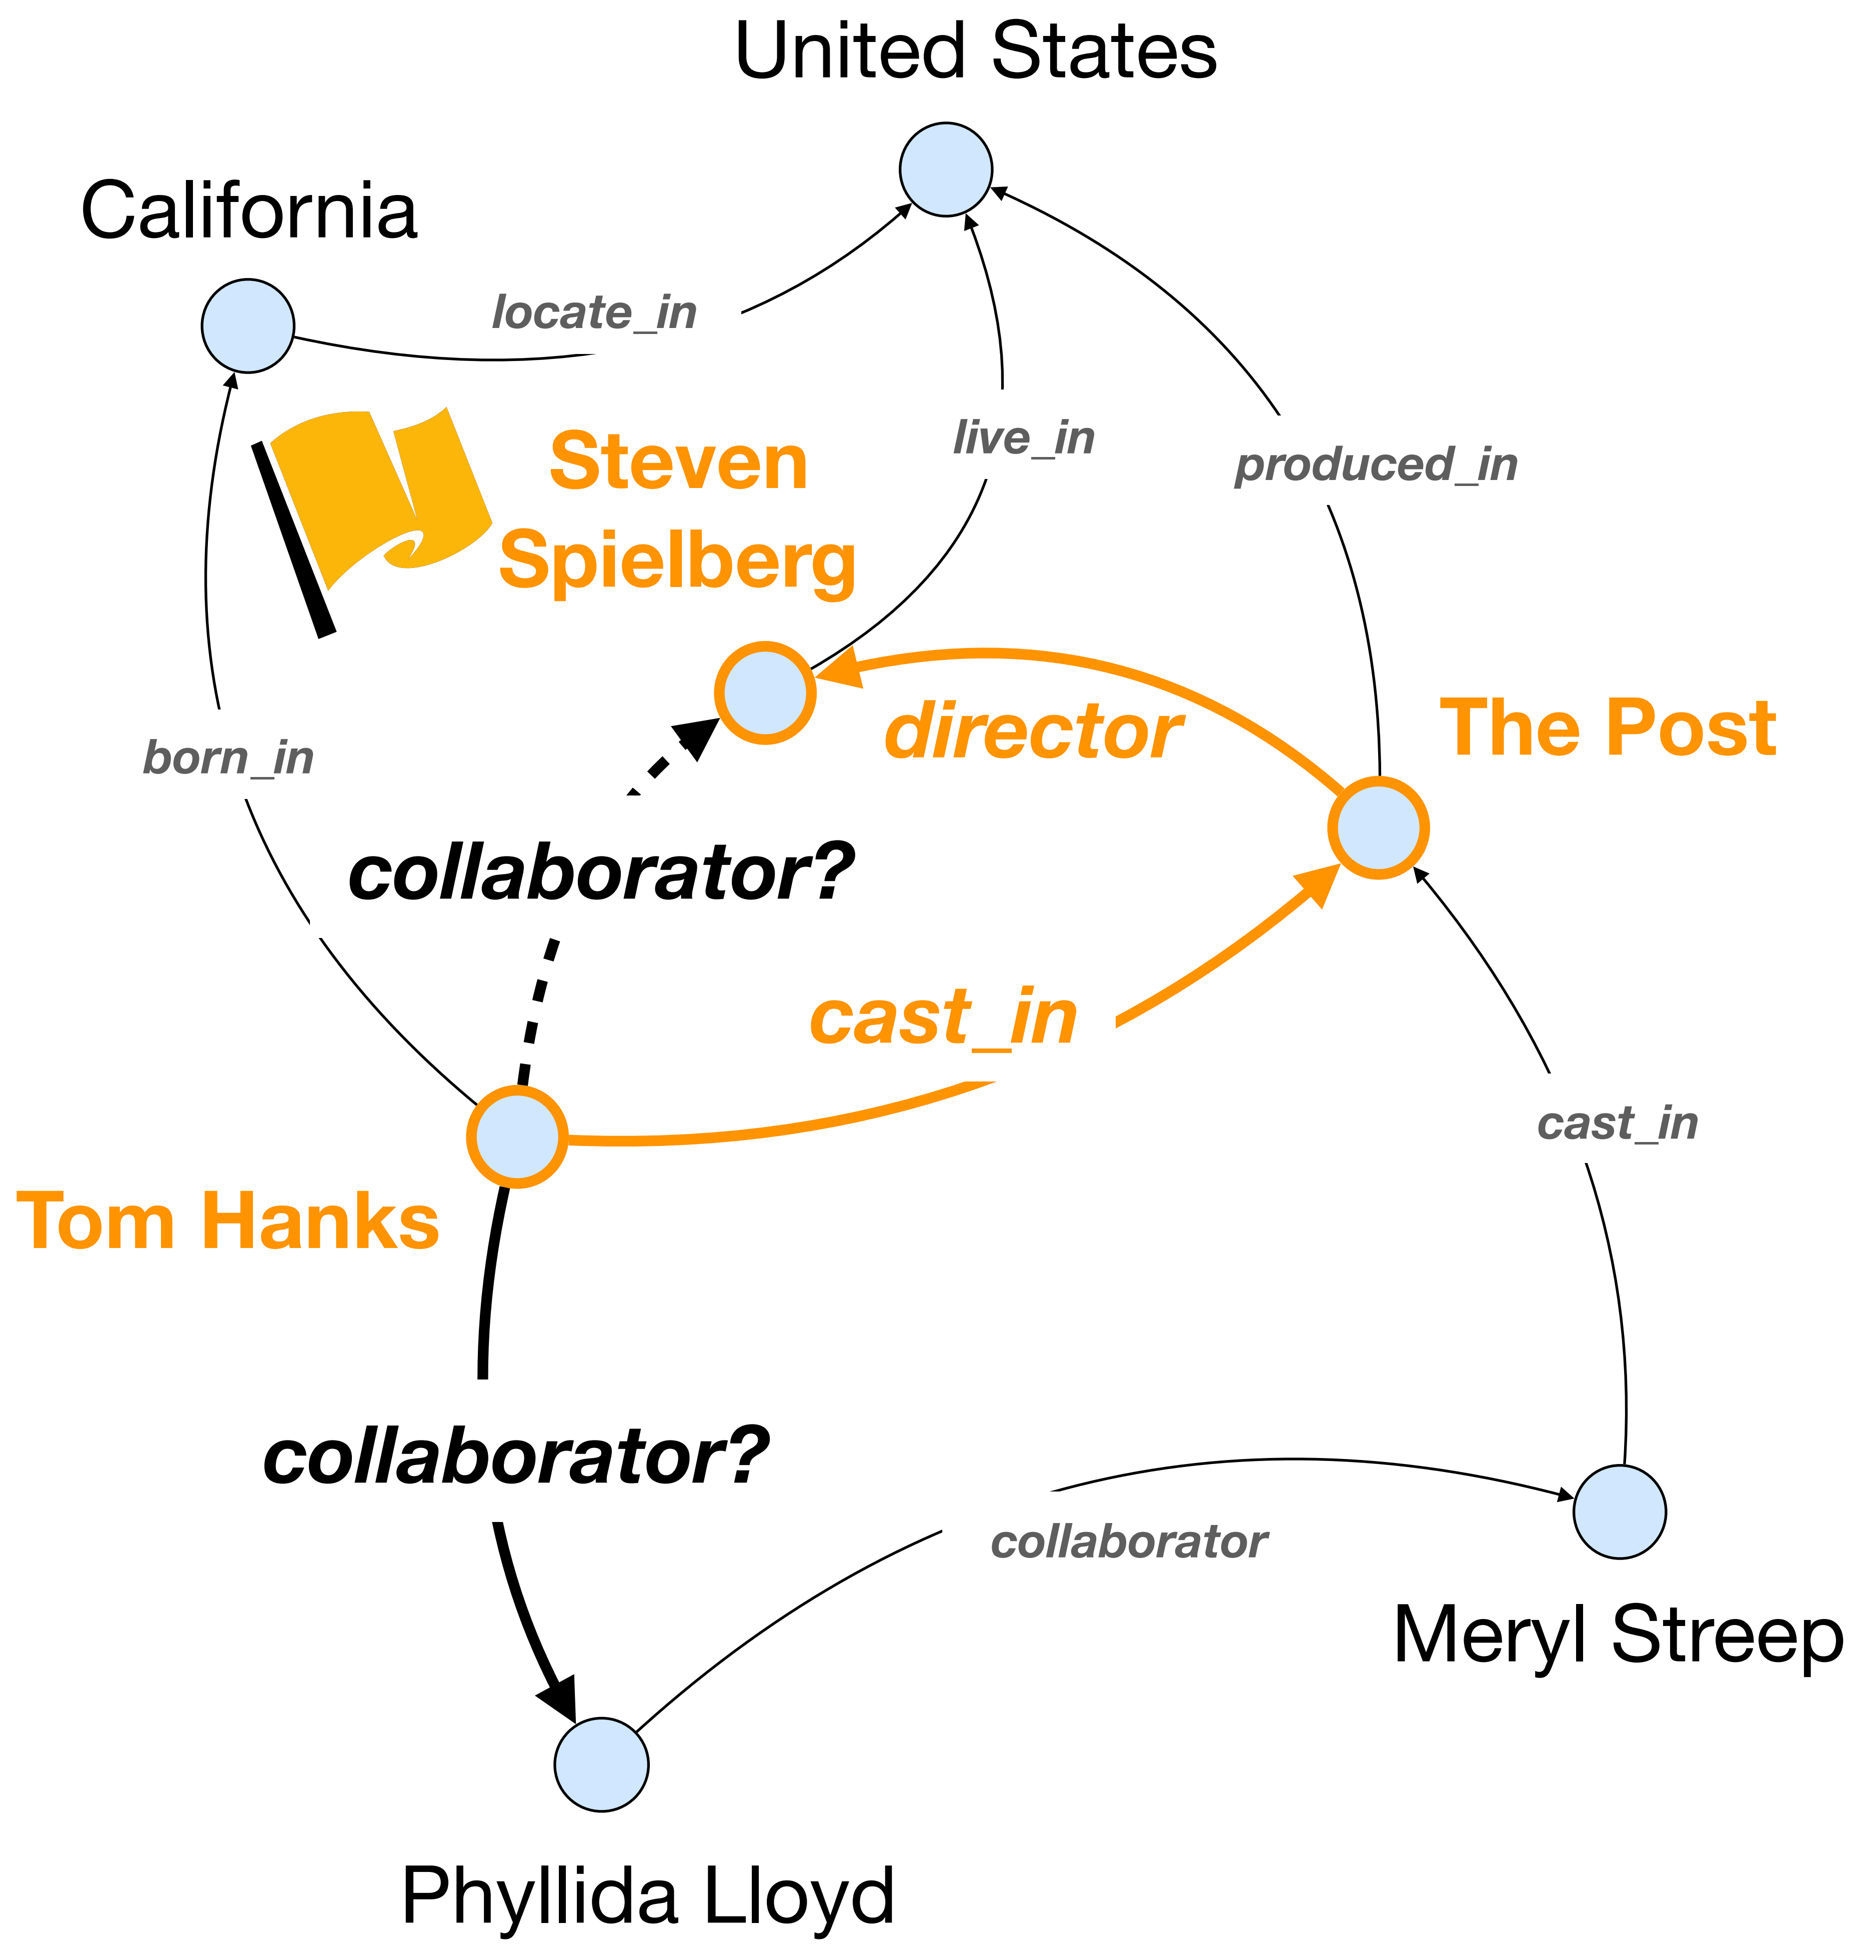

Reinforcement Learning Framework

Environment | State | Action | Transition | Reward

$e_s \xrightarrow{r_q}$

Reinforcement Learning Framework

Environment | **State** | Action | Transition | Reward

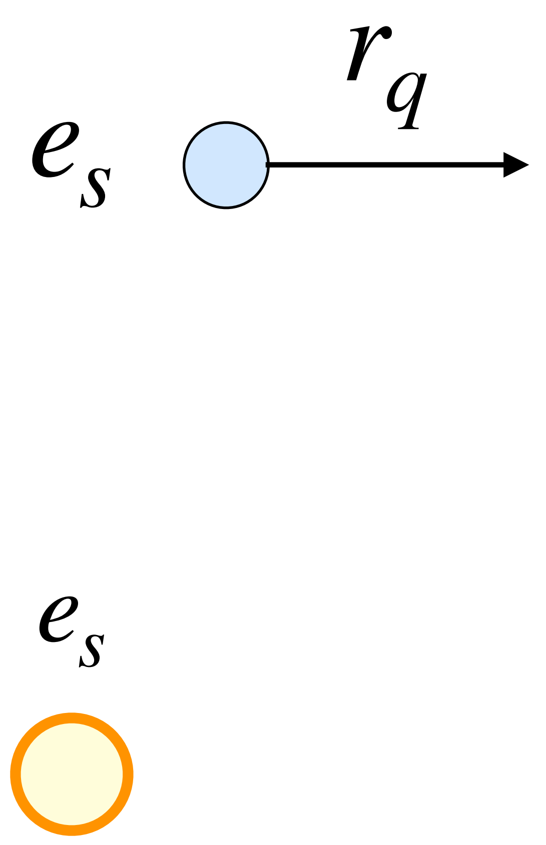

Reinforcement Learning Framework

Environment | State | **Action** | Transition | Reward

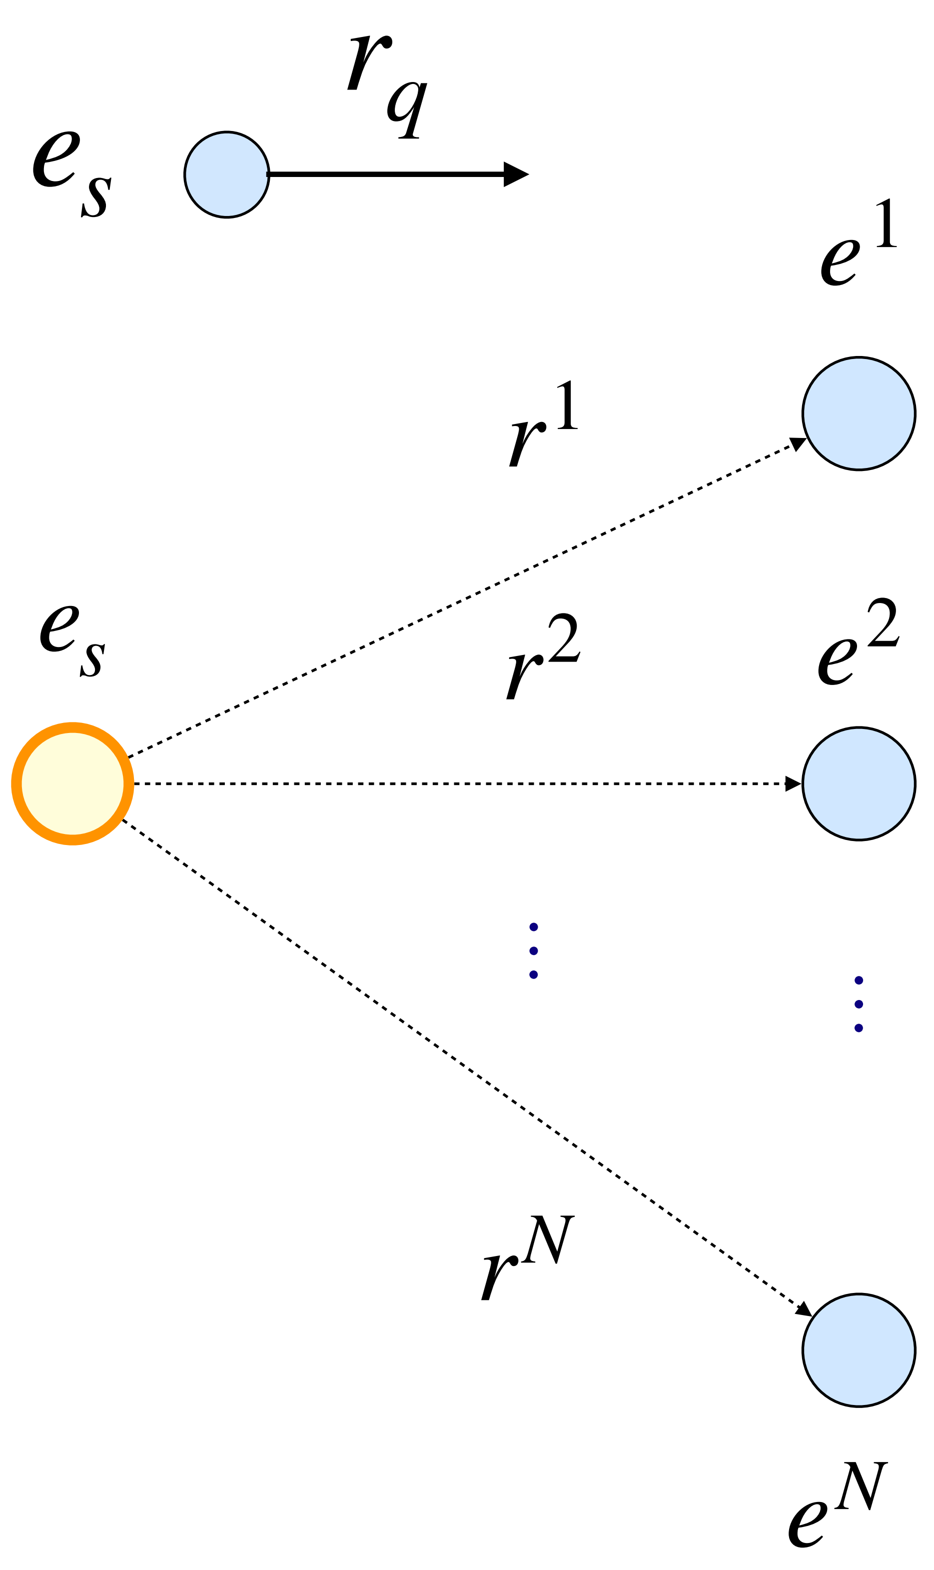

A_t

Reinforcement Learning Framework

Environment | State | **Action** | Transition | Reward

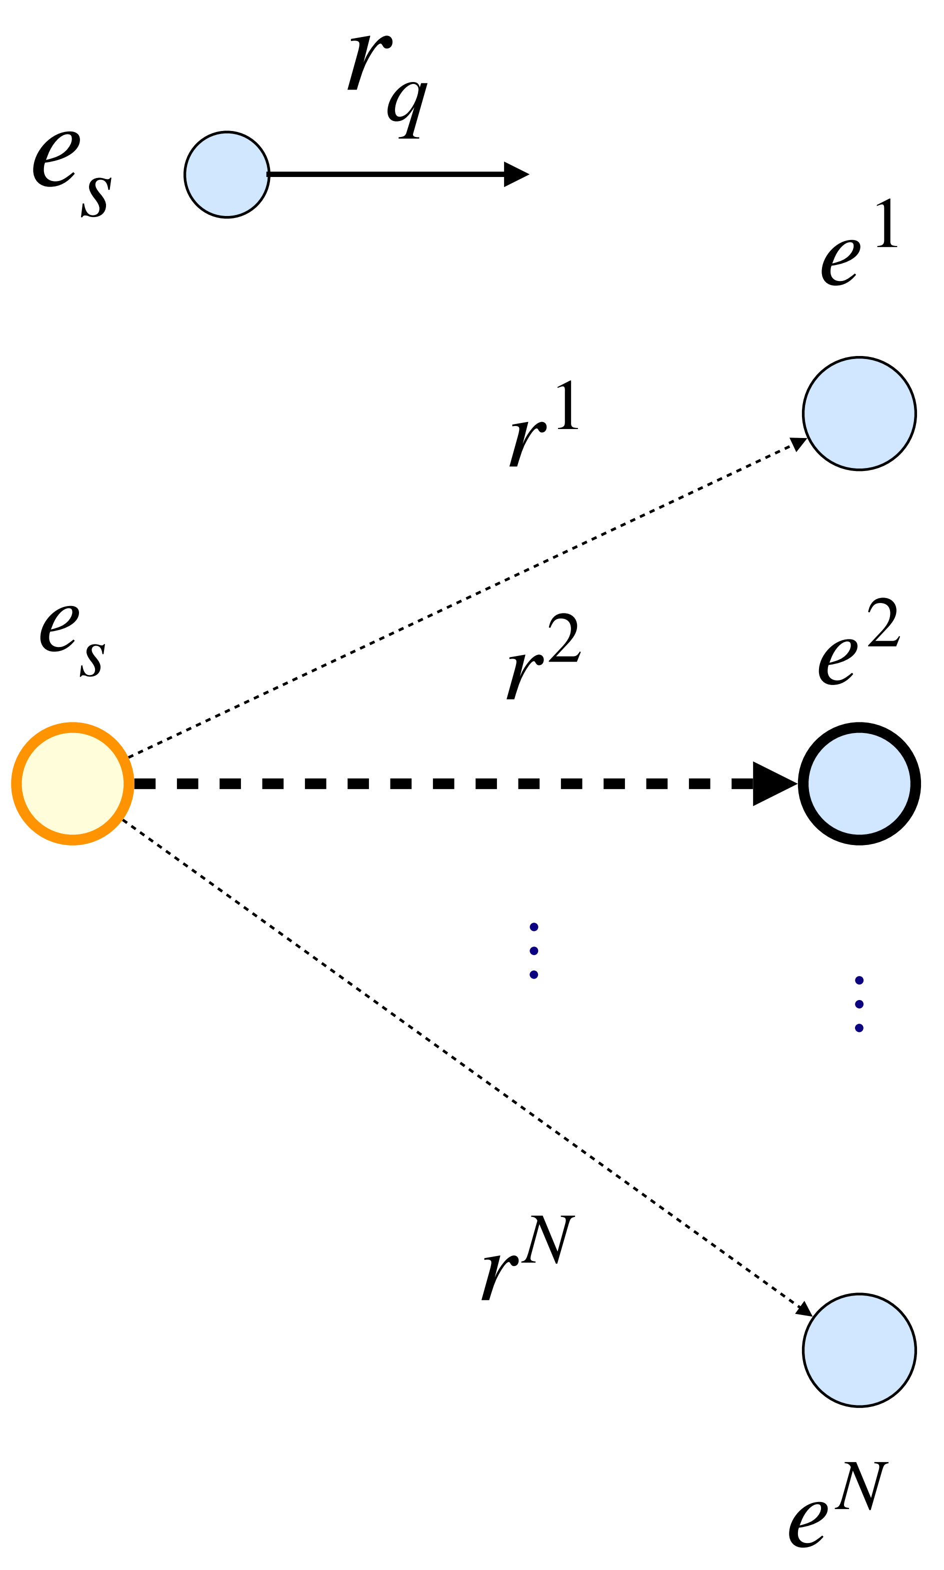

A_t

Action space: All nodes with incoming edges from e_s in the graph

Reinforcement Learning Framework

Environment | State | Action | **Transition** | Reward

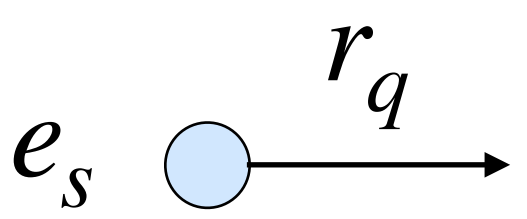

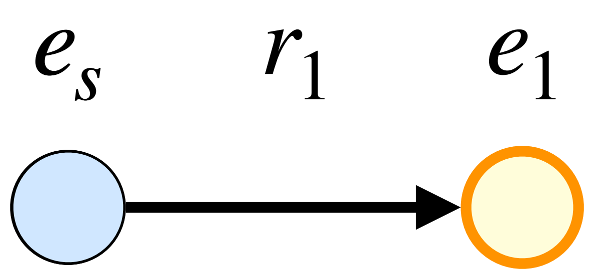

Reinforcement Learning Framework

Environment | State | Action | **Transition** | Reward

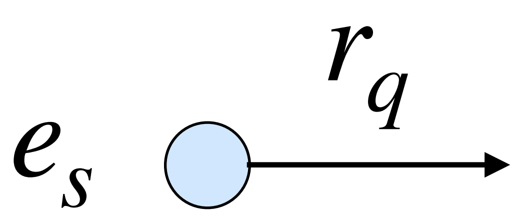

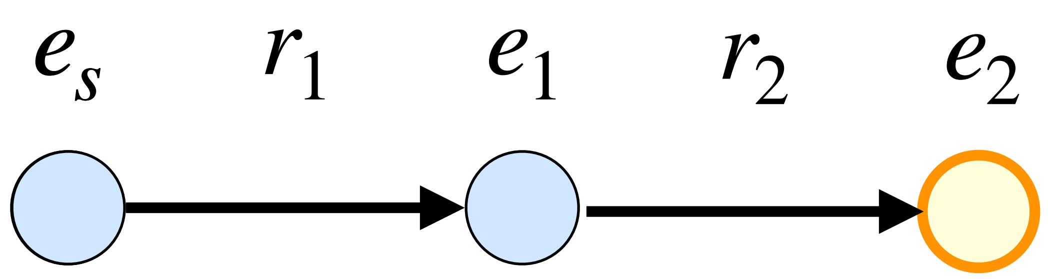

Reinforcement Learning Framework

Environment | State | Action | **Transition** | Reward

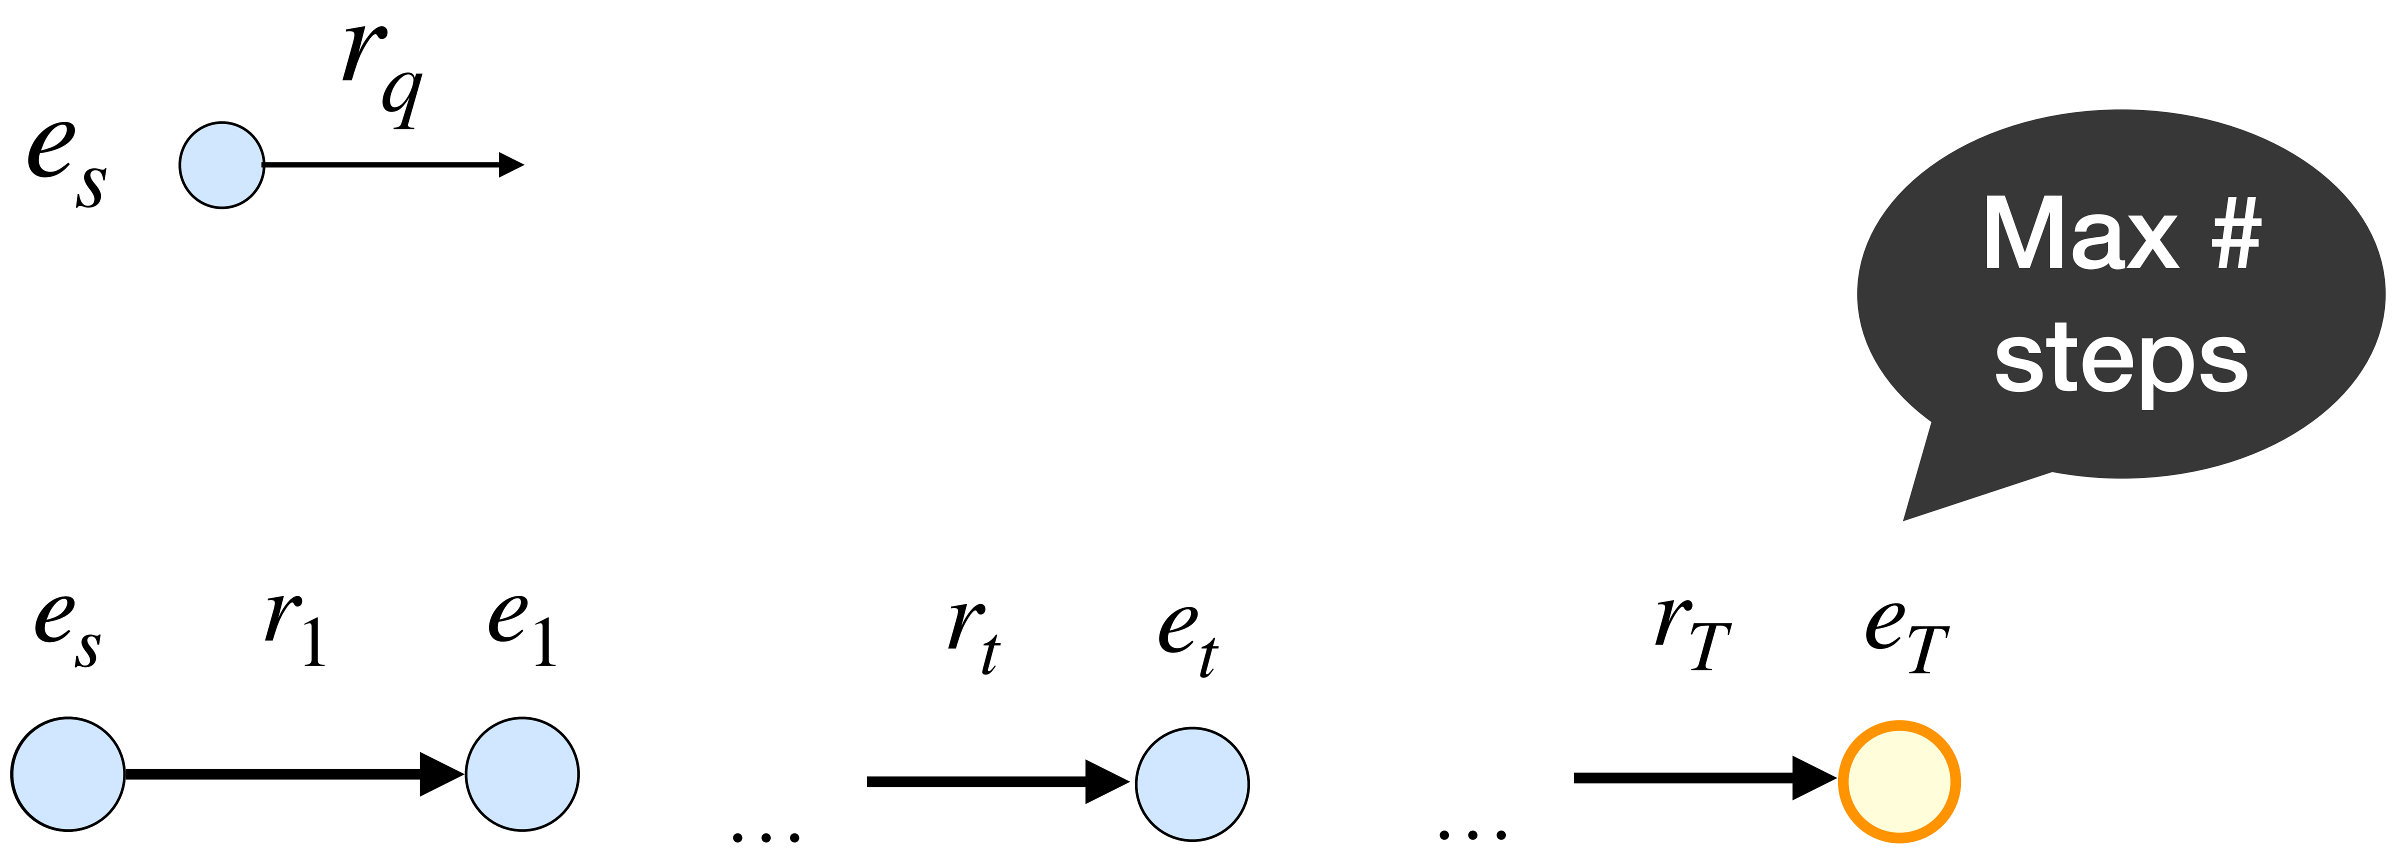

Reinforcement Learning Framework

Environment | State | Action | **Transition** | Reward

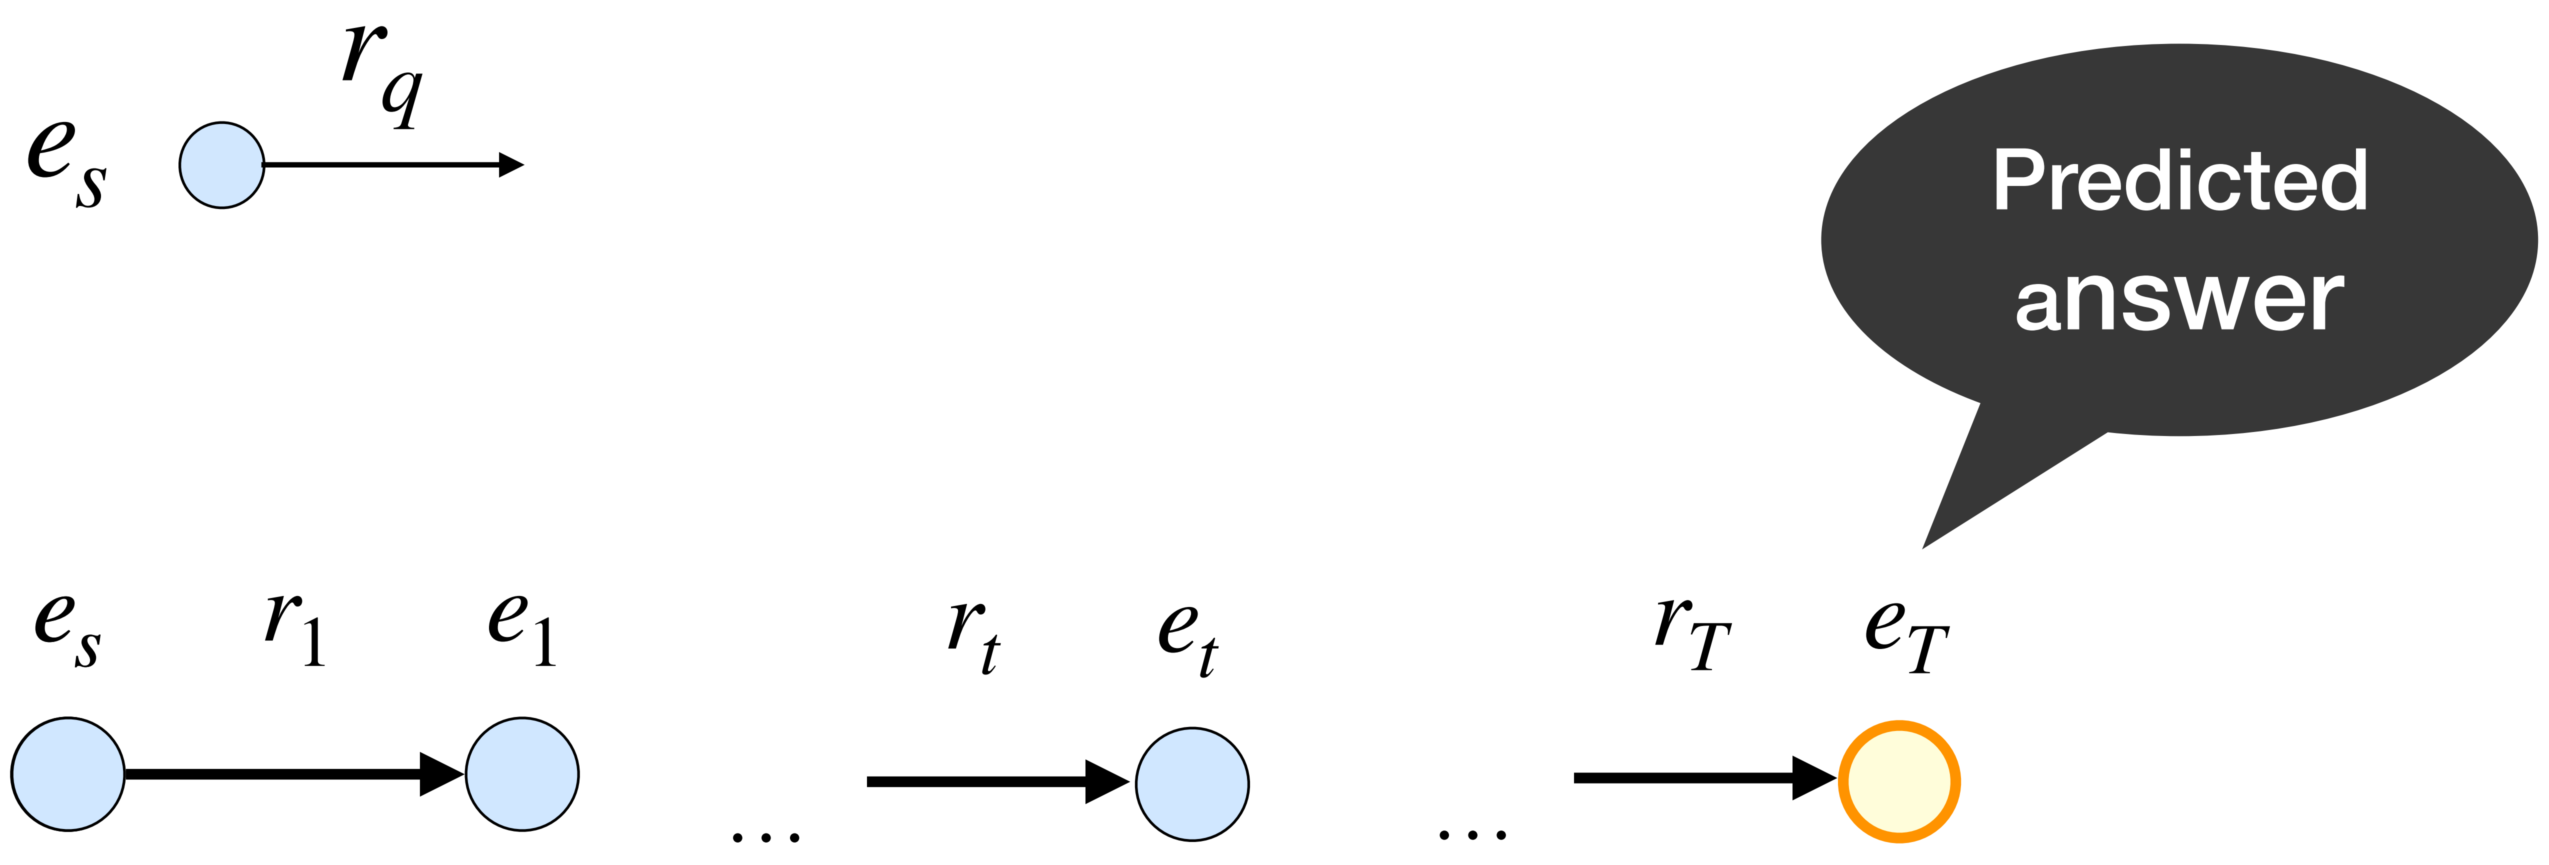

Reinforcement Learning Framework

Environment | State | Action | Transition | **Reward**

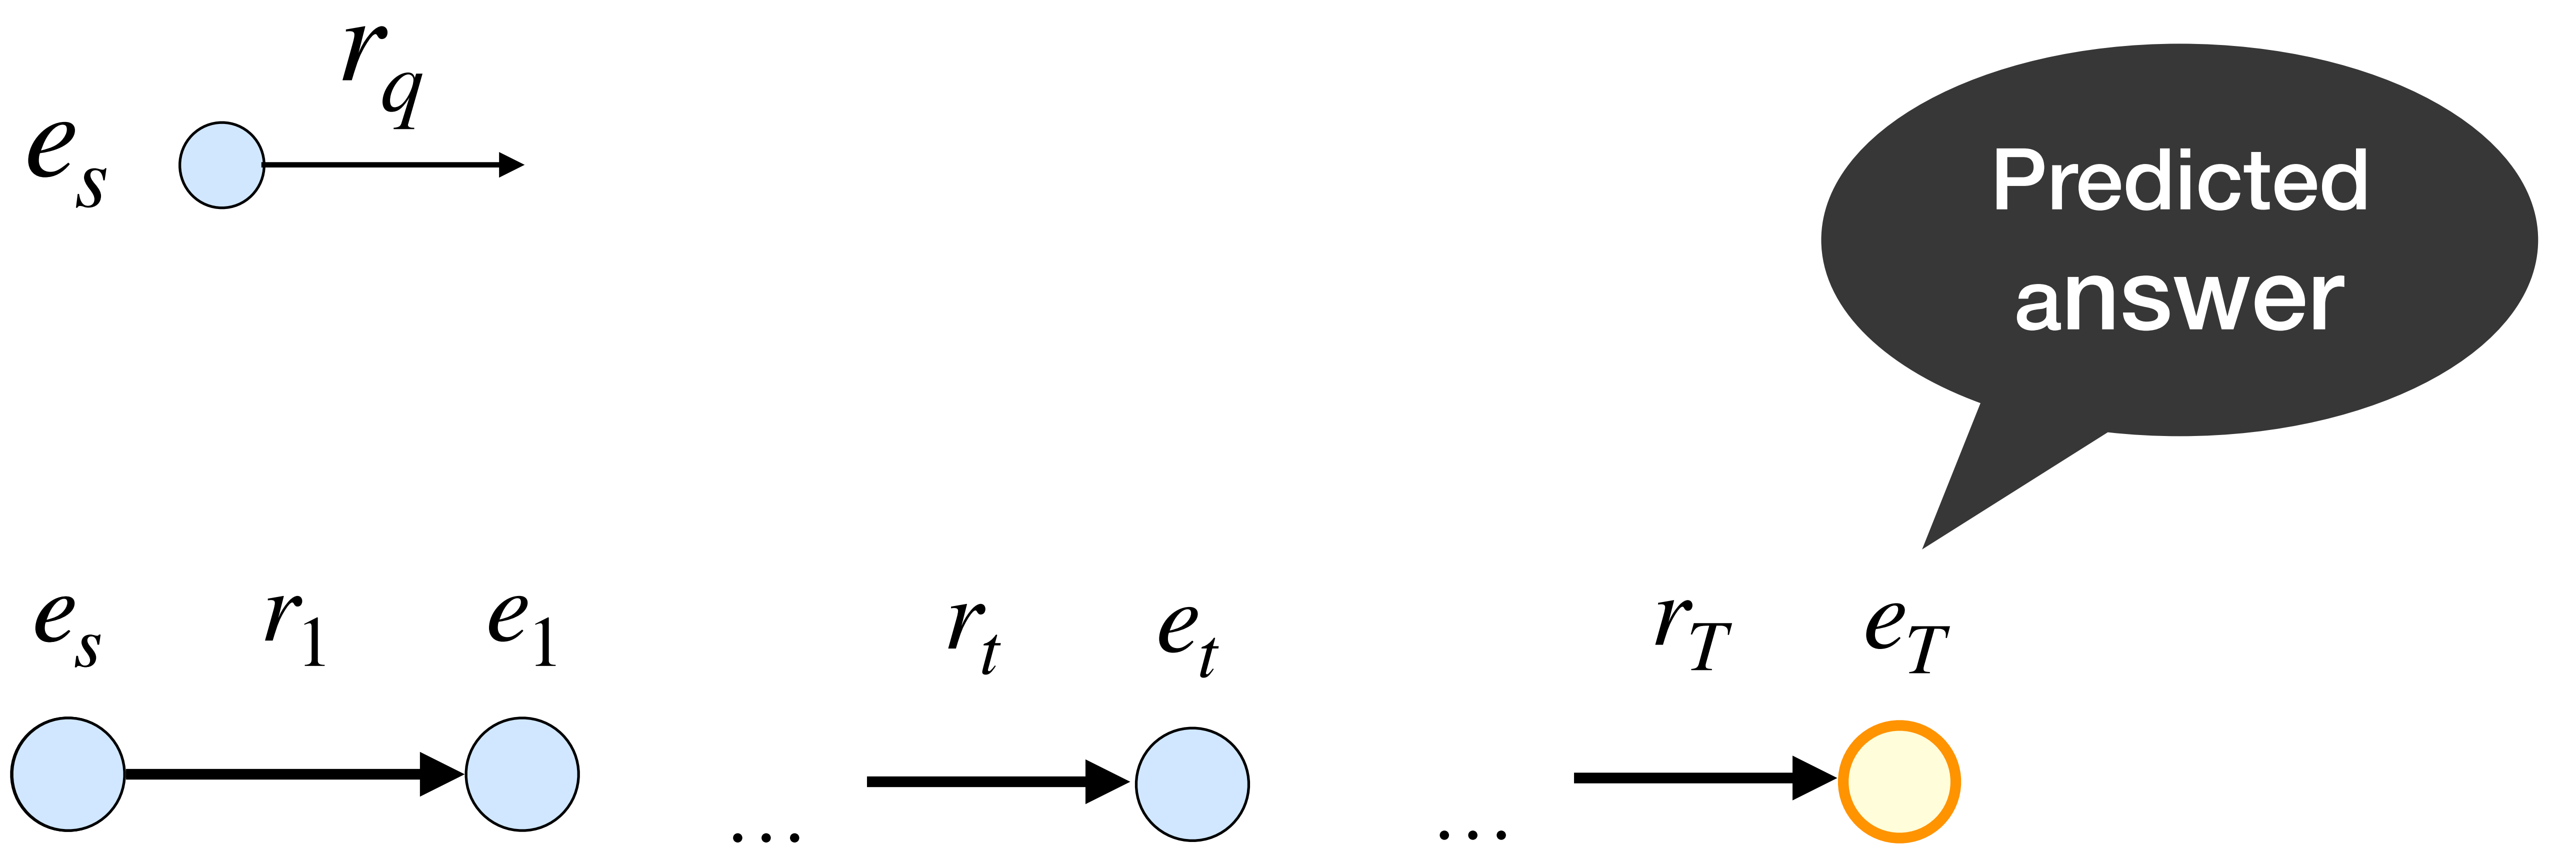

$$\boxed{R_b(s_T) = \mathbf{1}\{(e_s, r_q, e_T) \in G\}}$$

Reinforcement Learning Framework

Environment | State | Action | Transition | **Reward**

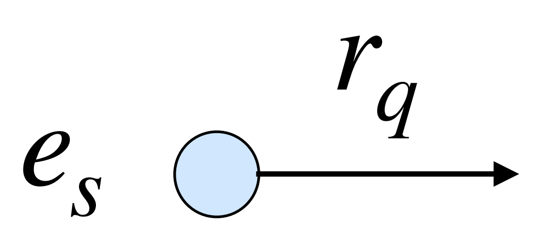

Policy gradient

Learn ***which action to choose*** given a state

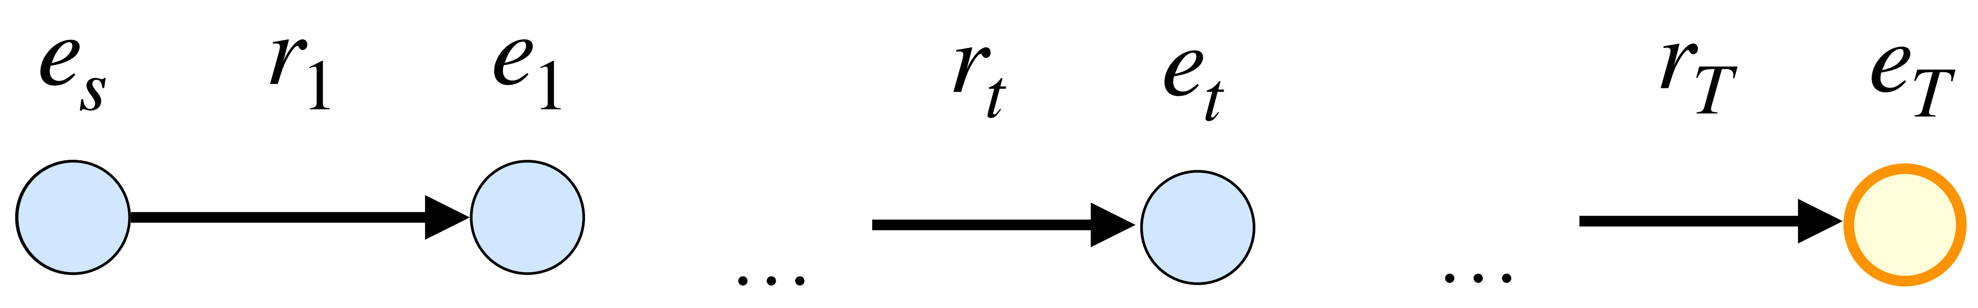

$$R_b(s_T) = \mathbf{1}\{(e_s, r_q, e_T) \in G\}$$

Policy Gradient

Policy function $\pi_\Theta(a_t \mid s_t)$

Probability of choosing an action given the current state

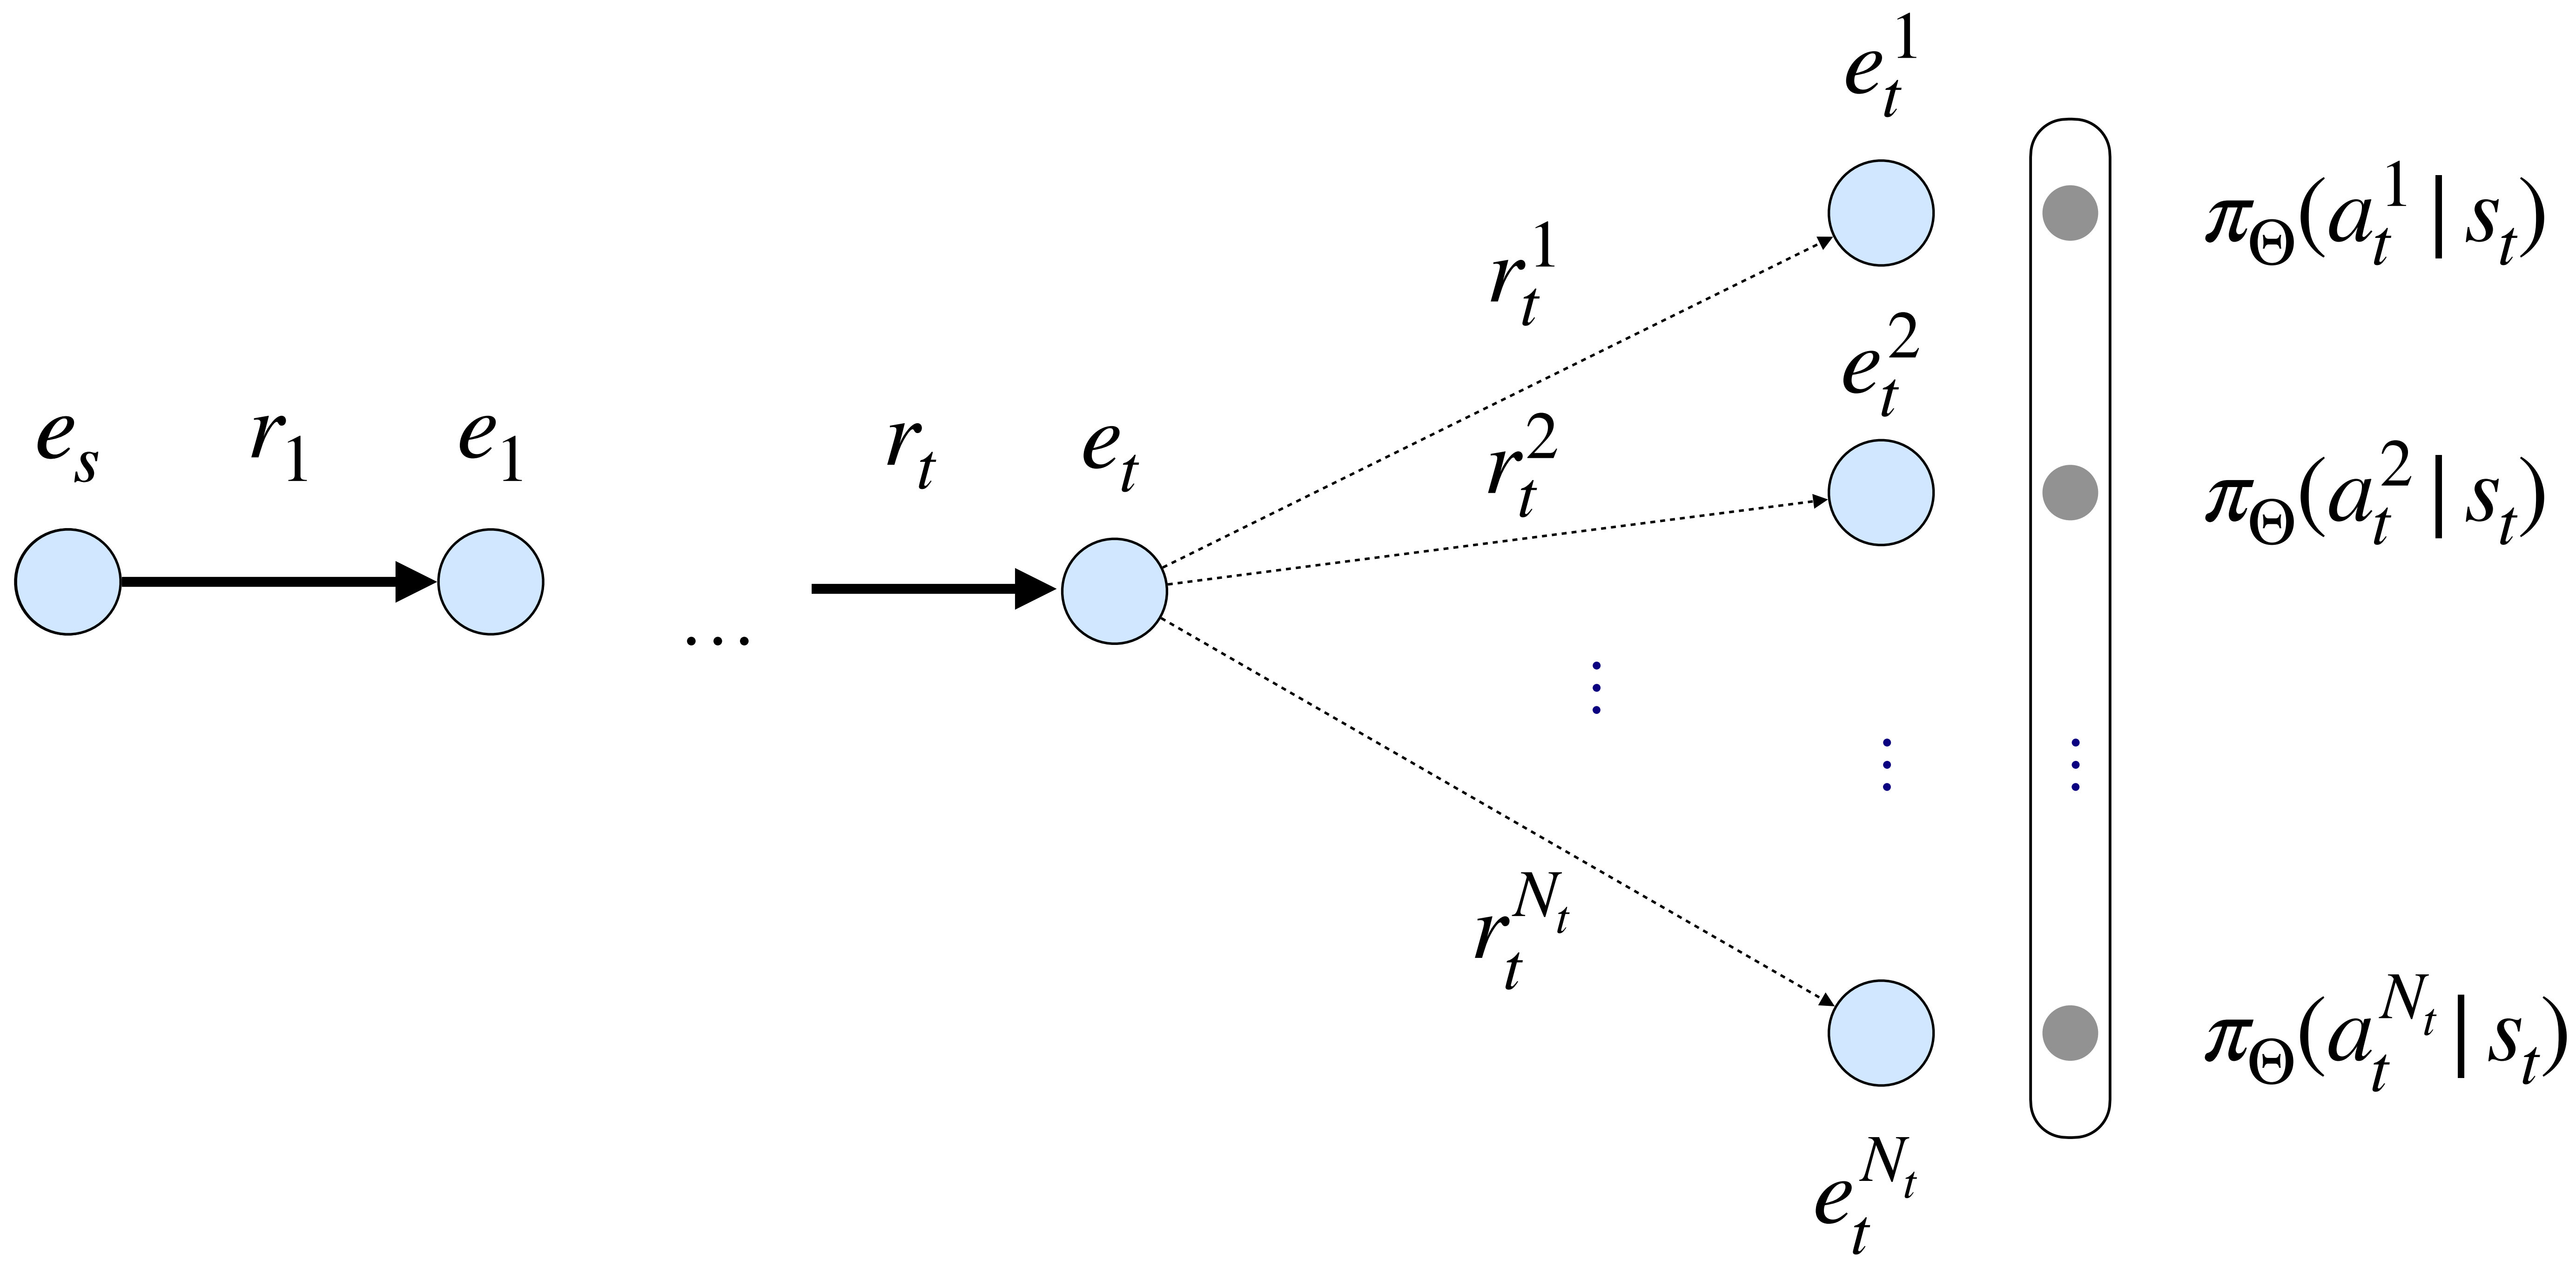

$a_t^i = (r_t^i, e_t^i)$

Policy Gradient

Policy function $\pi_\Theta(a_t \mid s_t)$

Design choice: include search history in the state representation

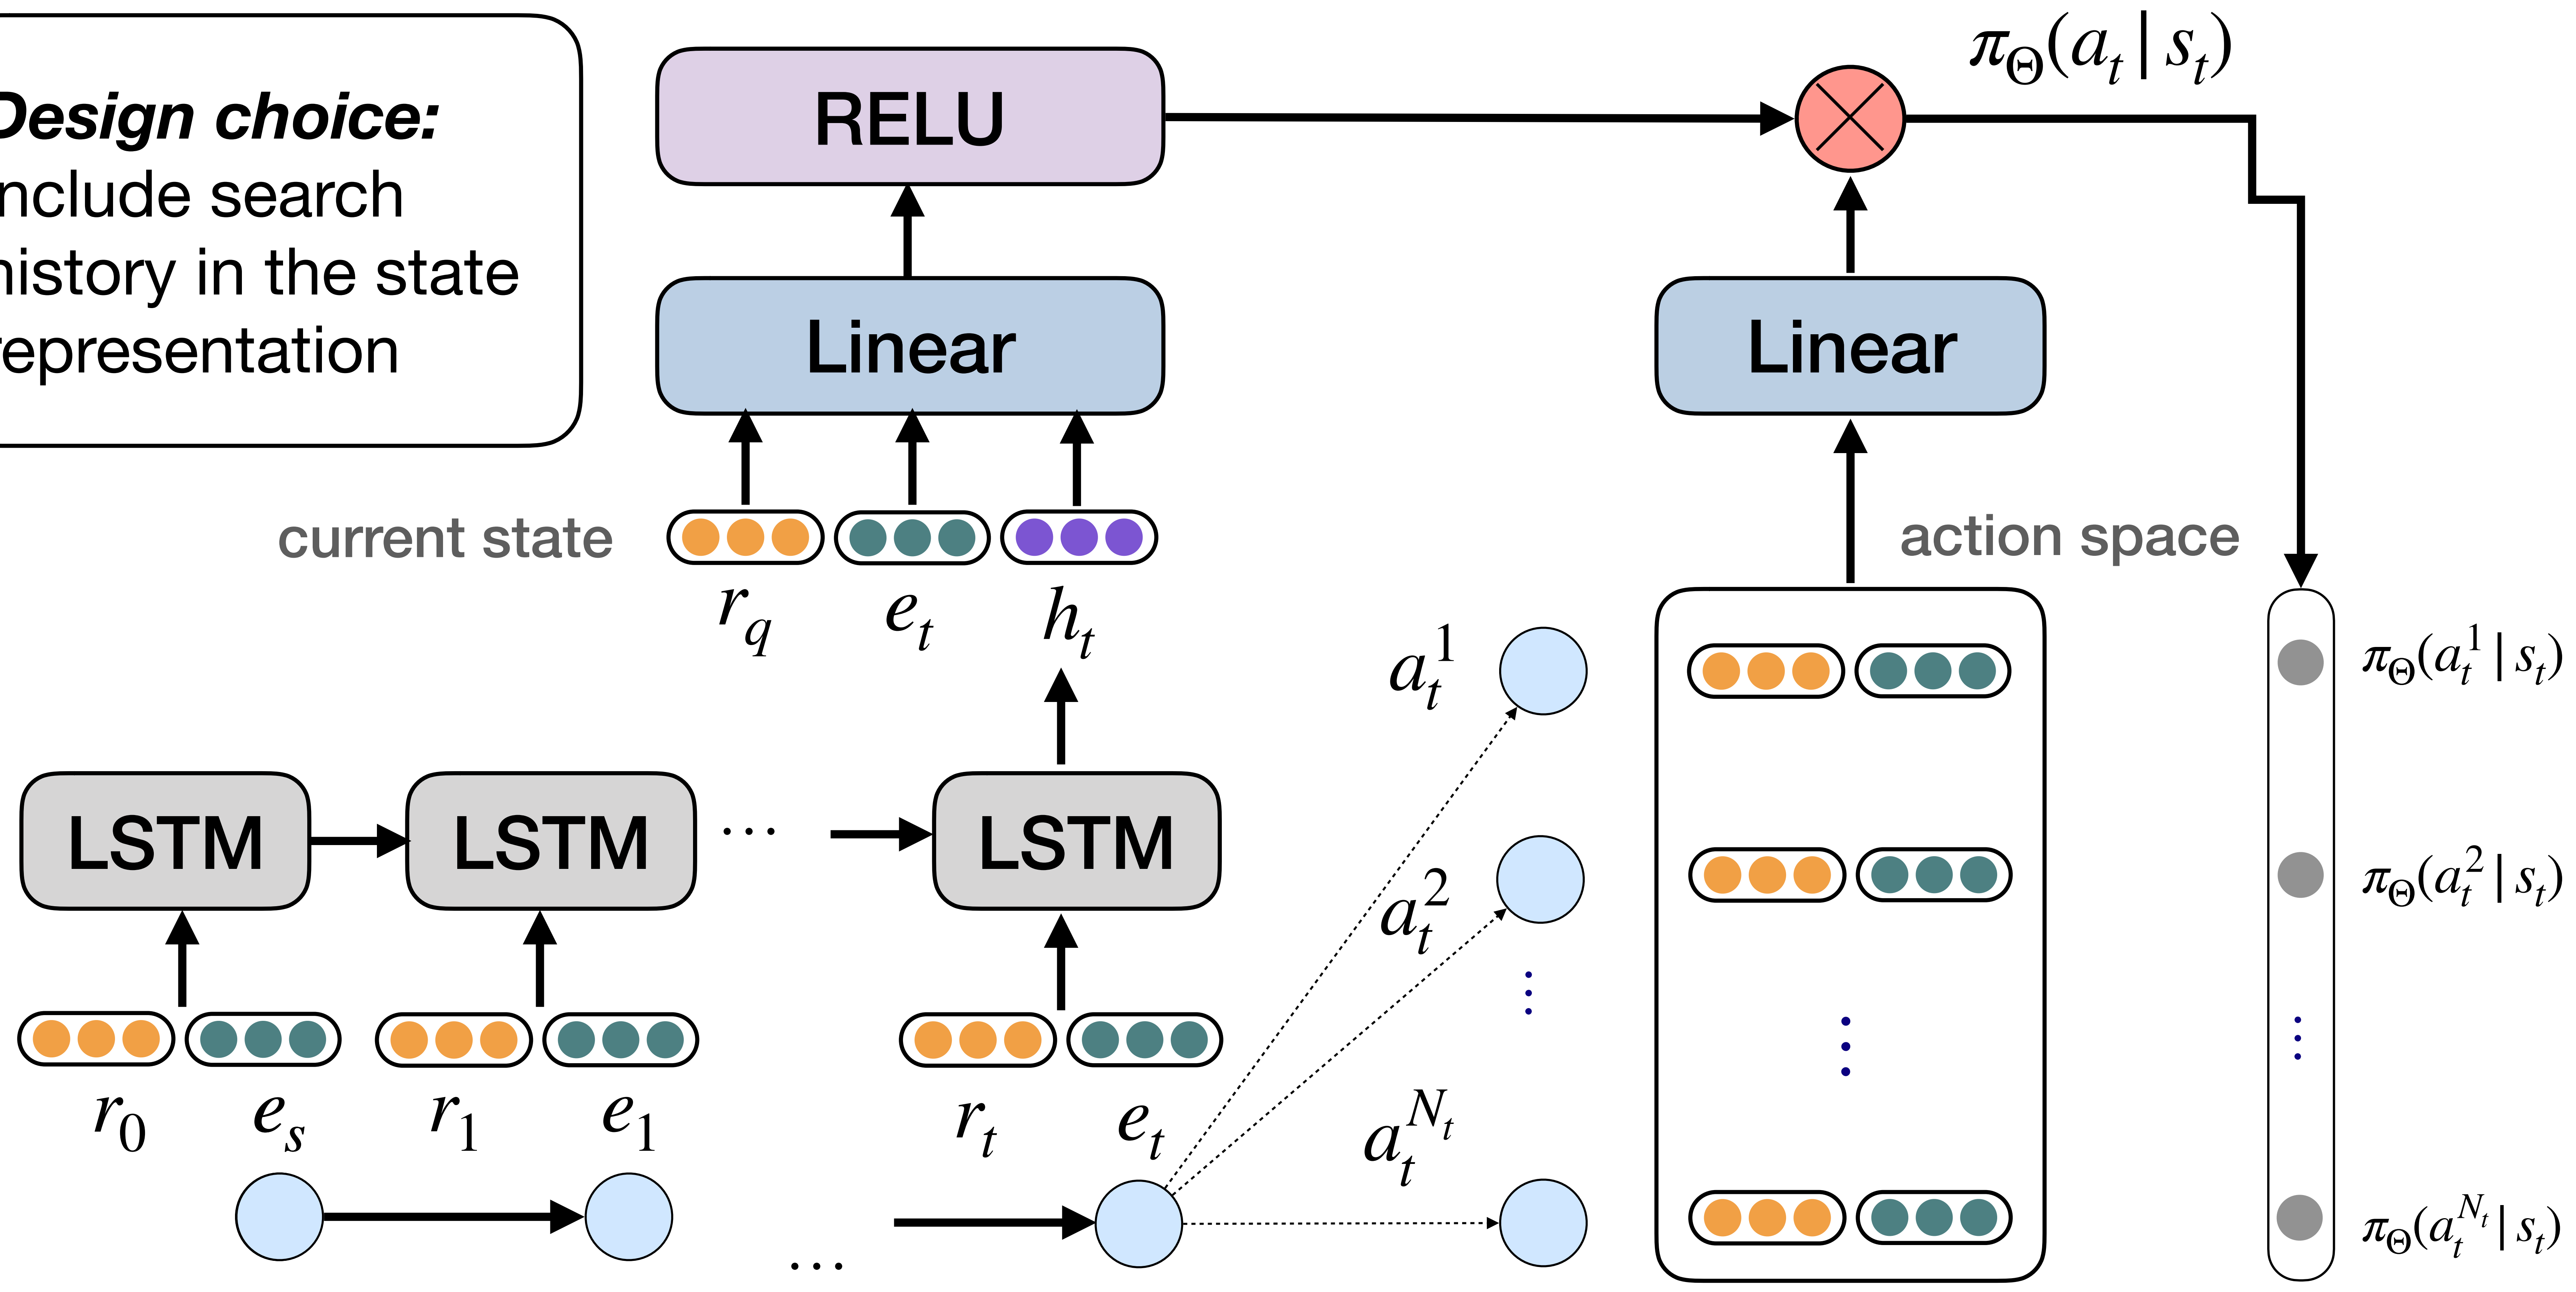

REINFORCE Training

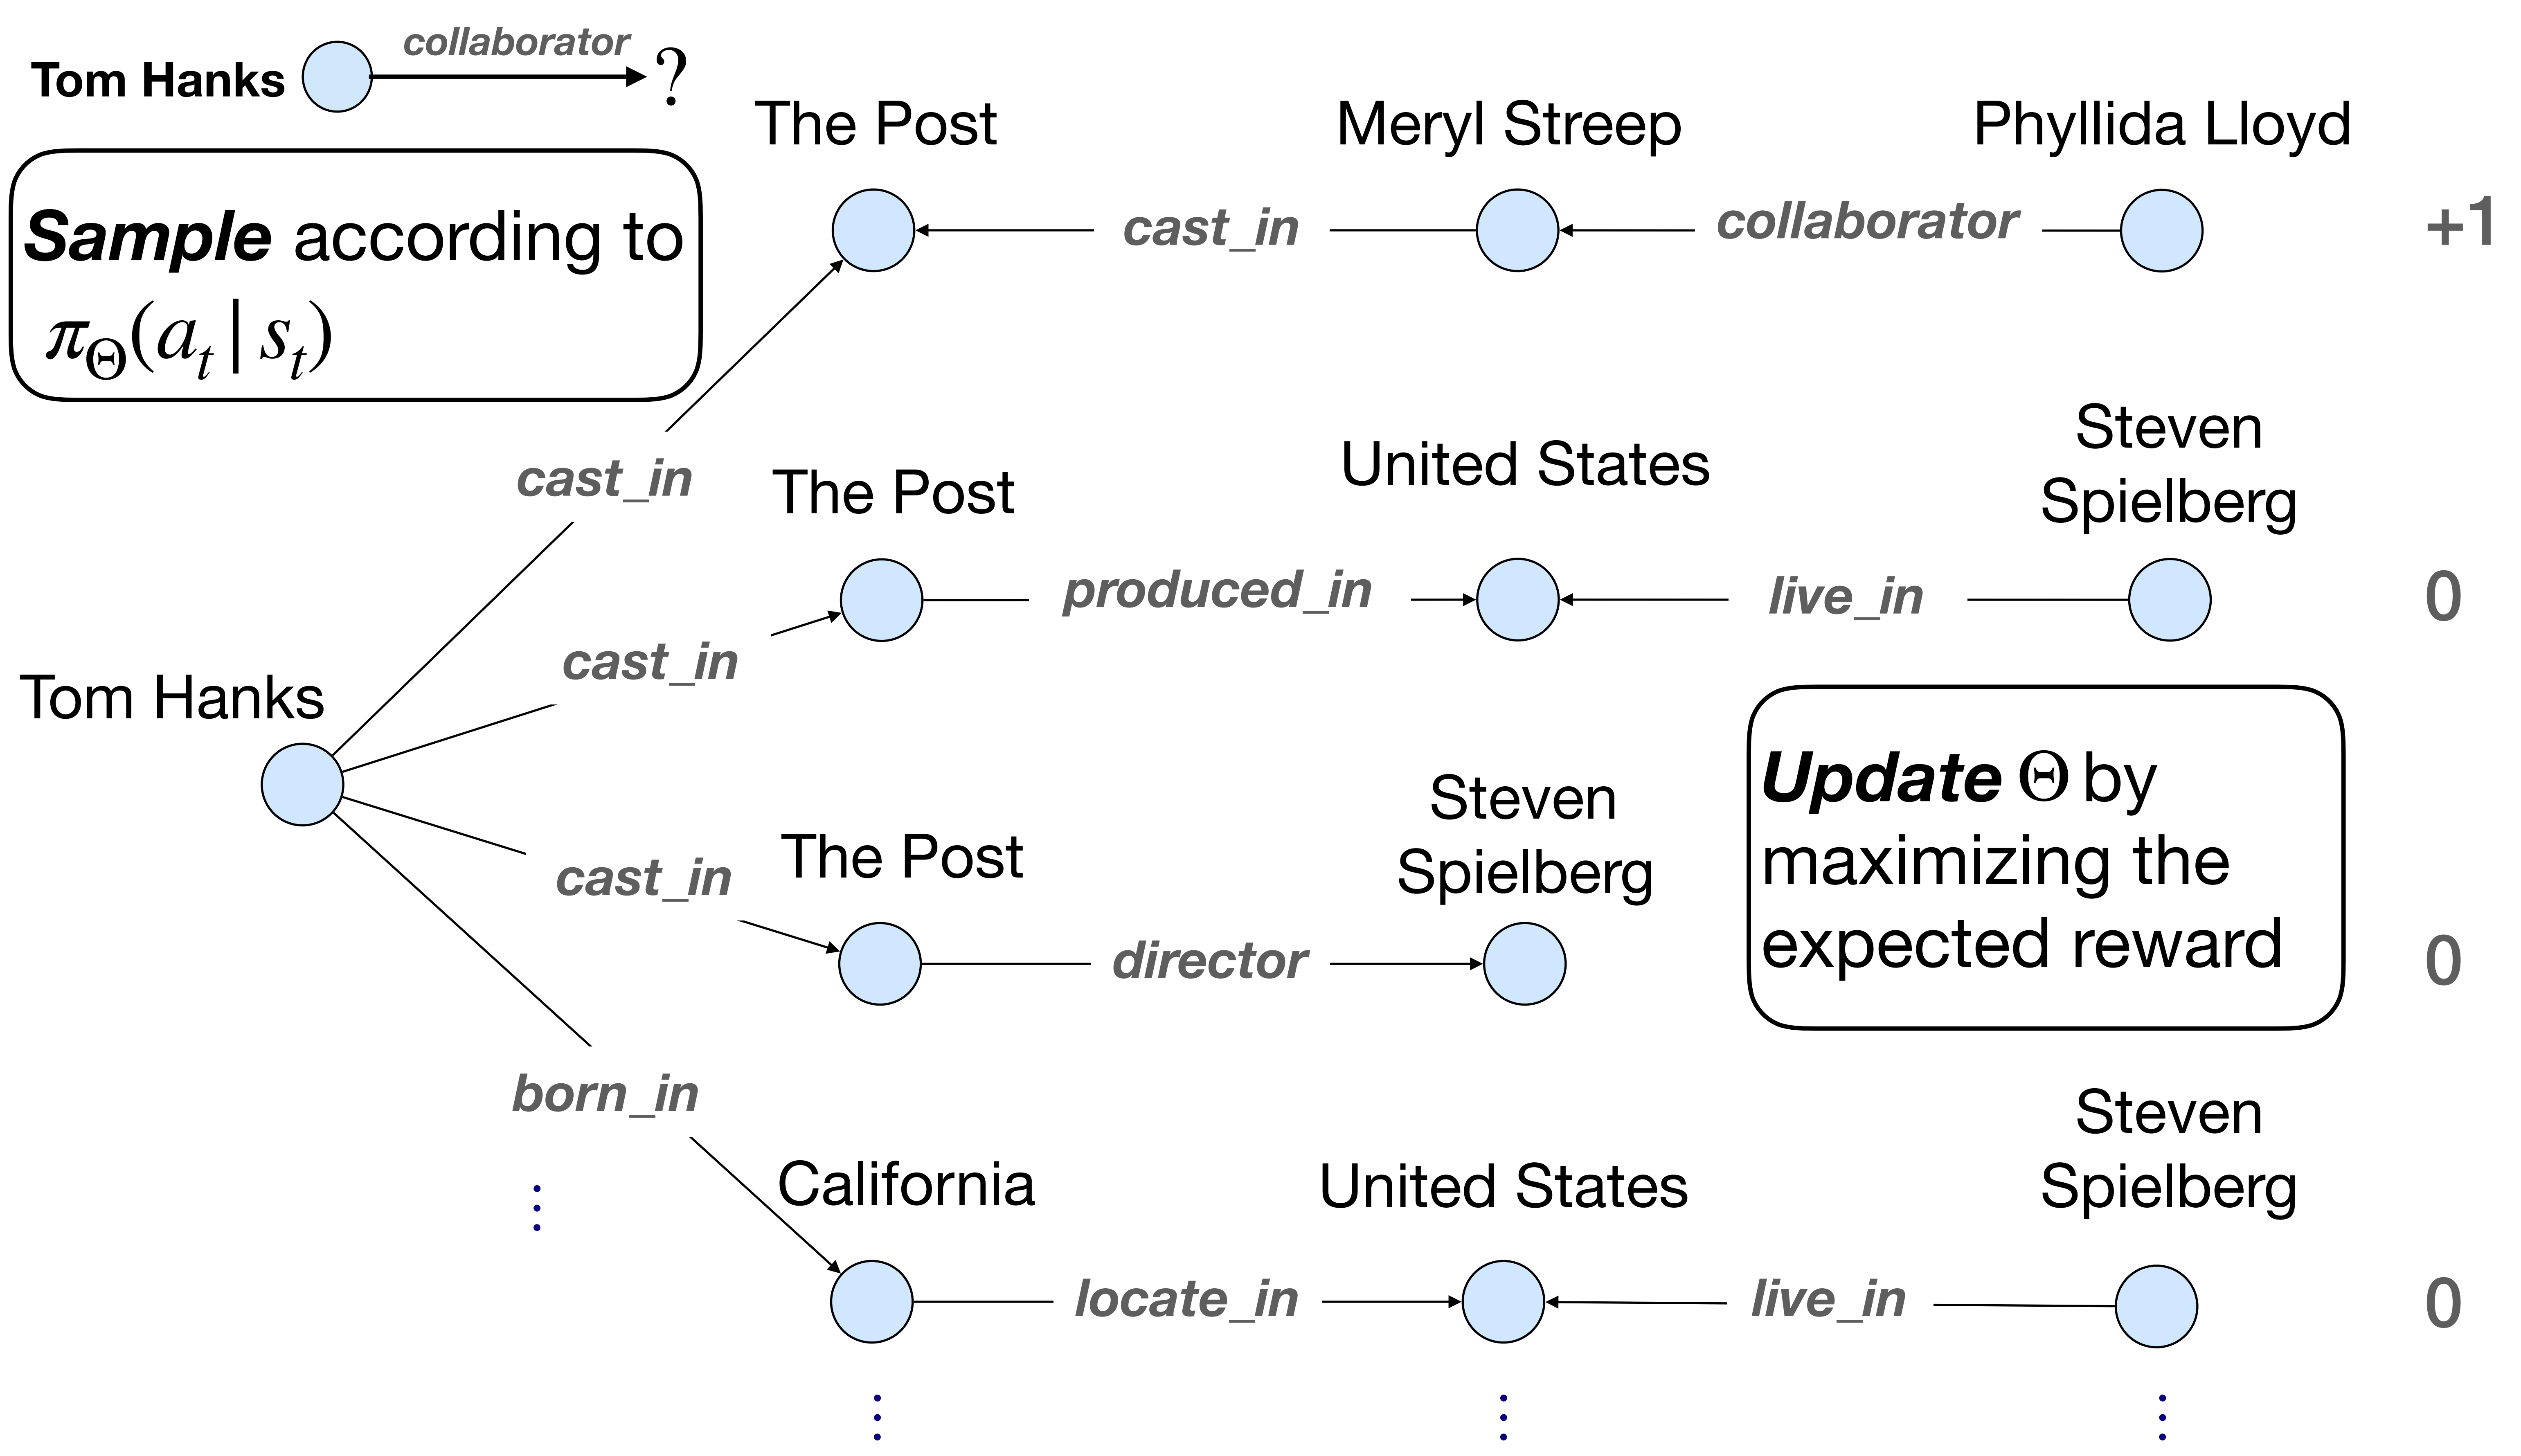

REINFORCE (Williams, 1992)

REINFORCE Training

Issues: ***False-negative*** & ***nearly-correct*** entities ➡ true- negatives

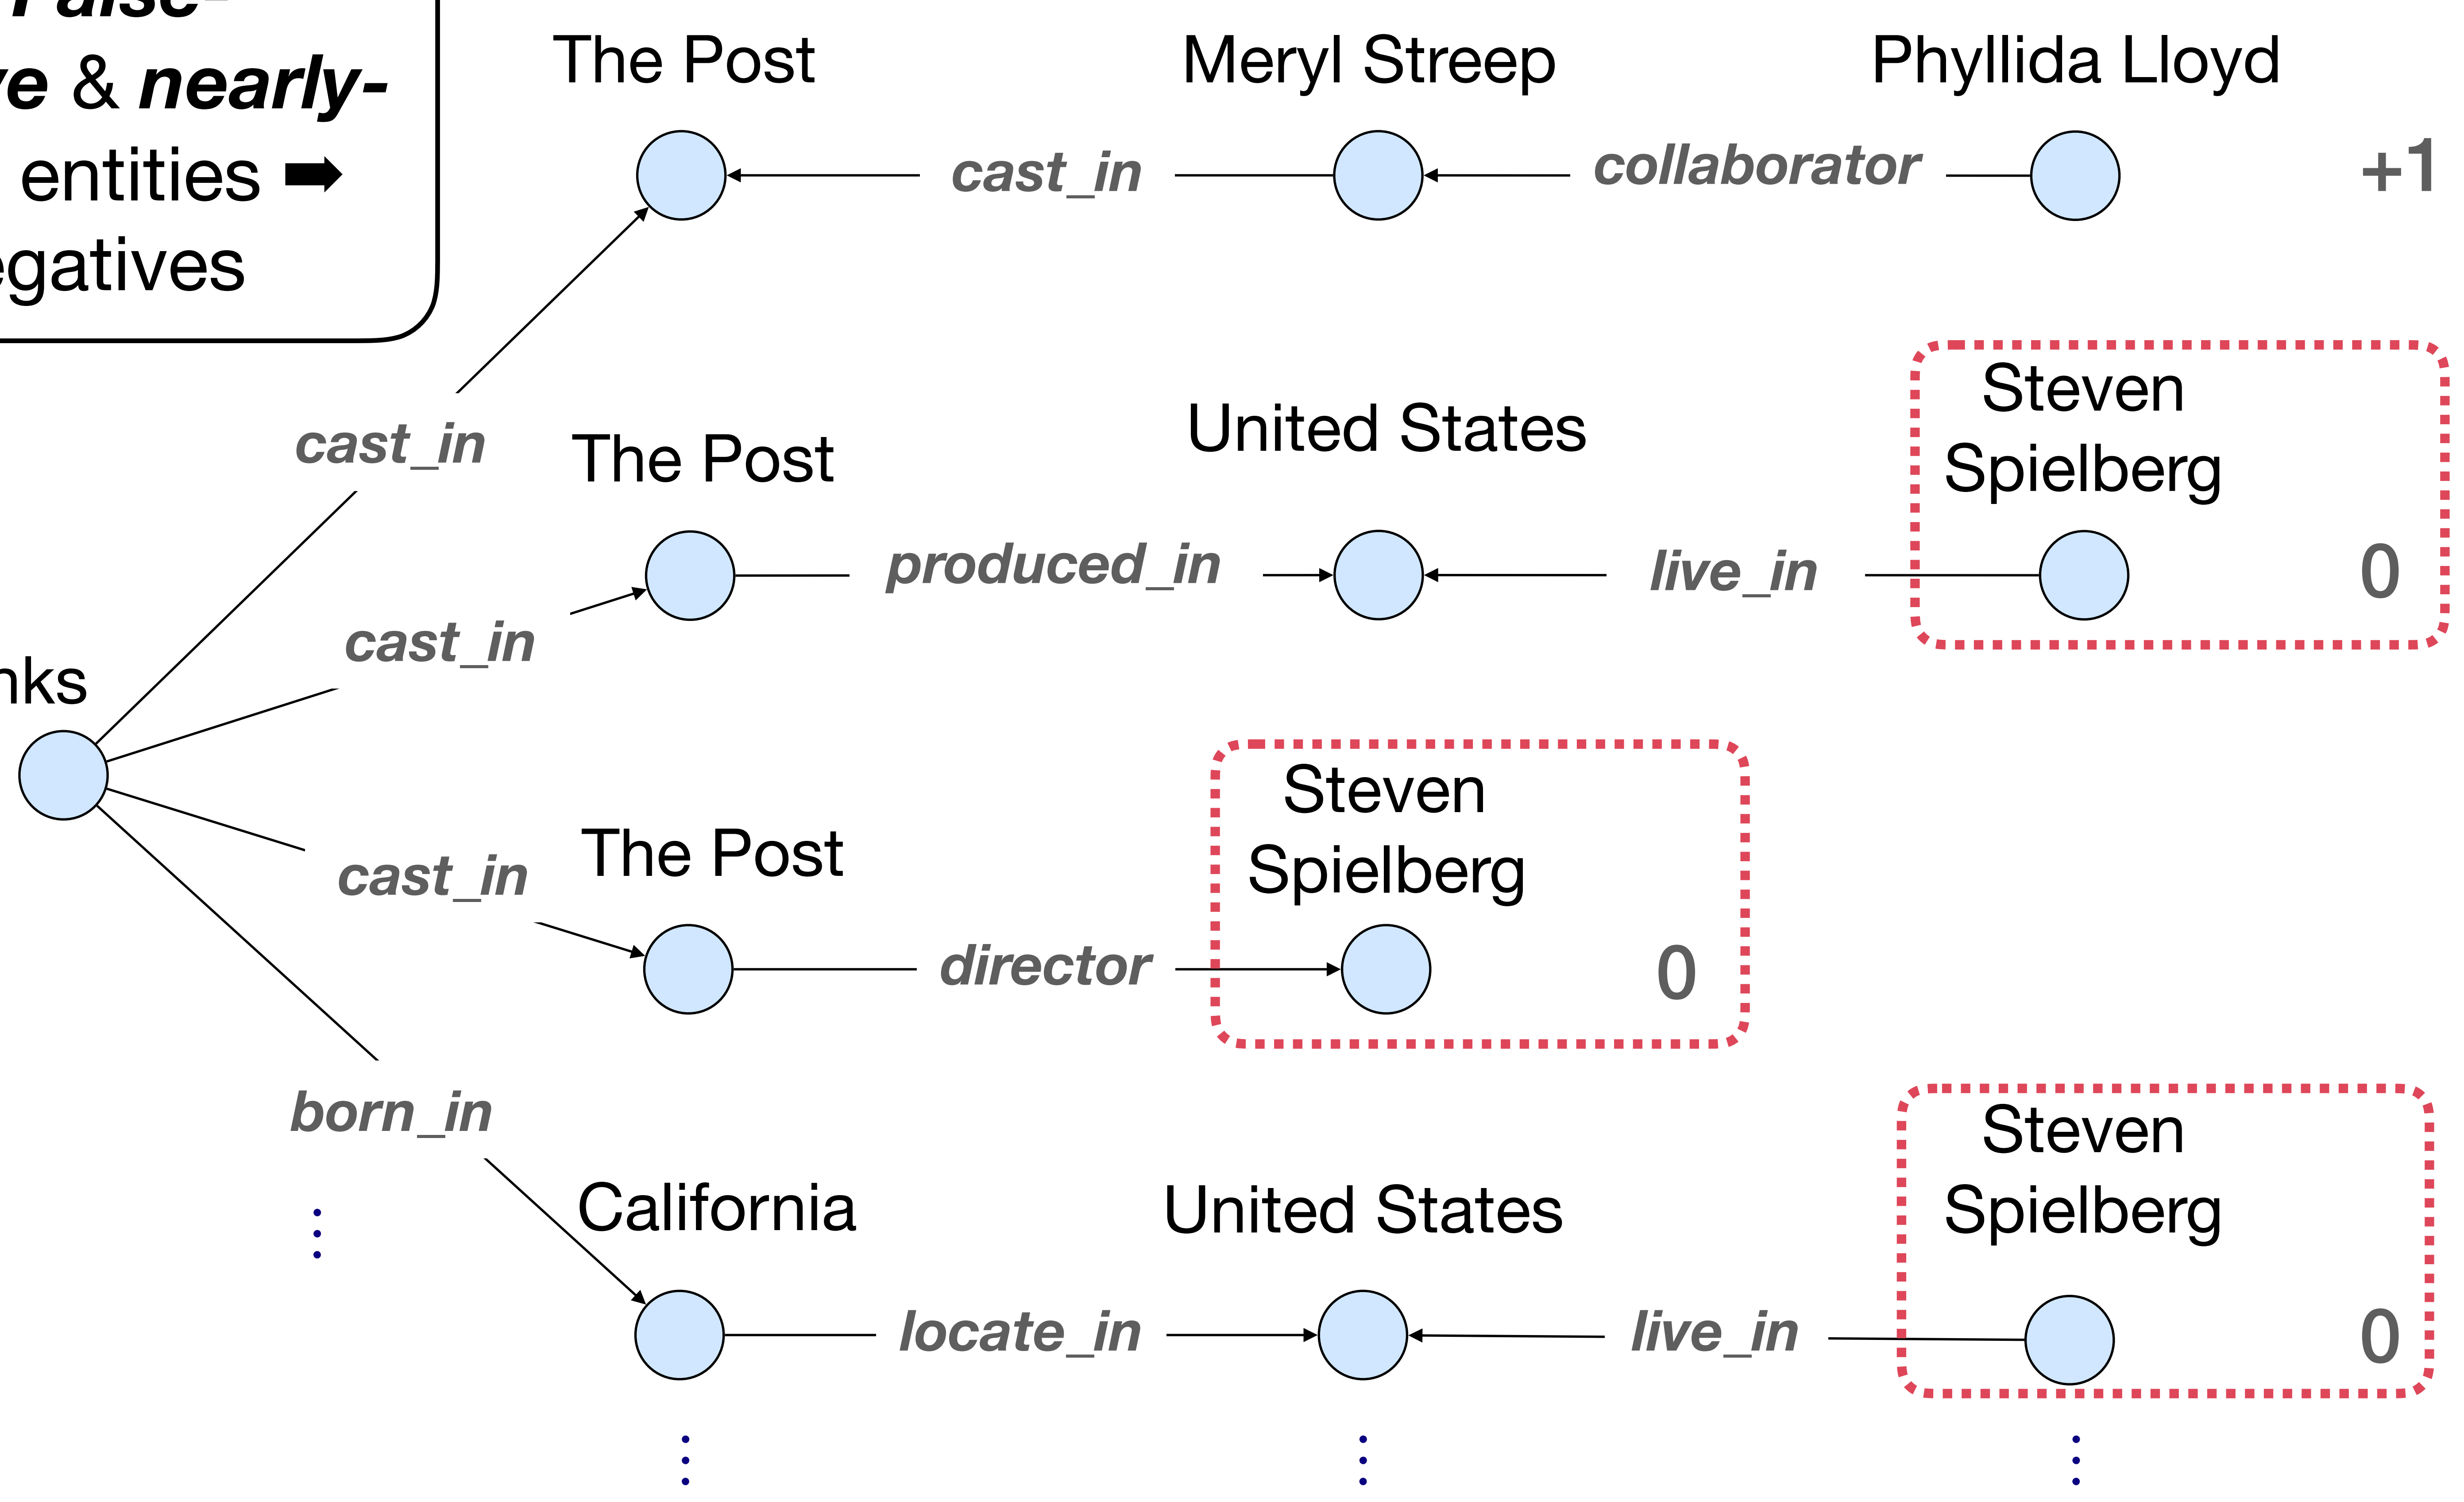

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Reward Shaping

Unobserved facts

Soft correctness

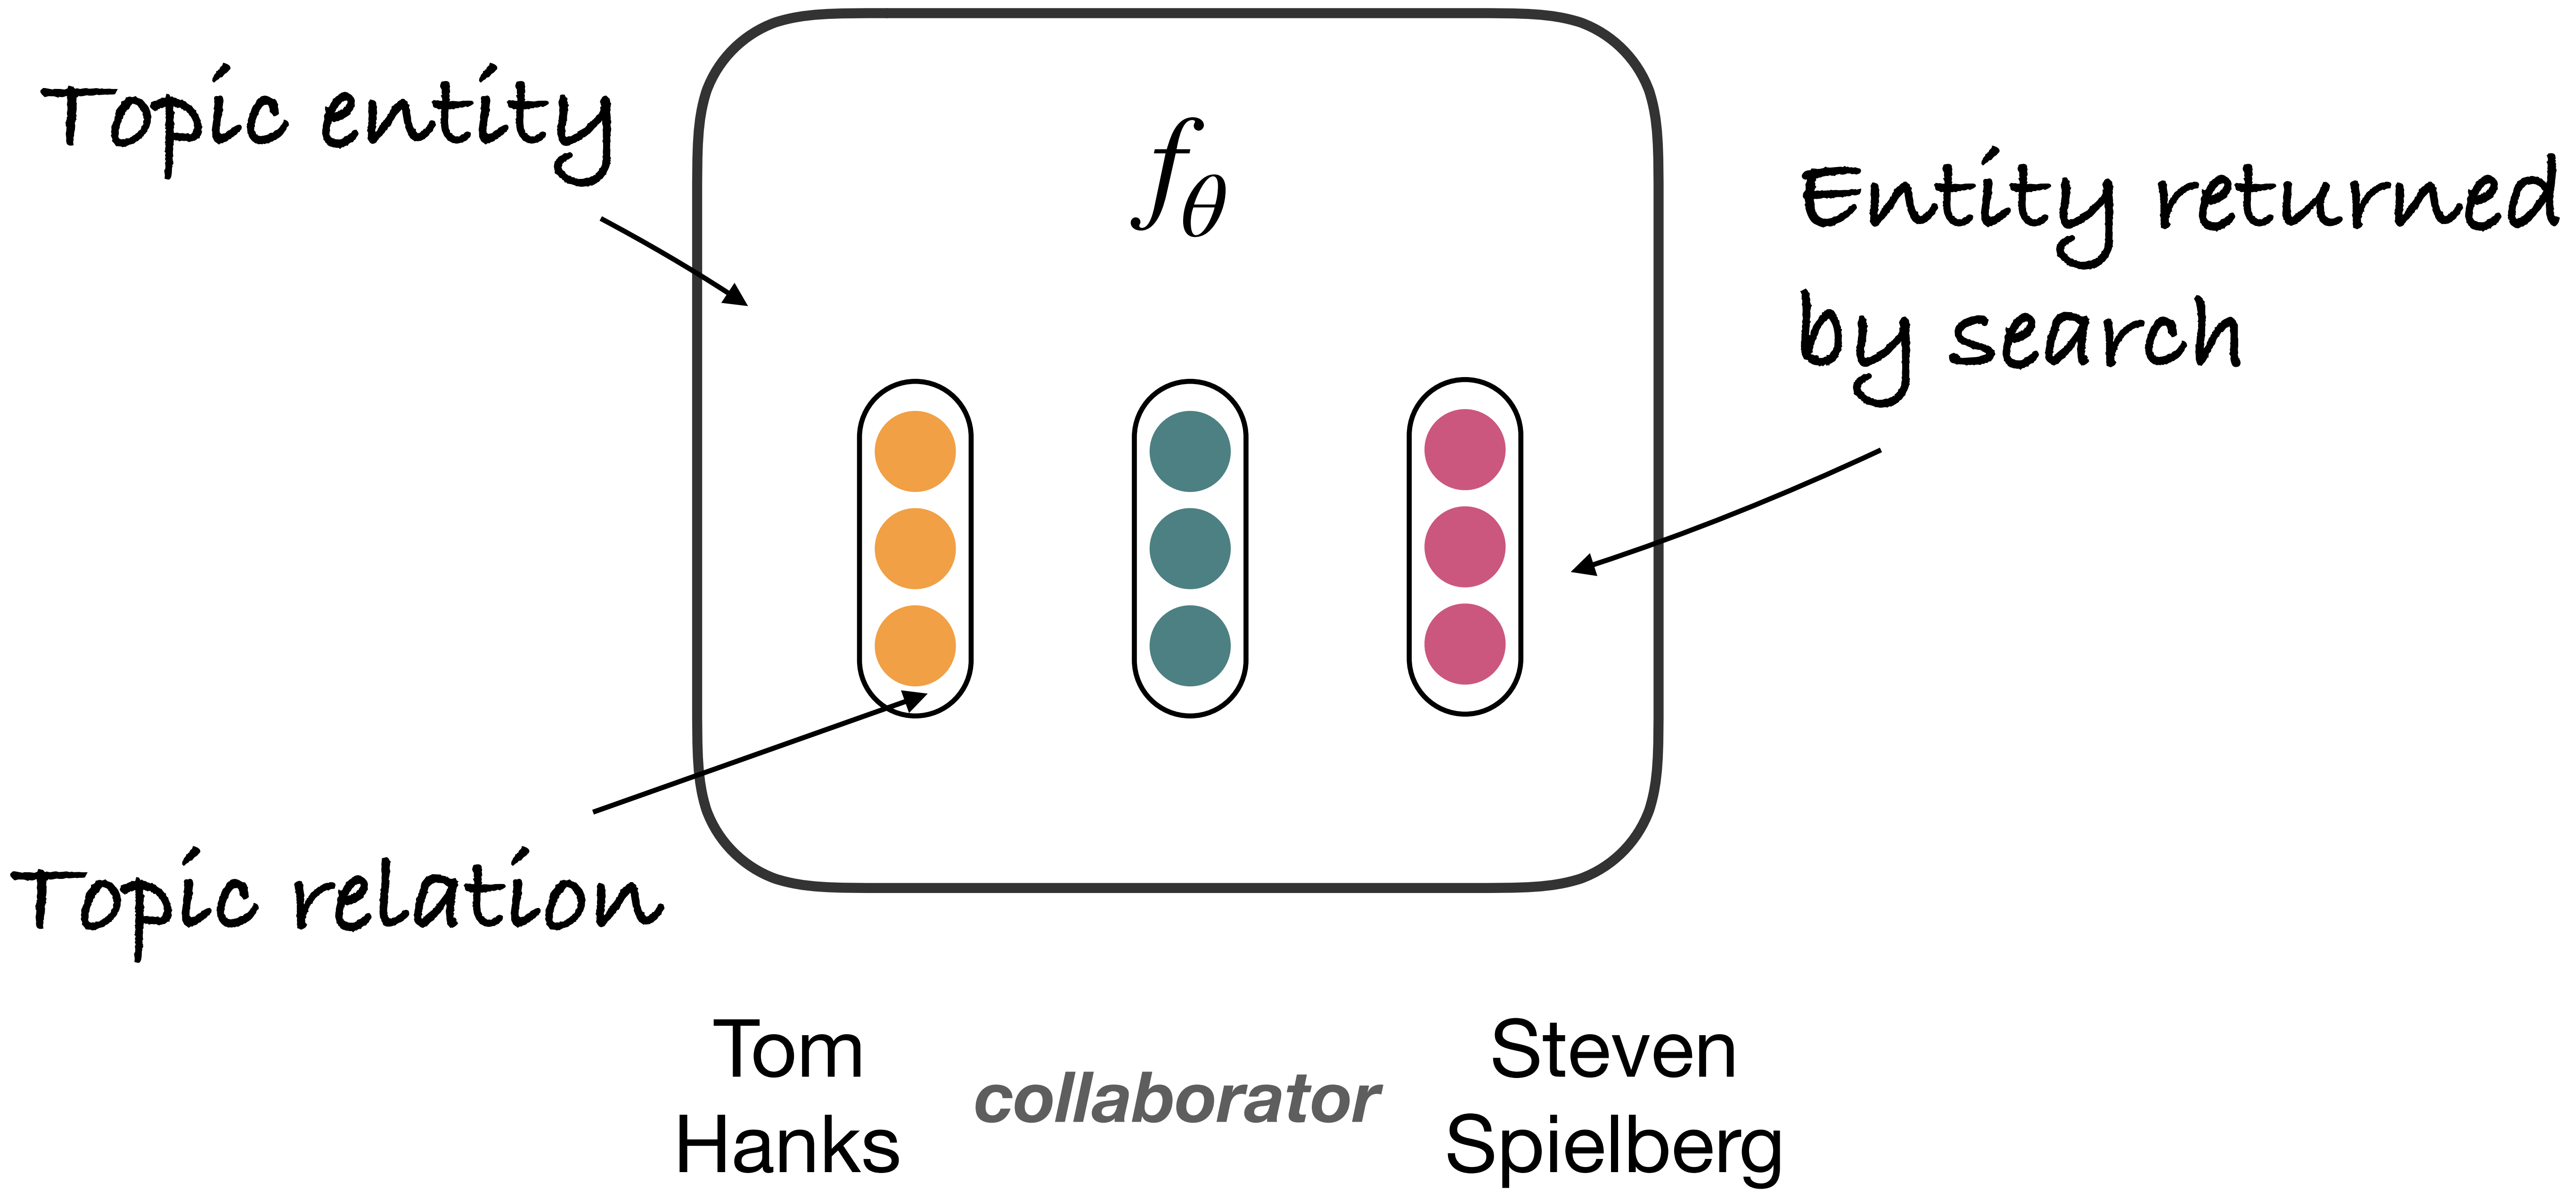

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Reward Shaping

Unobserved facts

Soft correctness

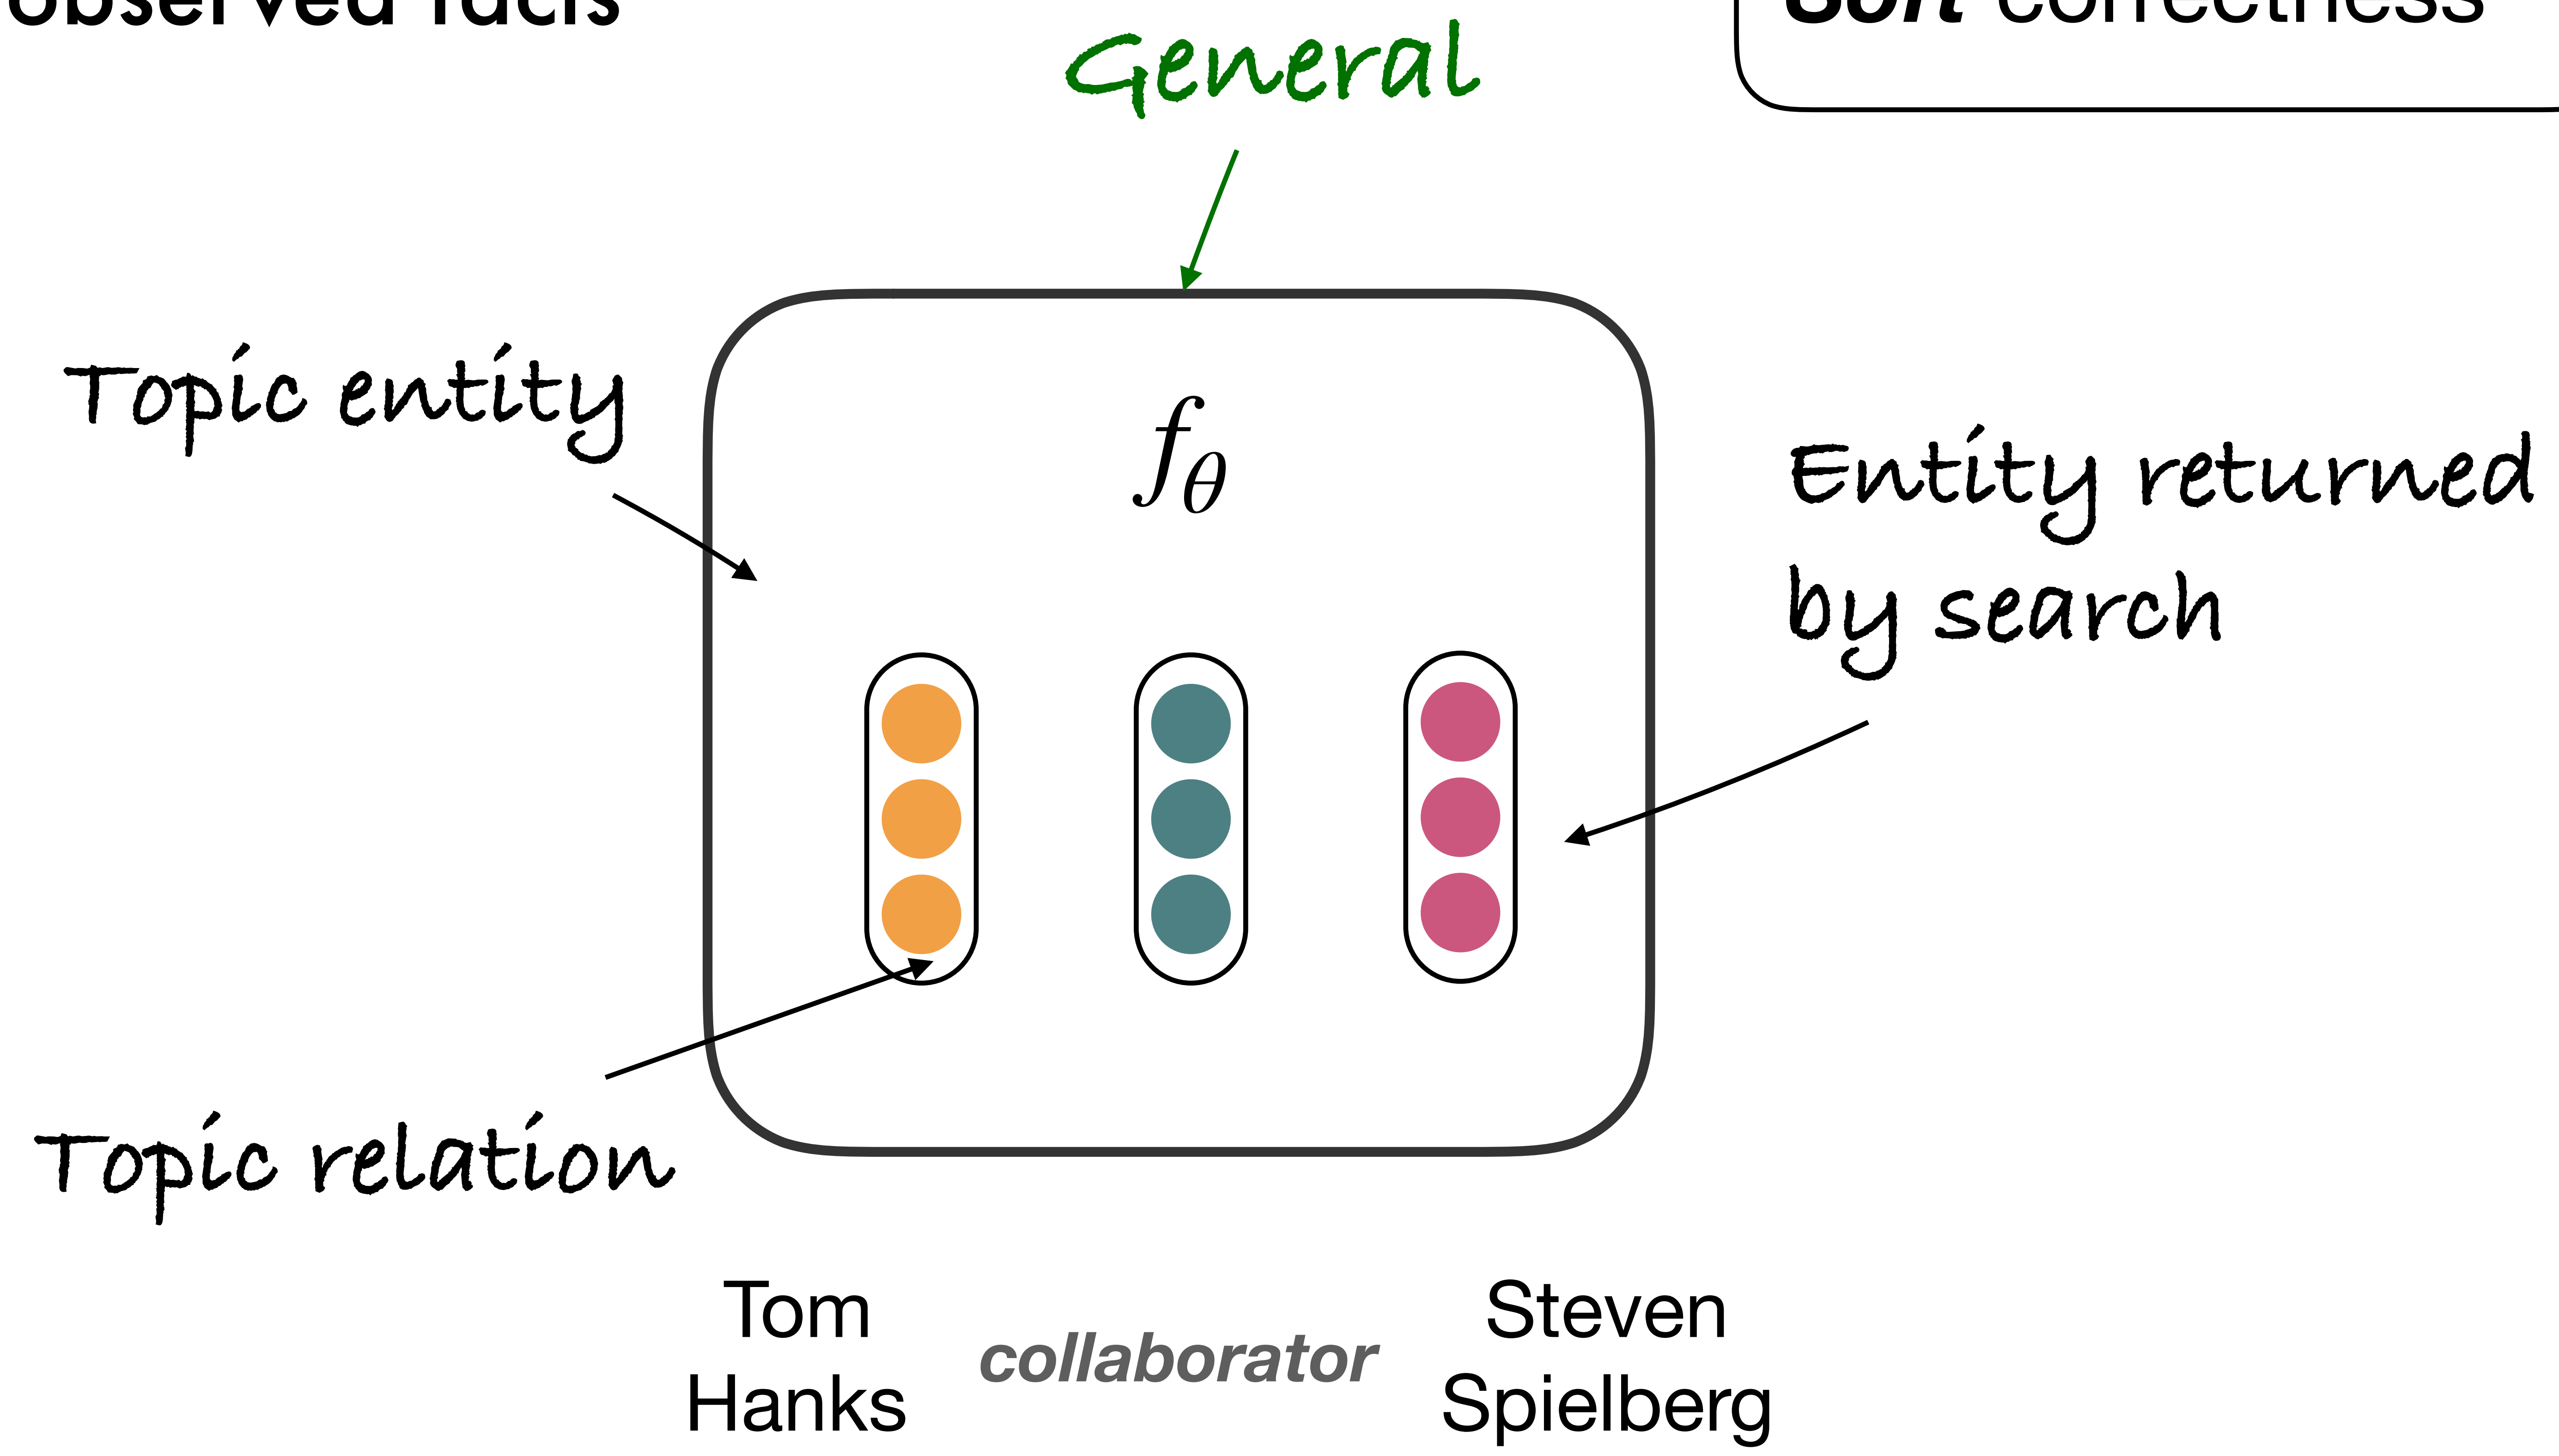

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Reward Shaping

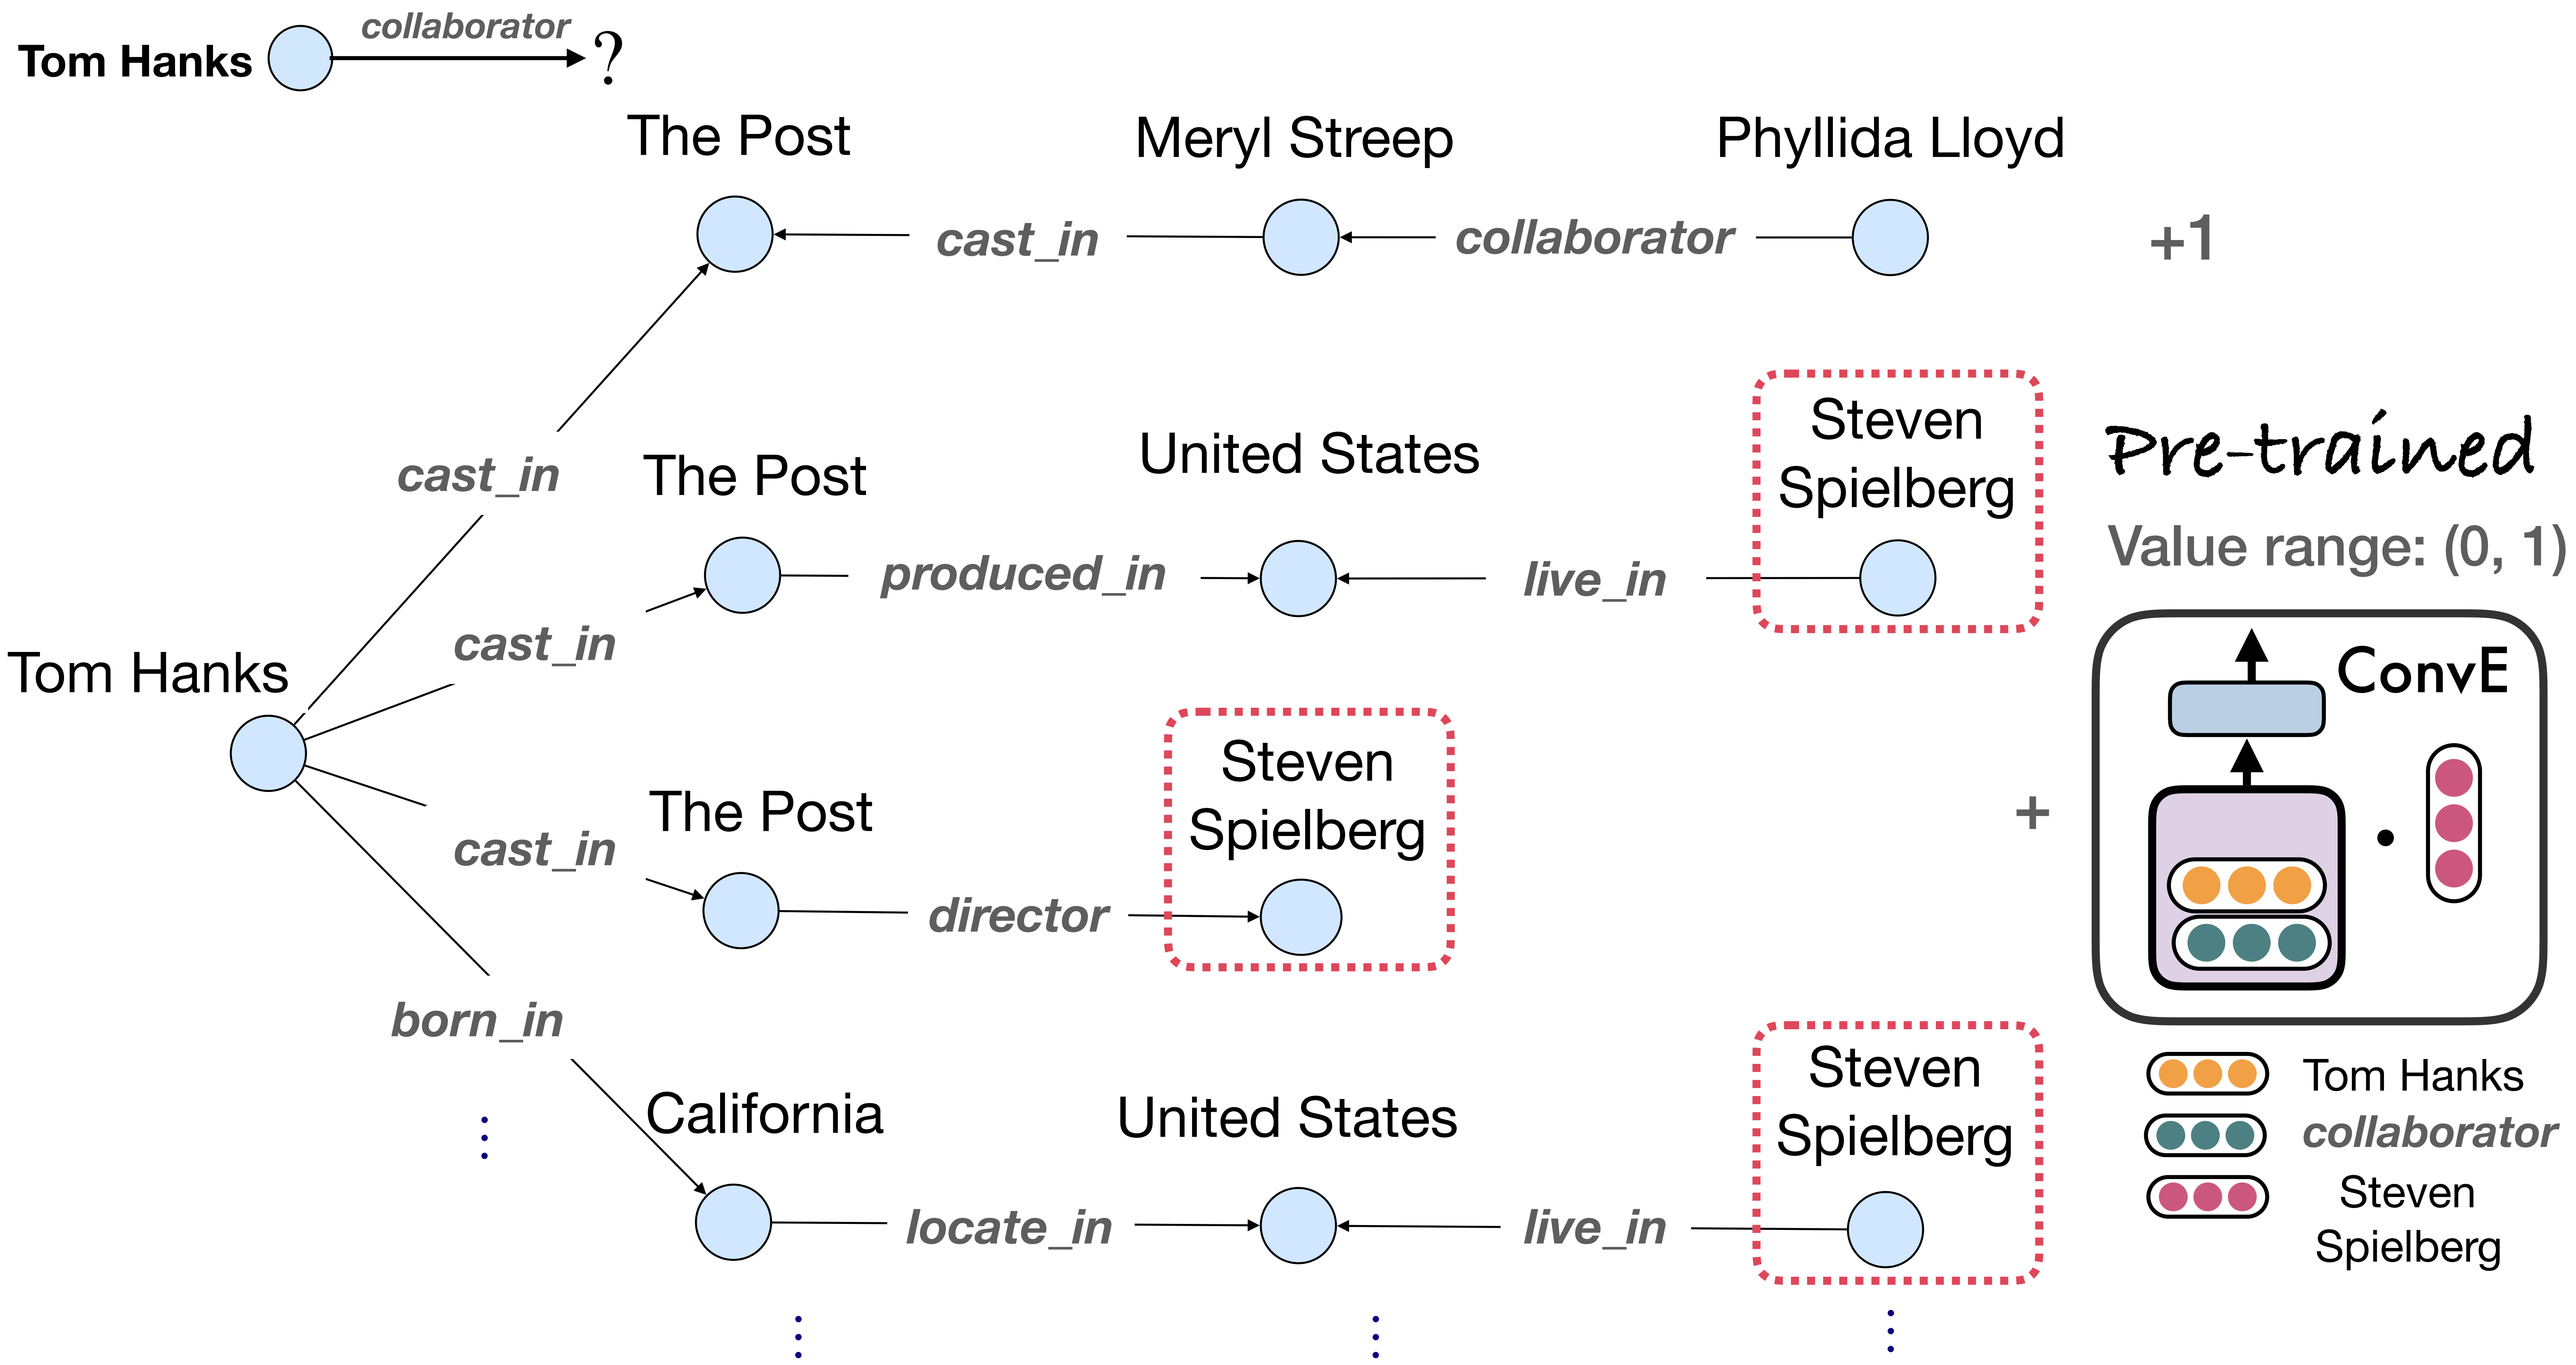

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Experiment Setup

KG Benchmarks

Name	# Ent.	# Rel.	# Fact	# Degree Avg	# Degree Median
Kinship	104	25	8,544	85.15	82
UMLS	135	46	5,216	38.63	28
FB15k-237	14,505	237	272,115	19.74	14
WN18RR	40,945	11	86,835	2.19	2
NELL-995	75,492	200	154,213	4.07	1

Decreasing connectivity

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Main Results

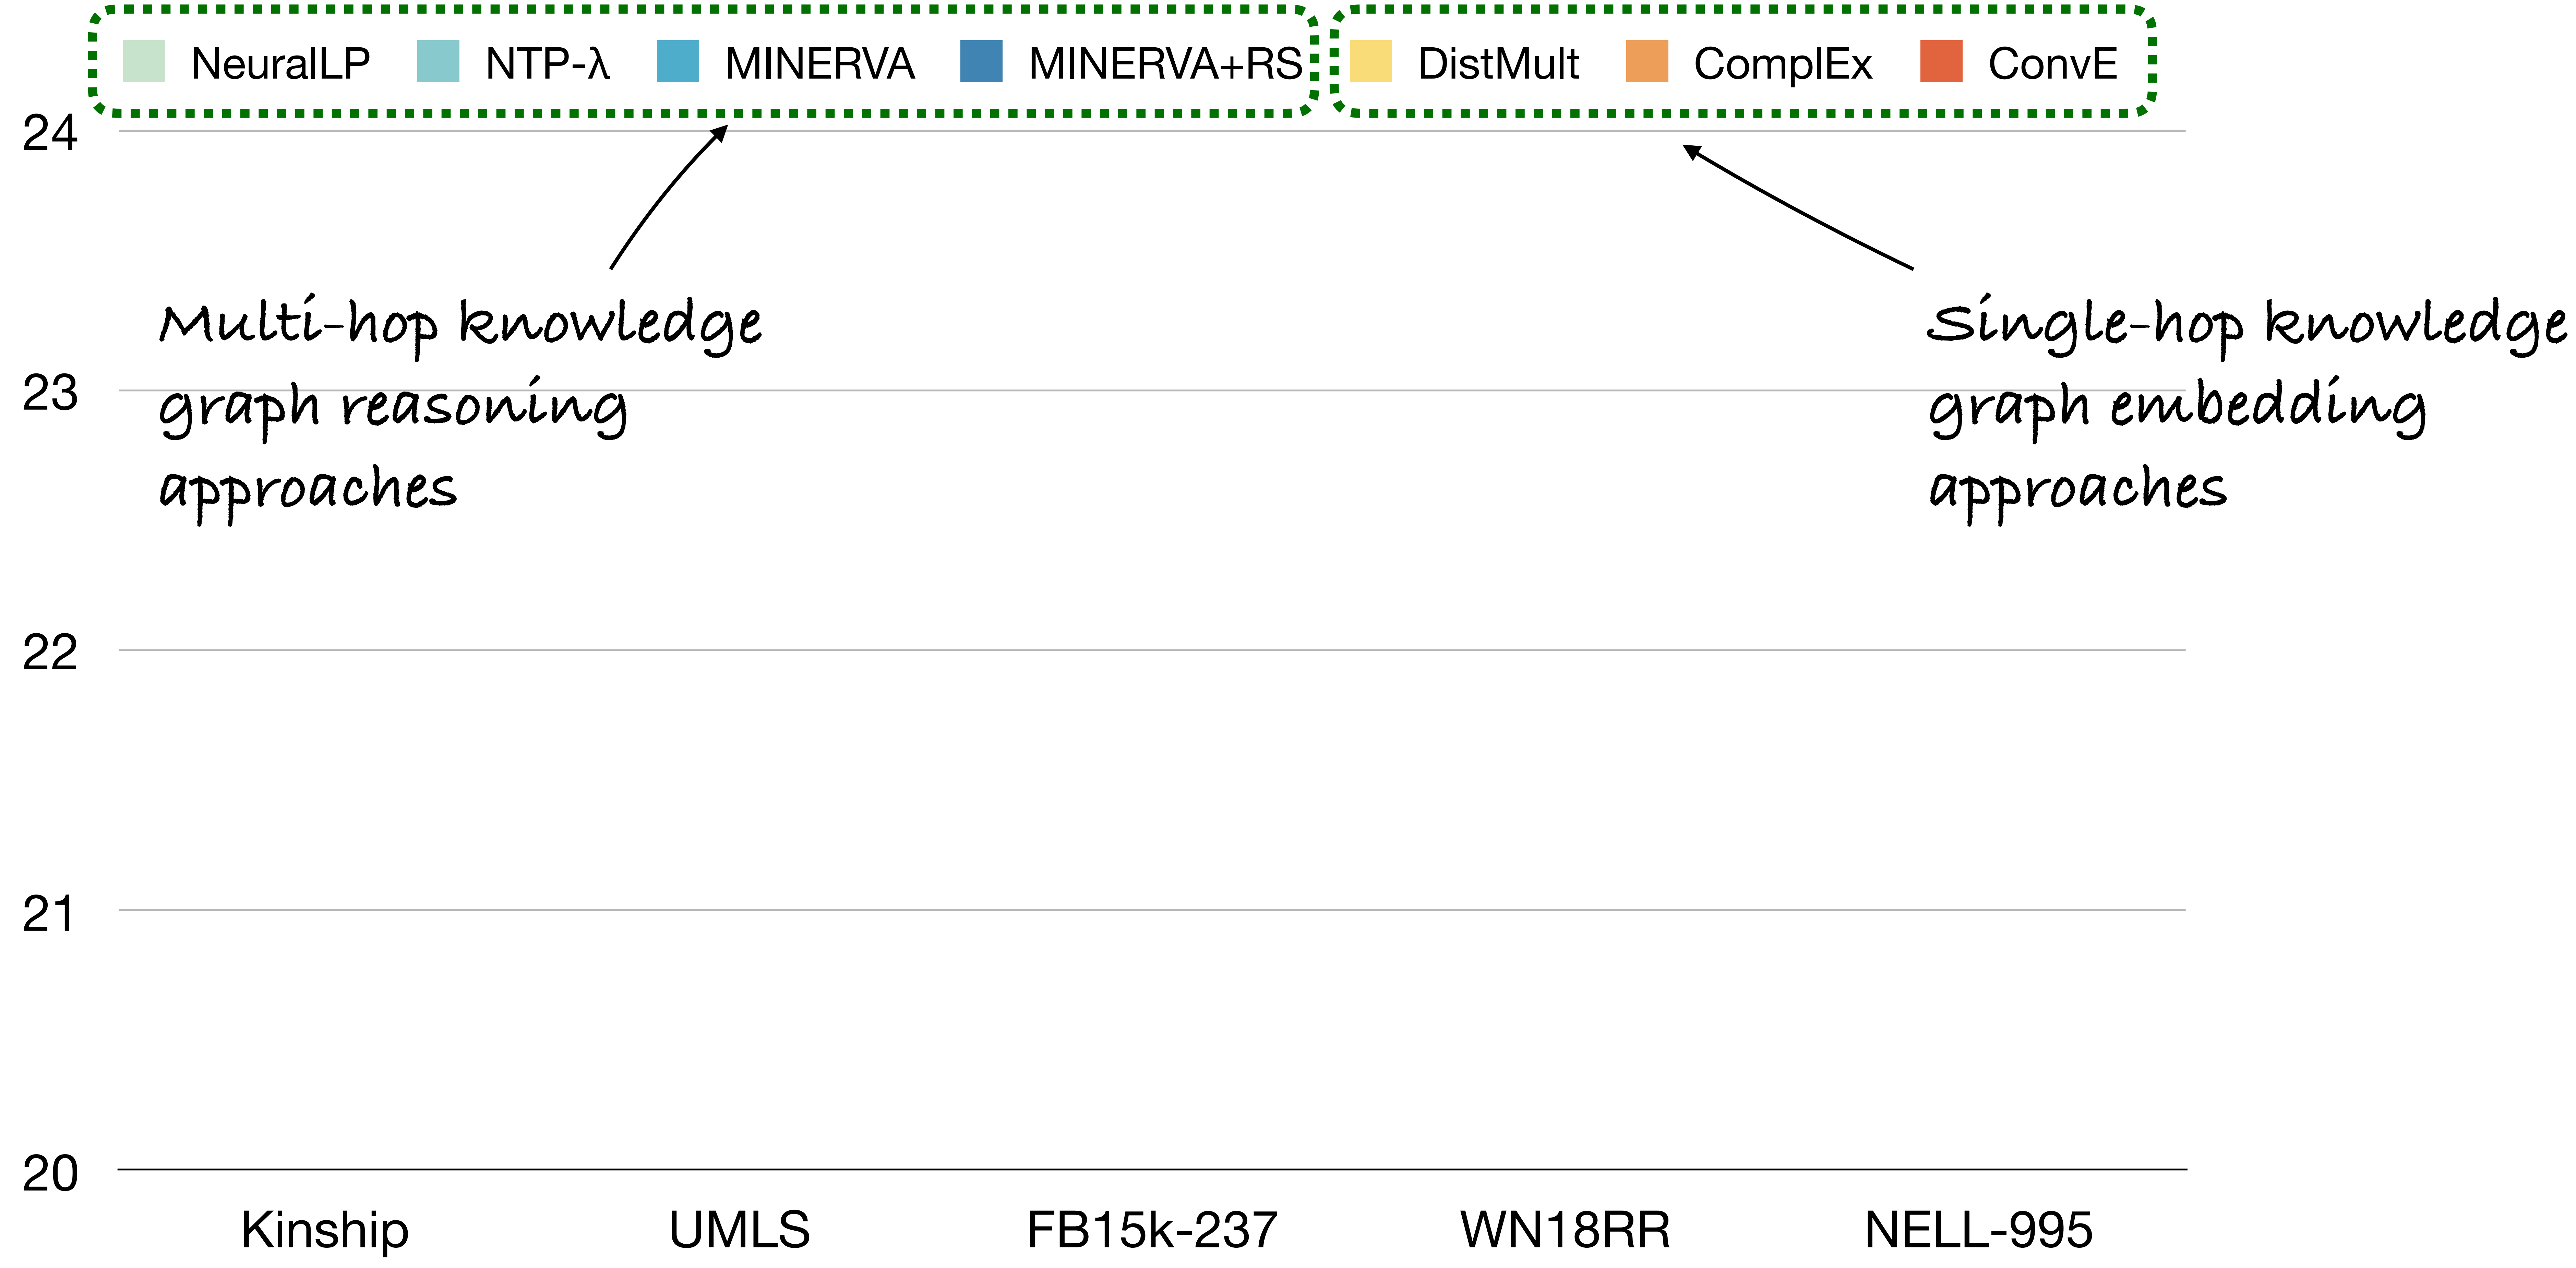

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Main Results

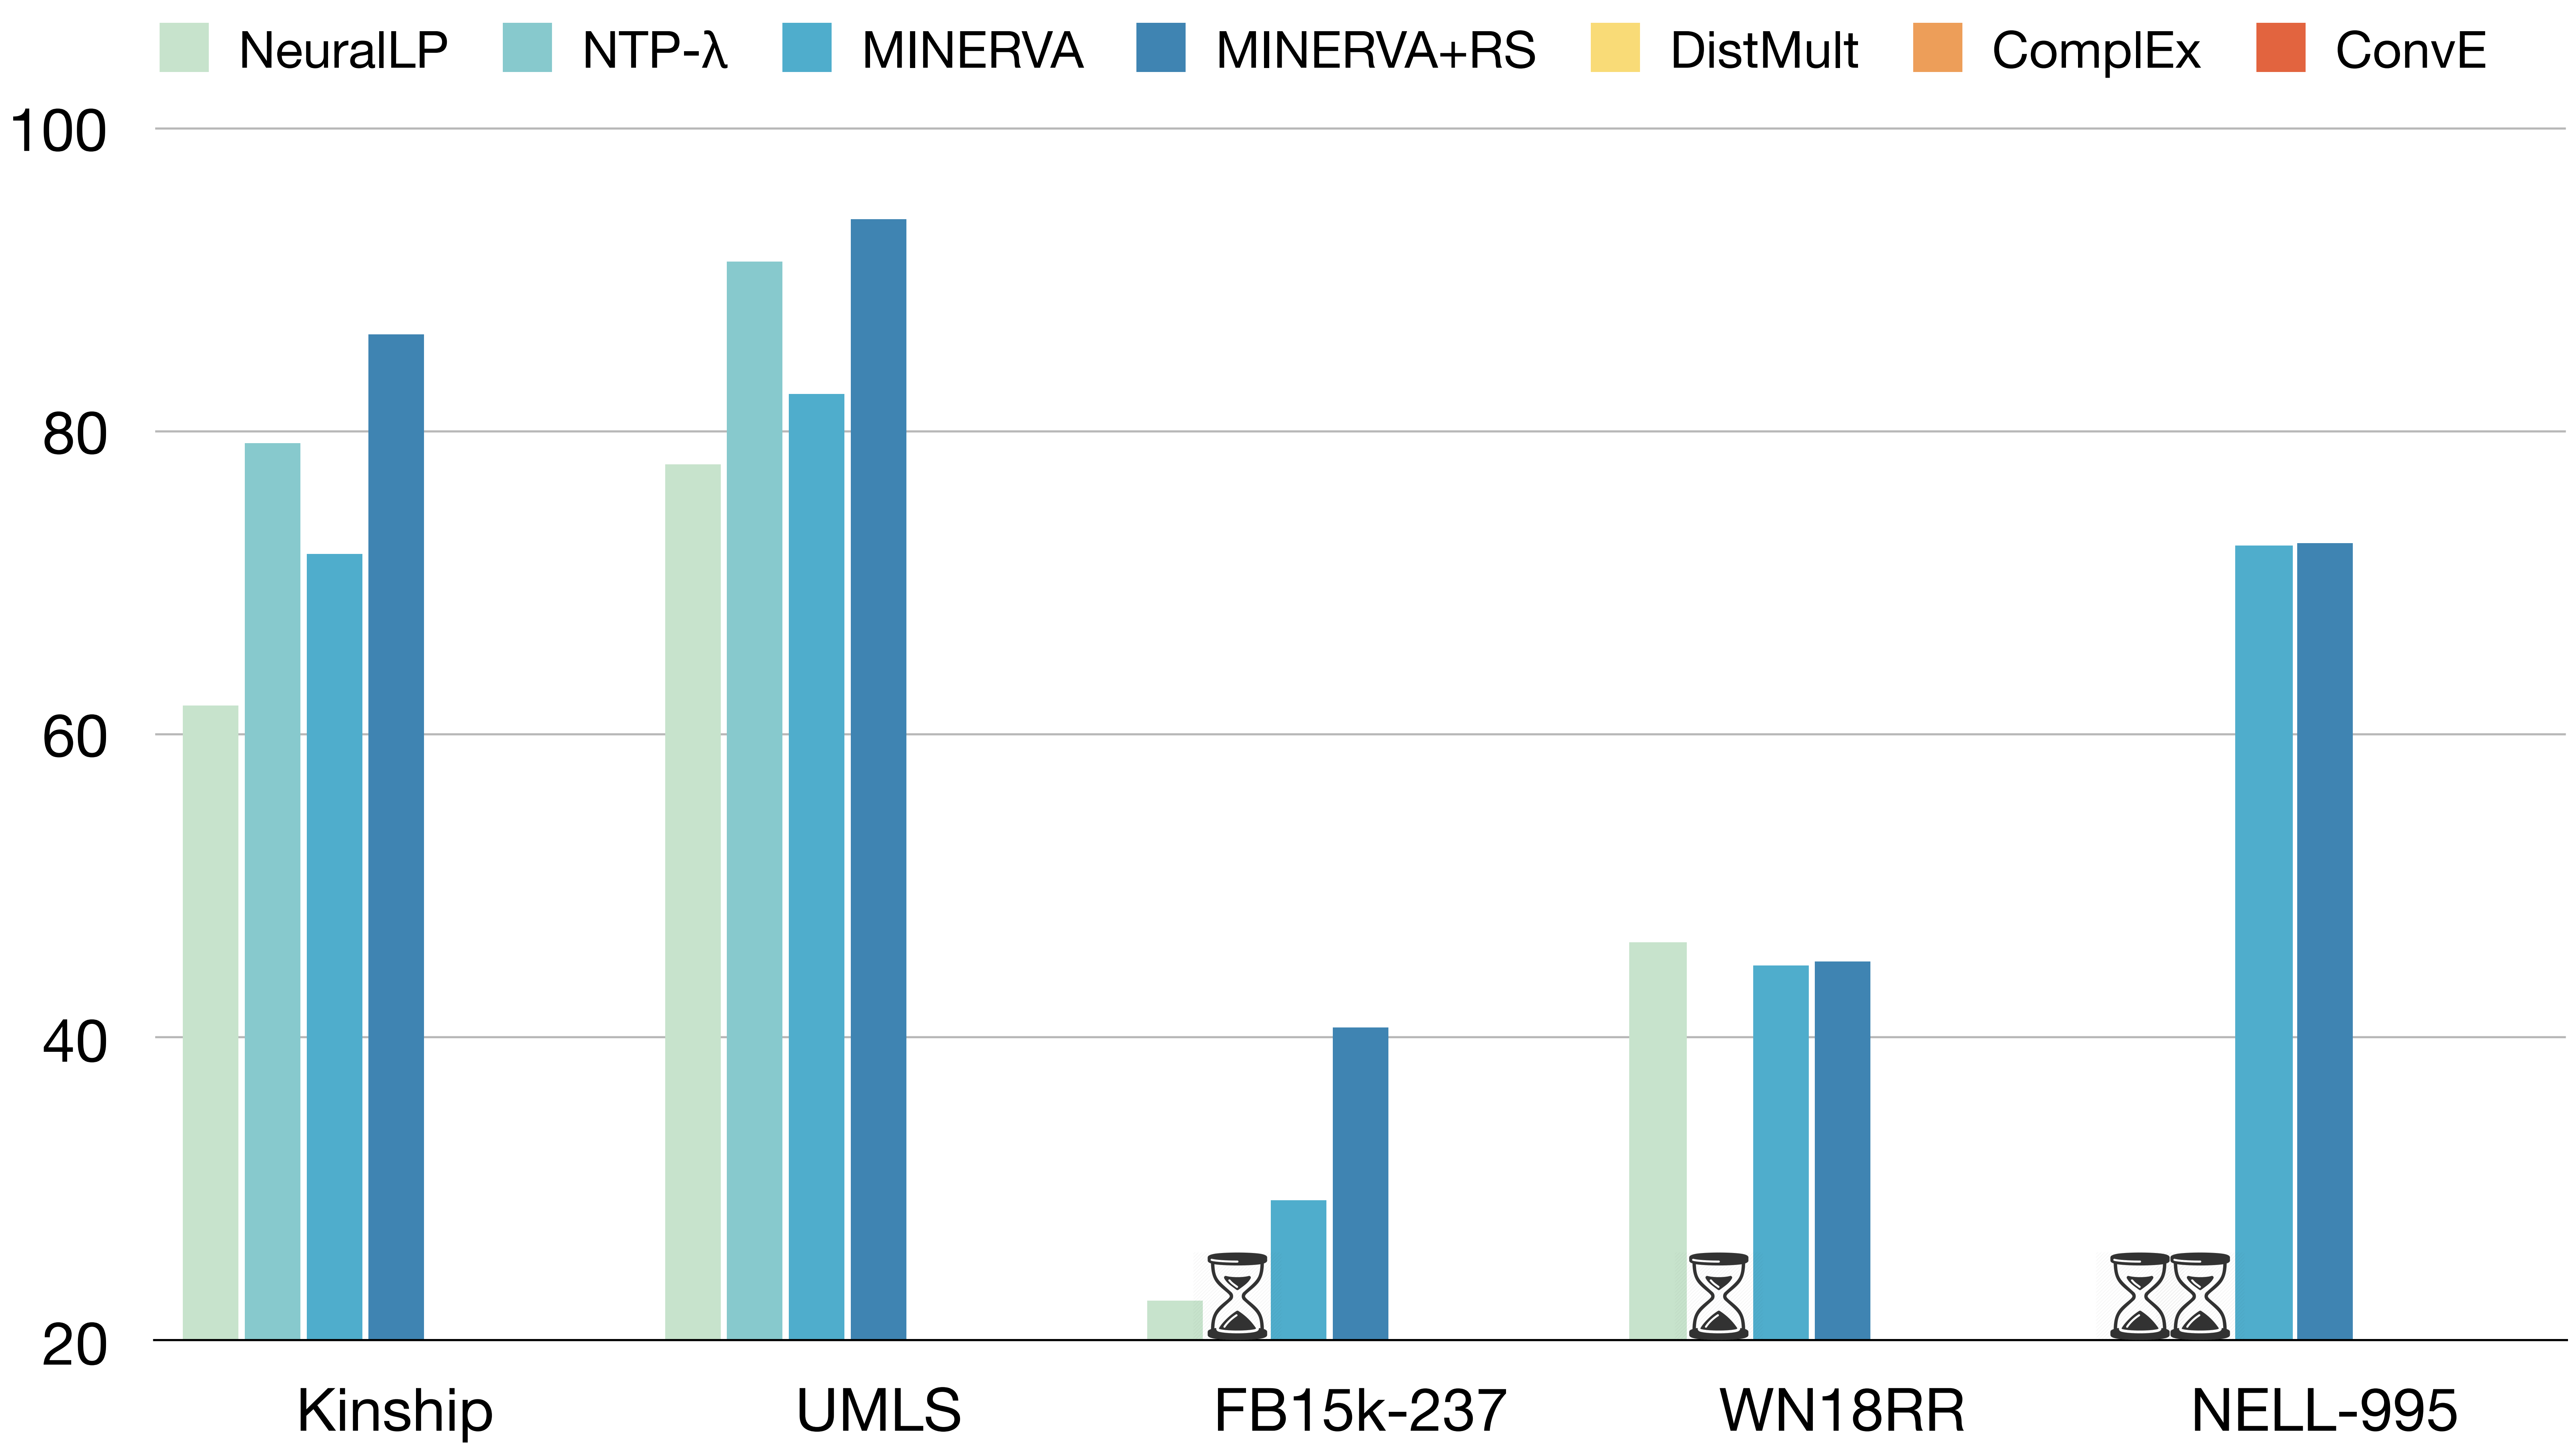

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Main Results

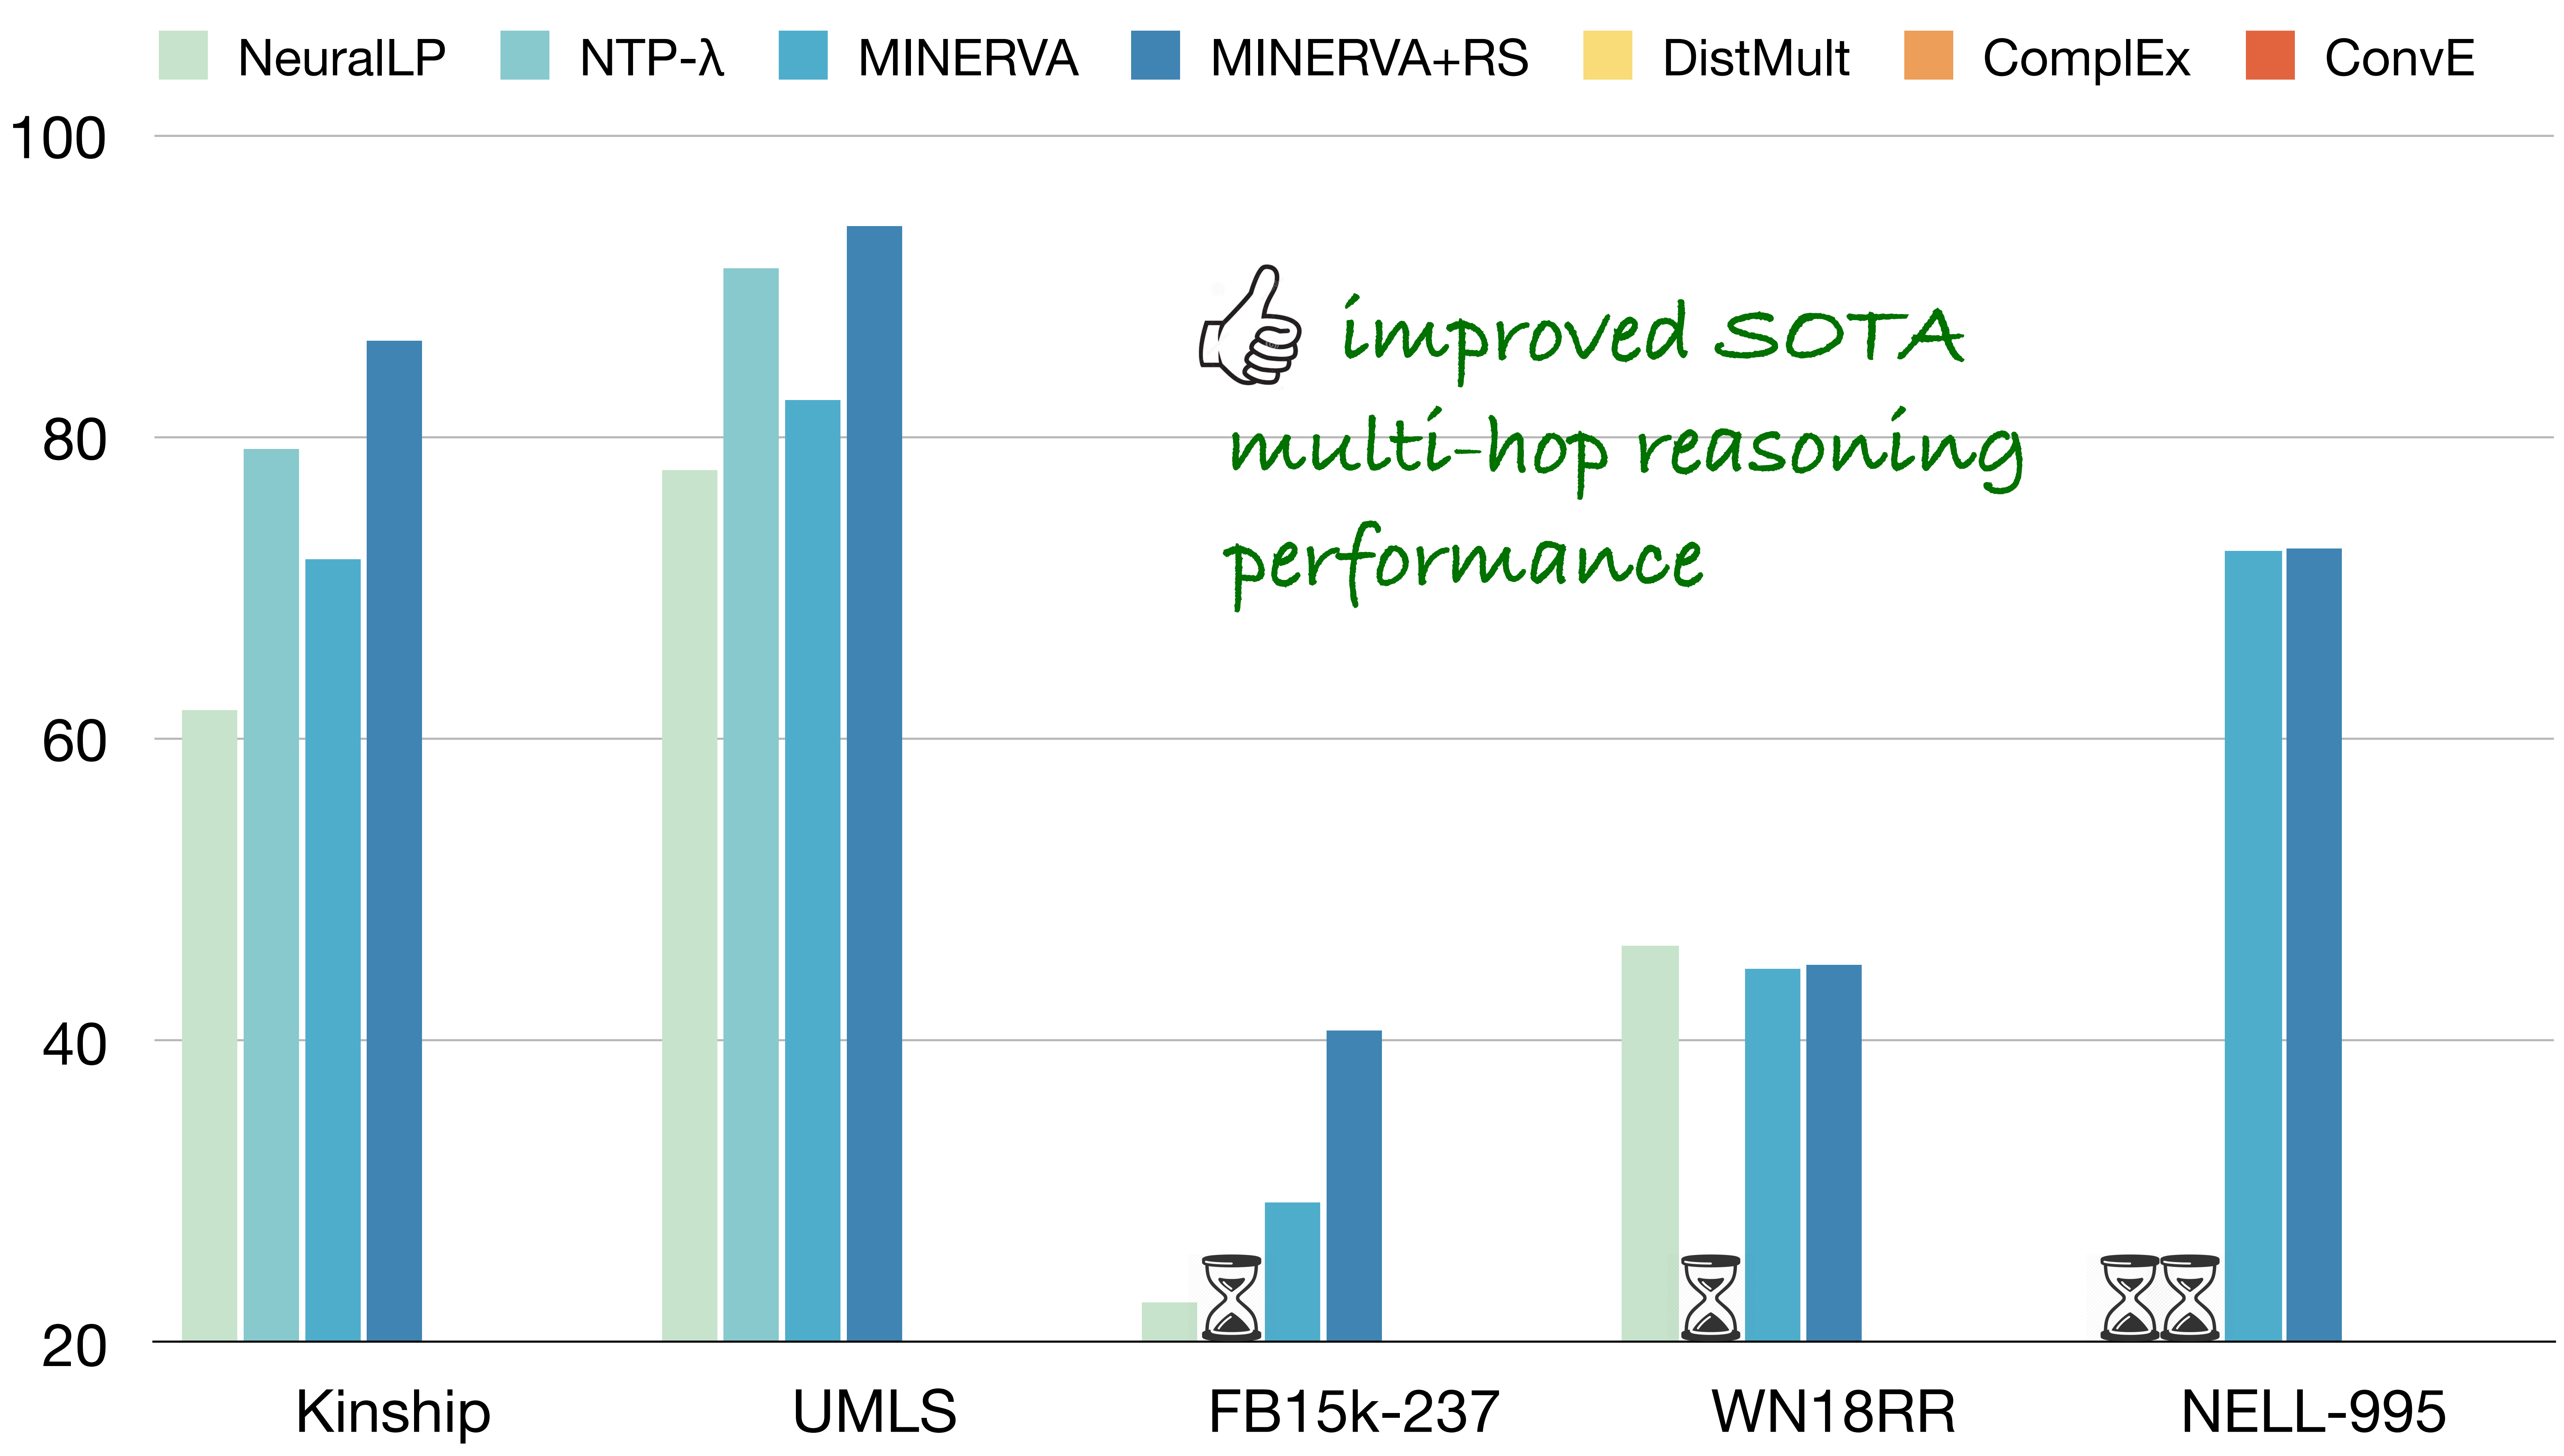

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Main Results

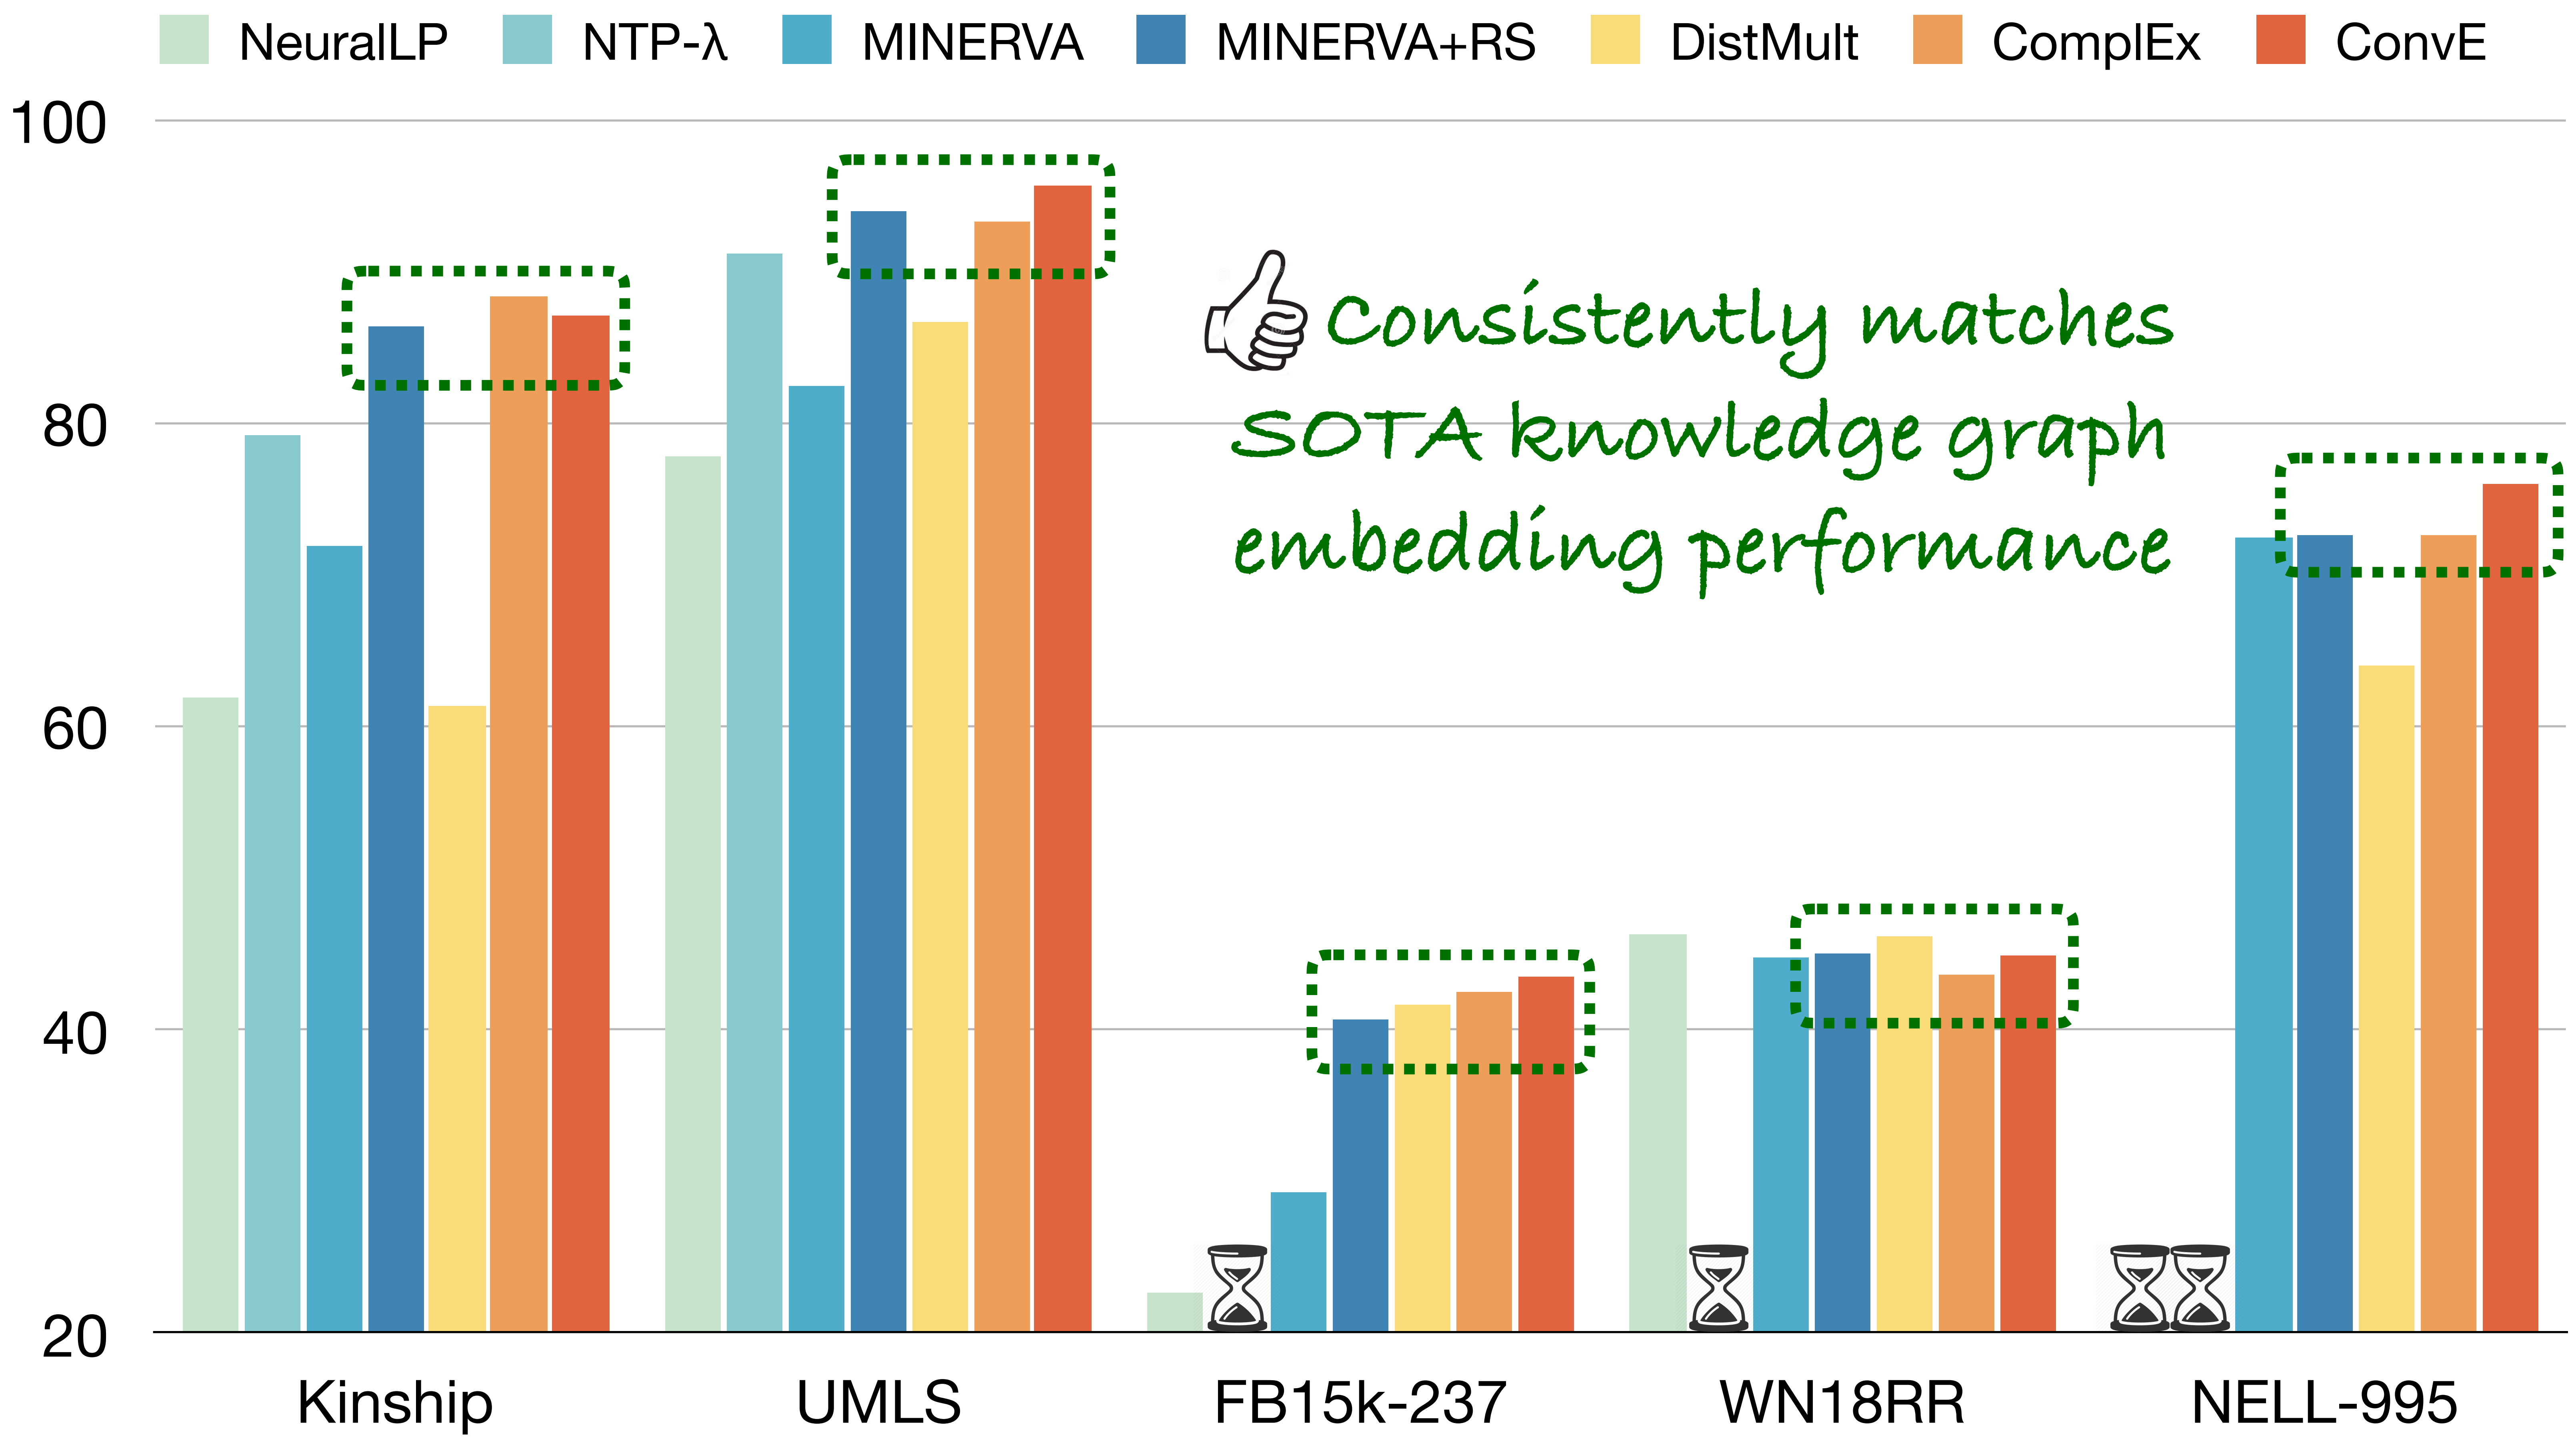

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Interpretable Results

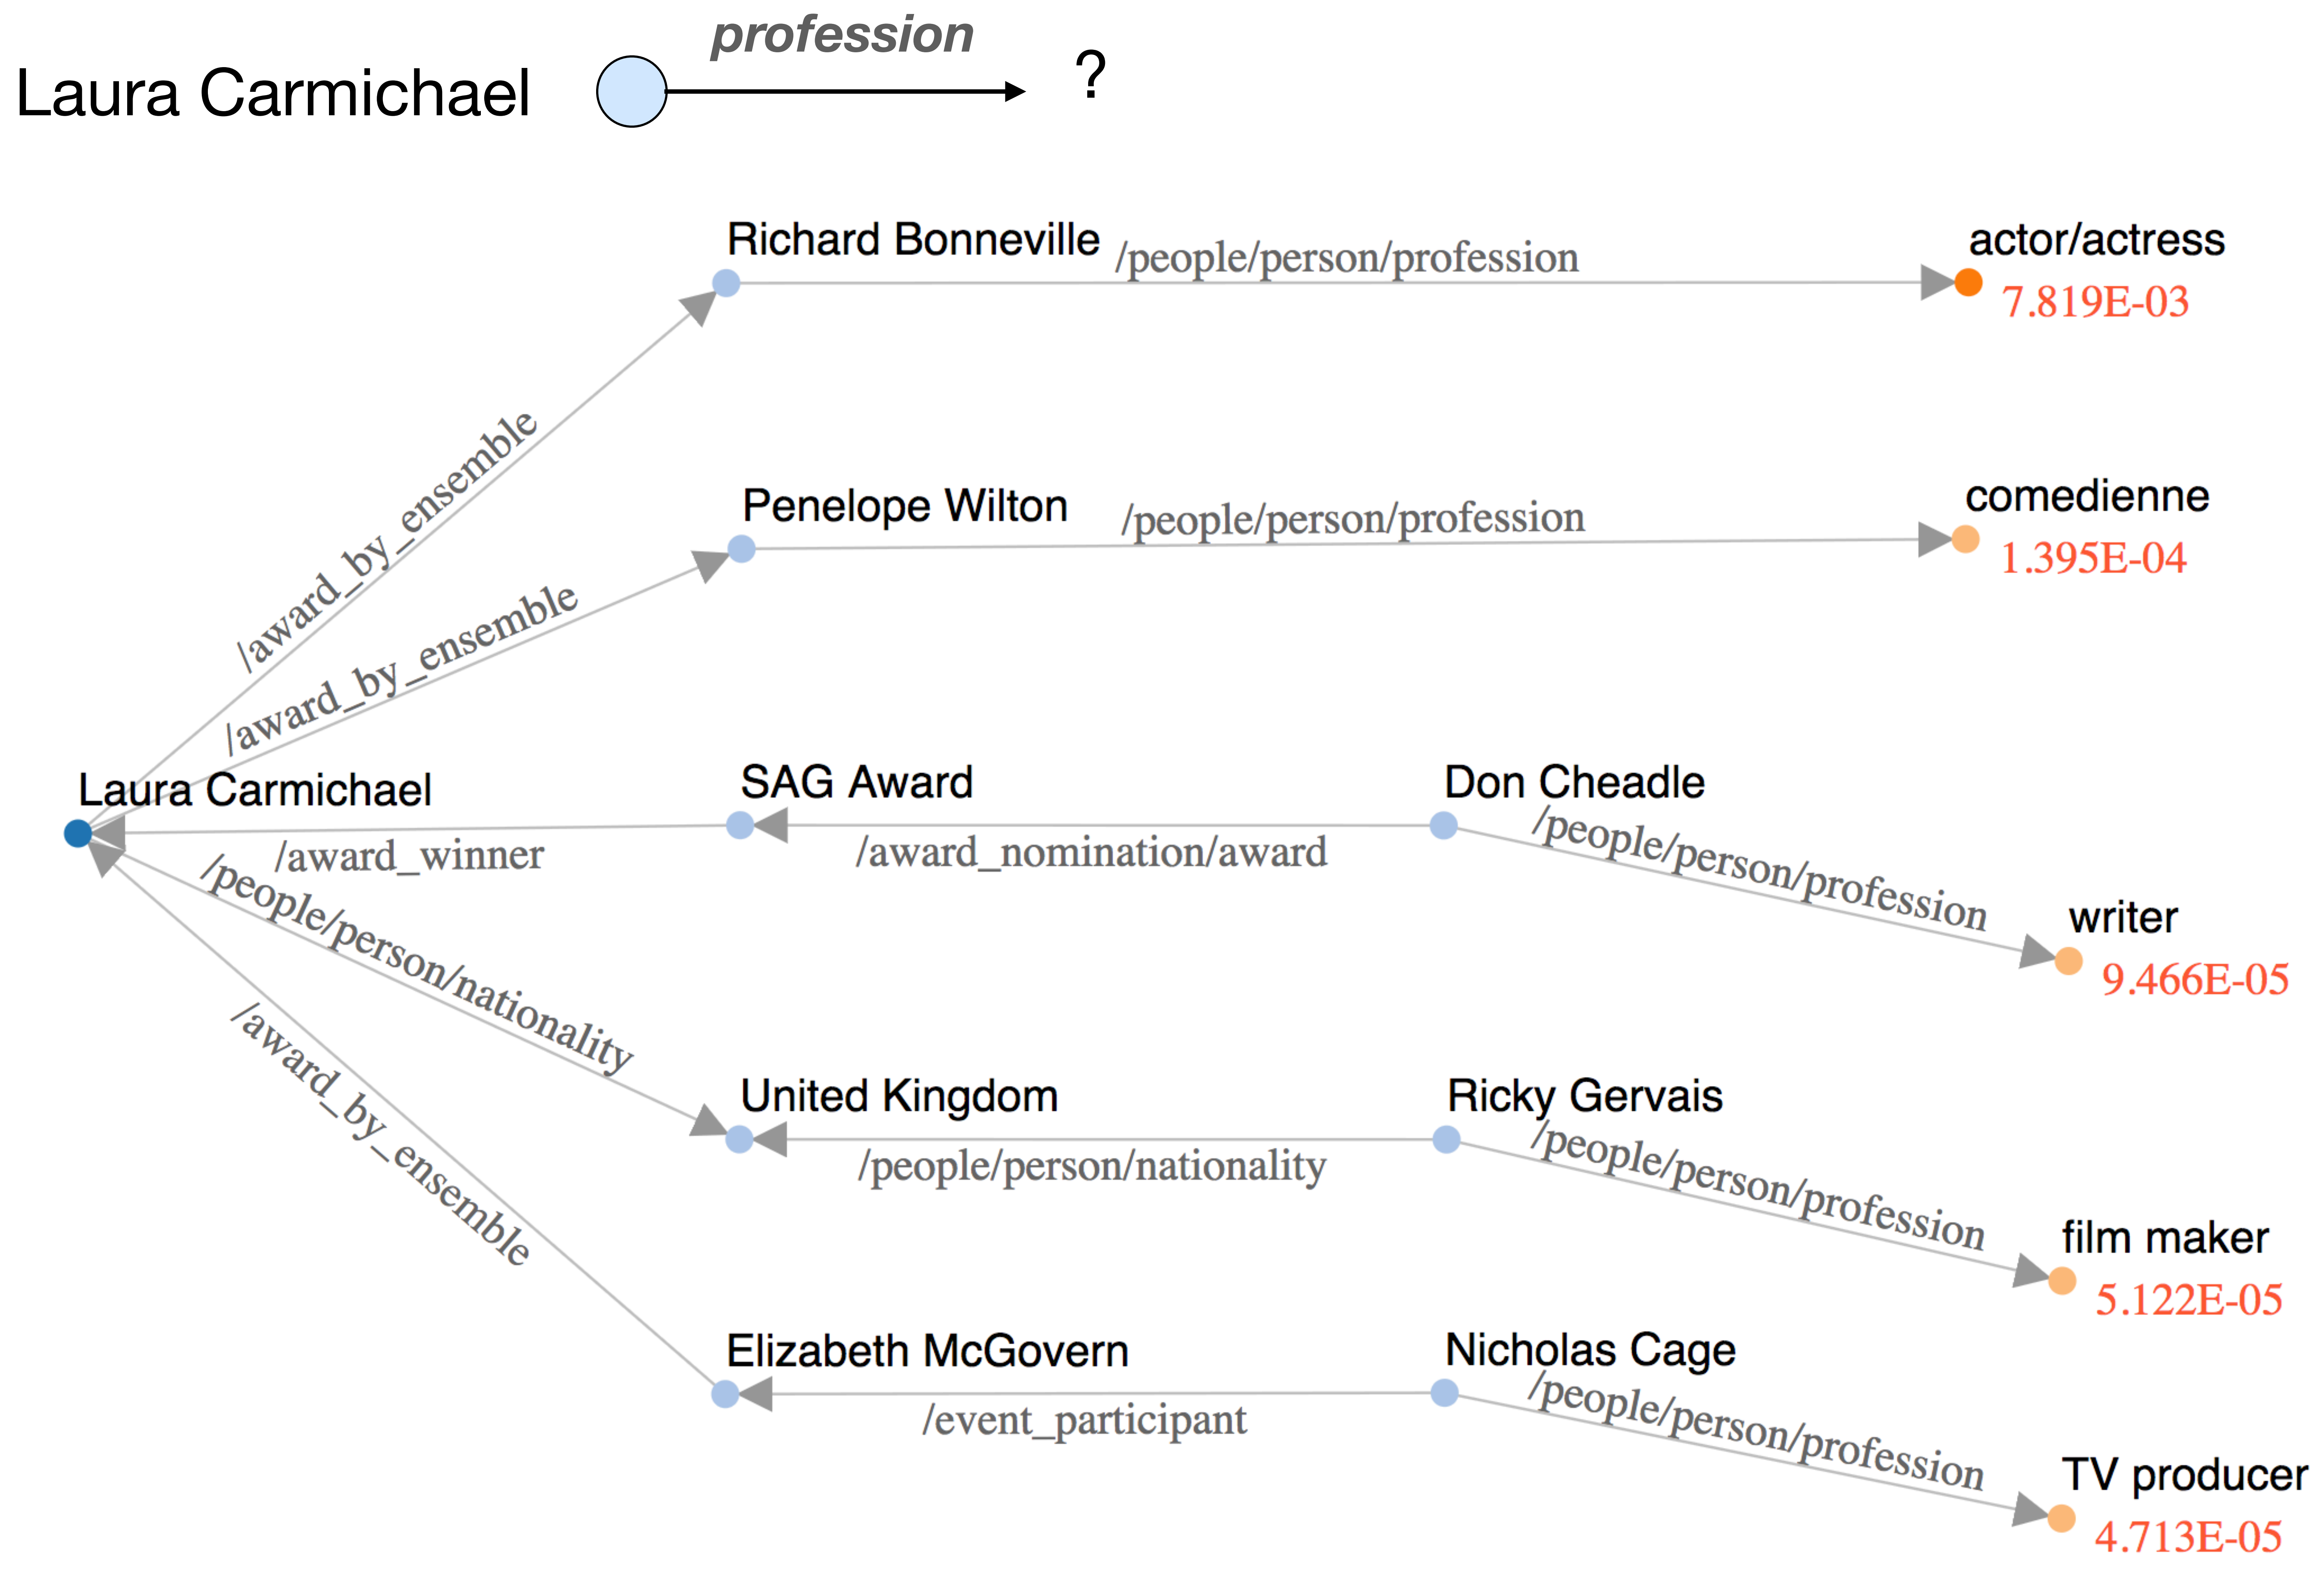

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Interpretable Results

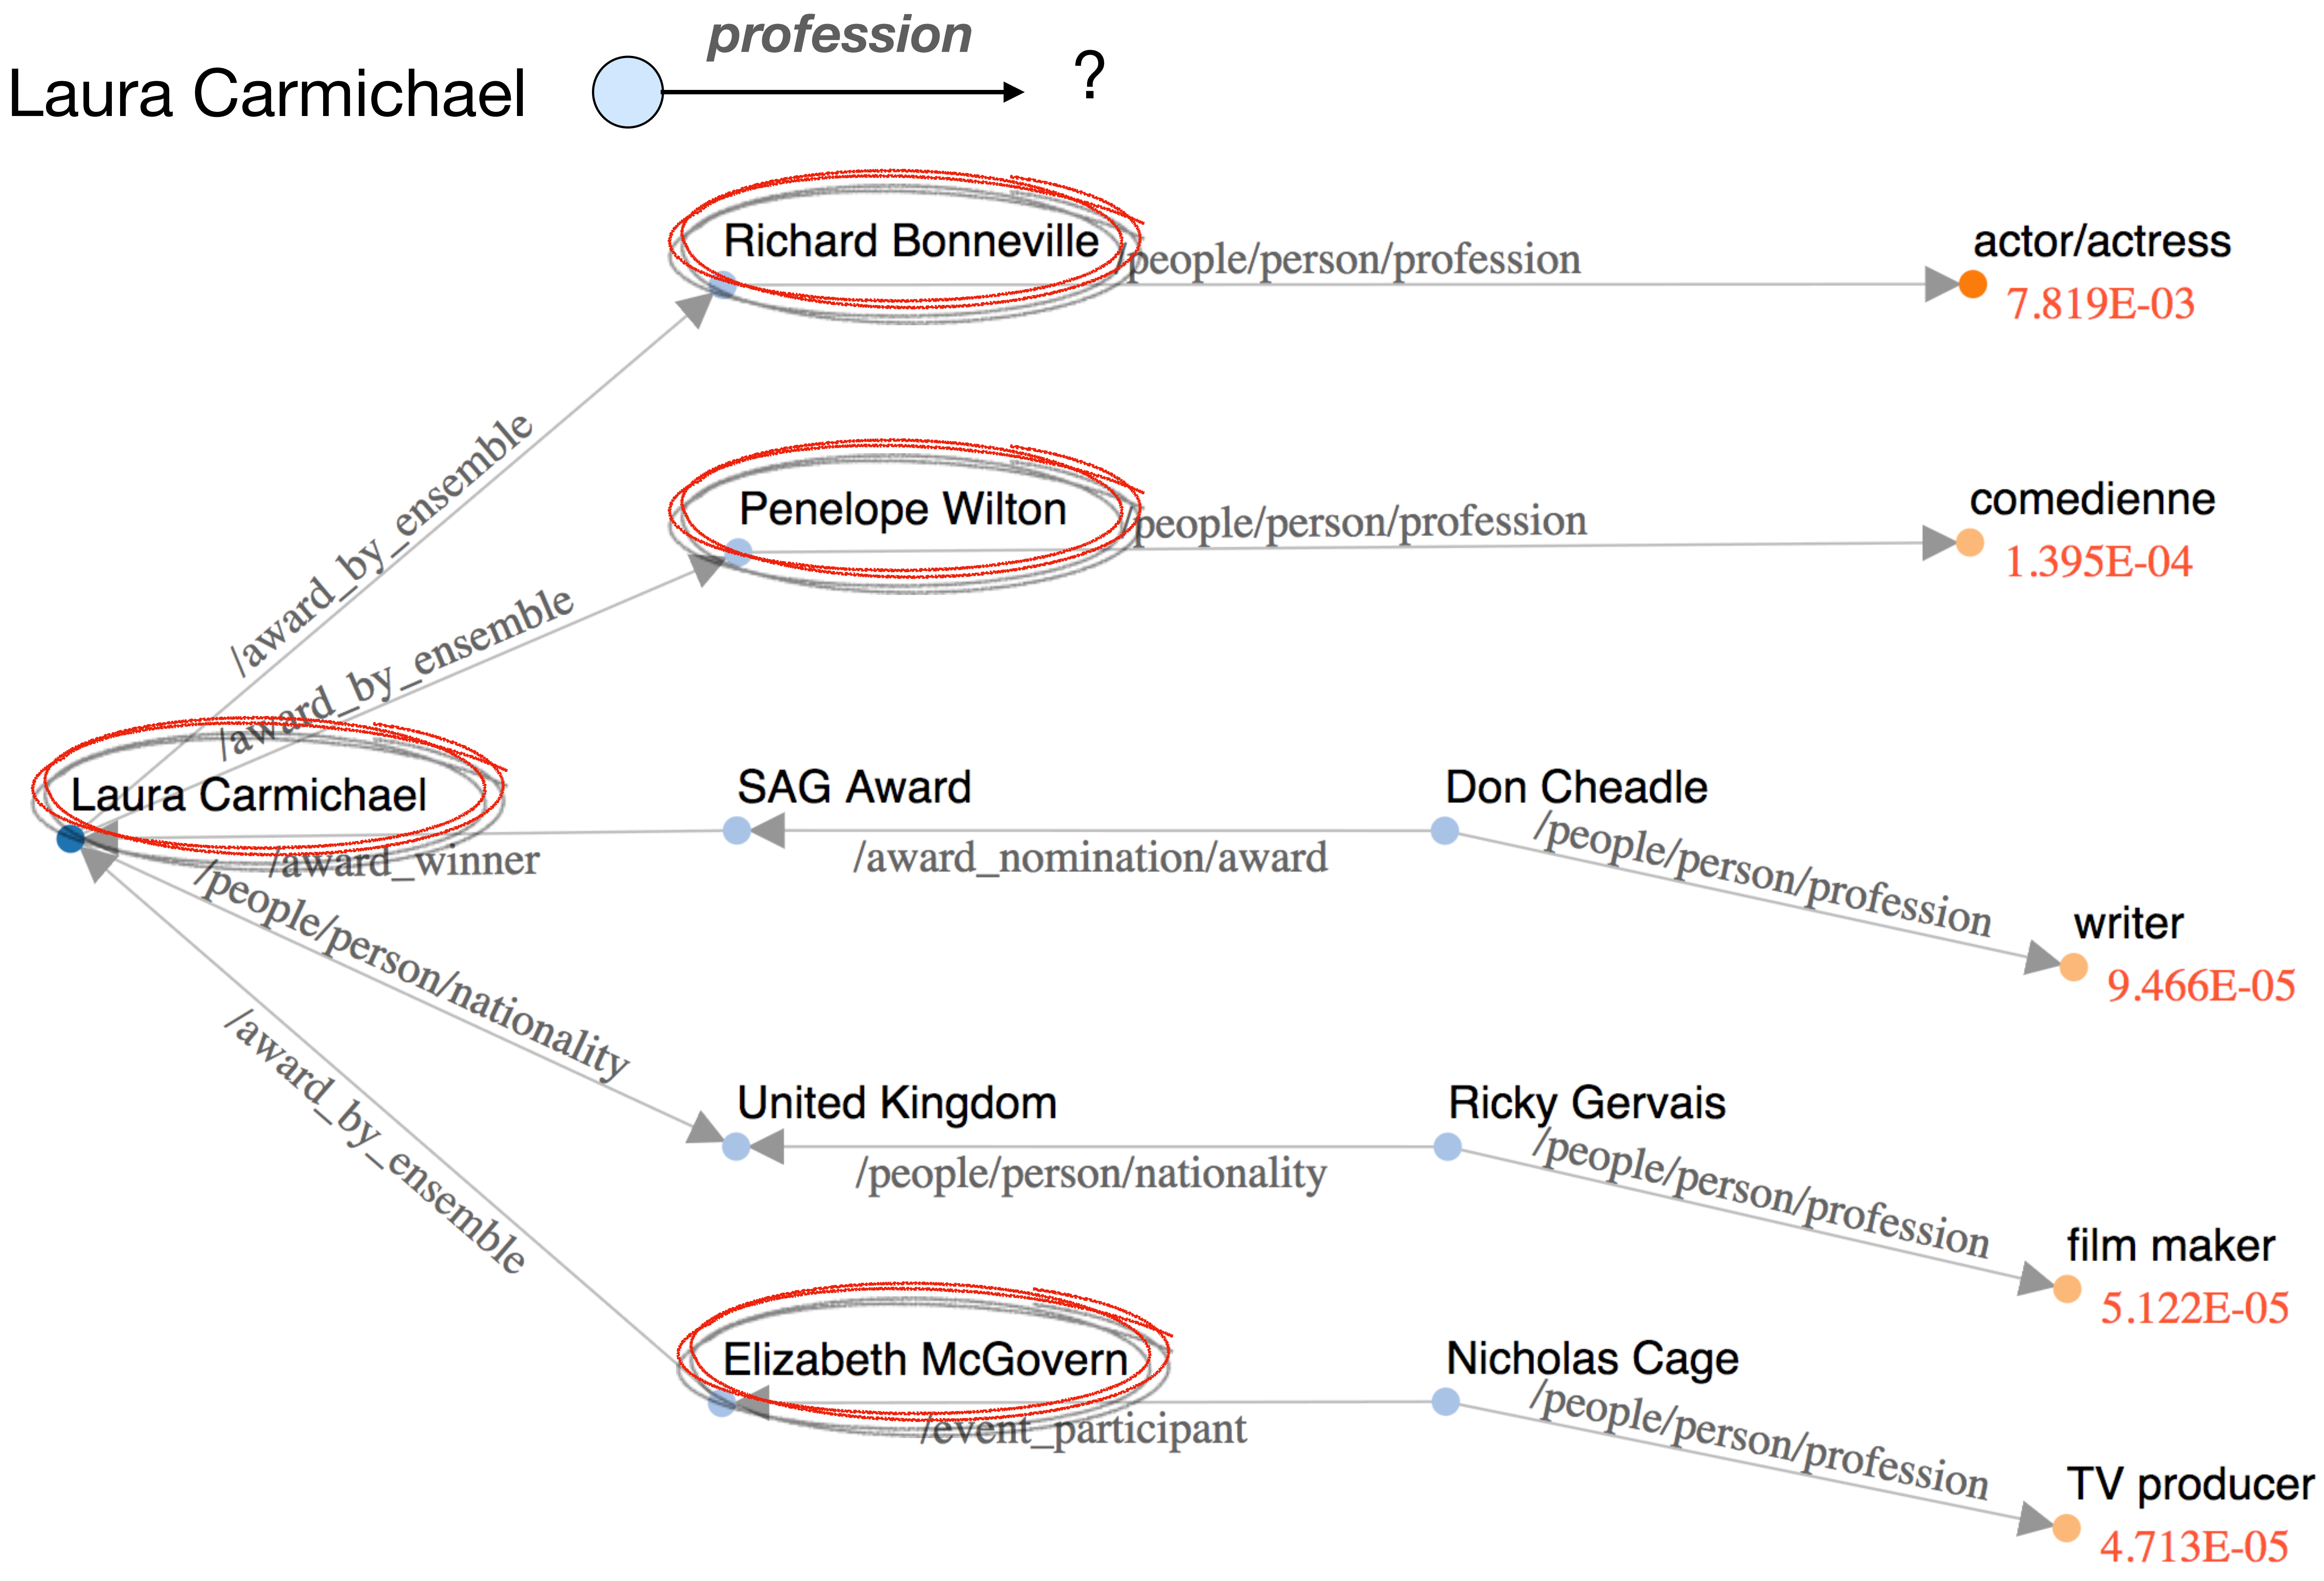

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Interpretable Results

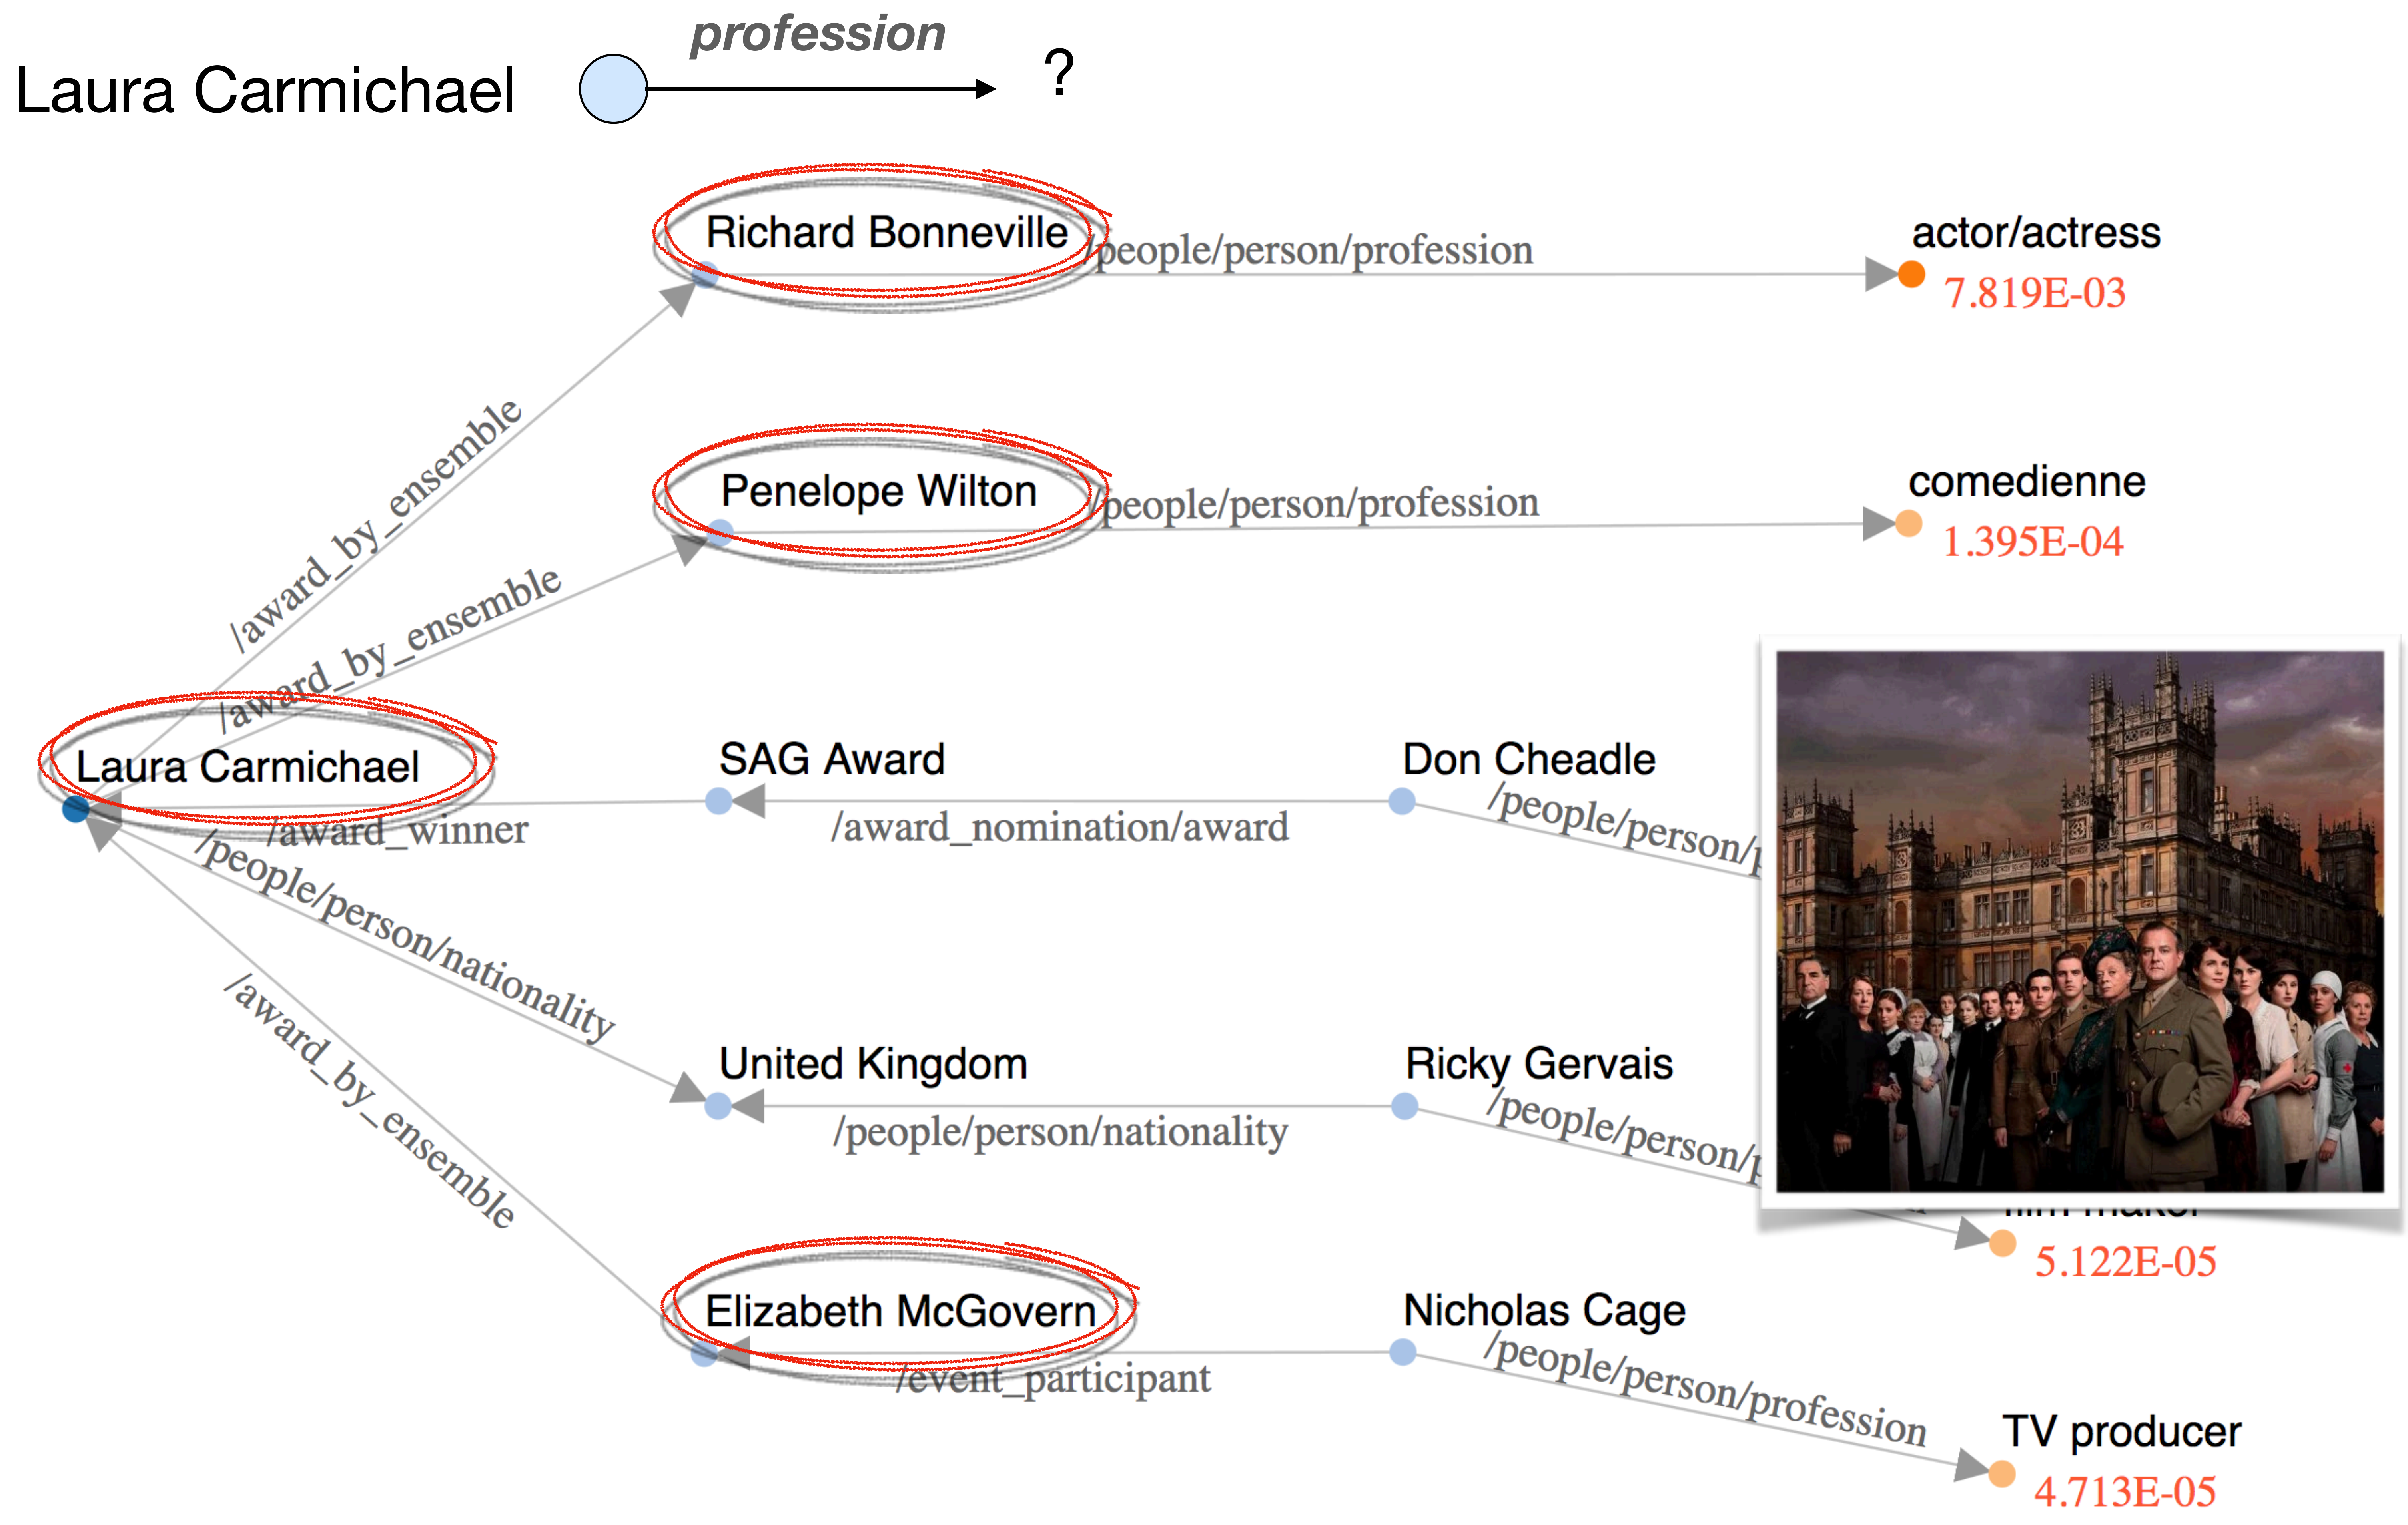

Multi-Hop Knowledge Graph Reasoning with Reward Shaping (Lin et. al. 2018)

Take-aways

- Practical knowledge graphs are incomplete and require automatic completion
- KG embeddings are effective approaches for recovering missing facts, but lack interpretability
- Multi-hop KG inference is interpretable but less accurate
- Multi-hop KG inference with embedding-based reward shaping combines the best of both approaches
- The method can potentially be extended to multi-hop reasoning over joint KG and text to better support downstream AI systems