Self-Supervised Natural Language Processing

William Wang

Assistant Professor of Computer Science

UC SANTA BARBARA

NLP and State-of-the-Arts

- Language modeling style pre-training took over in 2018.
  - $P(w_n|w_1,w_2...w_{n-1})$
  - Elmo (Peters et al., NAACL Best Paper 2018)
  - BERT (Devlin et al., NAACL Best Paper 2019)

- Wait, but isn't that Tomas Mikolov's 2010 thesis already told us about the effectiveness of RNNLM?
- And we know about LMs for more than 20 years, so what's new?
- One thing we didn't know very well about is the transfer learning capabilities of pre-trained LMs.

Connections to Low-Level Self-Supervised Learning in Vision

- The idea of language modeling, i.e., predicting the next word using previous words, or remove a word and use context to predict the target word…

is fairly similar to

- Randomly remove a patch, and predict the relative position of the missing patch (Doersch et al., ICCV 2015).
- Remove an object from an image, and use context to predict the object.
- Remove key frames in videos, and use context to predict the missing frame.

But How do We Get to Task-Level Self-Supervised Learning like AlphaGo Zero?

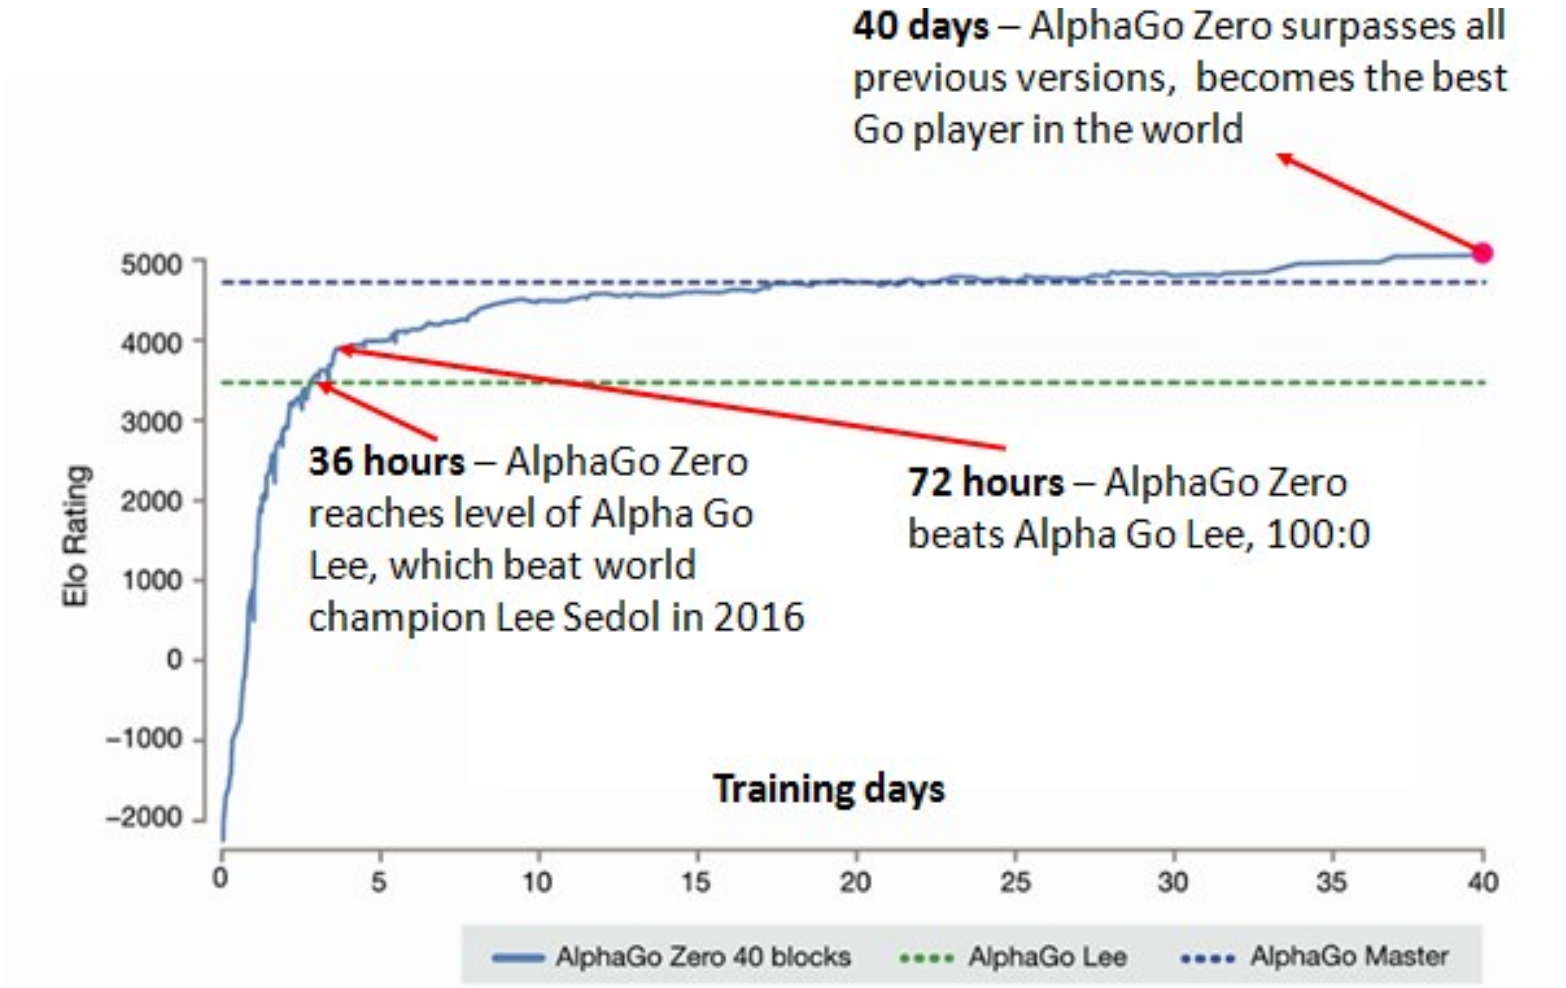

Three Areas of Machine Learning

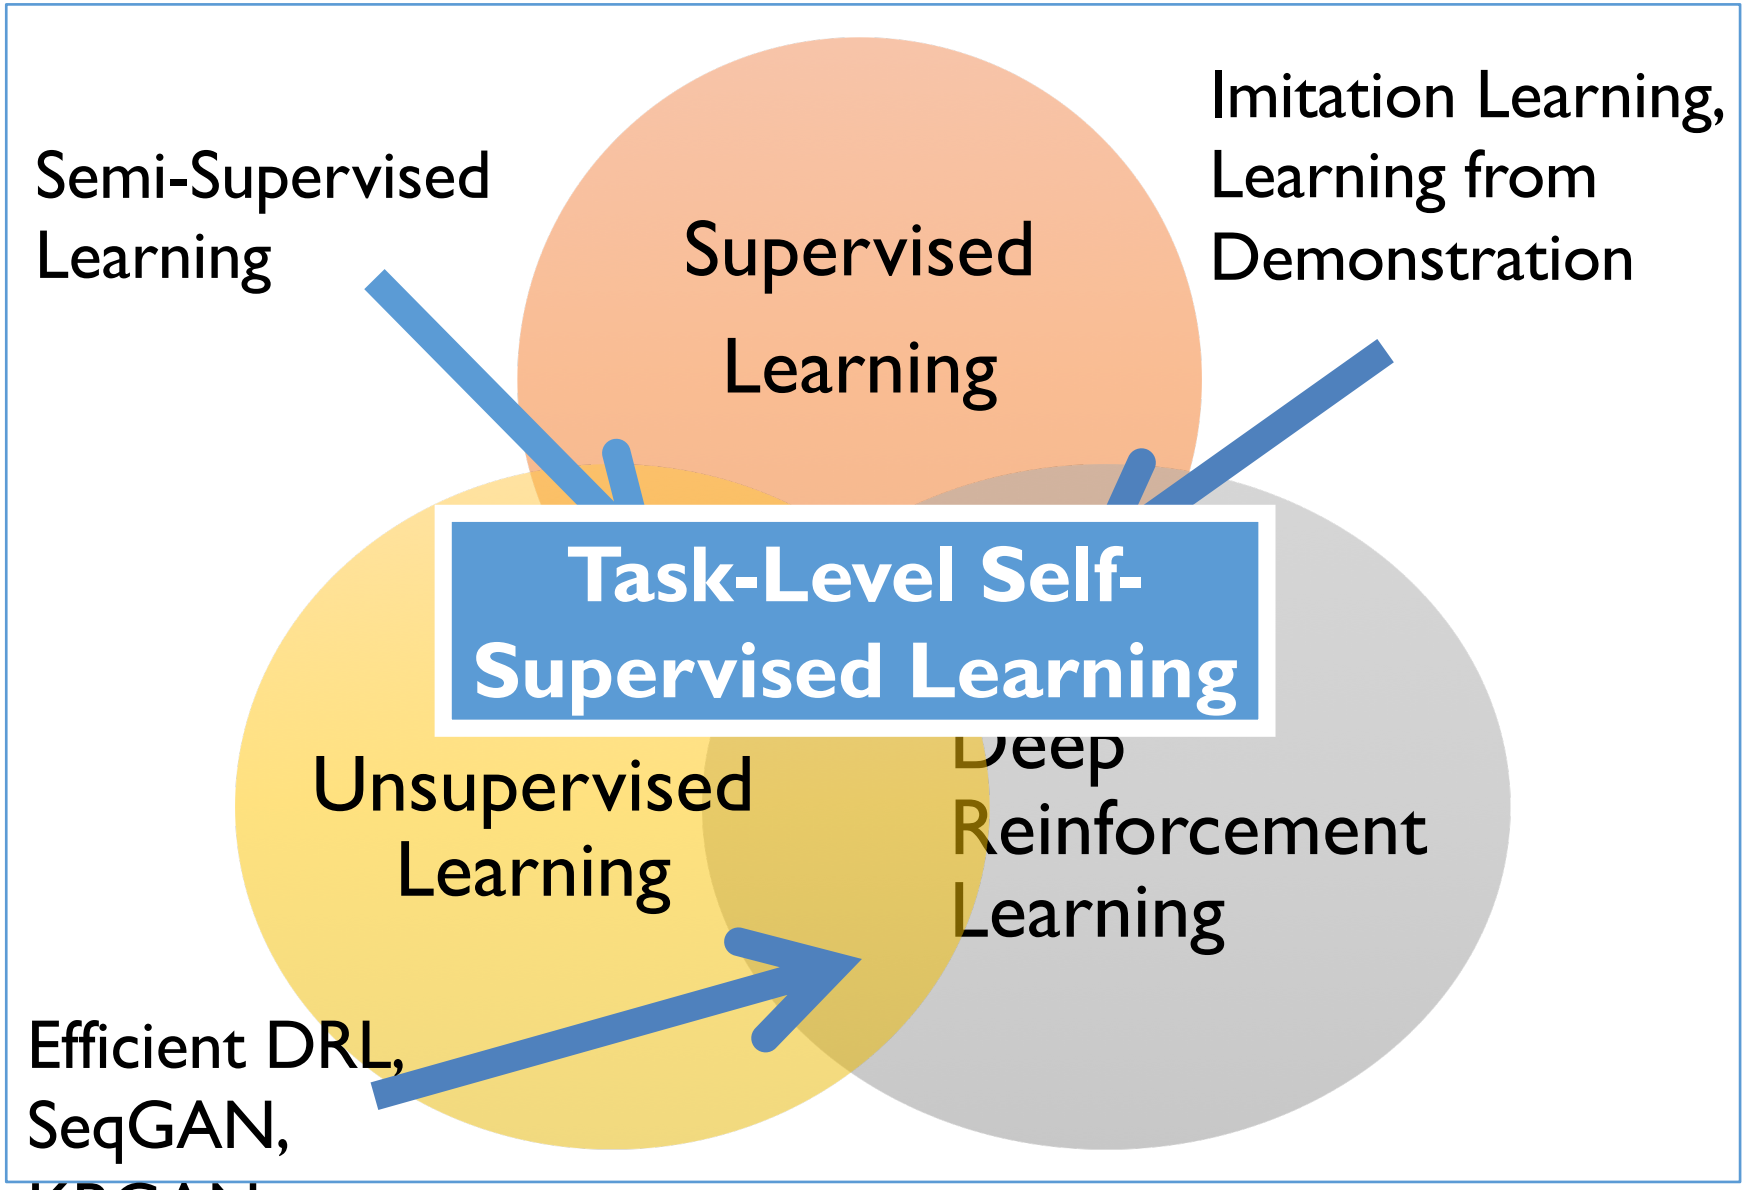

Our Work

- Unlike games, we show that modeling the **complex objective** and **diverse** answers is the key to task-level SSL in Lang + Vision.

- Combined with pre-trained KB embedding models, we show how to reason with DRL for **explainable** knowledge graph reasoning.

Outline

- Motivation
- Inverse RL for Visual Story Telling
- Learning to Reason with DeepPath
- Conclusion
- Other Research Interests

No Metrics are Perfect:

From Optimizing End Metrics (e.g., BLEU/ROUGE) to Reward Learning

(Wang, Chen et al., ACL 2018)

Existing Automatic Evaluation Metrics for Language Generation

- Input: generation candidate and human reference(s).
- Output: a score.

- Metrics:
  - BLEU: precision-driven n-gram overlap.
  - ROUGE: recall-driven n-gram overlap.
  - METEOR: weighted f1 n-gram overlap.
  - CIDEr: TF-IDF + cosine similarity.

Pop Quiz: assuming reasonable references, what is the METEOR score of this sample output?

"We had a great time to have a lot of the. They were to be a of the. They were to be in the. The and it were to be the. The, and it were to be the."

Average METEOR score: 40.2

(SOTA model: 35.0)

How about this one?



"I had a great time at the restaurant today. The food was delicious. I had a lot of food. I had a great time."

BLEU-4 score: 0

No Metrics Are Perfect: Adversarial Reward Learning (ACL 2018)

- Task: visual storytelling (generate a story from a sequence of images in a photo album).
- Difficulty: how to quantify a good story?
- Idea: given a policy, learn the reward function.

No Metrics Are Perfect: Adversarial Reward Learning (Wang, Chen et al., ACL 2018)

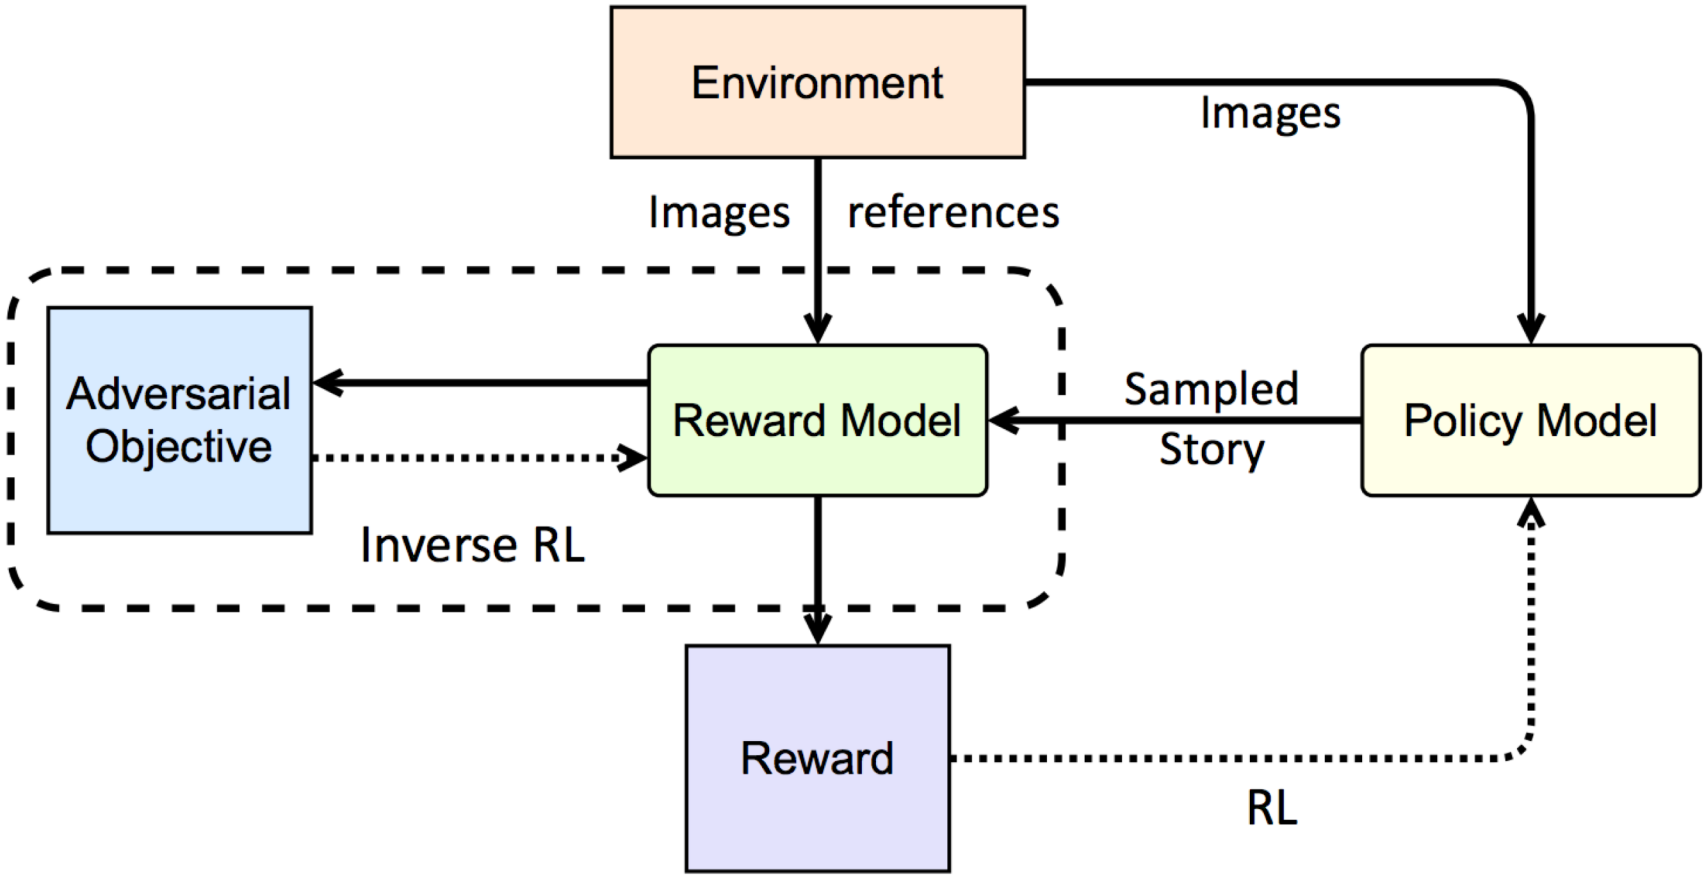

Baseline: MIXER (Ranzato et al., ICLR 2016)

- Optimize the cross-entropy loss and the BLEU score directly using REINFORCE (Williams, 1992).

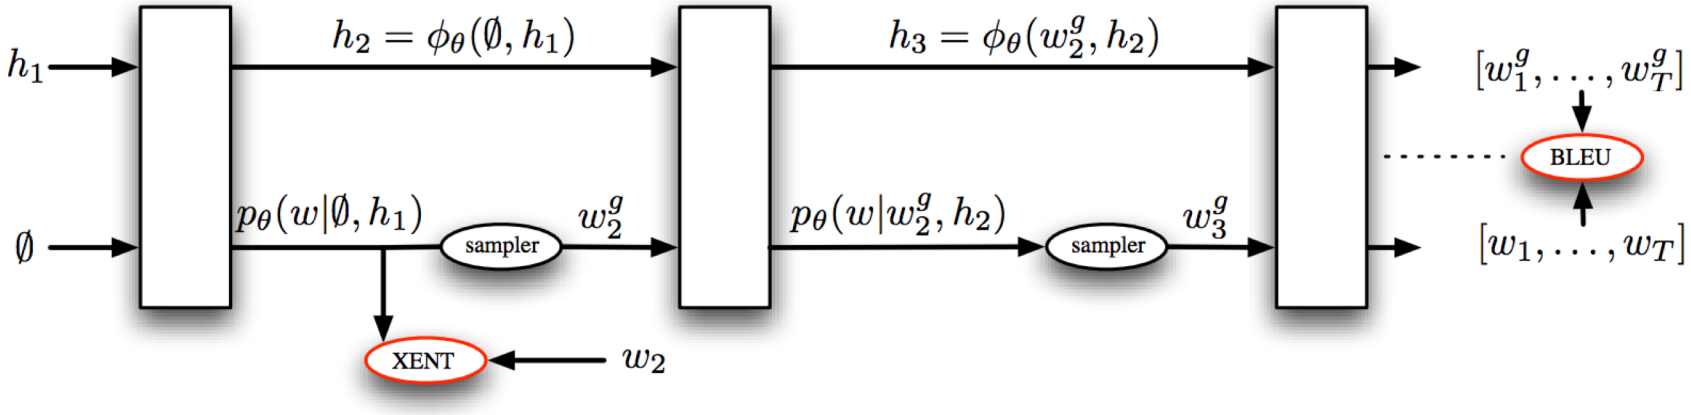

AREL Storytelling Evaluation

- Dataset: VIST (Huang et al., 2016).

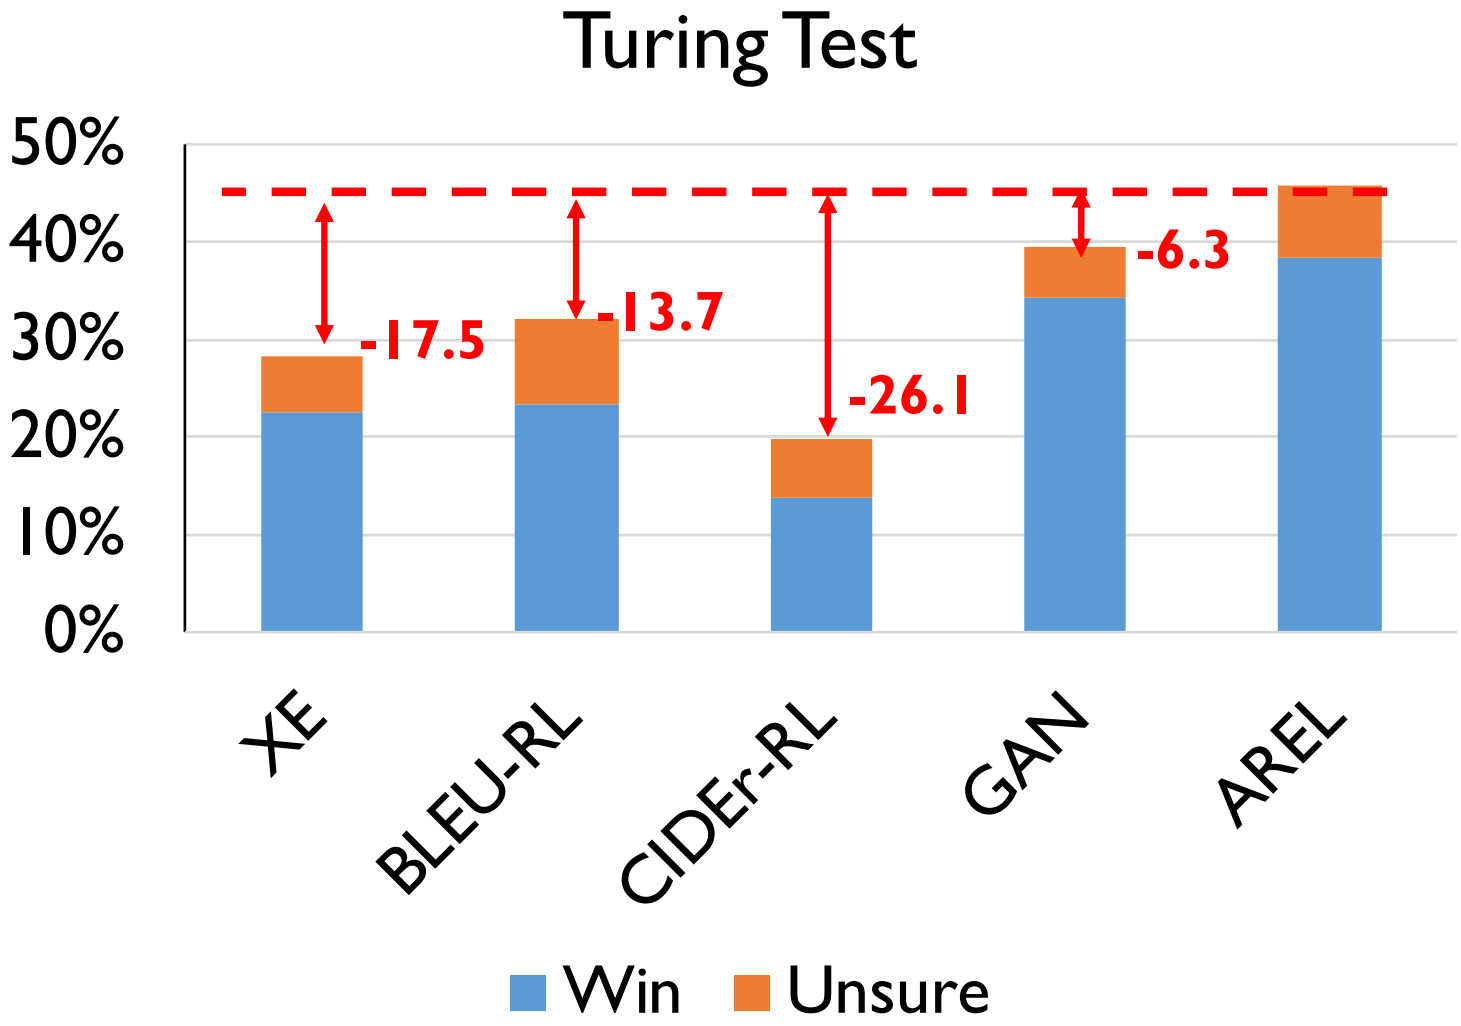

When will IRL work?

- When the optimization target is complex.
- There are no easy formulations of the reward.
- If you can clearly define the reward, don't use IRL and it will not work.

Outline

- Motivation
- Reinforced Semi-Supervised Learning
- Inverse RL for Visual Story Telling
- Learning to Reason with DeepPath
- Conclusion
- Other Research Interests and Goals

Reasoning:

Can DRL learn interpretable reasoning paths for knowledge base completion?

(Xiong et al., EMNLP 2017)

Reasoning on Knowledge Graph

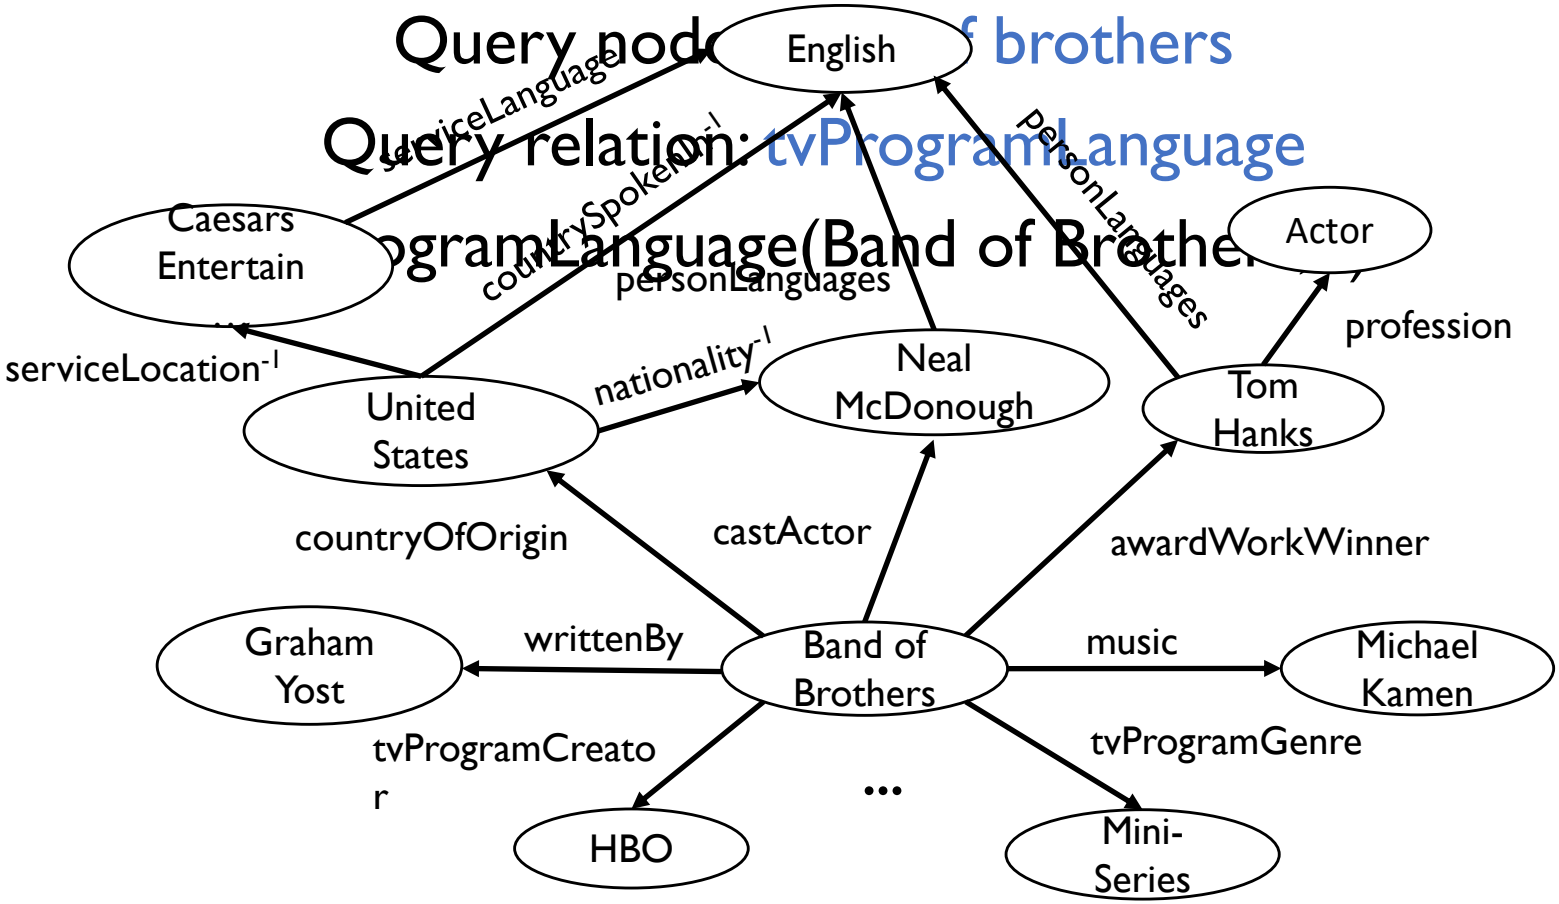

DeepPath: DRL for KG Reasoning (Xiong et al., EMNLP 2017)

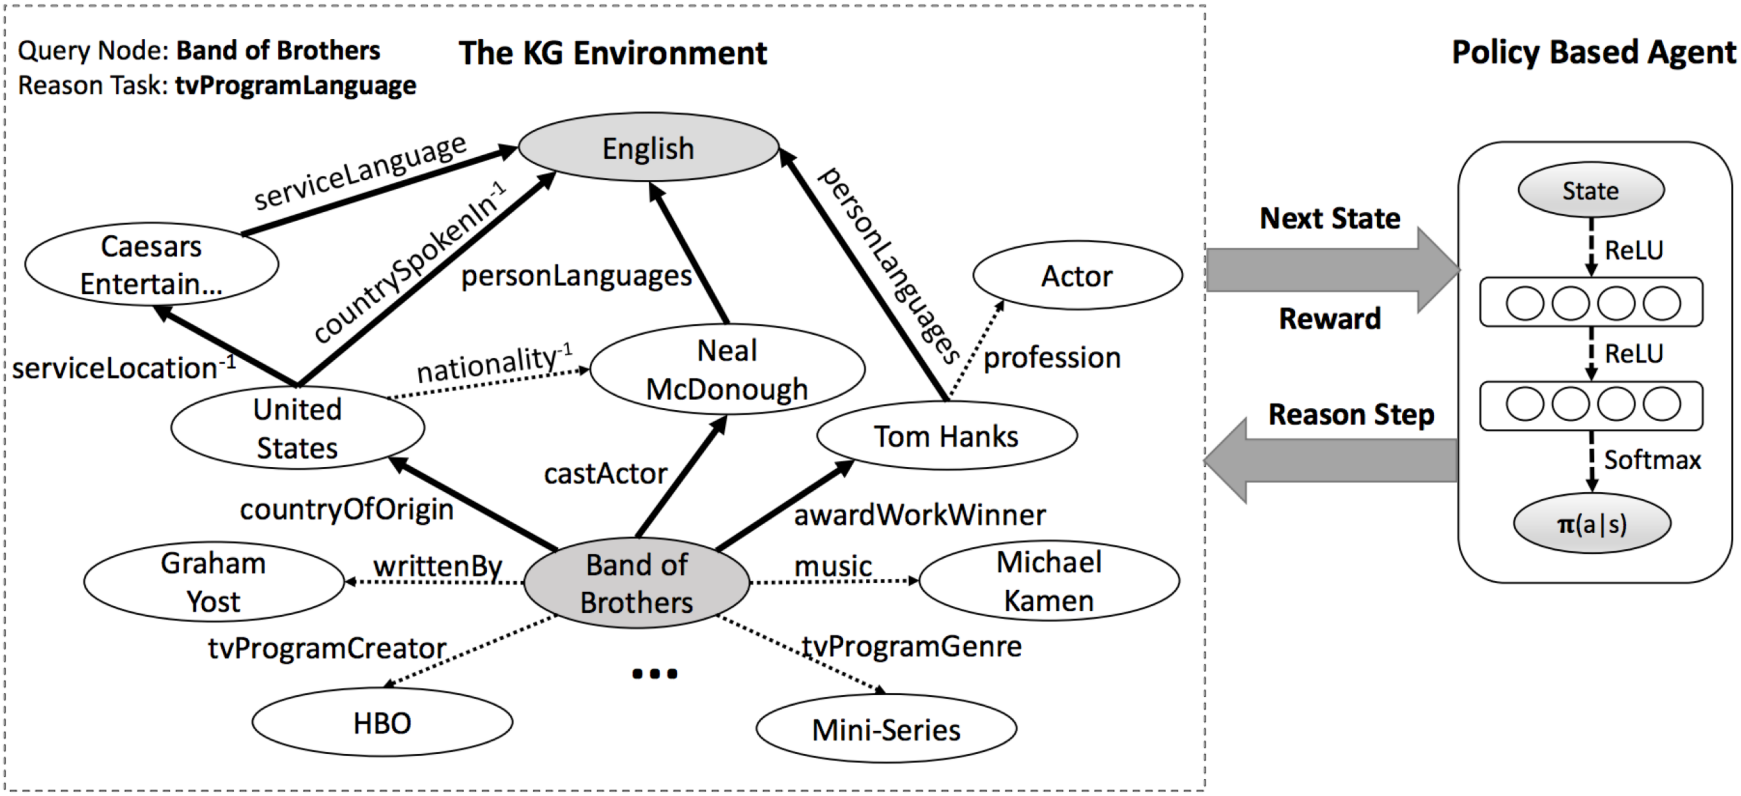

Components of MDP

- 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

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

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

Training with Policy Gradient

- Pre-training with BFS-selected paths.
- Monte-Carlo Policy Gradient (REINFORCE, William, 1992)

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

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

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

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

UCSB

Conclusion

- We show that inverse reinforcement learning is a secret weapon for addressing the diverse and complex nature of language generation problems.
- We introduce an interpretable path-based knowledge graph reasoning framework DeepPath.

Self-Supervised Learning Beyond Word-Level Representation Learning (Wang et al., ACL 2019)

Masked Paragraph	Last week, I went to attend a one-day meeting. I booked the flight in advanced. [masked sentence] The earliest next flight will be a few days later. I had to use the online discussion instead.
Candidate Sentences	But the flight was cancelled due to the weather. But I lost my passport. The meeting was cancelled. The weather is good today.

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

https://arxiv.org/abs/1711.04090

Free Label!

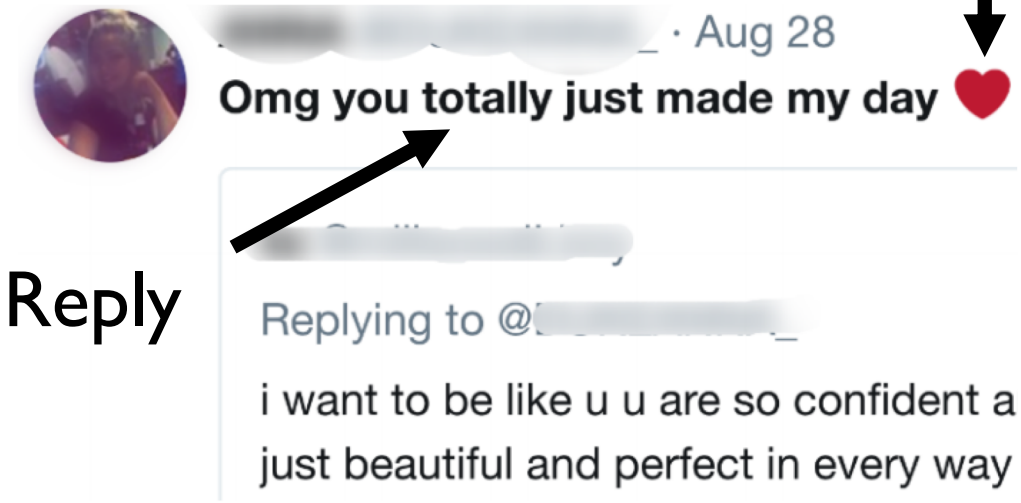

Reply

Original Post

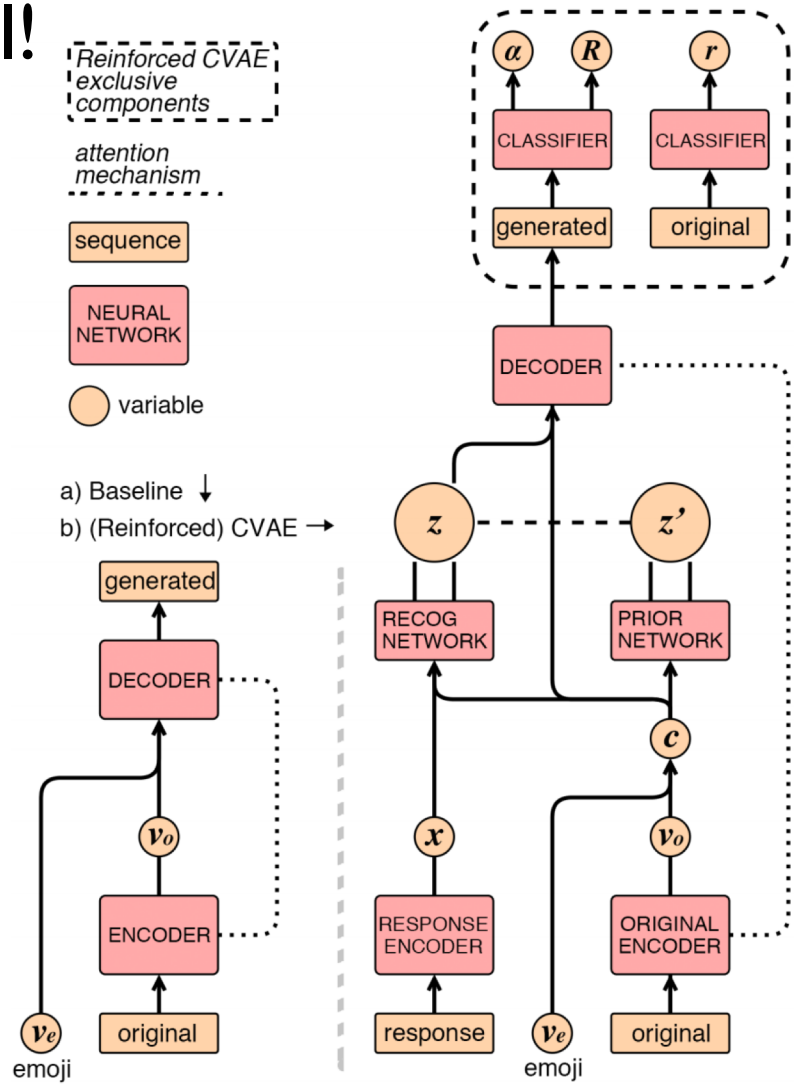

Controlling Emotions for RC-VAE Generated Sentences

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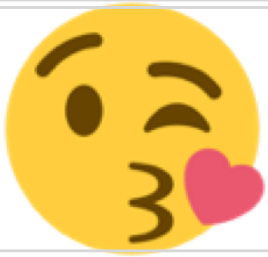	
Generated by Seq2Seq Baseline	**i 'm sorry you 're going to be missed it**	**i 'm sorry for your loss**
Generated By MojiTalk	**hope you are okay hun !**	**hi jason , i 'll be praying for you**

Other Research Highlights: Interdisciplinary and Socially Responsible Data Science

Language, Vision, Speech, & Dialog

- Image and Video Captioning and Summarization
  - Zero-Shot Video Captioning (AAAI 2019)
  - Hierarchical RL for video captioning (CVPR 2018)
  - Local and global cross-modal attention for multimodal video captioning w. speech, language, vision (NAACL 2018)
  - Multimodal summarization of events (NAACL 2016)
  - Generating memes and humor (NAACL 2015)
- Cross-lingual transfer learning for dialog tracking (EMNLP 2018)
- Semantically-Conditioned Dialogue Generation via Disentangled Self-Attention (ACL 2019)

Deep Reinforcement Learning for Language + Vision + Robotics

- Combining modal-based and model free RL for vision-language navigation (ECCV 2018)
- Scheduled Policy Optimization: Imitation Learning + RL (IJCAI 2018)
- Self-supervised imitation learning for vision-language navigation (CVPR 2019 Oral, Best Student Paper)

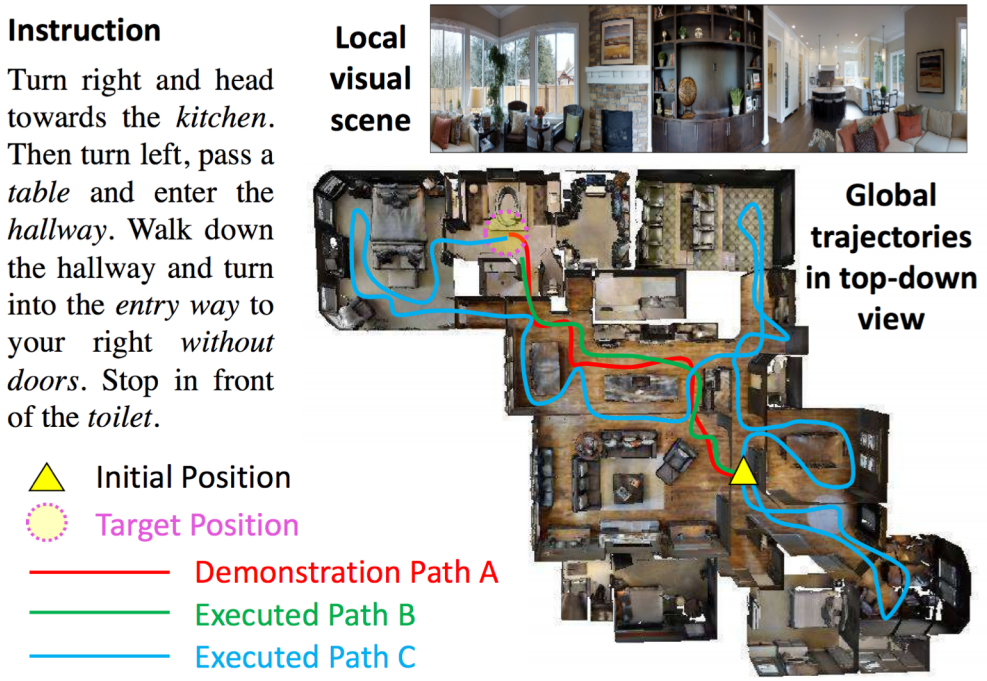

Socially Responsible NLP

- Hate speech analysis, detection, and intervention
  - Hate speech intervention (EMNLP 2019)
  - Deciphering hate symbols (NAACL 2019)
  - Representation learning for detection (NAACL 2018)
  - Linguistic analysis of hate speech (ICWSM 2018)
  - Hate group and topic detection (EMNLP 2018)

- Fake news detection and diffusion
  - LIAR benchmark dataset (ACL 2017, 200+ citations)
  - Predicting the spread of misinformation

Machine Learning / NLP + X

- ML/NLP + Politics
  - Political ideology detection (NAACL 2018)
- ML/NLP + Statistics + Finance
  - Semiparametric Gaussian Copula Regression Model for Predicting Financial Risks (ACL 2014)
- ML/NLP + Math
  - Riemannian Normalizing Flows and Wasserstein VAE (NAACL 2019)
- ML/NLP + Marketing
  - Structured learning for computational branding analytics (EMNLP 2013)
- ML/NLP + Law + History
  - Mixed-effect models for studying legal opinions (ACL 2012)

Open Challenges

- VaTeX Vision-and-Language Dataset (ICCV 2019, Oral)

Open Challenges

- TweetQA: social media question answering.

Passage: *Oh man just read about Paul Walkers death. So young. Ugggh makes me sick especially when it's caused by an accident. God bless his soul. – Jay Sean (@jaysean) December 1, 2013*

Q: *why is sean torn over the actor's death?*
A: *walker was young*

Open Challenges

- TabFact: Language and Semi-Structured Reasoning

United States House of Representatives Elections, 1972

District	Incumbent	Party	Result	Candidates
California 3	John E. Moss	democratic	re-elected	John E. Moss (d) 69.9% John Rakus (r) 30.1%
California 5	Phillip Burton	democratic	re-elected	Phillip Burton (d) 81.8% Edlo E. Powell (r) 18.2%
California 8	George Paul Miller	democratic	lost renomination democratic hold	Pete Stark (d) 52.9% Lew M. Warden , Jr. (r) 47.1%
California 14	Jerome R. Waldie	republican	re-elected	Jerome R. Waldie (d) 77.6% Floyd E. Sims (r) 22.4%
California 15	John J. Mcfall	republican	re-elected	John J. Mcfall (d) unopposed

John E. Moss and Phillip Burton are both re-elected in the house of representative election in 1972.

Acknowledgment

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

Thank you!

- UCSB NLP Group: nlp.cs.ucsb.edu

- AREL: https://github.com/eric-xw/AREL
- DeepPath: https://github.com/xwhan/DeepPath/
- Walk the Block: https://github.com/xwhan/walk_the_blocks
- Cross-Lingual Dialog State Tracking: https://github.com/wenhuchen/Cross-Lingual-NBT
- MojiTalk: https://github.com/claude-zhou/MojiTalk

UCSB