Knowledge Graph Reasoning: Recent Advances

William Wang

Department of Computer Science

UC SANTA BARBARA

CIPS Summer School 2018

Beijing, China

UCSB

Agenda

- Motivation
- Path-Based Reasoning
- Embedding-Based Reasoning
- Bridging Path-Based and Embedding-Based Reasoning: DeepPath, MINERVA, and DIVA
- Conclusions
- Other Research Activities at UCSB NLP

UCSB

Knowledge Graphs are Not Complete

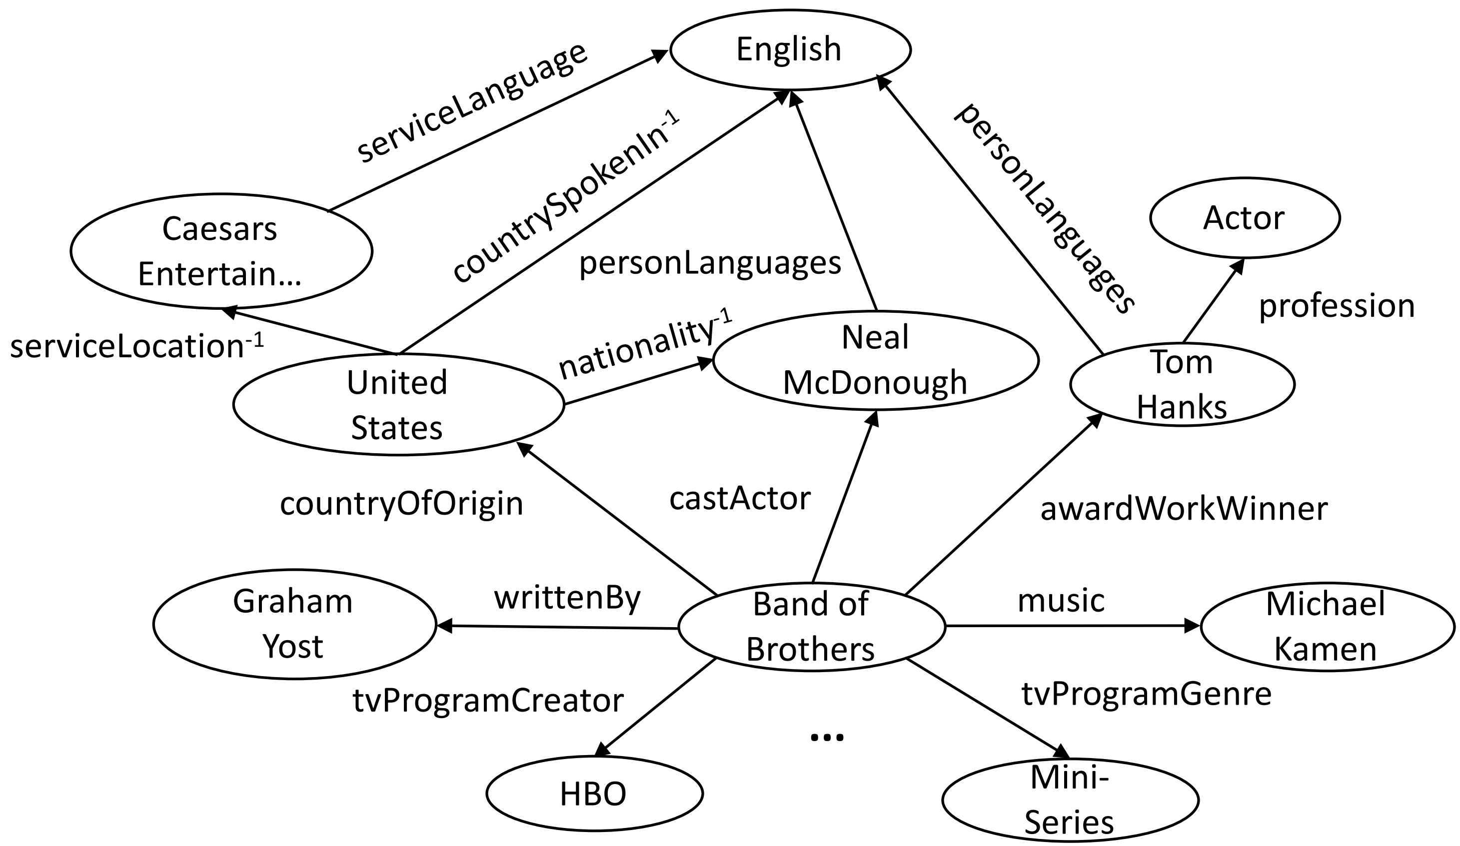

UCSB

Benefits of Knowledge Graph

- Support various applications
  - Structured Search
  - Question Answering
  - Dialogue Systems
  - Relation Extraction
  - Summarization
- Knowledge Graphs can be constructed via information extraction from text, but…
  - There will be a lot of missing links.
  - Goal: complete the knowledge graph.

UCSB

Reasoning on Knowledge Graph

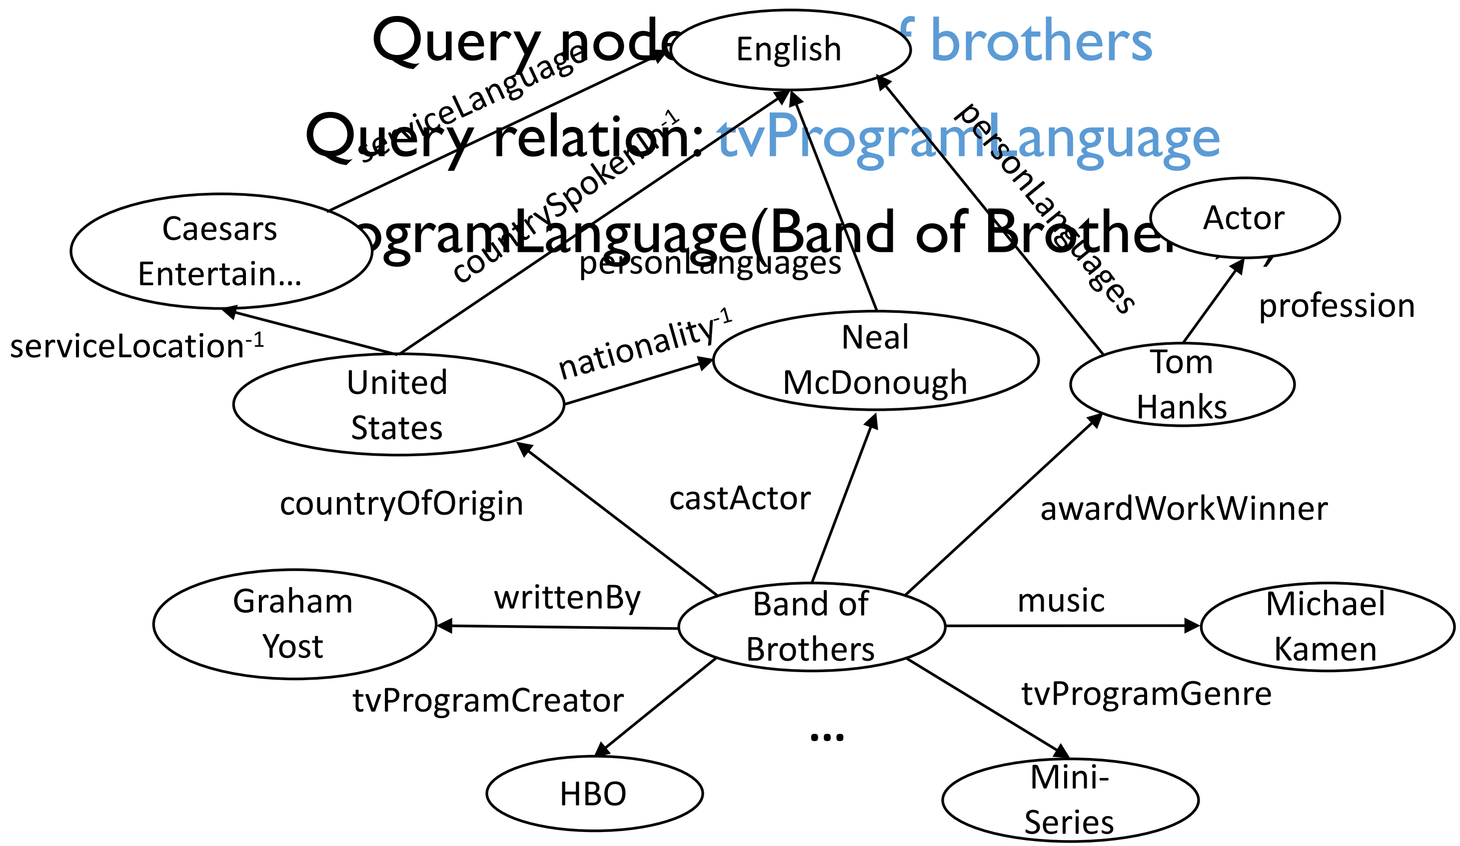

UCSB

KB Reasoning Tasks

- Predicting the missing link.
  - Given e1 and e2, predict the relation r.
- Predicting the missing entity.
  - Given e1 and relation r, predict the missing entity e2.
- Fact Prediction.
  - Given a triple, predict whether it is true or false.

UCSB

Related Work

- **Path-based methods**
  - Path-Ranking Algorithm, Lao et al. 2011
  - ProPPR, Wang et al, 2013 (My PhD thesis)
  - Subgraph Feature Extraction, Gardner et al, 2015
  - RNN + PRA, Neelakantan et al, 2015
  - Chains of Reasoning, Das et al, 2017

Why do we need path-based methods?

It's accurate and explainable!

UCSB

Path-Ranking Algorithm (Lao et al., 2011)

- 1. Run random walk with restarts to derive many paths.
- 2. Use supervised training to rank different paths.

teamHomeStadium

$$c \xrightarrow{\text{teamPlaysInCity}} c \xrightarrow{\text{cityStadiums}} c$$

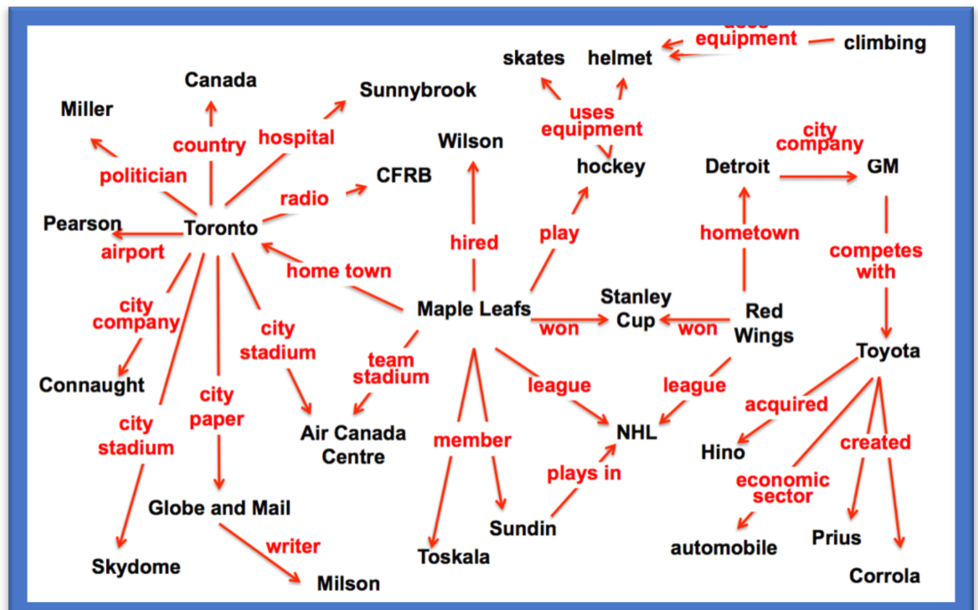

ProPPR (Wang et al., 2013;2015)

- ProPPR generalizes PRA with recursive probabilistic logic programs.
- You may use other relations to jointly infer this target relation.

```
about(X,Z):- handLabeled(X,Z)                # base
about(X,Z):- sim(X,Y),about(Y,Z)             # prop
sim(X,Y):- link(X,Y)                         # sim,link
sim(X,Y):-
    hasWord(X,W),hasWord(Y,W),
    linkedBy(X,Y,W)                          # sim,word
linkedBy(X,Y,W):- true                       # by(W)
```

UCSB

Chain of Reasoning (Das et al, 2017)

- 1. Use PRA to derive the path.
- 2. Use RNNs to perform reasoning of the target relation.

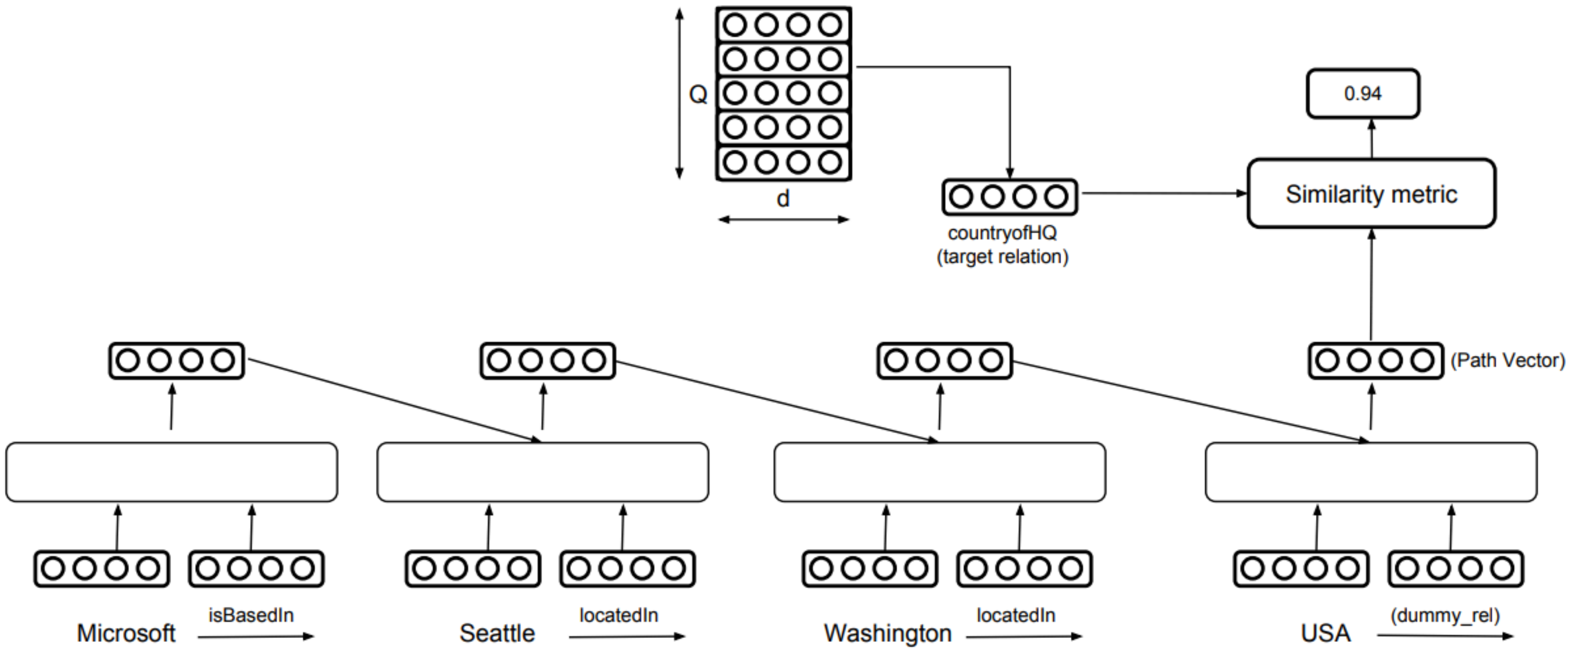

UCSB

Related Work

- **Embedding-based method**
  - RESCAL, Nickel et al, 2011
  - TransE, Bordes et al, 2013
  - Neural Tensor Network, Socher et al, 2013
  - TransR/CTransR, Lin et al, 2015
  - Complex Embeddings, Trouillon et al, 2016

Embedding methods allow us to compare, and find similar entities in the vector space.

UCSB

Bridging Path-Based and Embedding-Based Reasoning with Deep Reinforcement Learning: DeepPath (Xiong et al., 2017)

UCSB

RL for KB Reasoning: DeepPath (Xiong et al., 2017)

- Learning the paths with RL, instead of using random walks with restart
- Model the path finding as a MDP
- Train a RL agent to find paths
- Represent the KG with pretrained KG embeddings
- Use the learned paths as logical formulas

UCSB

Machine Learning

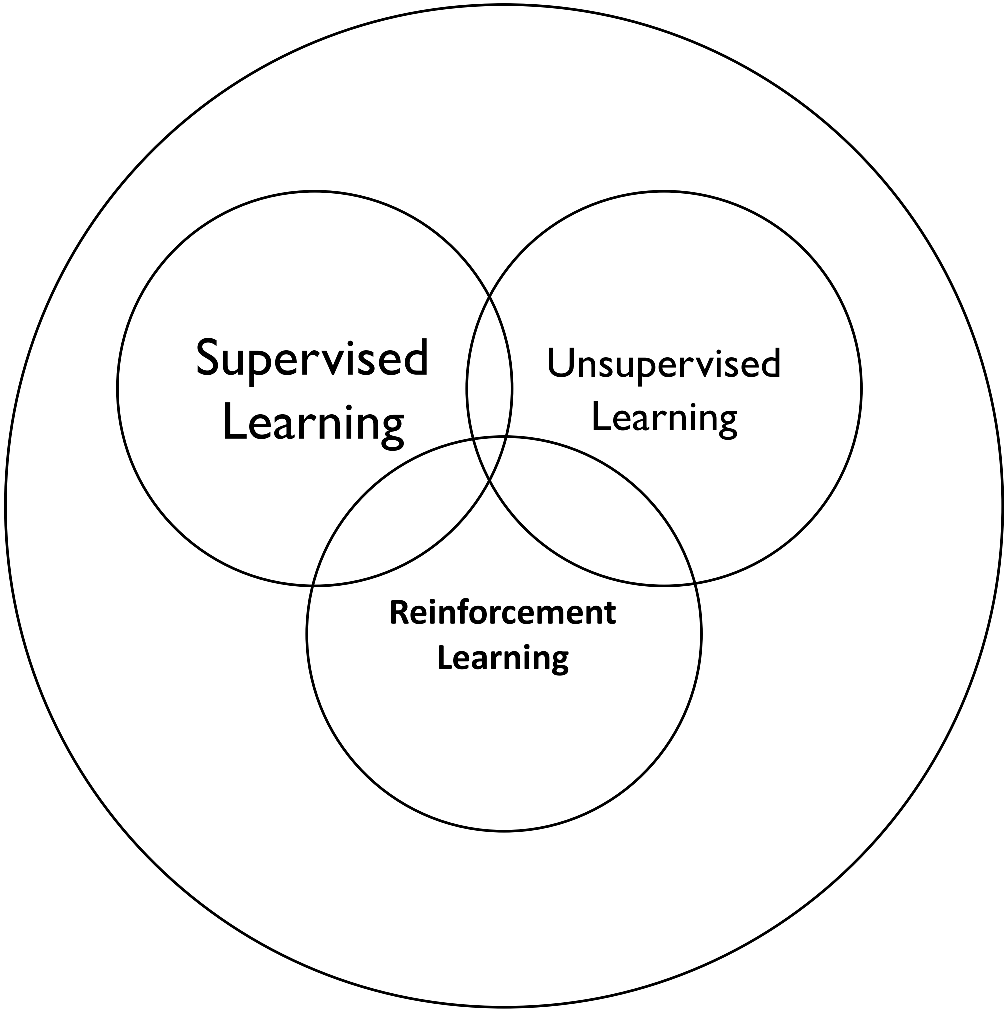

UCSB

Supervised v.s. Reinforcement

Supervised Learning

- Training based on supervisor/label/annotation
- Feedback is instantaneous
- Not much temporal aspects

Reinforcement Learning

- Training only based on reward signal
- Feedback is delayed
- Time matters
- Agent actions affect subsequent exploration

UCSB

Reinforcement Learning

- RL is a general purpose framework for **decision making**
  - RL is for an *agent* with the capacity to *act*
  - Each *action* influences the agent's future *state*
  - Success is measured by a scalar *reward* signal
  - Goal: *select actions to maximize future reward*

UCSB

Reinforcement Learning

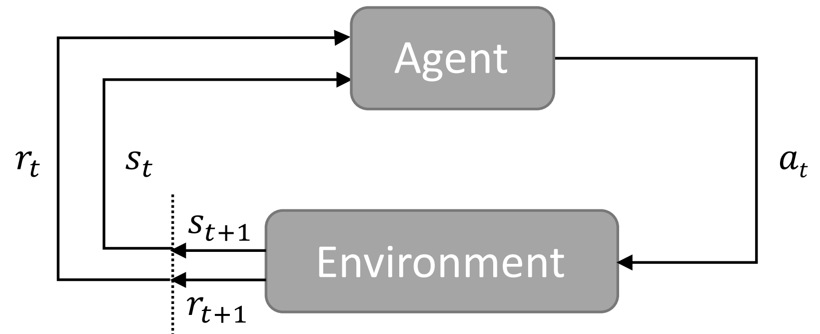

Agent → Multi-layer neural nets $\psi(s_t)$

Environment → KG modeled as a MDP

DeepPath: RL for KG Reasoning

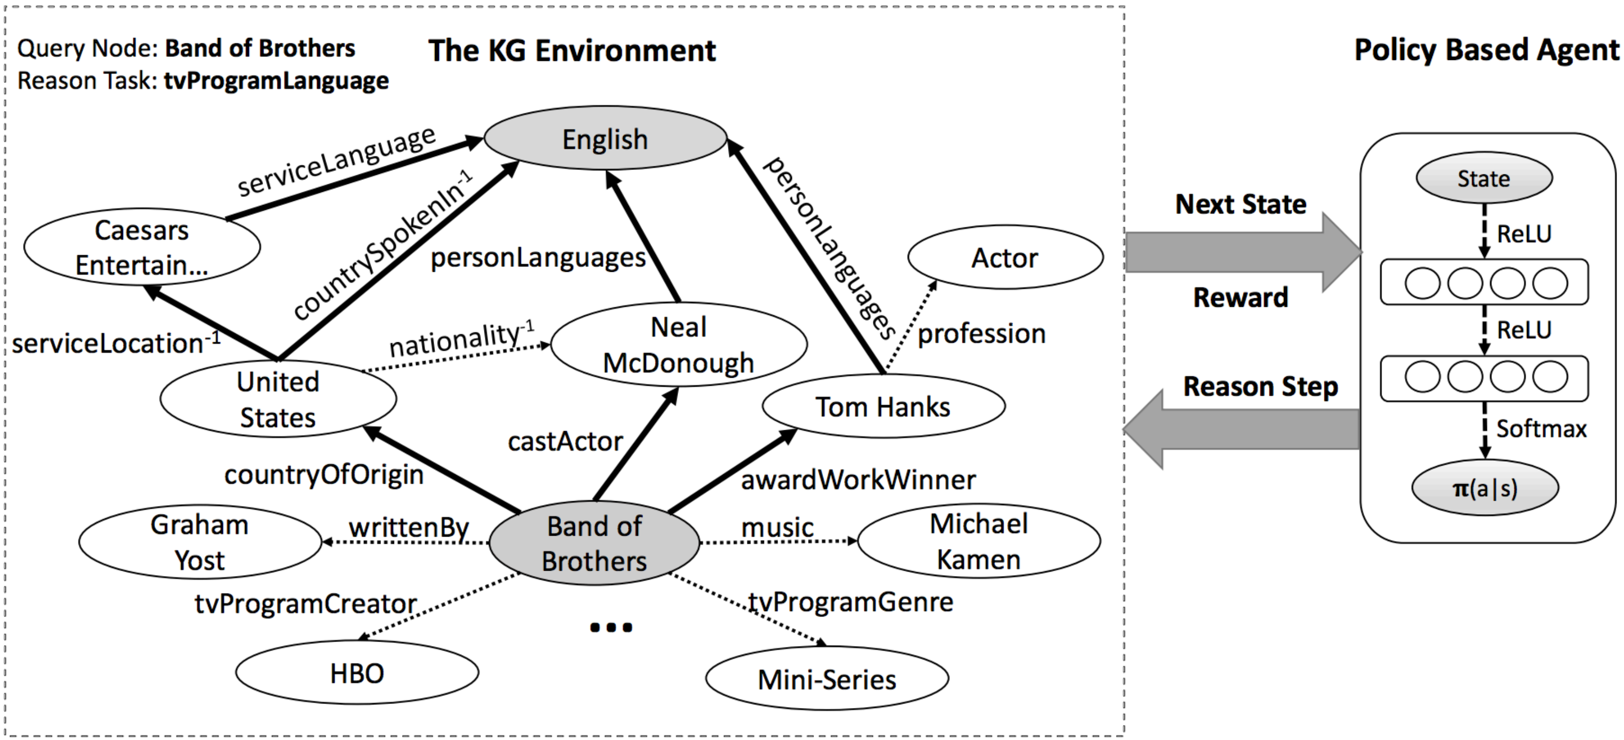

UCSB

Components of MDP

- Markov decision process $< S, A, P, R >$
  - S: continuous states represented with embeddings
  - A: action space (relations)
  - $P(S_{t+1} = s' | S_t = s, A_t = a)$: transition probability
  - $R(s, a)$: reward received for each taken step

- With pretrained KG embeddings
  - $s_t = e_t \oplus (e_{target} - e_t)$
  - $A = \{r_1, r_2, \ldots, r_n\}$, all relations in the KG

UCSB

Reward Functions

- Global Accuracy

$$r_{\text{GLOBAL}} = \begin{cases} +1, & \text{if the path reaches } e_{target} \\ -1, & \text{otherwise} \end{cases}$$

- Path Efficiency

$$r_{\text{EFFICIENCY}} = \frac{1}{length(p)}$$

- Path Diversity

$$r_{\text{DIVERSITY}} = -\frac{1}{|F|} \sum_{i=1}^{|F|} cos(\mathbf{p}, \mathbf{p}_i)$$

UCSB

Training with Policy Gradient

- Monte-Carlo Policy Gradient (REINFORCE, William, 1992)

$$\nabla_\theta J(\theta) = \sum_t \sum_{a \in \mathcal{A}} \pi(a|s_t;\theta) \nabla_\theta \log \pi(a|s_t;\theta) R(s_t, a_t)$$
$$\approx \nabla_\theta \sum_t \log \pi(a = r_t|s_t;\theta) R(s_t, a_t)$$

$$R(s_t, a_t) = \lambda_1 r_{global} + \lambda_2 r_{efficiency} + \lambda_3 r_{diversity}$$

Challenge

- Typical RL problems
  - Atari games (Mnih et al., 2015): 4~18 valid actions
  - AlphaGo (Silver et al. 2016): ~250 valid actions
  - Knowledge Graph reasoning: >= 400 actions

Issue:

- large action (search) space -> poor convergence properties

UCSB

Supervised (Imitation) Policy Learning

- Use randomized BFS to retrieve a few paths
- Do imitation learning using the retrieved paths
- All the paths are assigned with +1 reward

$$\nabla_\theta J(\theta) = \sum_t \sum_{a \in \mathcal{A}} \pi(a|s_t;\theta) \nabla_\theta \log \pi(a|s_t;\theta)$$

$$\approx \nabla_\theta \sum_t \log \pi(a = r_t | s_t; \theta)$$

UCSB

Datasets and Preprocessing

Dataset	# of Entities	# of Relations	# of Triples	# of Tasks
FB15k-237	14,505	237	310,116	20
NELL-995	75,492	200	154,213	12

FB15k-237: Sampled from FB15k (Bordes et al., 2013), redundant relations removes

NELL-995: Sampled from the 995th iteration of NELL system (Carlson et al., 2010b)

- **Dataset processing**
  - Remove useless relations: *haswikipediaurl*, *generalizations, etc*
  - Add inverse relation links to the knowledge graph
  - Remove the triples with task relations

UCSB

Effect of Supervised Policy Learning

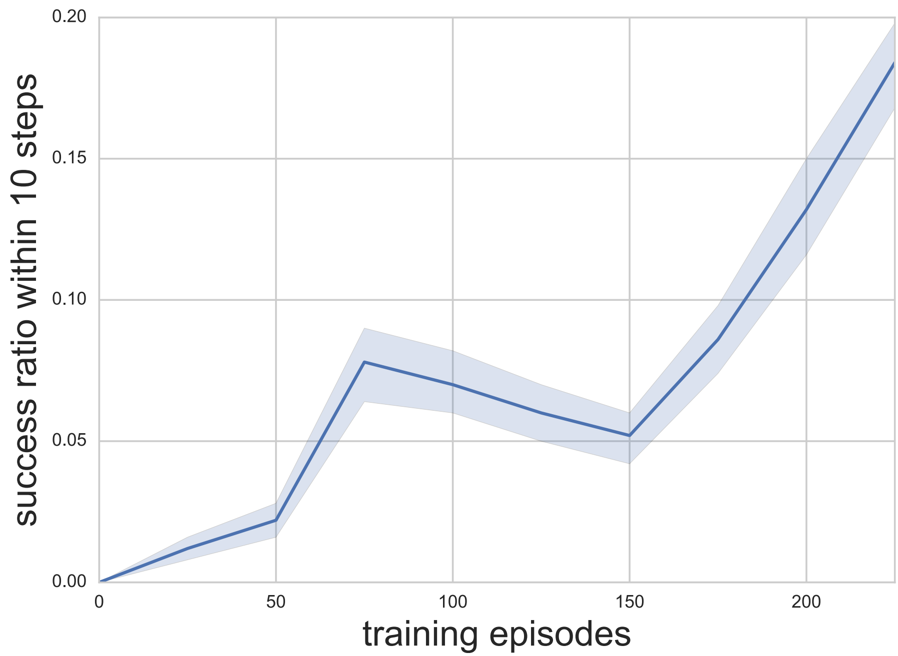

- **x-axis:** number of training epochs
- **y-axis:** success ratio (probability of reaching the target) on test set

-> Re-train the agent using reward functions

UCSB

Inference Using Learned Paths

- Path as logical formula
  - **FilmCountry:** actionFilm^{-1} -> personNationality
  - **PersonNationality:** placeOfBirth -> locationContains^{-1}
  - etc …
- Bi-directional path-constrained search
  - Check whether the formulas hold for entity pairs

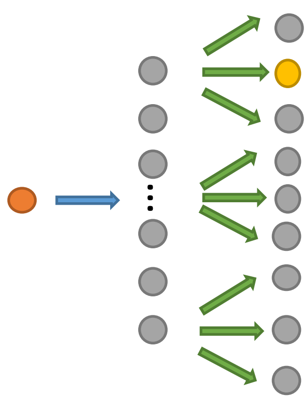

Uni-directional search

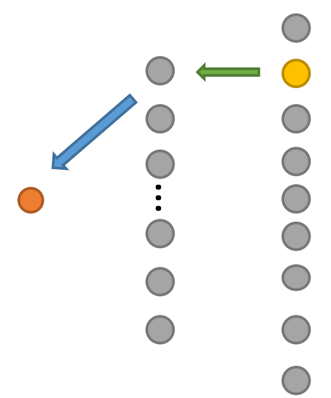

bi-directional search

UCSB

Link Prediction Result

Tasks	PRA	Ours	TransE	TransR
worksFor	0.681	**0.711**	0.677	0.692
atheletPlaysForTeam	**0.987**	0.955	0.896	0.784
athletePlaysInLeague	0.841	**0.960**	0.773	0.912
athleteHomeStadium	0.859	**0.890**	0.718	0.722
teamPlaysSports	0.791	0.738	0.761	**0.814**
orgHirePerson	0.599	**0.742**	0.719	0.737
personLeadsOrg	0.700	**0.795**	0.751	0.772
…				
Overall	0.675	**0.796**	0.737	0.789

Mean average precision on NELL-995

UCSB

Qualitative Analysis

Path length distributions

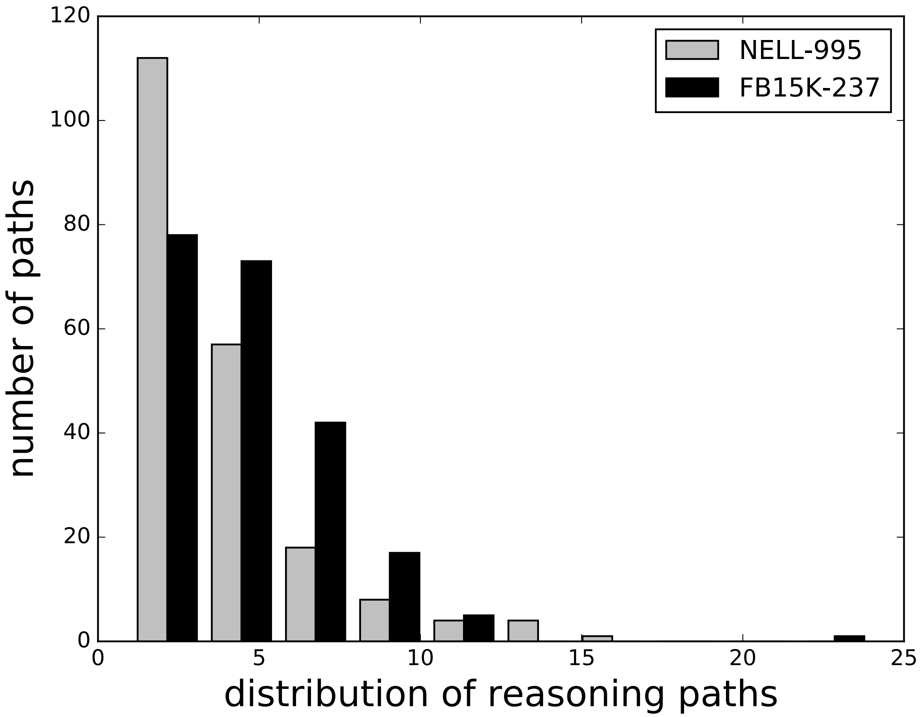

UCSB

Qualitative Analysis

Example Paths

personNationality:
- placeOfBirth -> locationContains^{-1}
- peoplePlaceLived -> locationContains^{-1}
- peopleMariage -> locationOfCeremony -> locationContains^{-1}

tvProgramLanguage:
- tvCountryOfOrigin -> countryOfficialLanguage
- tvCountryOfOrigin -> filmReleaseRegion-1 -> filmLanguage
- tvCastActor -> personLanguage

athletePlaysForTeam:
- athleteHomeStadium -> teamHomeStadium^{-1}
- athletePlaysSports -> teamPlaysSports^{-1}
- atheleteLedSportsTeam

Bridging Path-Finding and Reasoning w. Variational Inference (teaser): DIVA (Chen et al., NAACL 2018)

UCSB

DIVA: Variational KB Reasoning (Chen et al., NAACL 2018)

- Inferring latent paths connecting entity nodes.

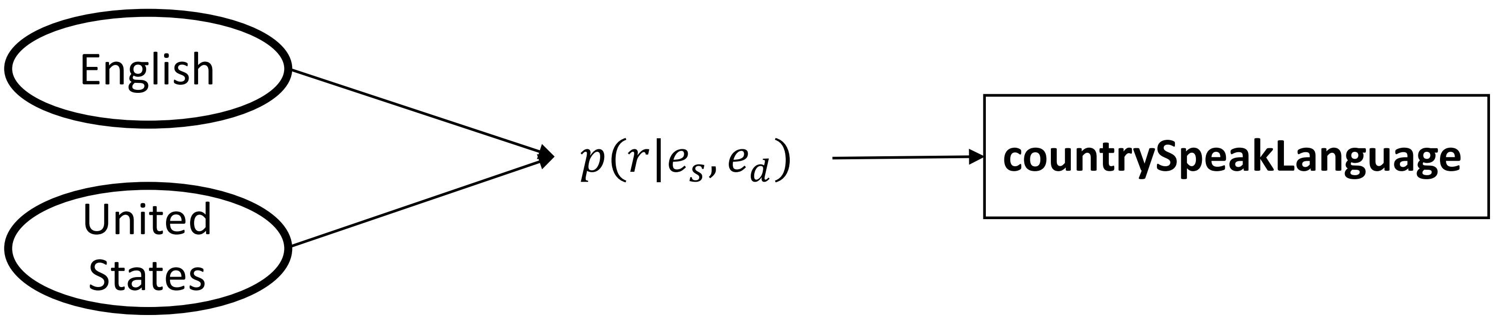

Condition (e_s, e_d)

Observed Variable r

$$\bar{p} = argmax_p \log p(r|e_s, e_d)$$

UCSB

DIVA: Variational KB Reasoning (Chen et al., NAACL 2018)

- Inferring latent paths connecting entity nodes by parameterizing likelihood (path reasoning) and prior (path finding) with neural network modules.

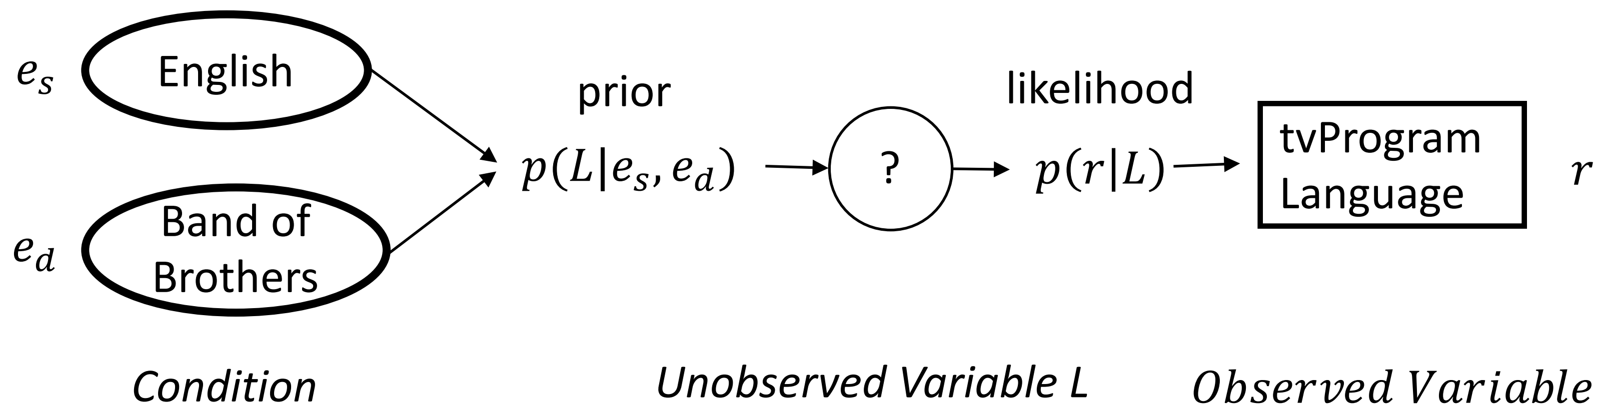

$$p = argmax_p\, p(r|e_s, e_d) = argmax_p \log \int_L^{\infty} p(r|L)p(L|e_s, e_d)$$

UCSB

DIVA: Variational KB Reasoning (Chen et al., NAACL 2018)

- Marginal likelihood $\log \int_L p(r|L) p(L|e_s, e_d)$ is intractable
- We resort to Variational Bayes by introduce a posterior distribution $q(L|e_s, e_d, r)$

$$\log \mathrm{p}(\mathrm{r}|e_s, e_d) \rightarrow \geq \rightarrow ELBO$$

$$ELBO = \mathbb{E}_{q(L|e_s, e_d, r)}[\log p(r|L)] - KL(q(L|e_s, e_d, r)||p(L|e_s, e_d))$$

DP Kingma et al. 2013

Parameterization – Path-finder

- Approximate posterior $q_{\varphi}(L|e_s, e_d, r)$ and prior $p_{\beta}\ (L|e_s, e_d)$: parameterize with RNN

Transition Probability: $p(a_{\tau+1}, e_{\tau+1}|a_{1:\tau}, e_{1:\tau})$

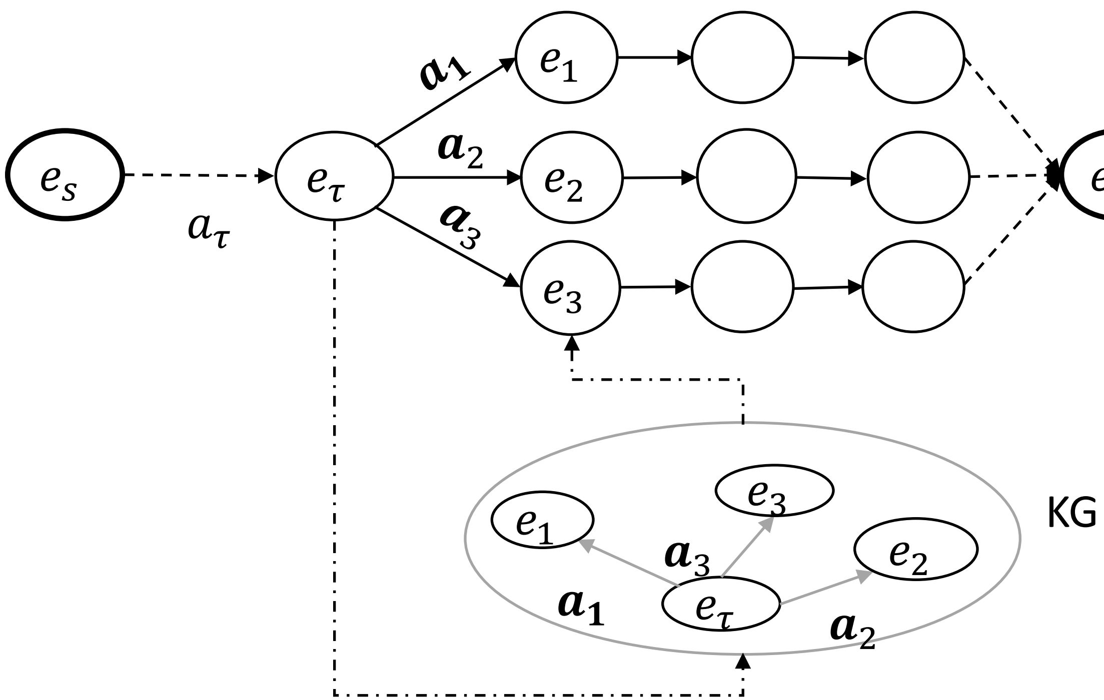

UCSB

Parameterization – Path Reasoner

- Likelihood $p_\theta\ (r|L)$: parameterize with CNN

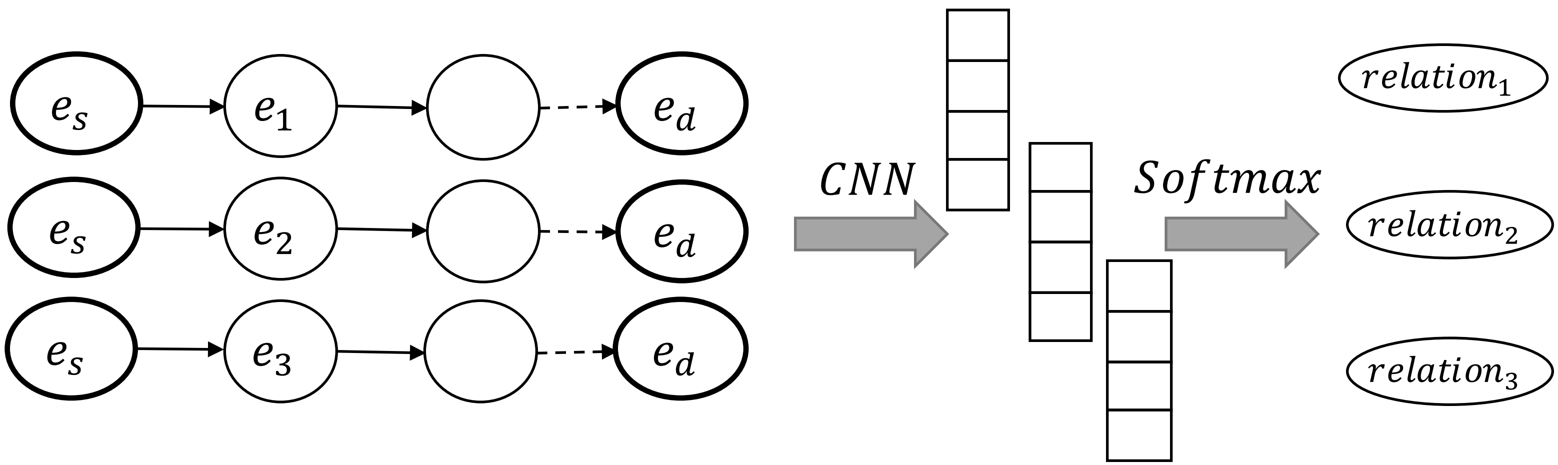

DIVA: Variational KB Reasoning (Chen et al., NAACL 2018)

- Training

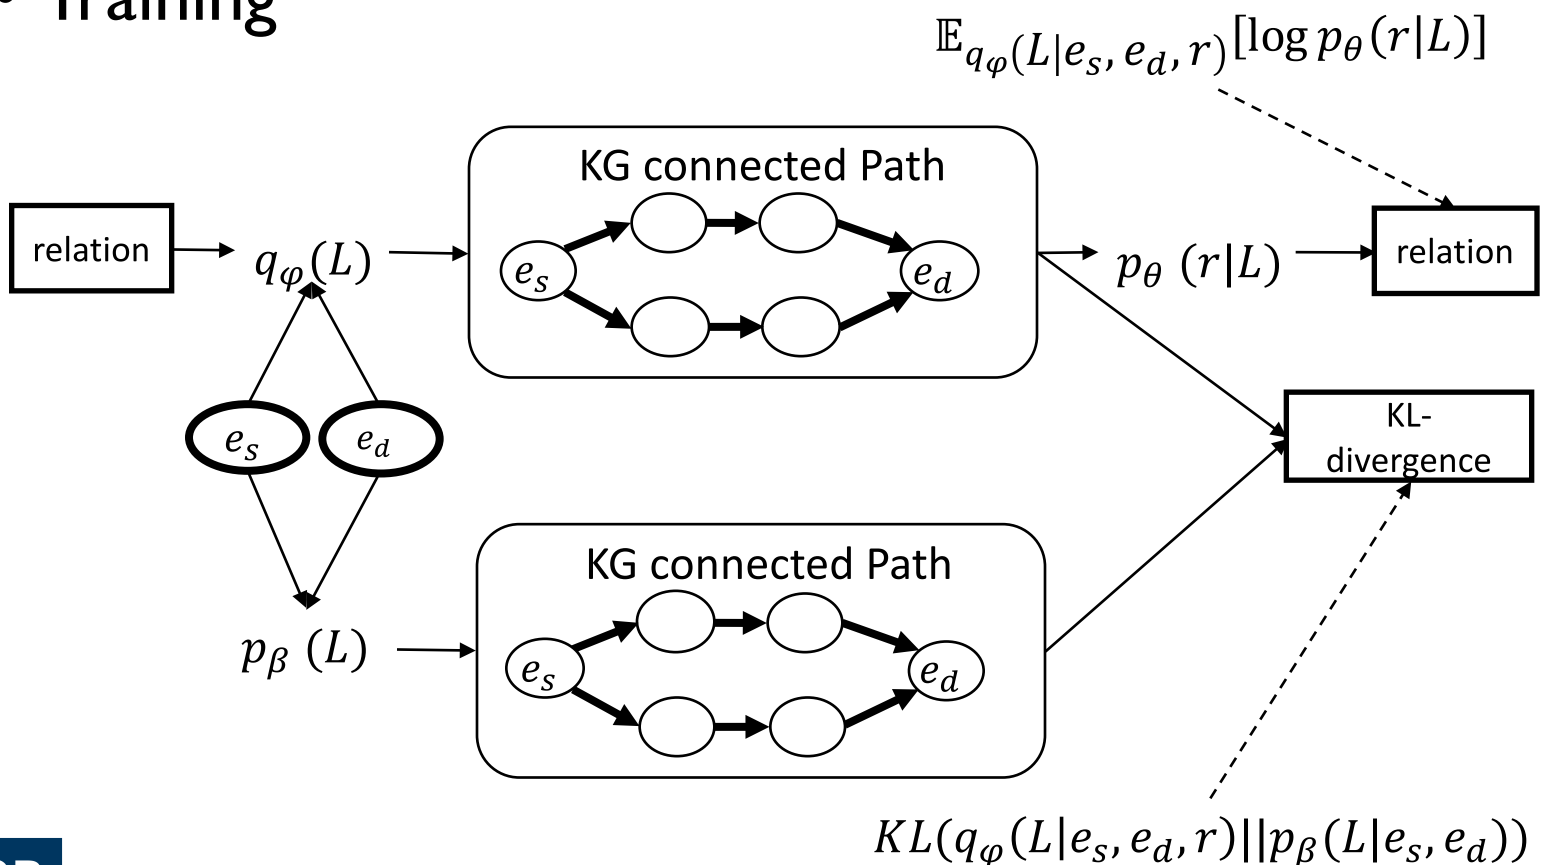

UCSB

$posterior$: q_φ, likelihood: $p_\theta\ (r|L)$, prior: $p_\beta(L|e_s,e_d)$

DIVA: Variational KB Reasoning (Chen et al., NAACL 2018)

- Testing

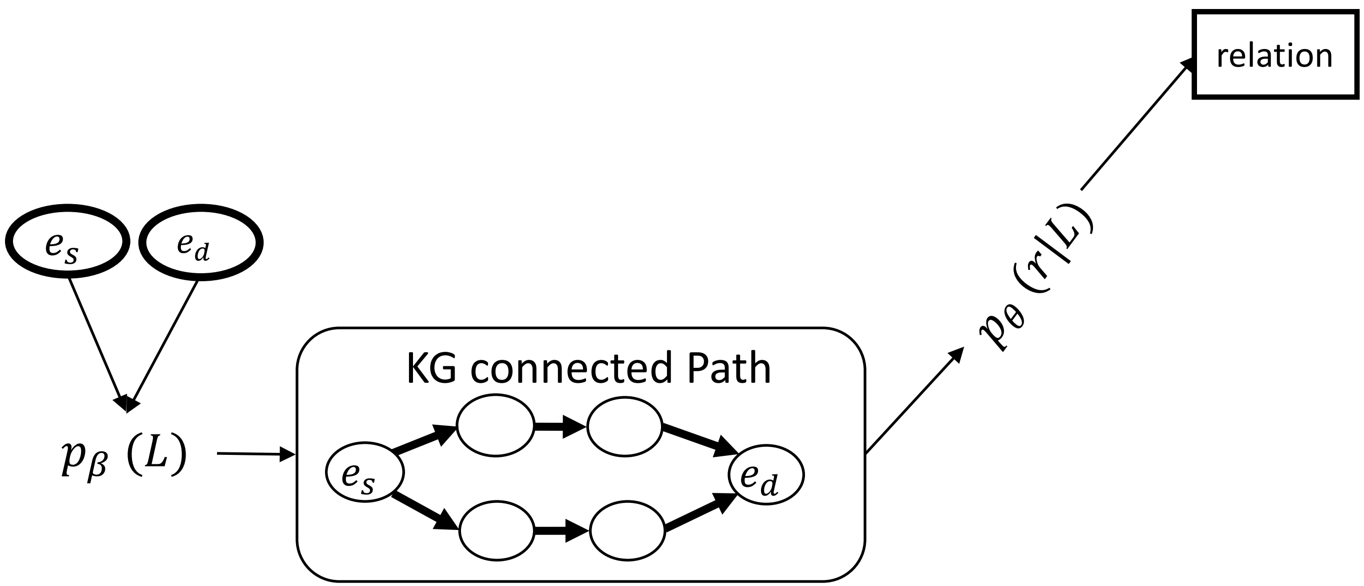

UCSB

$posterior$: q_φ, likelihood: $p_\theta\ (r|L)$, prior: $p_\beta(L|e_s, e_d)$

Conclusions

- Embedding-based methods are very scalable and robust.
- Path-based methods are more interpretable.
- There are some recent efforts in unifying embedding and path-based approaches.
- DIVA integrates path-finding and reasoning in a principled variational inference framework.

UCSB

UCSB NLP

- Natural Language Processing
  - Information Extraction: **relation extraction, and distant supervision**.
  - Summarization: **abstractive summarization.**
  - Social Media: **non-standard English expressions.**
  - Language & Vision: **action/relation detection, and video captioning**.
  - Spoken Language Processing: **task-oriented neural dialogue systems**.
- Machine Learning
  - Statistical Relational Learning: **neural symbolic reasoning.**
  - Deep Learning: **sequence-to-sequence models.**
  - Structure Learning: **learning the structures for neural models.**
  - Reinforcement Learning: **efficient and effective methods for DRL and NLP.**
- Artificial Intelligence
  - Knowledge Representation & Reasoning: **beyond Freebase/OpenIE.**
  - Knowledge Graphs: **construction, completion, and reasoning.**

Other Research Activities at UCSB's NLP Group

Natural Language Generation

Reinforced Conditional Variational Autoencoder for Generating Emotional Sentences (Zhou and Wang, ACL 2018)

https://arxiv.org/abs/1711.04090

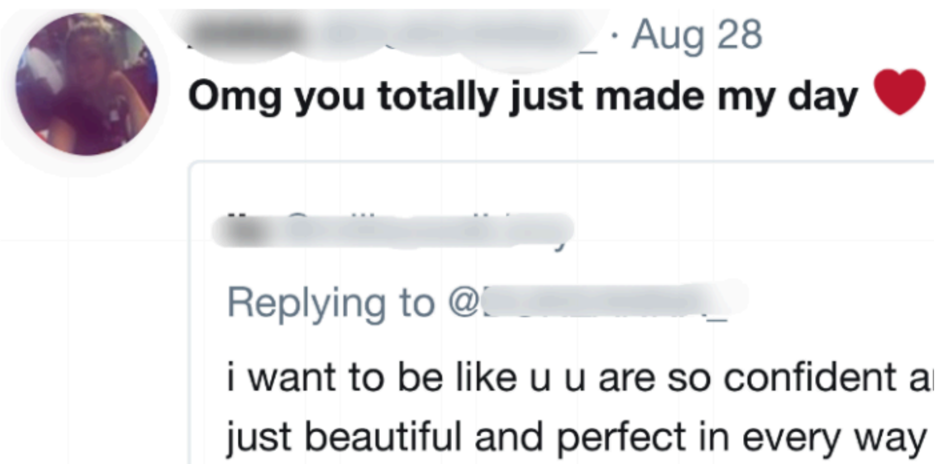

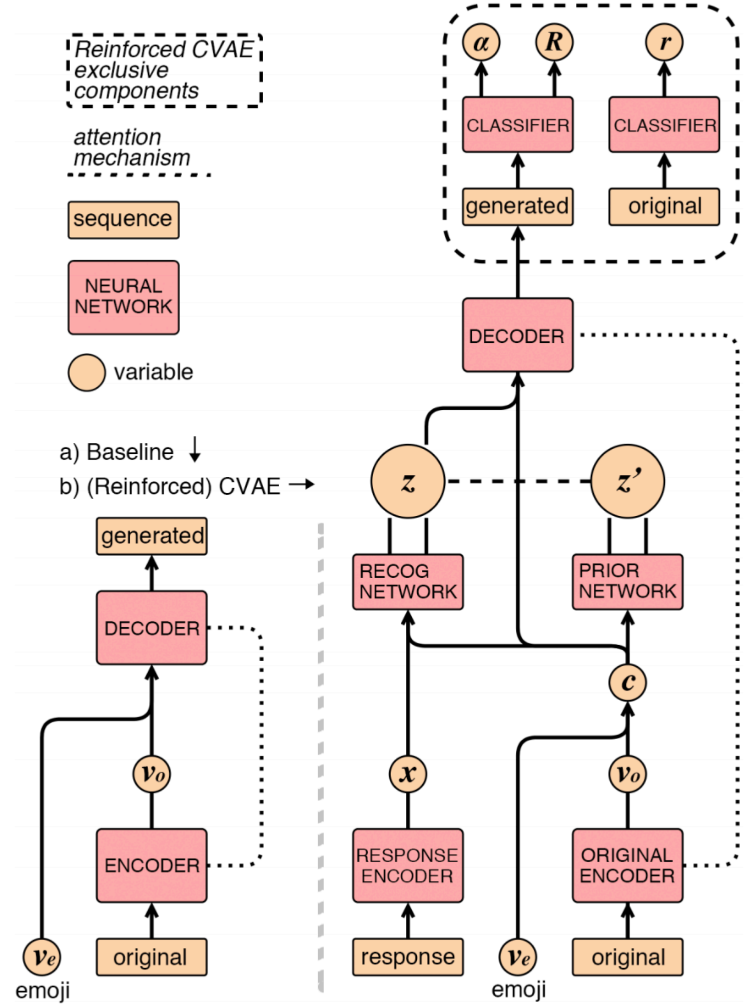

Controlling Emotions for RC-VAE Generated Sentences

User's Input	**sorry guys , was gunna stream tonight but i 'm still feeling sick**	
Designated Emojis	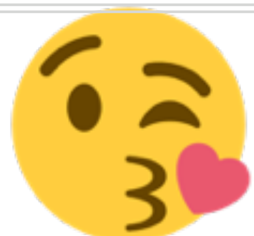	
Generated by Seq2Seq Baseline	**i 'm sorry you 're going to be missed it**	**i 'm sorry for your loss**
Generated By MojiTalk	**hope you are okay hun !**	**hi jason , i 'll be praying for you**

Hierarchical Deep Reinforcement Learning for Video Captioning (Wang et al., CVPR 2018)

Caption #1: A woman offers her dog some food.

Caption #2: A woman is eating and sharing food with her dog.

Caption #3: A woman is sharing a snack with a dog.

Caption: A person sits on a bed and puts a laptop into a bag. The person stands up, puts the bag on one shoulder, and walks out of the room.

Deep Multimodal Video Captioning (Wang et al., NAACL 2018)

Ground Truth: A girl is singing.
A girl sings to a song.
Video Only: A woman is talking in a room.
Video + Audio: A girl is singing a song.

Adversarial Reward Learning for Visual Storytelling (Wang et al., ACL 2018)

XE-ss	We took a trip to the mountains.	There were many different kinds of different kinds.	We had a great time.	He was a great time.	It was a beautiful day.
AREL	The family decided to take a trip to the countryside.	There were so many different kinds of things to see.	The family decided to go on a hike.	I had a great time.	At the end of the day, we were able to take a picture of the beautiful scenery.
Human-created Story	We went on a hike yesterday.	There were a lot of strange plants there.	I had a great time.	We drank a lot of water while we were hiking.	The view was spectacular.

Figure 6: Qualitative comparison example with XE-ss. The direct comparison votes (AREL:XE-ss:Tie) were 5:0:0 on Relevance, 4:0:1 on Expressiveness, and 5:0:0 on Concreteness.

Computational Social Science

Learning to Generate Slang Explanations (Ke Ni, IJCNLP 2017)

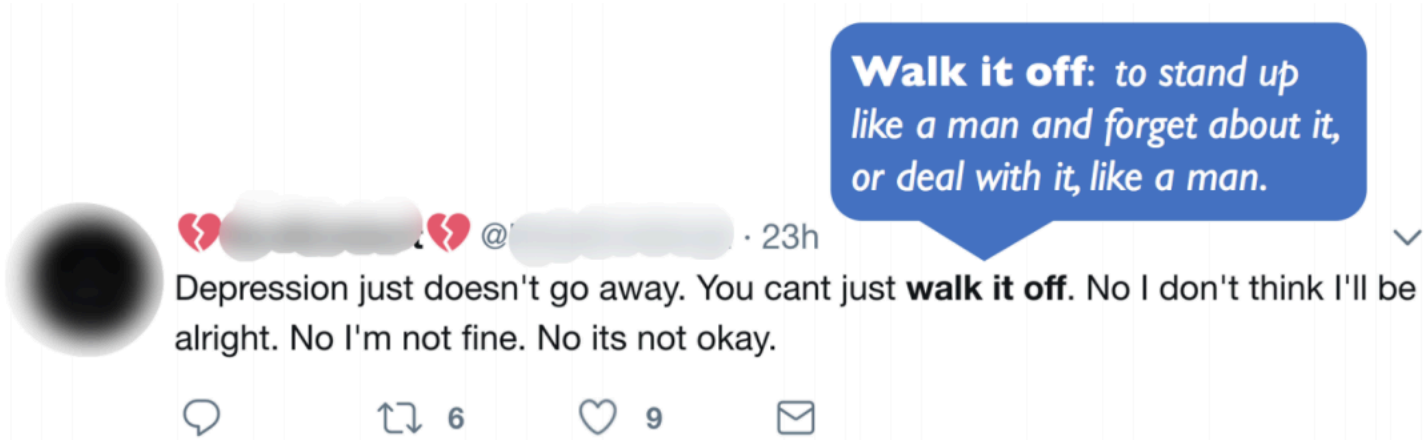

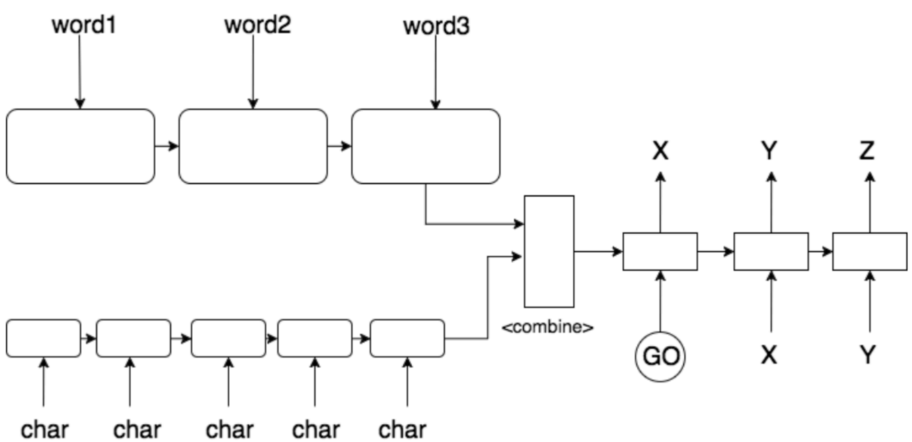

Automatic Generation of Slang Words (Kulkarni and Wang, NAACL 2018)

Word	**Derived From**	**Type**
dink	double income no kids	alphabetism
lambortini	lamborghini + martini	blend
diamat	dialectical + materialism	blend
tude	attitude	clipping (fore)
brill	brilliant	clipping (back)
teenie-weenie	teenie	reduplicative
yik-yak	yik	reduplicative

Leveraging Intra-Speaker and Inter-Speaker Representation Learning for Hate Speech Detection (Qian et al., NAACL 2018)

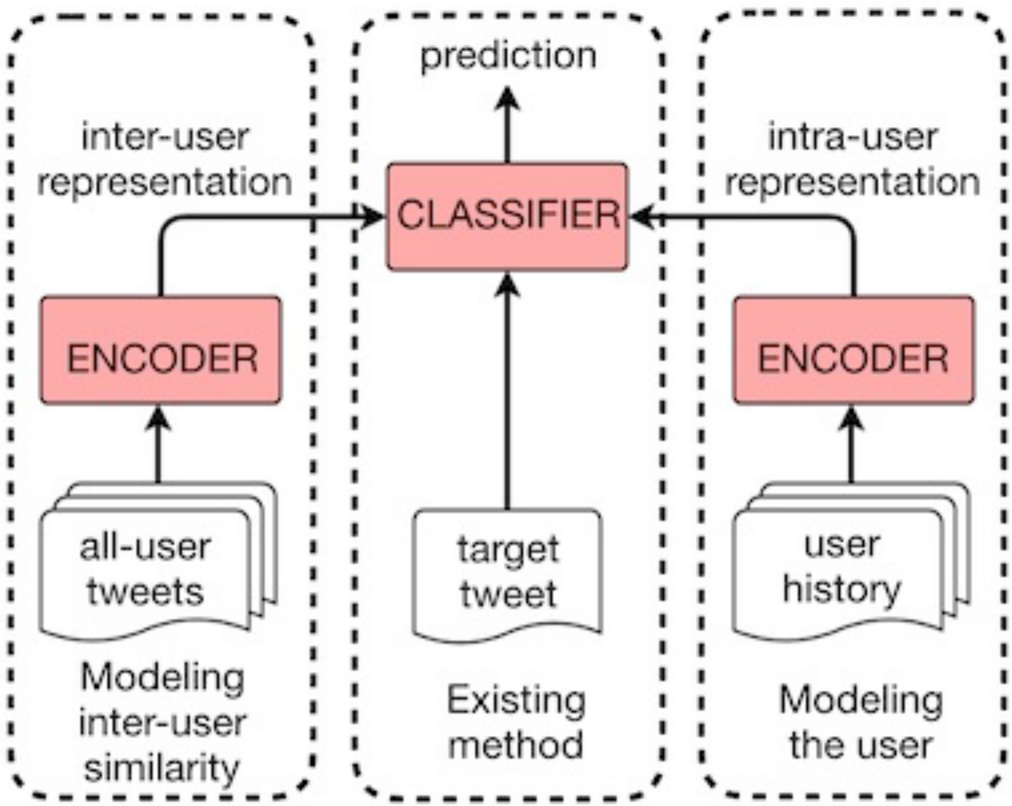

Deep Reinforcement Learning

Scheduled Policy Optimization (Xiong et al., IJCAI 2018)

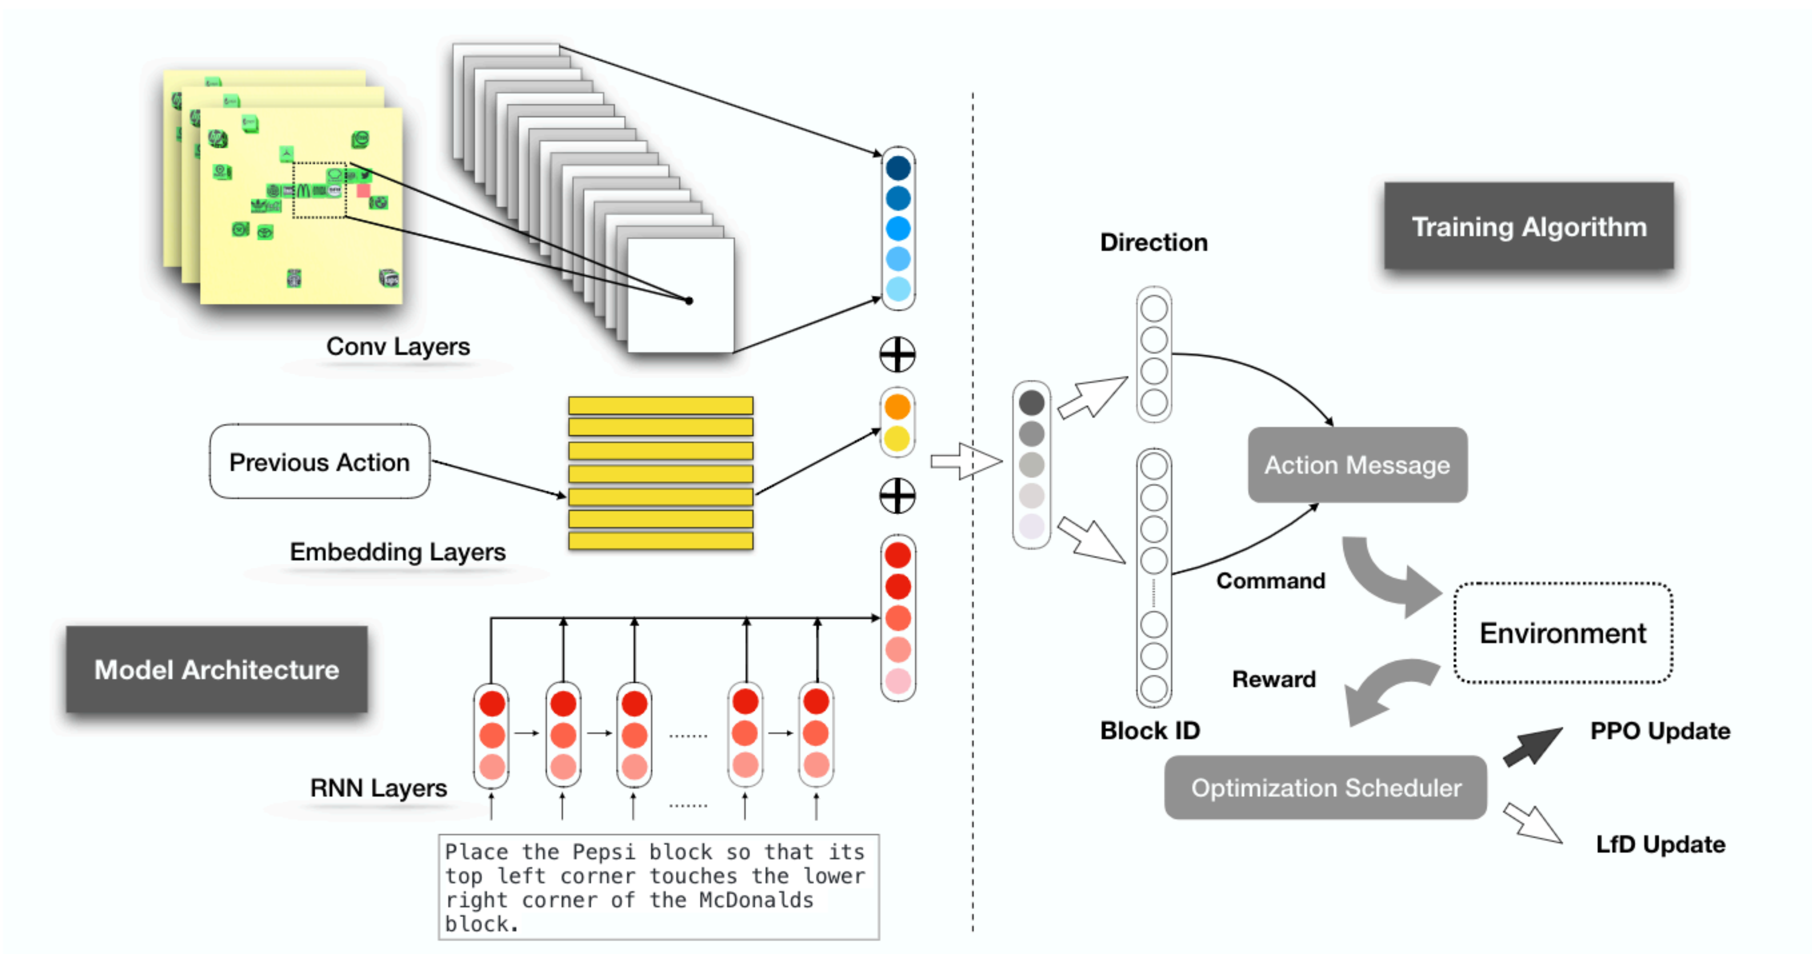

Combining Model-Free and Model Based Deep Reinforcement Learning (Wang et al., ECCV 2018)

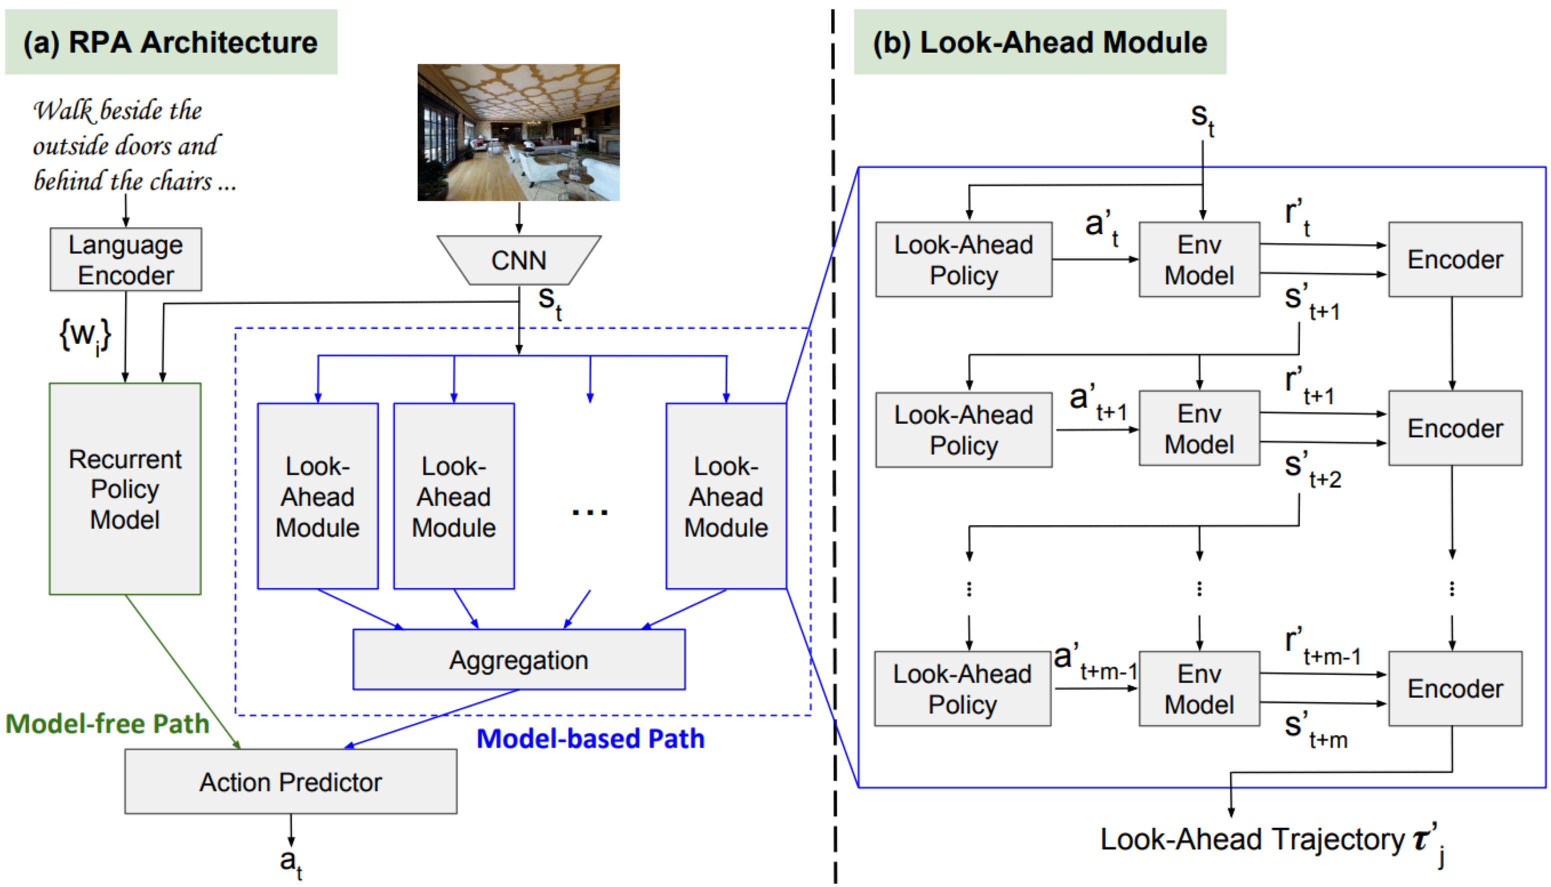

Fig. 2: The overview of our method.

Structure Learning

Learning Activation Functions in Deep Neural Networks (Conner Vercellino, NIPS MetaLearning)

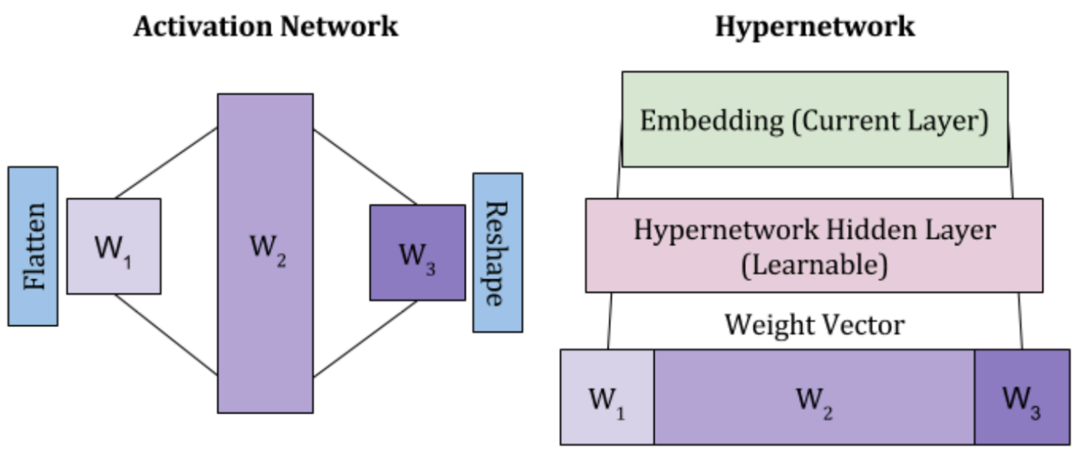

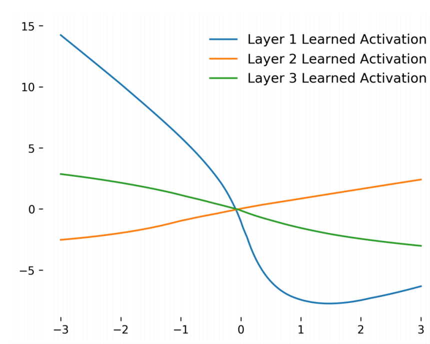

Acknowledgment

Sponsors: Adobe, Amazon, ByteDance, DARPA, Facebook, Google, IBM, LogMeIn, NVIDIA, Tencent.

Sponsors: Adobe, Amazon, ByteDance, DARPA, Facebook, Google, IBM, LogMeIn, NVIDIA, Tencent.

Thanks! Questions?

nlp.cs.ucsb.edu

DeepPath Source code:
https://github.com/xwhan/DeepPath
KBGAN Source code:
https://github.com/cai-lw/KBGAN
Scheduled Policy Optimization:
https://github.com/xwhan/walk_the_blocks
ProPPR Source code:
https://github.com/TeamCohen/ProPPR
AREL Source code:
https://github.com/littlekobe/AREL