NS4NLP: Neuro-Symbolic methods for Natural Language Processing

Maria Pacheco **Dan Goldwasser**

Purdue University

{pachecog, dgoldwas}@purdue.edu

IJCAI 2021 Tutorial

Aug. 2021

Introduction

Neuro (and/vs.) Symbolic Approaches

- **The Neuro-Symbolic distinction is often characterized as**
  - **System 1:** thinking fast, *typically associated with learned neural models*
  - **System 2:** thinking slow, *typically associated with symbolic reasoning*
- **Many dimensions to be considered:**
  - What is the ”right” representation?
  - Learning vs. Reasoning
  - Explainability
  - Modularity
  - ..
- **A fascinating, *albeit never-ending*, debate in AI..**

***Not* our main focus today!**

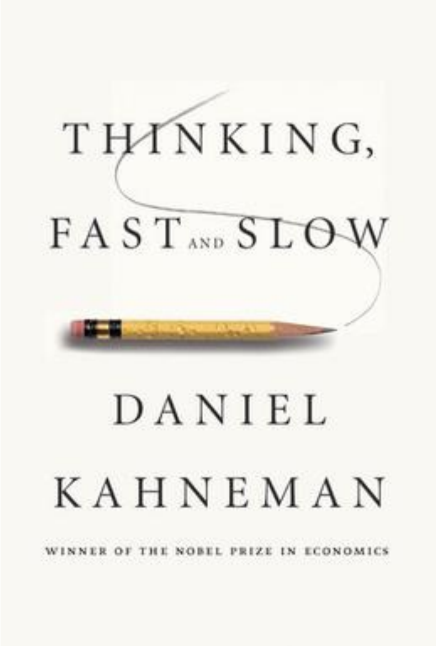

Neuro-Symbolic Approaches for NLP

- **Ignoring the last 8-10 years, NLP is symbolic.**
  - Not surprising, *words are symbols!*
  - Their combination results in symbolic linguistic structures
- **The last 8-10 years ignored that NLP was symbolic**
  - Not surprising, *it just works so well!*
- **On-going debate**: *do symbolic structures still matter? Does inference?*
  - At what level of granularity? How should it be used?
  - N-S work in NLP focuses on NLI, QA, and grounded language applications
- **This tutorial**: *Neuro-Symbolic methods for opinion analysis*
  - **New domain** for N-S: what are the relevant symbols?
  - **Interaction** using N-S methods: record human insights, abstracting over the text

This Tutorial in one Slide

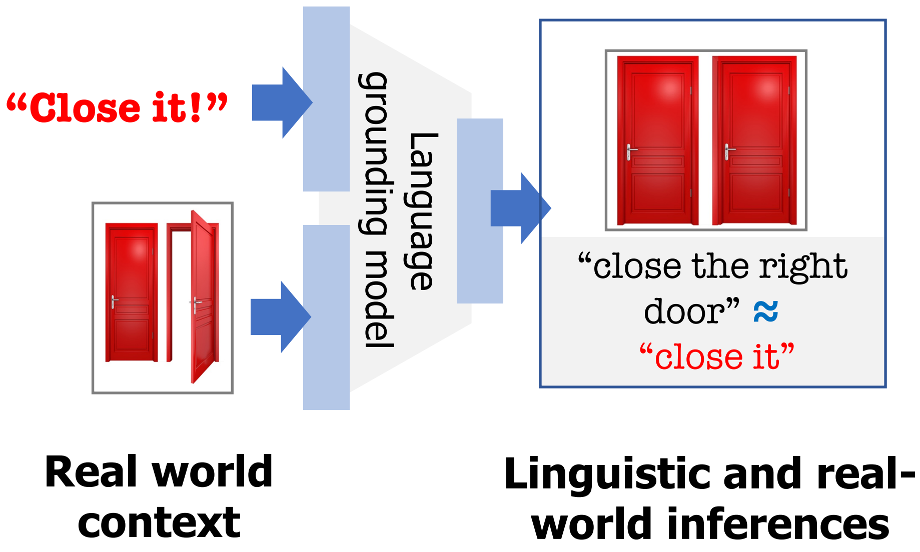

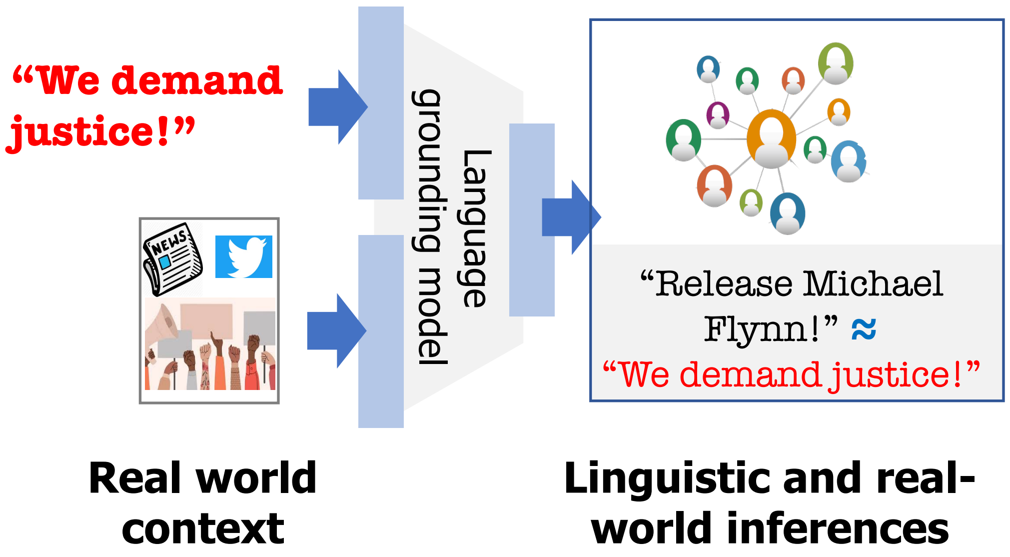

- What are the inferences needed?
- What symbols (entities, properties, relations) are they defined over?

- How about now?
- Can we ask ***people*** to help us?

.. and now this tutorial in a few more slides!

Symbolic vs. Distributed Representations

- **Symbolic Representation**

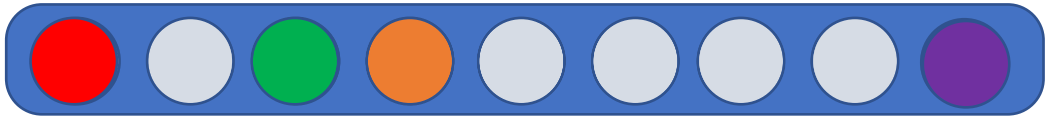

Game Play
Megan_Rapinoe
Ian_McKellen

- **Distributed Representation**

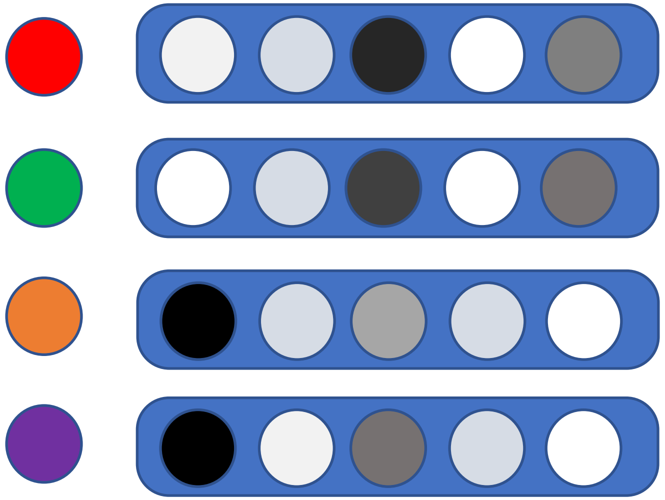

Representing Context and Structure

Megan Rapinoe's **play** vs. *Ian McKellen's* **play**

- Symbolic Language Models can capture context, but at a high cost
  - Assume a 10-gram, using a vocabulary of 50K words

97,656,250,000,000,000,000,000,000,000,000,000,000,000,000,000 #parameters

- Clearly, overly simplistic analysis ignoring many language modeling lessons
  - *But on the other hand, GPT-3 suddenly seems very economical..*

00's - early '10's NLP:
Instead of relying on words directly, we can exploit linguistic structure and devise better representations!

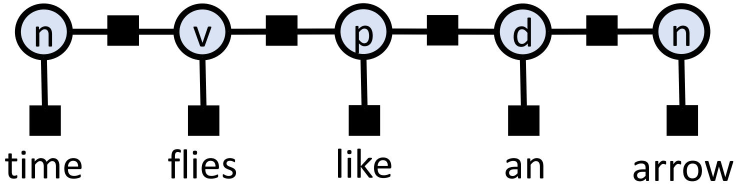

Gormley and Eisner ACL'15 BP Tutorial

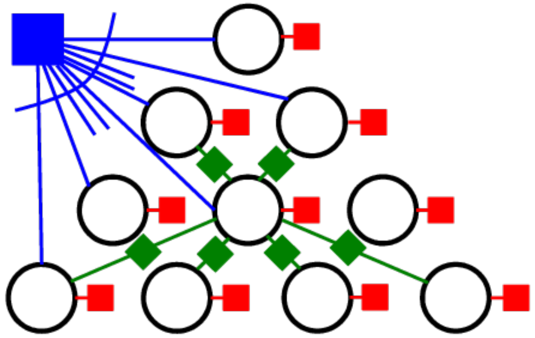

Burkett & Klein (2012)

Symbolic Representation of Meaning

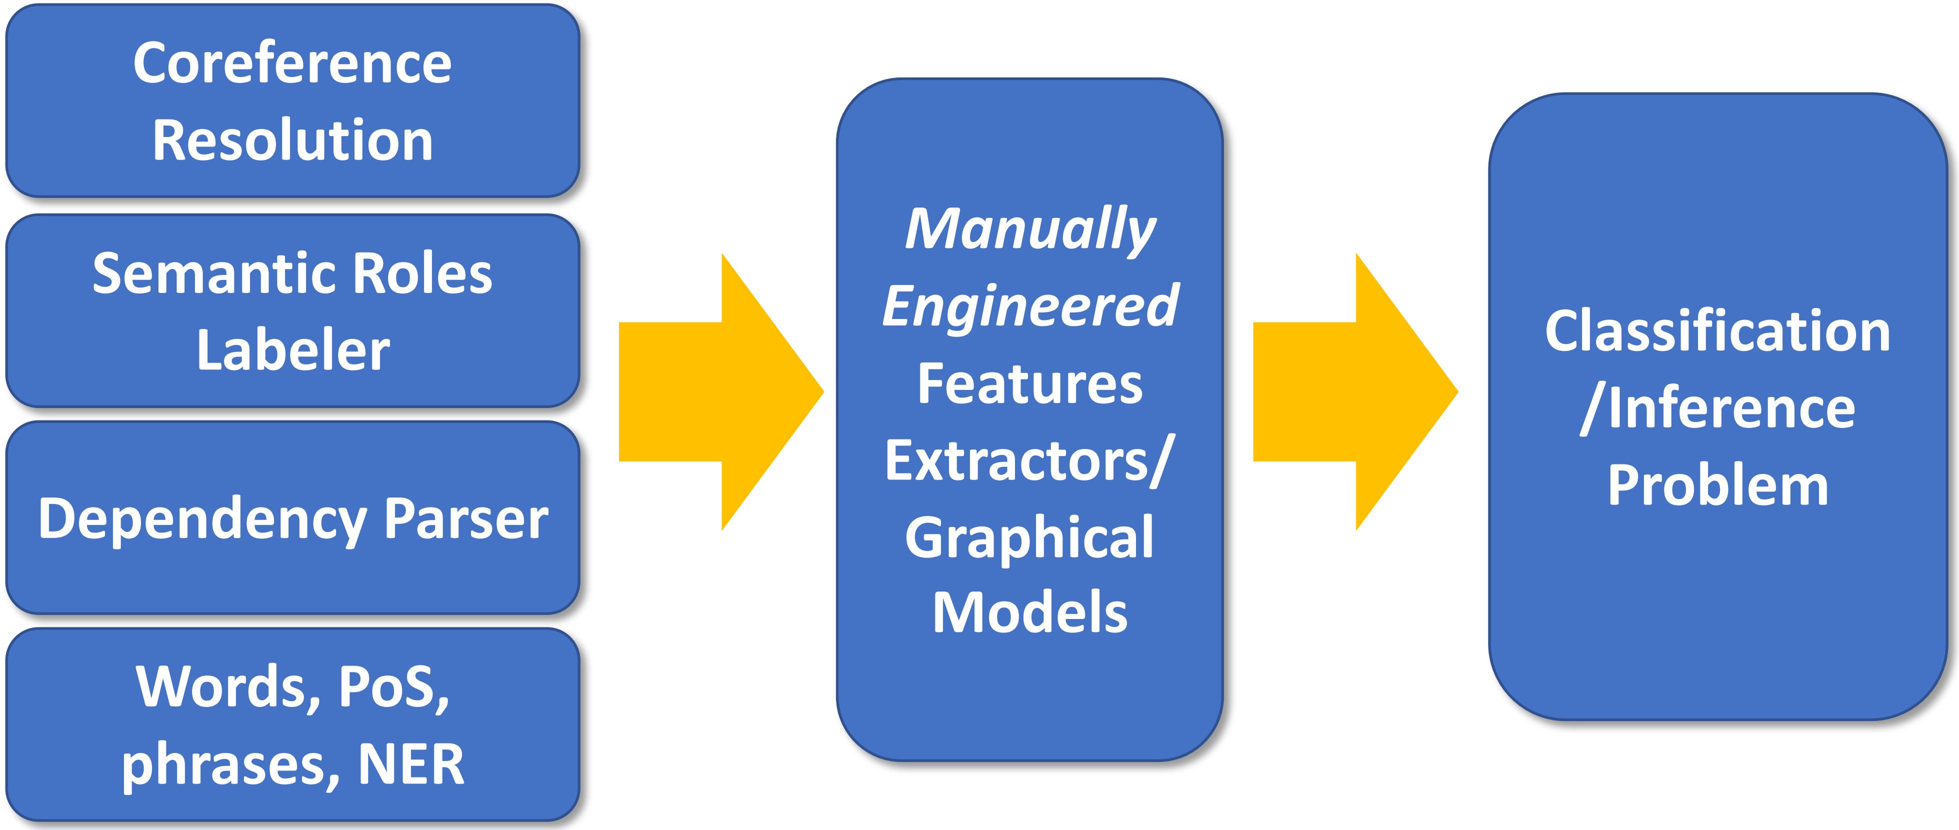

00’s - early ‘10’s NLP: *Instead of relying on words directly, we can exploit linguistic structure and devise better representations!*

Representing Context and Structure

This phone is sick!

This person is sick!

On the **Neural side,** *Contextualized Word Embedding* methods (*Sesame Street*) are very good at disambiguating word usage.

- Words are represented as dense low-dimensional vectors
- Words-in-context are represented as a composition of neighboring words (up to a fixed window)
- Both initial vectors and composition function trained over **massive** amounts of data, by predicting their context
- Representation can then be specialized for specific tasks

Representing Context and Structure

Known as *contextualized language models*

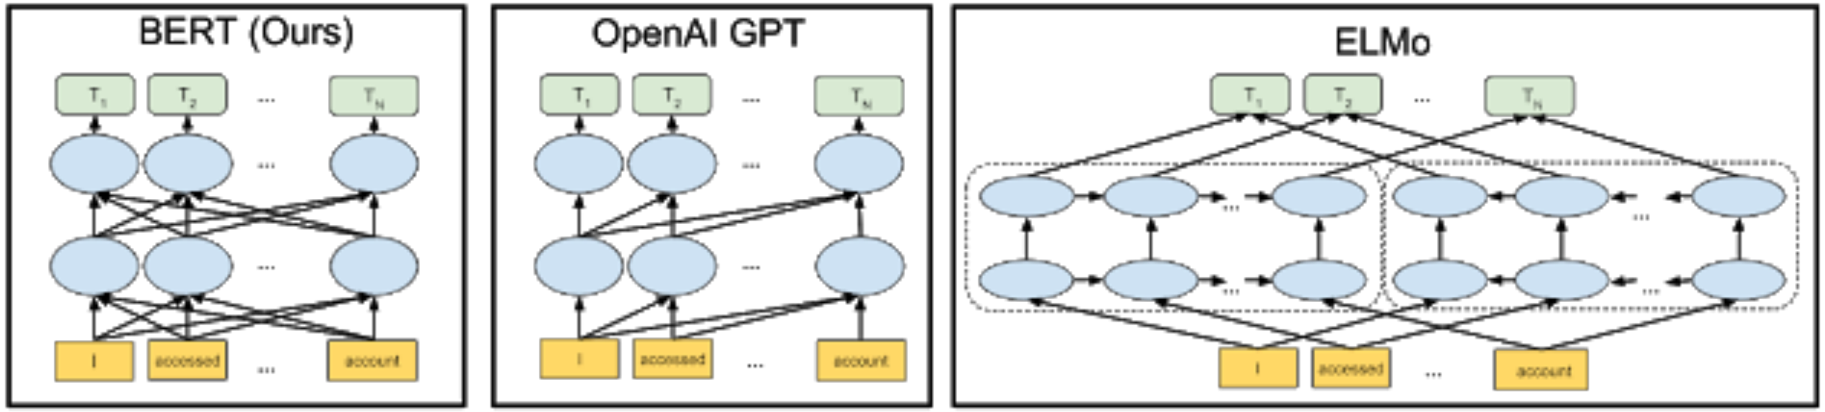

Devlin et-al. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding” 2019

What does BERT learn?

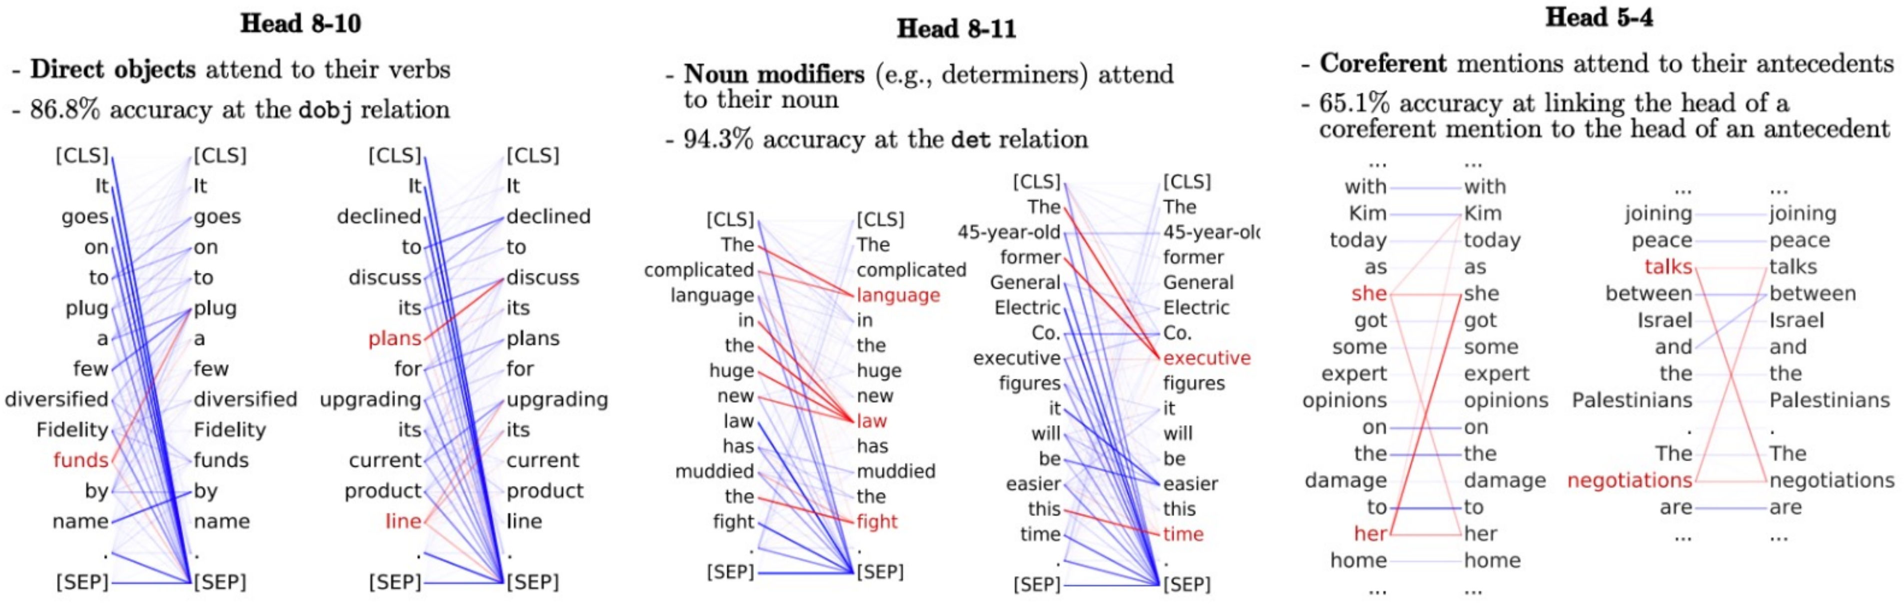

Linguistic structure emerges without direct supervision

Emergent linguistic structure in artificial neural networks trained by self-supervision. PNAS Manning et-al, 2020

Using BERT for Reasoning Tasks

- **BERT-based *near-human performance* on Winograd Schema**

World Knowledge and Commonsense inferences reflected in coref decisions

		Twin sentences	Options (**answer**)
✓ (1)	a	The trophy doesn't fit into the brown suitcase because **it**'s too *large*.	**trophy** / suitcase
	b	The trophy doesn't fit into the brown suitcase because **it**'s too *small*.	trophy / **suitcase**
✓ (2)	a	Ann asked Mary what time the library closes, *because* **she** had forgotten.	**Ann** / Mary
	b	Ann asked Mary what time the library closes, *but* **she** had forgotten.	Ann / **Mary**
✗ (3)	a	The tree fell down and crashed through the roof of my house. Now, I have to get **it** *removed*.	**tree** / roof
	b	The tree fell down and crashed through the roof of my house. Now, I have to get **it** *repaired*.	tree / **roof**
✗ (4)	a	The lions ate the zebras because **they** are *predators*.	**lions** / zebras
	b	The lions ate the zebras because **they** are *meaty*.	lions / **zebras**

WINOGRANDE: An Adversarial Winograd Schema Challenge at Scale. Sakaguchi et-al, AAAI'20

Can "*thinking-slow*" tasks be accomplished with "*thinking-fast*" systems?

Not a panacea (McCoy et al ACL'19, others), often relies on simple heuristics when learning complex decisions

Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference. McCoy et-al, ACL'19

Challenges and Opportunities

- **Purely neural models face challenges**
  - Energy efficiency
  - Data efficiency
  - Explainability and Human Interaction
  - Domain Transfer
  - Reasoning beyond surface level patterns
- **Neuro-Symbolic methods: *Best of both worlds!***

Neural methods can train expressive models using massive datasets to identify patterns in raw data

Symbolic methods can map neural representations to symbols and reason over higher level patterns

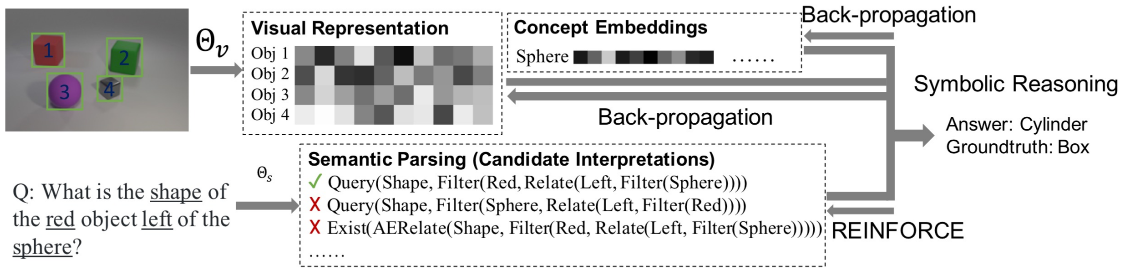

The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision. Mao et-al ICLR2019

New

NS 4 NLP Domains

and 4 interaction

Learning from Static Data *vs.* Interaction

- Machine learning, an *alternative* view: it is all about ***communicating knowledge* from humans to machines**
- Currently very *inefficient* : communication using labeled examples
- Very(!) successful.. **But**
  - Costly annotation, especially in complex domains
  - End-2-end approach is monolithic and hard to reuse
  - Can be very sensitive to changes in the input (concept drift)

Learning from Static Data *vs.* Interaction

- **Learning from interaction**
  - Natural supervision!
  - Often, easy & cheap
  - *Broadly applicable:*
    - Clickthrough data
    - Behavioral data
    - Expected behavior/answers
    - User emotional tone

- Very(!) successful.. **But**
  - Still, not a great account for *knowledge communication*

Learning from Interaction

My teacher said that lightning is dangerous.
Knives are also dangerous, they are sharp.
<u>Lightning is sharp too!</u>

- **The bandwidth of interaction can be much greater!**
- Humans are great "reasoning machines"
  - Learn by matching new data to previously acquired concepts
- We can characterize the learning problem using intermediate concepts which can be shared across many learning problems
- Learning as a form of **knowledge communication**
  - Intermediate concepts can be viewed as a shared vocabulary supporting interaction between machine learner and human teacher
  - Human teacher can "debug" the learner's internal representation

Learning from Interaction

- **Working definition**: communicate human's rationale about the task, via intermediate judgements and explanations, sub-goals or steps.

Learning from Explanations

"a vehicle is suspicious if it is in working condition and has camouflage color"

NL

SuspiciousTruck(x) ← vehicle(x) **AND** works(x) **AND** color(x,CAMOUFLAGE)

Human-readable Symbolic Representation

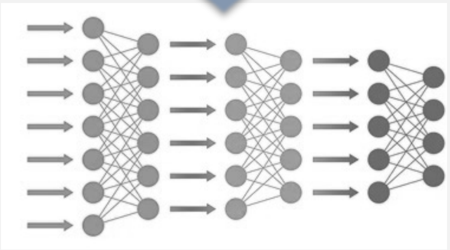

Neural Representation (Machine-representation)

Interaction over the symbolic representation

- Communicate knowledge by capturing dependencies between concepts.
- Debug concept by interaction

Open Problems

- Where do relevant concepts come from?
- How to ground concepts in raw data?
- How can the symbolic representation be "**compiled**" into a classifier?

Advice (post-training)

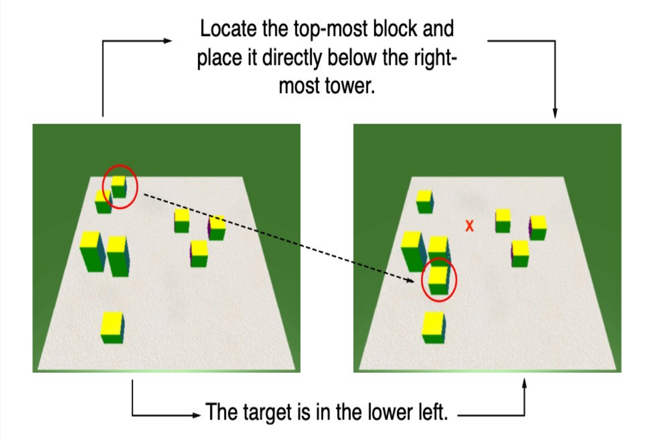

Improving natural language interaction with robots using advice. Mehta and Goldwasser NAACL'19

Tackling Fake News Detection by Interactively Learning Representations using Graph Neural Networks. Mehta Goldwasser. InterNLP '21

Understanding Opinions and Political Discourse

- **Natural fit** for N-S and interactive learning:
  - **Text + Context:** lots of text coupled with behavior
  - **Very dynamic**: *a moving target for supervised learning approaches*
  - **Explanations rely on complex concepts:**
    - Ideology, interests, arguments, many more!

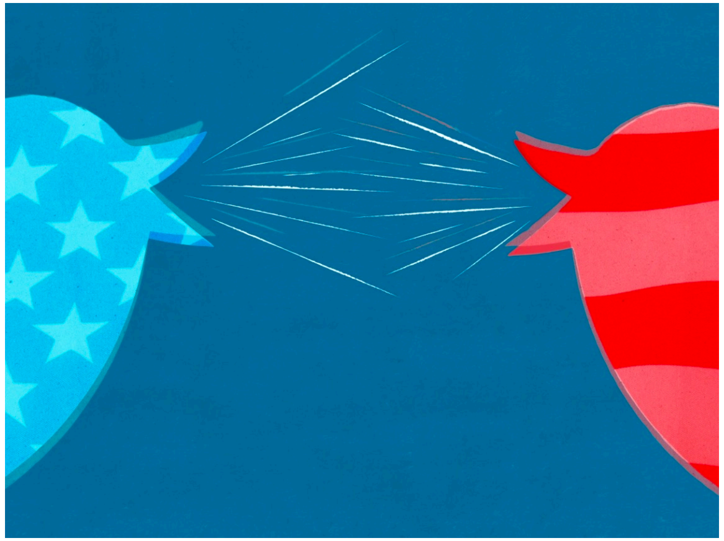

"if you talk about healthcare as a human right then…"

"..probably voted in favor of Obamacare"

Tweet(x) author(x,y) HasFrame(x,fairness) HasTopic(x,healthcare) ➔ VotedFor(y,Obamacare)

PURDUE UNIVERSITY

Beyond Linguistic Context!

Understanding the real-world context of text can help disambiguate it!

E.g., transformers are very good at disambiguating word usage, **but**...

This movie is sick!

Explanations can also consider the social context of the text!

"if the author is a Trump supporter, then.."

".. article likely to oppose impeachment"

"if the author follows OAN, then.."

".. author likely to support Trump"

We also have other reasons for looking at this domain

NLP meets the *#world*

Social media is the primary space for public discourse

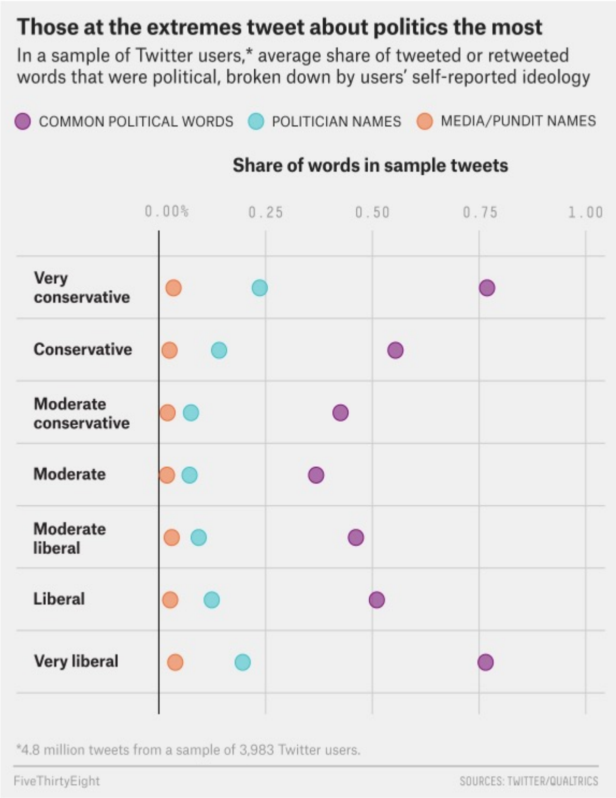

FiveThirtyEight

Politics Sports Science Podcasts Video

NOV. 1, 2017, AT 10:13 AM

Political Twitter Is No Place For Moderates

By Dan Hopkins

Americans Who Mainly Get Their News on Social Media Are Less Engaged, Less Knowledgeable

Those who rely on social media for news are less likely to get the facts right about the coronavirus and politics and more likely to hear some unproven claims

BY AMY MITCHELL, MARK JURKOWITZ, J. BAXTER OLIPHANT AND ELISA SHEARER

NLP meets the *#world*

Social media users are continuously exposed to extreme opinions and low-quality information

- NLP can help provide context for understanding opinions and alternative perspectives on the issues
- It can help analyze political messaging at scale and quantify its effectiveness
- It can help combat "information echo-chambers" and ensure all perspective are covered
- It can help identify extremism and harmful content on social media

Underlying issue – **explaining information!**

Tutorial Structure and Goals

Goals: *(1) social/political discourse analysis is a great opportunity for NS researchers. (2) Leverage human insights. (3) Use DRaiL.* Shameless promotion warning

- **Part 1: Introduce NS 4 NLP**
  - Challenges and opportunities for analyzing opinions and political discourse
- **Part 2: Technical Overview**
  - Statistical Relational Learning, Neuro-Symbolic frameworks
  - NLP Applications using Neuro-Symbolic methods

Break

- **Part 3: Neuro-Symbolic analysis for discourse**
  - Parsing the opinions landscape using DRaiL
- **Part 4: Symbolic explanations for opinions and behaviors**
  - Psychological states, Framing theory, Moral Foundation theory

Break

- **Part 5: DRaiL Demo**
  - Human interaction supporting automated opinion analysis

PURDUE UNIVERSITY

Technical Overview

Statistical Relational Learning

- Complex relational structure
  - Capture real world domains by encoding objects and relations between them
  - Express generic facts about these objects and relations
- Uncertainty
  - Move away from "all-or-nothing" logic and model the uncertainty

Combine the expressivity of logic with statistical methods

- Many approaches and frameworks
  - Extend logic to handle probabilities
  - Extend statistical models to capture relational structures

An Example: Markov Logic Networks

- Key idea: **add weights to first-order formulae**
  - Expresses the strength of the formula

$$\begin{array}{l|l} 8.2 & \forall x : smokes(x) \Rightarrow cancer(x) \\ 1.7 & \forall x, y : friends(x, y) \Rightarrow (smokes(x) \Leftrightarrow smokes(y)) \end{array}$$

- MLN is a set of pairs <Formula_i, weight_i>
- Describes an **undirected graphical model**
- Ground it in data and use it for **inference**

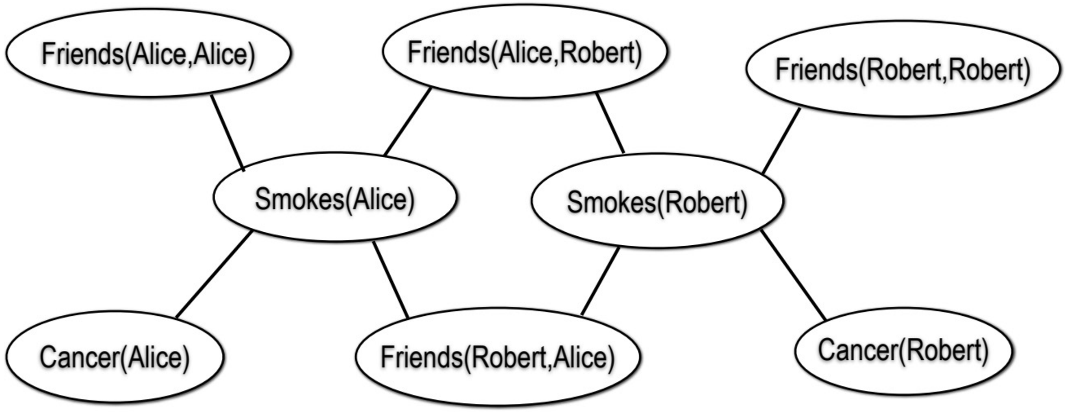

Richardson, M., Domingos, P. “Markov Logic Networks” 2006

An Example: Markov Logic Networks

- We can learn MLNs **from data**
- ***Parameter learning***
  - Assume the formulae are given, but not the weights
  - Max-margin, EM
- ***Structure learning***
  - Try to learn both the formulae and the weights
  - Inductive Logic Programming methods
- Standard algorithms for inference in graphical models
  - Gibbs Sampling, Belief Propagation

Richardson, M., Domingos, P. "Markov Logic Networks" 2006

Another Example: Probabilistic Soft Logic

- Same key idea, different formalism
  - ***Same***: weighted logical formulae
  - ***Different***: restricted to **horn-clauses**

- Atoms have **continuous** truth values in the [0,1] interval
- Inference finds values for atoms that best satisfy rules given evidence
- Most probable explanation is a **linear-programming problem**
  - Making inference efficient!

Bach S. et al. “Hinge-Loss Markov Fields and Probabilistic Soft Logic”, 2017

Statistical Relational Learning for NLP

- How to model high-level language problems in a framework like PSL?

1. Enumerating relevant keywords
(Johnson et al., 2017)

Unigram(T, U) -> HasLabel(T, L)
Unigram(T, U) & Bigram(T, B) -> HasLabel(T, L)

Retweets(T1,T2) & HasLabel(T1,L) -> HasLabel(T2, L)

2. **Using local classifiers as priors**
(Sridhar et al., 2015)

localLabel(T, L) -> HasLabel(T, L)
localAgree(T1, T2) -> Agree(T1, T2)

Agree(T1,T2) & HasLabel(T1, L) => HasLabel(T2,L)

Sridhar, D. et al. "Joint Models of Disagreement and Stance in Online Debate", 2015
Johson, K. et al. "Leveraging Behavioral and Social Information for Weakly Supervised Collective Classification of Political Discourse"

Distributed Representations for Relational Data

- ***Common theme***: Complex relational structure
- ***Alternative approach***: represent data using symbolic structures (graphs) and use NNets to learn distributed representations for it

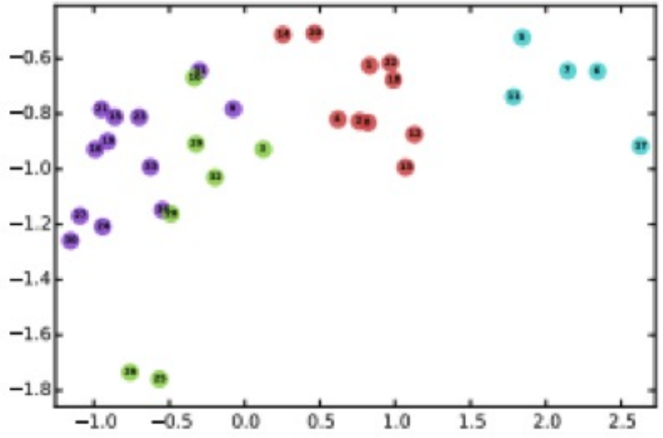

Distributed Reps: Node Embeddings

- ***Node embeddings*** (Perozzi et al., 2014; Tang et al., 2015; Grover and Leskovev, 2016)
  - Capture **adjacency information**
  - Similarity of two nodes $\propto$ graph distance and neighborhood overlap

- ***Node embeddings with textual properties*** (Pan et al. 2016; Xiao et al., 2017; Tu et al. 2017)
  - **Text** provides initial representation, **jointly learned** with adjacency info.

- ***Multi-relational node embeddings*** (Bordes et al., 2013; Wang et al. 2014; Trouillon et al., 2016)
  - Embed both nodes and **edge types**

Distributed Reps: Graph Embeddings

- ***Neural-module networks*** (Andreas et al., 2016; Johnson et al., 2017)
  - Translate input into **tree-structured functional program**
  - Input-specific NNets out of **discrete collection** of **specialized modules**

- ***Tree-structured NNets*** (Socher et. al, 2011; Tai et al., 2015)
  - Recursive NNets over **tree-structured inputs** (*more flexible than NMMs*)

- ***General Graph NNets*** (Kipf and Welling, 2017; Hamilton et al., 2017; Velickovic et. al, 2017)
  - **Contextualized representations of nodes** by recursively aggregating neighbors
  - Both single and multi-relational

At the Intersection: Neural-Symbolic Frameworks

- Combine **neural** and **symbolic** representations in a ***unified way***

- Approaches can be categorized as:
  - Lifted rules to specify compositional nets
  - Differentiable inference and rule induction from data
  - Deep representations + symbolic reasoning

Lifted Rules for Compositional Nets

- **RelNNs** map observed ground atoms, facts, and rules to neurons in a network
- Define composition functions over them

Input is symbolic!

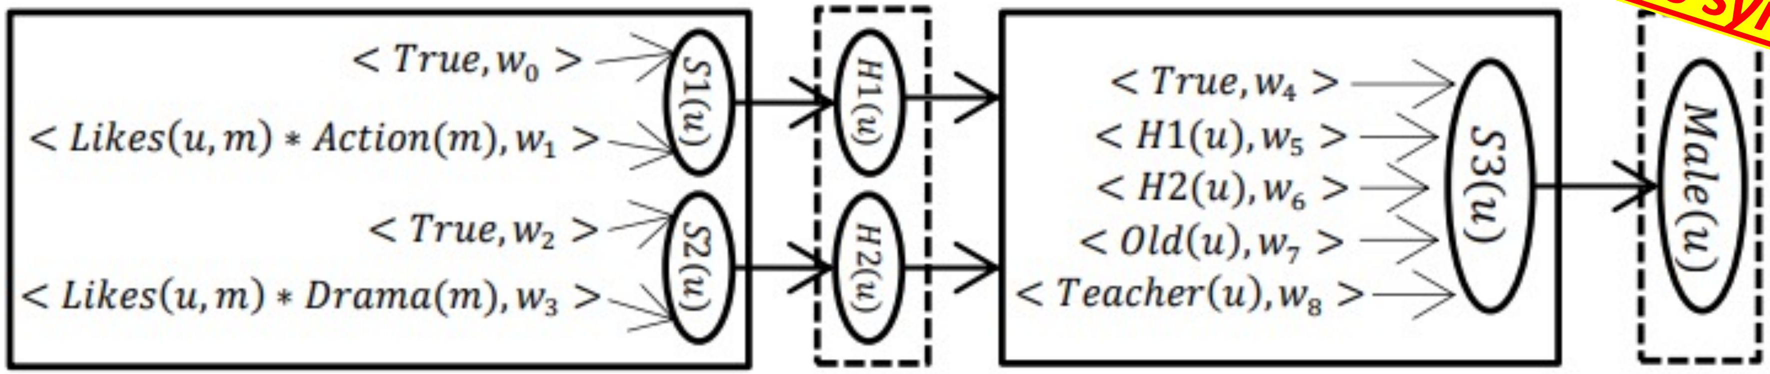

A RelNN model to predict the gender of users in a movie rating system with a layer-wise architecture

- Other systems that take a similar approach:
  - Lifted Relational Neural Networks

Kazemi S.M. and Pool D. "RelNN: A Deep Neural Model for Relational Learning", 2017
Sourek G., "Lifted Relational Neural Networks", 2015

Differentiable Inference

- ***TensorLog*** uses matrix multiplication for logic
- Entities are encoded as 1-hot vectors $v_i \in \{0,1\}^{|E|}$
- Relationships are encoded as adjacency matrices $M_R = \{0,1\}^{|E|x|E|}$
- Imitate logical rule inference for an entity $X = x$

$$R(Y,X) \Leftarrow P(Y,Z) \wedge Q(Z,X)$$

$$M_P * M_Q * v_x = s$$

- This can be generalized to rules of any length – and inference for x is defined as:

$$\alpha \;\; query(Y,X) \Leftarrow R_n(Y,Z_n) \wedge \;\dots \wedge R_1(Z_1,X)$$

$$score(y|x) = v_y^T s$$

$$s = \sum_l \alpha_l \prod_{k \in \beta_l} M_{R_k}$$

Cohen W., "TensorLog: A Differentiable Deductive Database", 2015

Differentiable Inference

- The learning problem for a query is then, where {x,y} are entity pairs that satisfy the query and $\{\alpha_l, \beta_l\}$ are to be learned

$$\max_{\{\alpha_l,\beta_l\}} \sum_{\{x,y\}} score(y|x)$$

Input is, again, symbolic

- Other systems that look at differentiable inference:
  - Logic Tensor Networks, Deep Logic Models

Cohen W., “TensorLog: A Differentiable Deductive Database”, 2015
Donadello I. “Logic Tensor Networks for Semantic Interpretation”, 2017

Rule Induction from Data

- ***Neural LP*** builds on TensorLog to learn rules
- Rewrite equation to address the problem of enumerating rules

$$\sum_{l} \alpha_l \prod_{k \in \beta_l} M_{R_k} \Rightarrow \prod_{t=1}^{T} \sum_{k}^{|R|} \alpha_t^k M_{R_k}$$

- Where T is the max length of rules and |R| the number of rels in the KB
- Assumption length = T is relaxed with recurrent formulation + attention
- Other systems that look at differentiable/end-to-end rule induction:
  - Neural Theorem Provers, DRUM, Neural Logic Machines

Yang F. et al, "Differentiable Learning of Logical Rules for Knowledge Base Reasoning", 2017
Rocktäschel T. and Riedel S. "End-to-End Differentiable Proving", 2017
Sadeghian A. et al, "DRUM: End-To-End Differentiable Rule Mining On Knowledge Graphs", 2019
Dong H. et al, "Neural Logic Machines", 2019

Neural Representations + Symbolic Reasoning

- ***ProbLog*** is a probabilistic logic programming framework

$$0.1 :: \texttt{burglary}. \quad 0.5 :: \texttt{hears\_alarm(mary)}.$$
$$0.2 :: \texttt{earthquake}. \quad 0.4 :: \texttt{hears\_alarm(john)}.$$

$$\texttt{alarm} :- \texttt{earthquake}.$$
$$\texttt{alarm} :- \texttt{burglary}.$$
$$\texttt{calls(X)} :- \texttt{alarm}, \texttt{hears\_alarm(X)}.$$

- ***DeepProbLog*** extends it to handle neural predicates

$$nn(m_q, \vec{t}, \vec{u}) :: q(\vec{t}, u_1); \ldots; q(\vec{t}, u_n) :- b_1, \ldots, b_m$$

$$nn(m_{\text{digit}}, \quad, [0, \ldots, 9]) :: \texttt{digit}(\quad, 0); \ldots; \texttt{digit}(\quad, 9).$$

De Raedt L. et al, “ProbLog: A probabilistic Prolog and its application in link discovery”, 2007
Manhaeve R. et al, “DeepProbLog: Neural Probabilistic Logic Programming”, 2018

DeepProbLog: Inference

- Step 1: Ground logic program w.r.t the query
- Step 2: Rewrites program into a formula in propositional logic
- Step 3: Compile formula into a Sentenial Decision Diagram
- Step 4: Evaluate the SDD bottom-up to calculate success probability
  - **When encountering a neural predicate -> Forward Pass!**

```
0.2::earthquake.
0.1::burglary.
alarm :- earthquake.
alarm :- burglary.
0.5::hears_alarm(mary).
calls(mary):-alarm, hears_alarm(mary).
```

(a) The ground program.

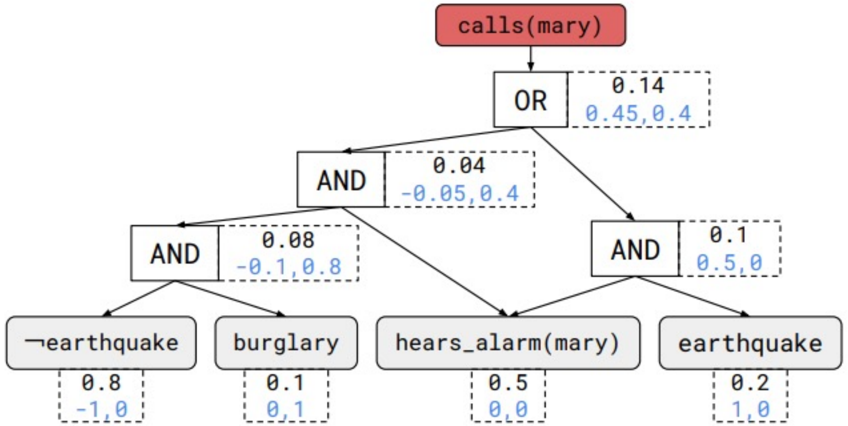

(b) SDD for query calls(mary).

Figure 1: Inference in ProbLog.

DeepProbLog: Learning

- Jointly learn parameters for probabilistic facts and NNets, loss based on query output

Learning based on a single query

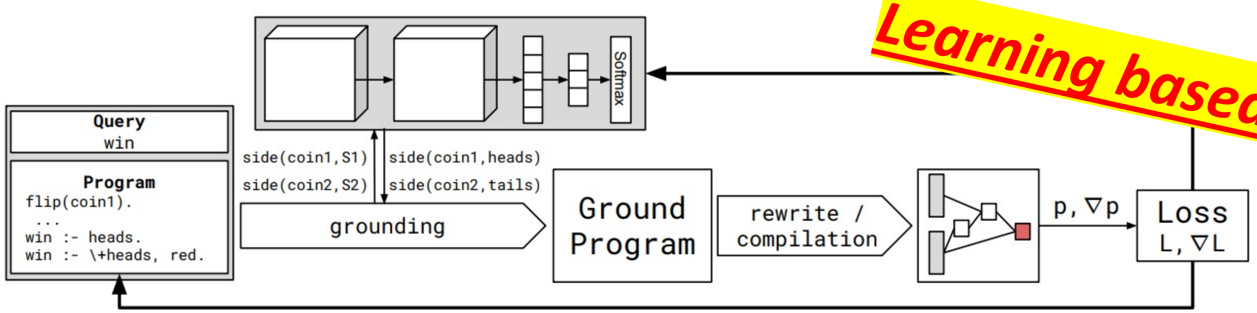

(a) The learning pipeline.

```
flip(coin1). flip(coin2).
nn(m_side,C,[heads,tails])::side(C,heads);side(C,tails).
t(0.5)::red;t(0.5)::blue.
heads :- flip(X), side(X,heads).
win :- heads.
win :- \+heads, red.
query(win).
```

(b) The DeepProbLog program.

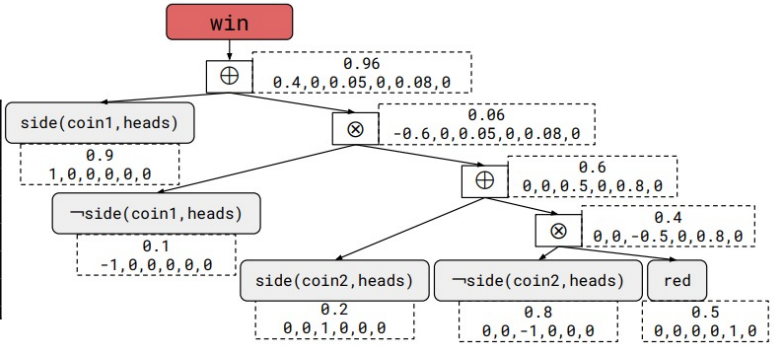

(c) SDD for query win.

Figure 2: Parameter learning in DeepProbLog.

Manhaeve R. et al, "DeepProbLog: Neural Probabilistic Logic Programming", 2018

Neural Representations + Symbolic Reasoning

- ***Deep Probabilistic Logic*** combines NNets with probabilistic logic for indirect supervision
- Label decisions modeled as latent variables
- Learning maximizes conditional likelihood of virtual evidence given input

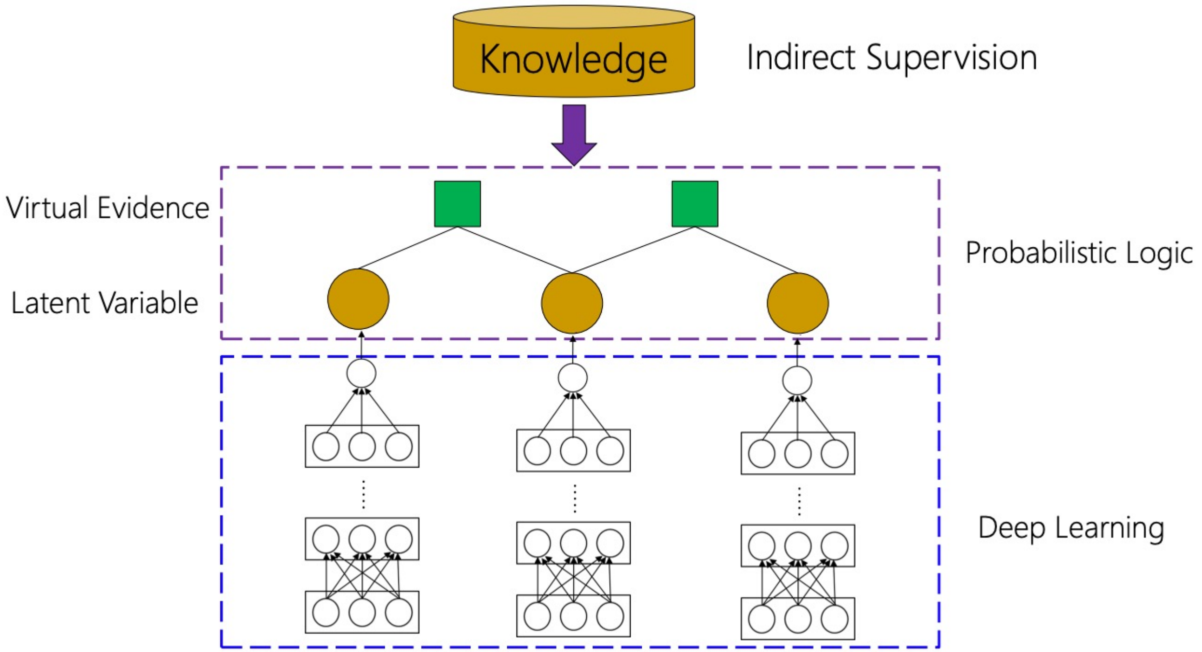

Wang H. and Poon H., "Deep Probabilistic Logic: A Unifying Framework for Indirect Supervision", 2018

Deep Probabilistic Logic

- We want to learn model $P(y|x)$ using a NNet with softmax on top
- Y is unobserved and learned using distance supervision
- We have a set of weak labeling functions $K = (\Phi_1, \dots, \Phi_V)$
- Dependencies between weak labeling functions and output
- Constraints on instances or model expectations can be introduced
- Learning is done using variational EM
- Approximate inference

Wang H. and Poon H., "Deep Probabilistic Logic: A Unifying Framework for Indirect Supervision", 2018

Deep Probabilistic Logic

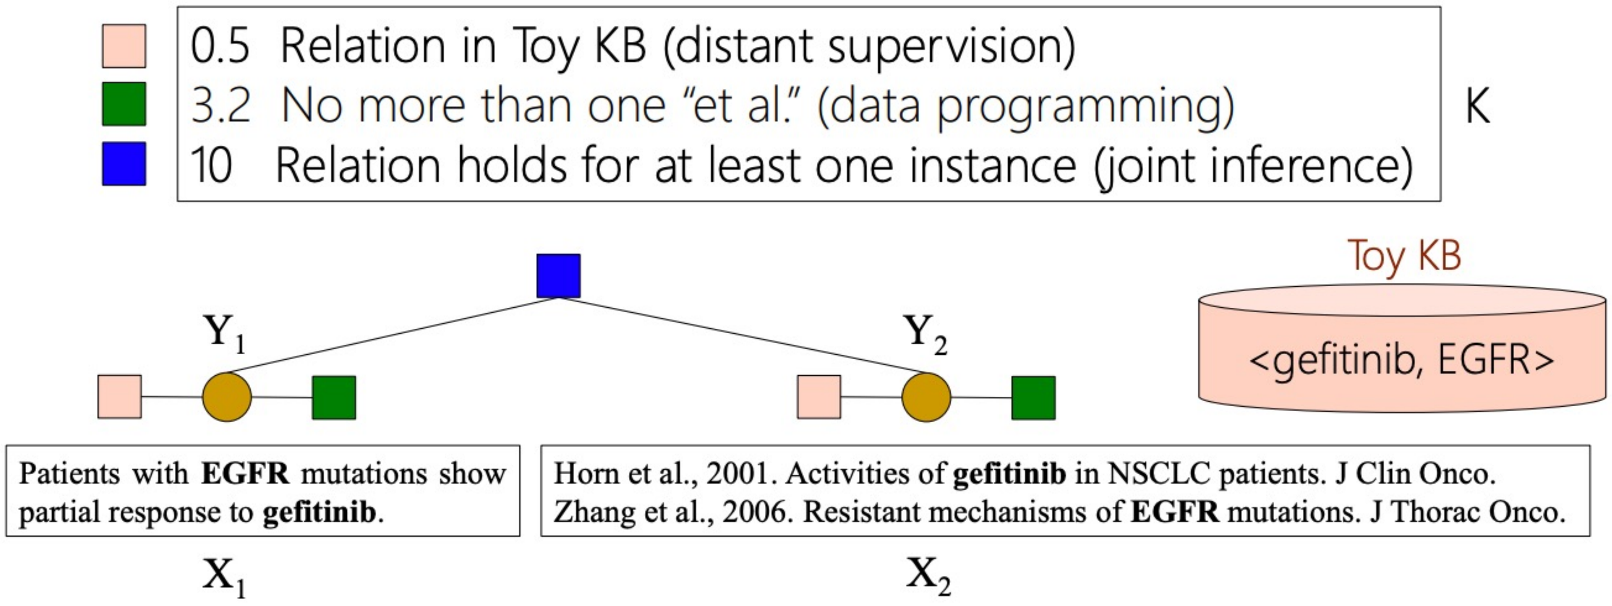

Y_1	Y_2	$P(K,Y \mid X) \propto$	$P(K, Y \mid X)$
T	T	$\exp(0.5\times2+3.2\times1+10\times1) = \exp(14.2)$	0.04
T	**F**	$\mathbf{\exp(0.5\times2+3.2\times2+10\times1) = \exp(17.4)}$	**0.94**
F	T	$\exp(0.5\times1+3.2\times1+10\times1) = \exp(13.7)$	0.02
F	F	$\exp(0.5\times0+3.2\times2+10\times0) = \exp(6.4)$	0

By combining distant supervision, data programming, and joint inference, DPL derives more accurate indirect supervision by inferring that the drug-gene relation likely holds in $\mathbf{X_1}$ but not in $\mathbf{X_2}$.

PURDUE UNIVERSITY

Wang H. and Poon H., "Deep Probabilistic Logic: A Unifying Framework for Indirect Supervision", 2018

NS Strategies in NLP Scenarios

NLP scenarios that have been at the center of recent neural-symbolic research

- Multi-Hop Reasoning for Question Answering
- Natural Language Grounding and Visual QA
- Common-sense Reasoning

Multi-Hop Reasoning for QA

Statements

- Socrates was born in Athens
- Athens belongs to Greece

Question

- Where was Socrates born?

Person born in city X & City located in country Y => Person born in country Y

Socrates was born in Athens

Socrates was born in Greece

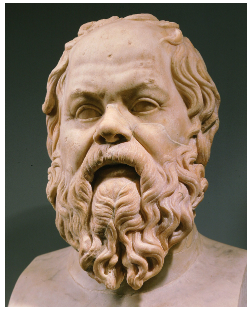

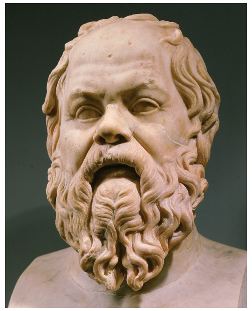

NNets + Logic Programming for Multi-Hop QA

- ***Prolog* -** Backward-chaining algorithm and unification procedure
  - **Horn-clauses** - $h\left(f_1^h, \ldots, f_n^h\right) \Leftarrow p1(f_1^1, \ldots, f_m^1) \wedge \ \ldots \wedge pB\left(f_1^B, \ldots, f_l^B\right)$
  - **Unification operator**: Given two atoms, find variable substitutions
    - $country(Greece, Socrates) \quad country(X, Y) \quad \{X/Greece, Y/Socrates\}$
  - **Backward chaining**: Given goal atom g
    - Check whether g is explicitly stated in KB
    - If not, find rules - unify g with heads of all available rules
    - If succeeds, resulting substitutions applied to body, and making each atom a subgoal
    - Recursively prove subgoals

Russel and Norvig "Artificial Intelligence, A Modern Approach", 2016

NNets + Logic Programming for Multi-Hop QA

- ***NLProlog*** - weak unification
  - Unify using differentiable **similarity fun.** with params θ
  - Comparing two atoms yields an aggregated **score**
  - Backward-chaining yields proofs with **proof scores**
  - Take **max over all scores** to get final proof success score

- ***Natural language statements*** as triples
  - (entity, entity, surface-pattern)
  - All elements are **embedded for computing similarities**
  - Resulting proof score is **end-to-end differentiable** w/r θ

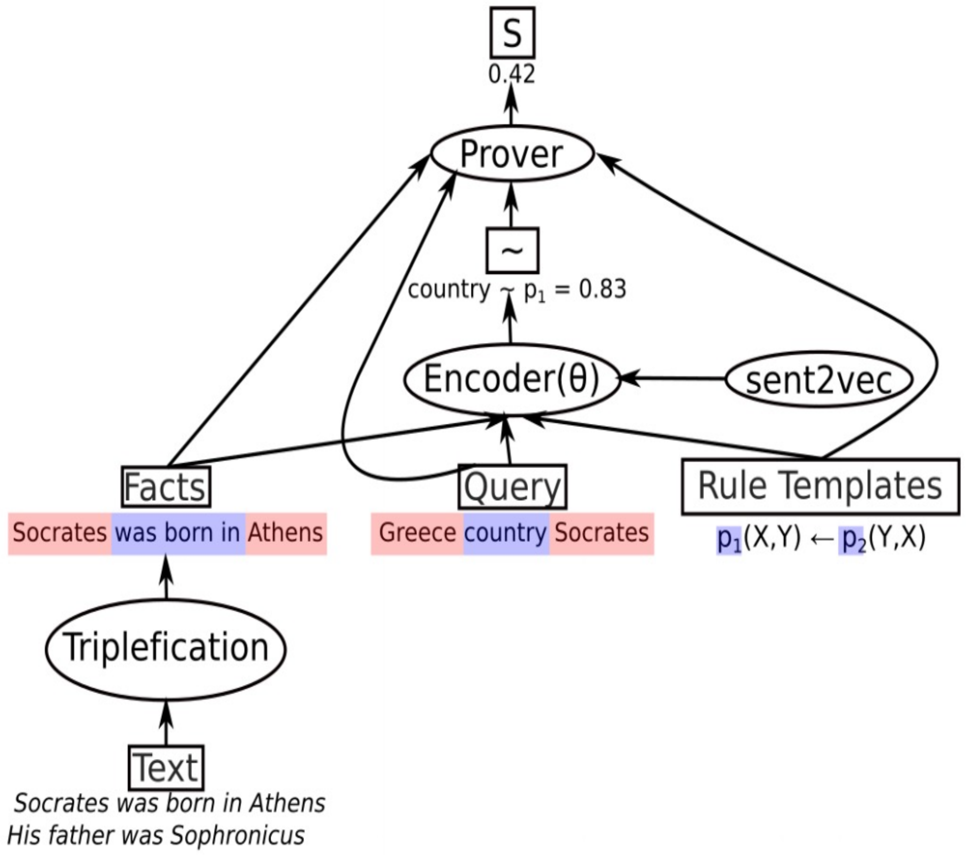

Weber L. et al., "NLProlog: Reasoning with Weak Unification for Question Answering in Natural Language " 2019

NNets + Logic Programming for Multi-Hop QA

- Rules that describe behaviors - two alternatives

  - ***Write down*** rules involving natural language patterns
  - Use rule templates to perform ***Inductive Logic Programming***
    - User defines the structure of rules as: $p1(X,Z) \Leftarrow p2(X,Y) \wedge p3(Y,Z)$
    - System instantiates multiple rules with randomly initialized embeddings for all p_i
    - Fine-tune using a downstream task

- Outcome
  - A system **competitive** with extractive neural QA approaches
  - Partially **interpretable model**, easier to debug

Weber L. et al., "NLProlog: Reasoning with Weak Unification for Question Answering in Natural Language " 2019

Grounding & Visual Question Answering

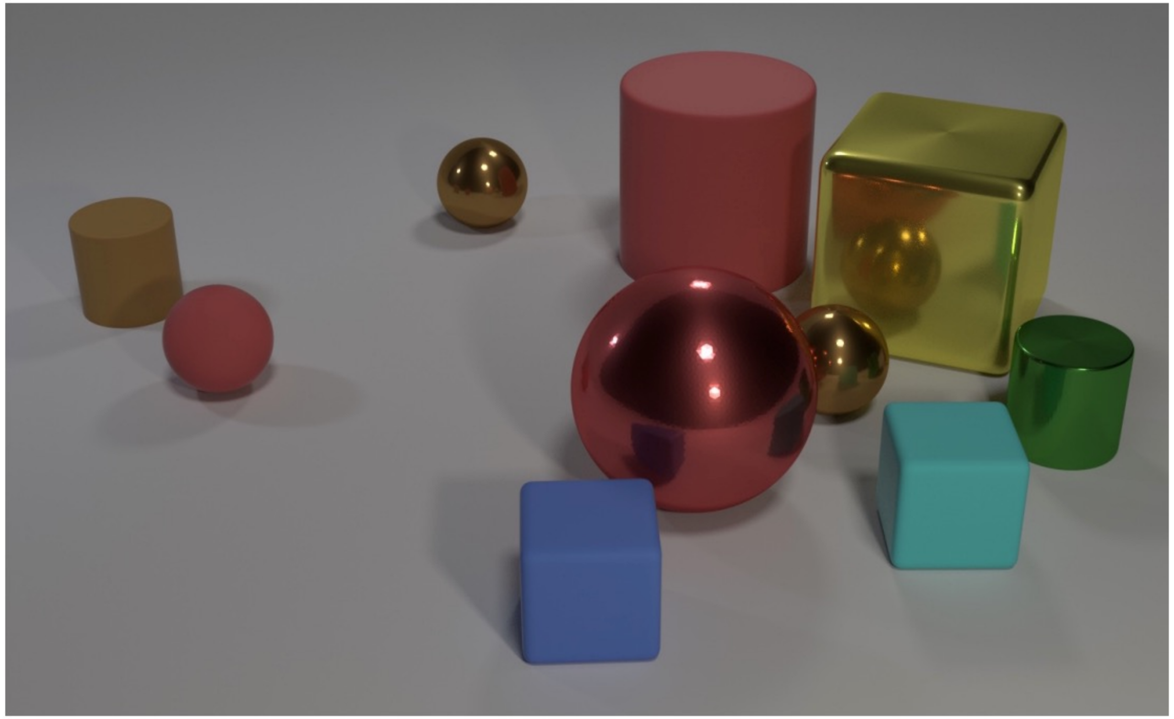

CLEVR dataset

There is a sphere which is the ***same size*** as the metal cube, is it made of the **same material** as the ***small*** red sphere?

Johnson, J. "CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning", 2017

NNets + Symbolic Reasoning for VQA

- **NS-VQA** disentangles vision and language understanding from reasoning
- NNets ***parse scenes*** from images, and ***generate programs*** for questions
- ***Symbolic program executor*** runs the program to obtain an answer
  - Direct mapping from vocabulary to python modules
- ***Learning:***
  - Scene parsing – standalone training
  - Small seed of labeled (question,program) examples, REINFORCE after execution

Yi K. et al., “Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding”

NNets + Symbolic Reasoning for VQA

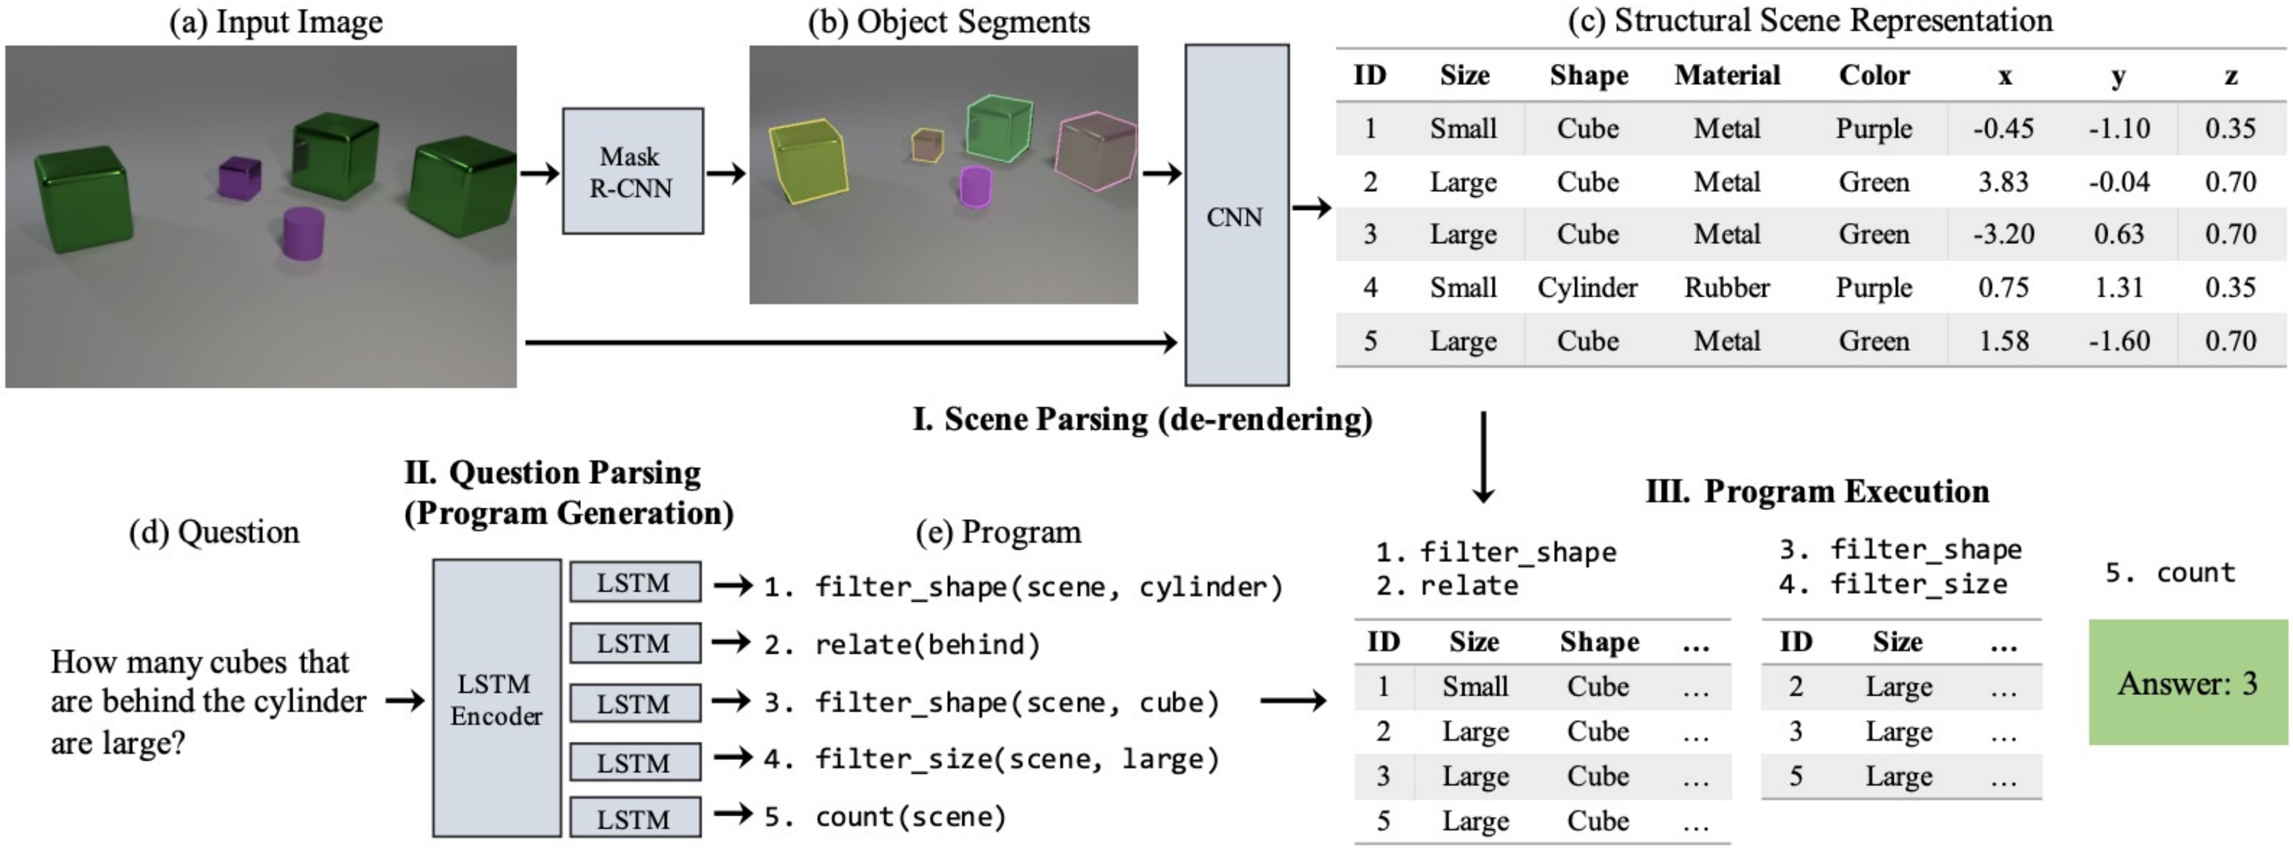

Yi K. et al., "Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding", 2018

NNets + Symbolic Reasoning for VQA

- ***NS-Concept Learner*** does not require visual supervision
- **Visual perception module** to construct object-based representation
  - NNet to generate object proposals
  - Embed attributes (color, shape) in the same embedding space

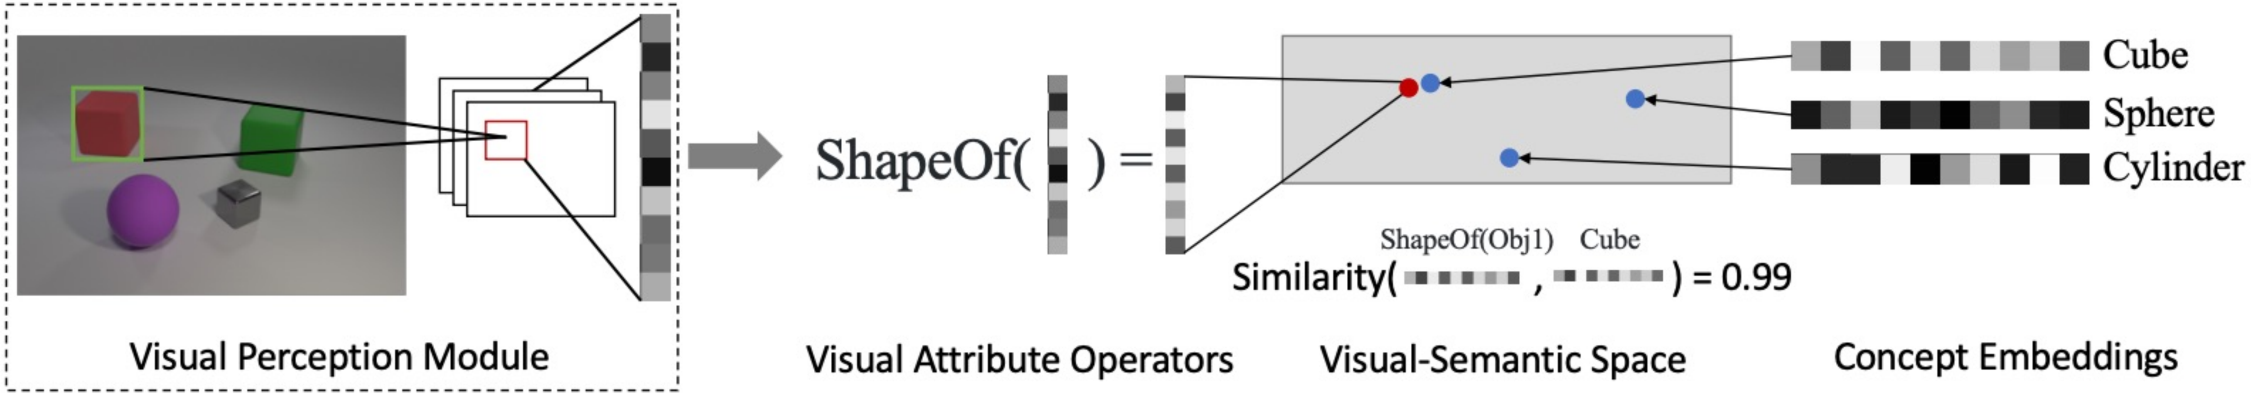

Mao J. et al., “The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words and Sentences from Natural Supervision”, 2019

NS-Concept Learner

- ***Semantic parsing module*** translates question into program
  - Encoder-decoder generates latent program
- ***Quasi-symbolic executor*** infers answer
  - Differentiable w.r.t visual representations
- Learn from (question,answer) pairs.
  - Visual parameters fully differentiable, REINFORCE to learn semantic parser

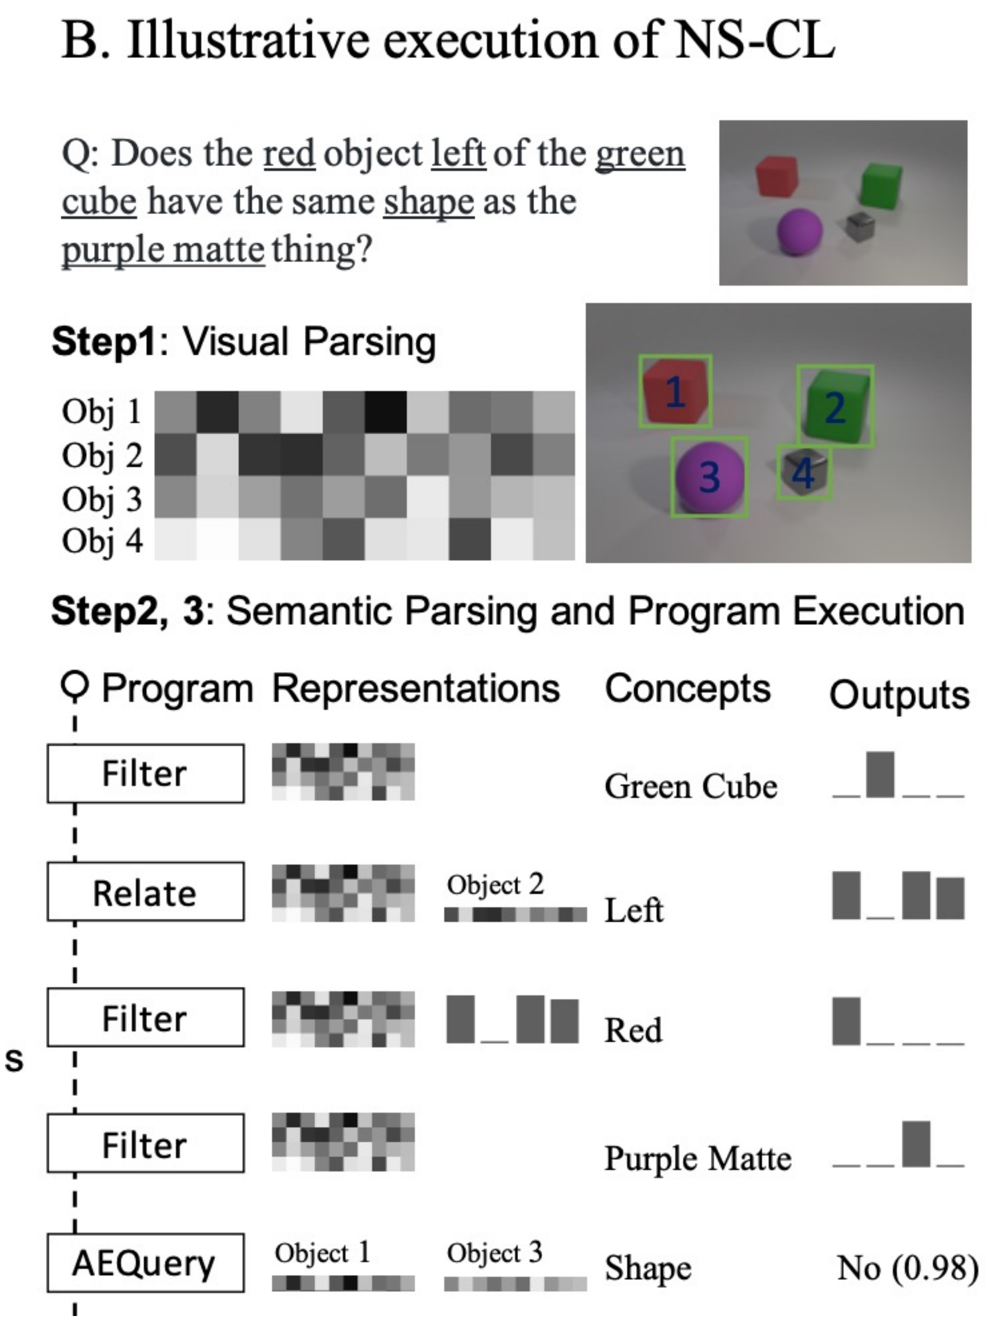

Mao J. et al., "The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words and Sentences from Natural Supervision", 2019

NS-Concept Learner

Curriculum learning approach splits learning in four stages

- Object-level visual concepts
- Relational questions
- Complex questions (perception modules fixed)
- Joint fine-tuning

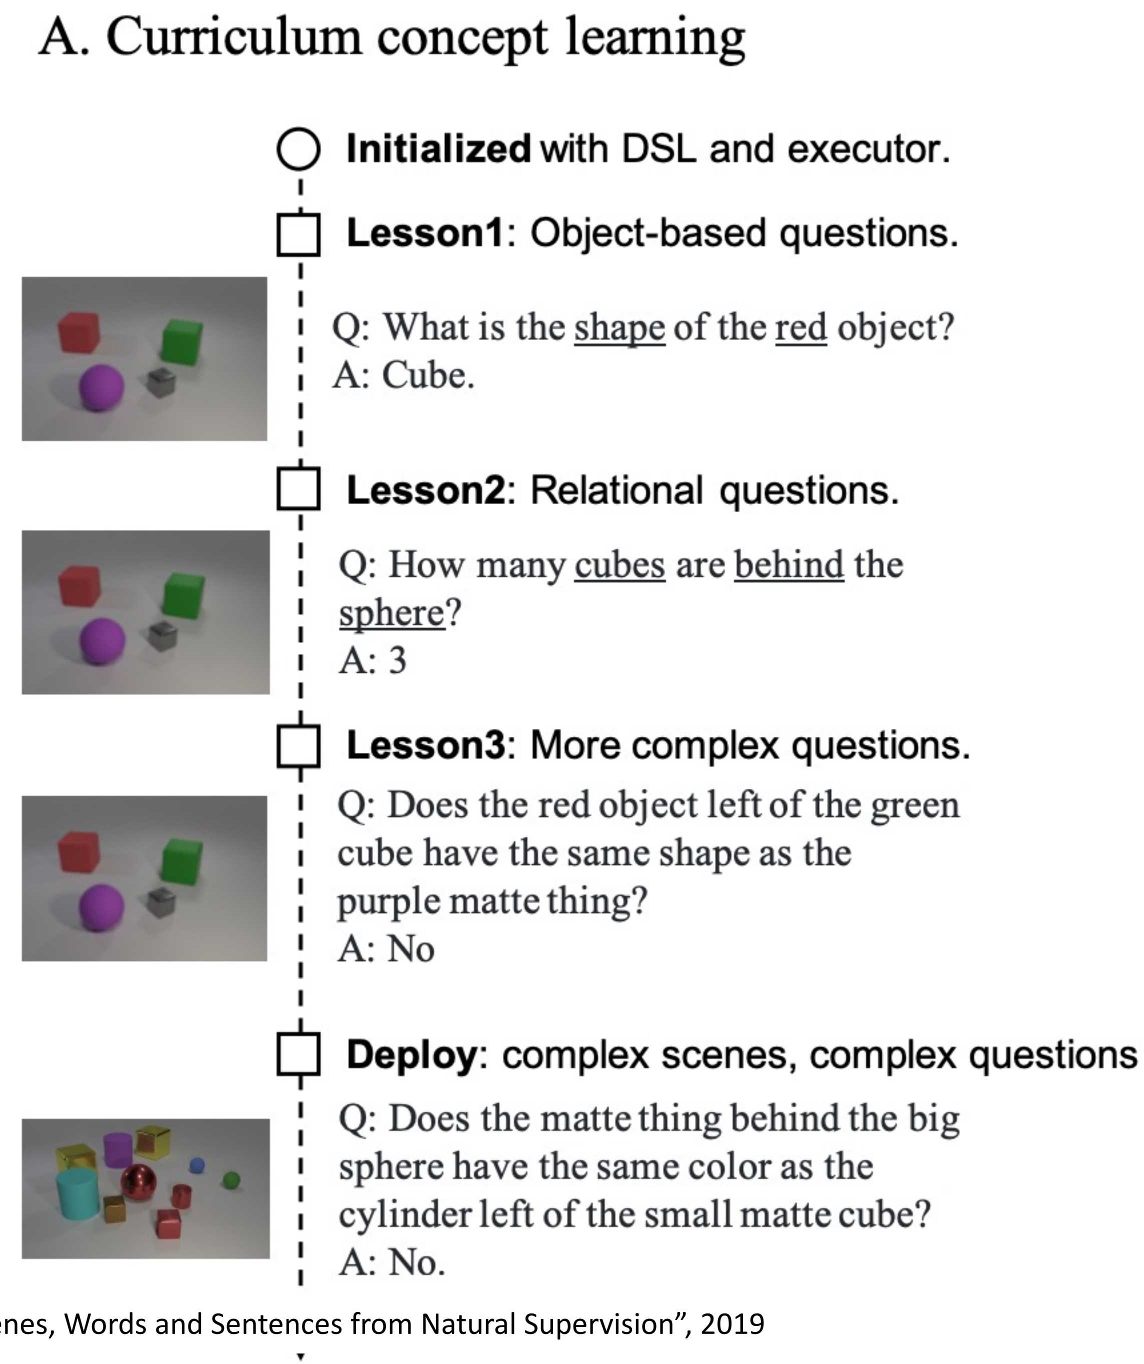

PURDUE UNIVERSITY

Mao J. et al., "The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words and Sentences from Natural Supervision", 2019

Commonsense Reasoning

- Understanding narratives requires **reasoning about implicit word knowledge**

They went to the club

They had to get dressed

They likely had drinks

NS Graph Construction for Commonsense QA

- Recent work augments DL models with CS KGs
- Dynamically generate knowledge that is ***contextually relevant***
- COMET (Bosselut et al, 2019) is a Transformer model for CS KGs, and can generate CS inferences for entities

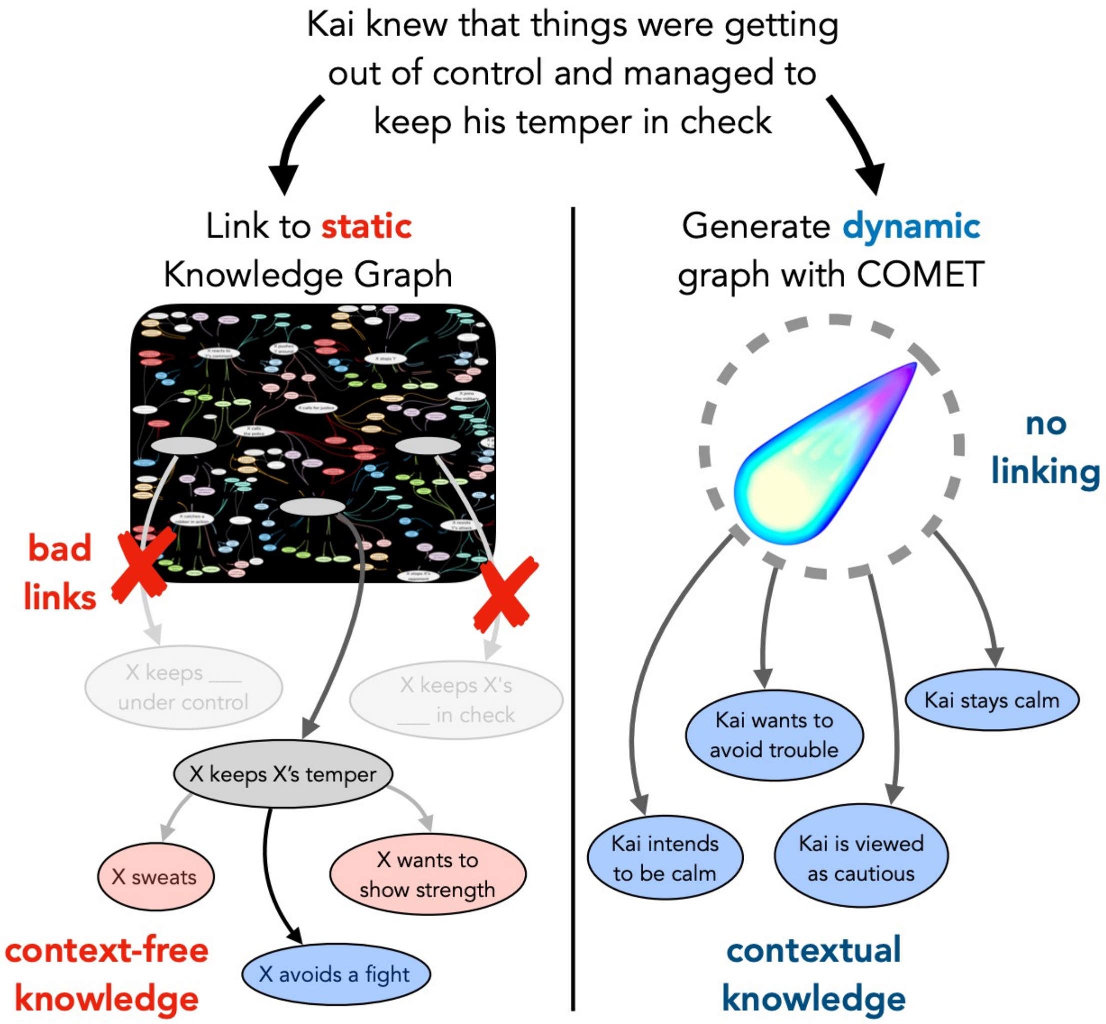

Bosselut A. et al., "COMET: Commonsense Transformer for Automatic Knowledge Graph Construction", 2020
Bosselut A. et al., "Dynamic Neuro-Symbolic Knowledge Graph Construction for Zero-shot Commonsense Question Answering", 2020

NS Graph Construction for Commonsense QA

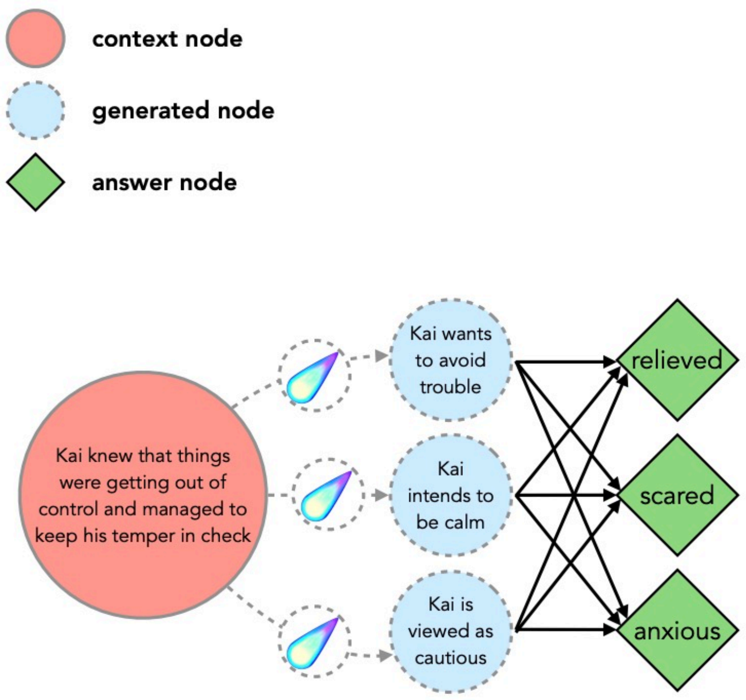

(a) COMeT receives the context c and generates new commonsense inferences to construct a local graph of knowledge about the situation (Section 2).

- Generate intermediate nodes by concatenating the context node with relation types in a KG (1-hop)
- Markovian assumption can generalize this to L-hops, recursively
- COMET also generates a score for each step
- Resulting factor graph can be reasoned over

Bosselut A. et al., "Dynamic Neuro-Symbolic Knowledge Graph Construction for Zero-shot Commonsense Question Answering", 2020

NS Graph Construction for Common-Sense QA

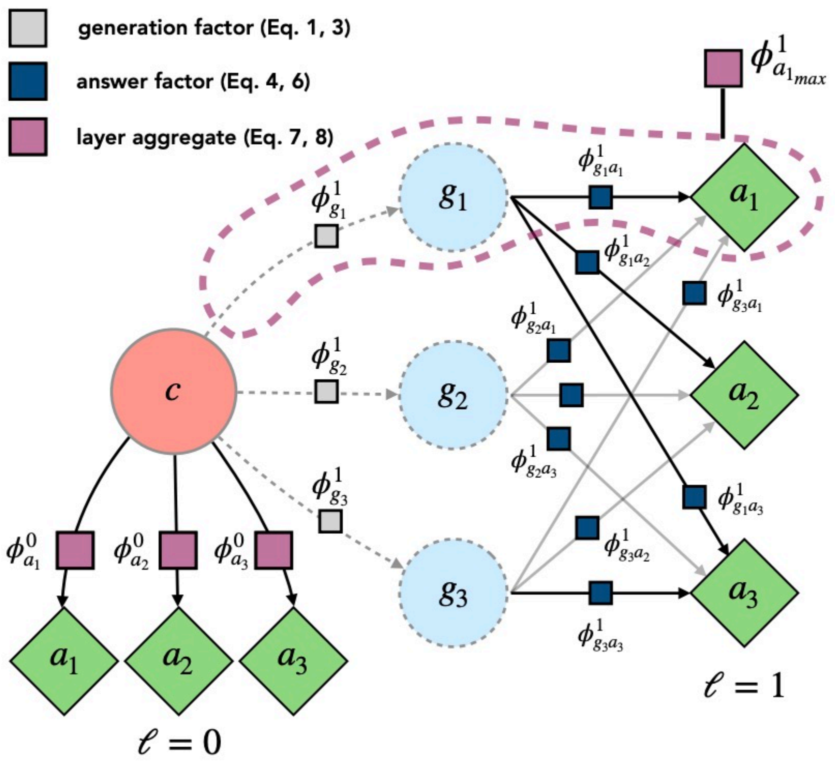

(b) Our inference algorithms reason over the graph by aggregating commonsense paths to answer questions about the situation (Section 3).

- Each path to an answer has a score
- Most likely answer can be found by marginalizing over all paths to the answers at layer L

Bosselut A. et al., “Dynamic Neuro-Symbolic Knowledge Graph Construction for Zero-shot Commonsense Question Answering”, 2020

Break!

NS 4 Discourse

Explanations: *Parsing the Landscape of Opinions and Perspectives*

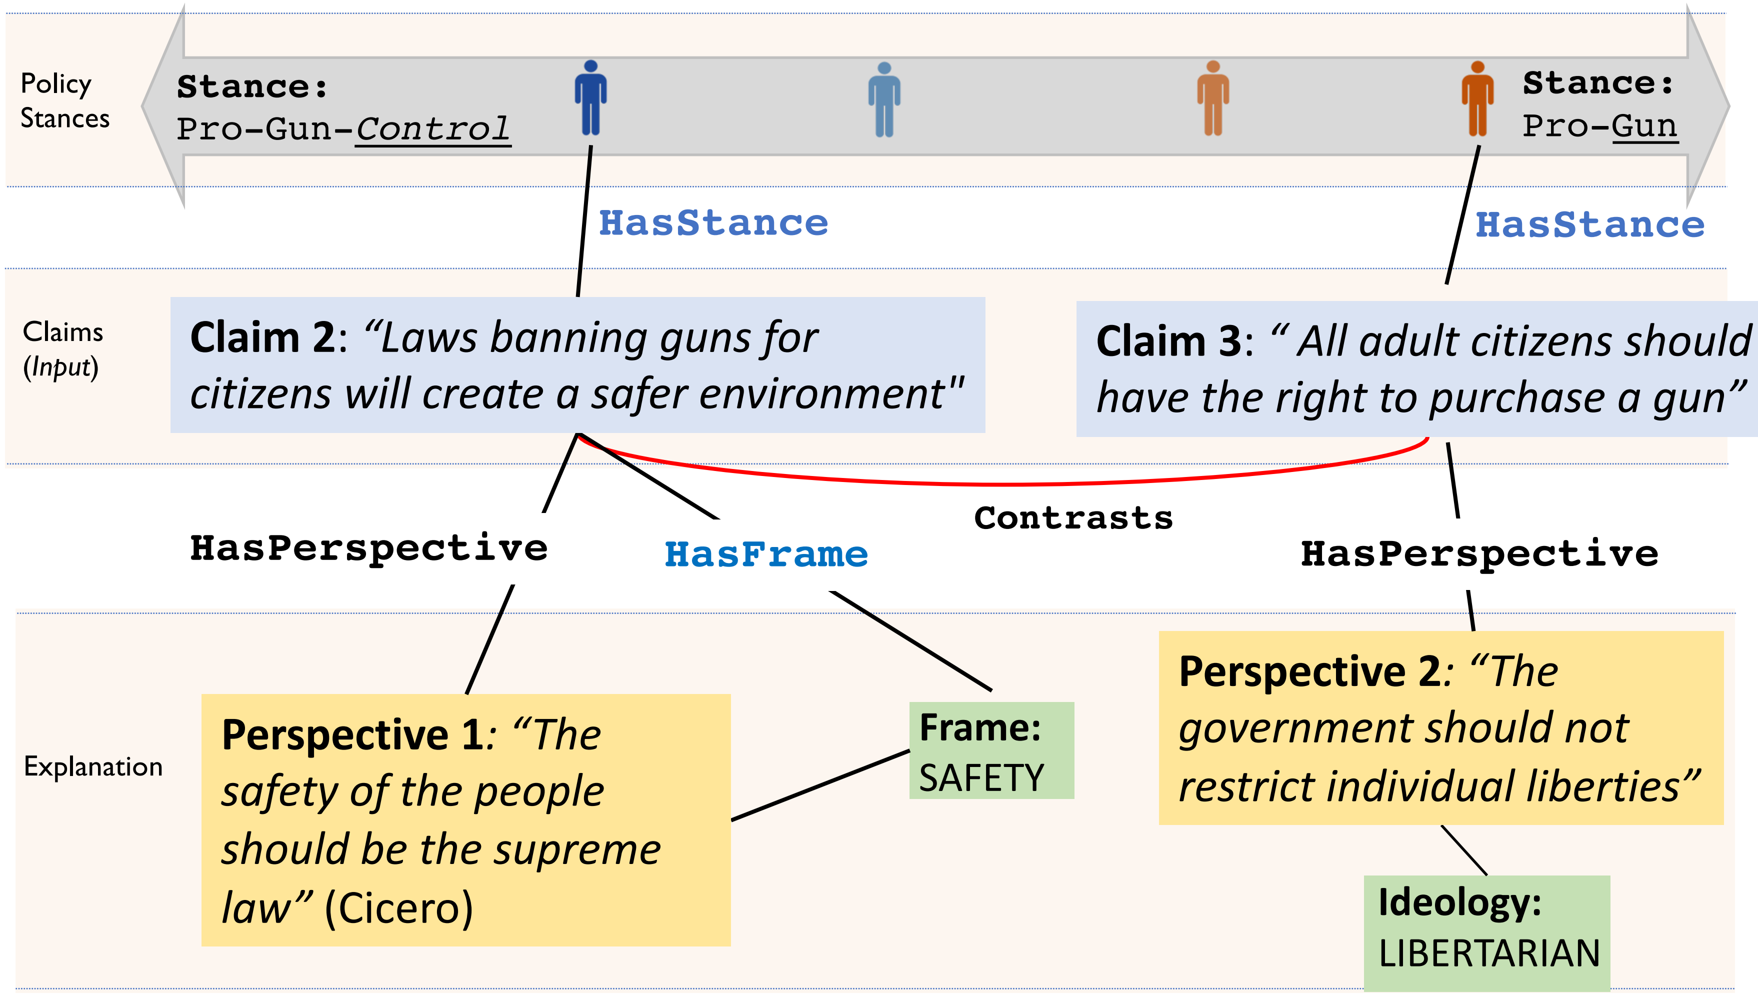

Capturing Symbolic Dependencies

Stance:
Pro-Gun-*Control*

HasStance

Stance:
Pro-Gun

HasStance

Claims
(*Input*)

Claim 2: *"Laws banning guns for citizens will create a safer environment"*

Claim 3: *" All adult citizens should have the right to purchase a gun"*

Contrasts

HasFrame

Identifying framing dimensions explains the stance prediction

Frame:
SAFETY

Enforcing dependencies between predictions can help constrain the output space

From Classification to Inference

- **Model relevant context as latent variables** which can be learned through human interactions
- **Replace classification with inference**: many decisions that should agree with each other, to support the decision

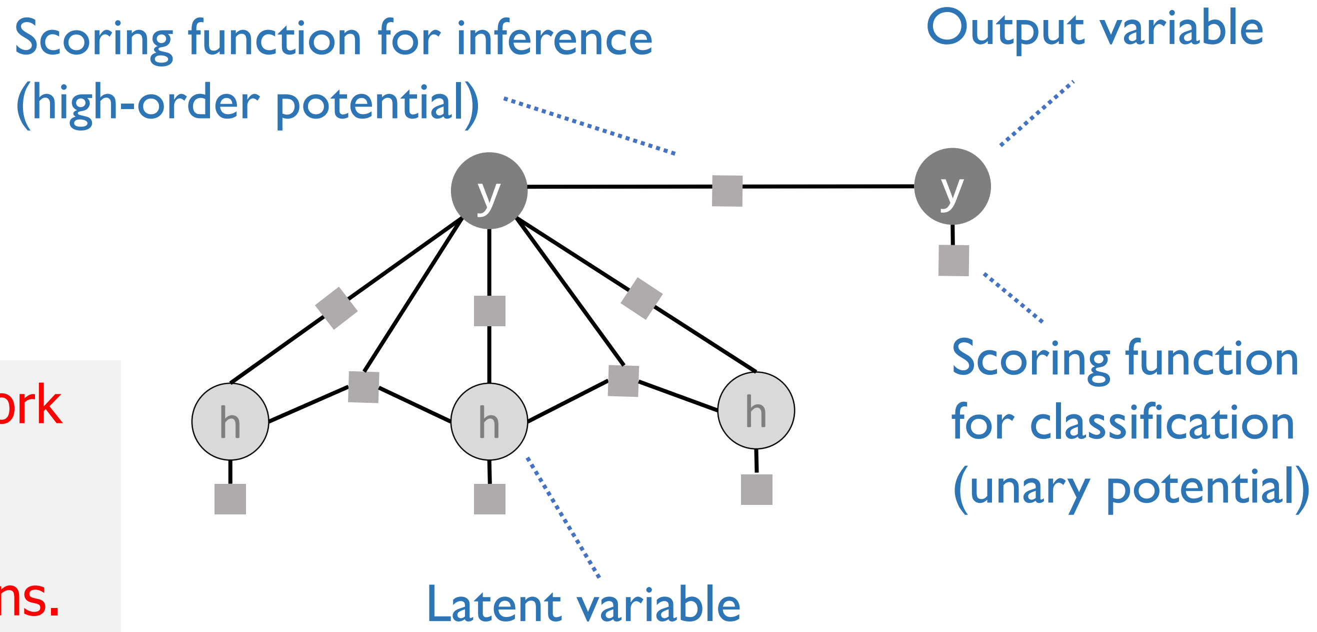

Very general framework for capturing social behavior and commonsense patterns.

Declarative Modeling

- Often easier to think about structure in a declarative way
- Define **entities** and **relations** and **probabilistic rules**

Agree(doc1, doc2) ∧ Label(doc1, Pos) → Label (doc2, Pos)

$$\arg\max_{\boldsymbol{y}\in\{0,1\}^n} P(\boldsymbol{y}|\boldsymbol{x}) \equiv \arg\max_{\boldsymbol{y}\in\{0,1\}^n} \boxed{\sum_{\psi_{r,t}\in\Psi} w_r\ \psi_r(\boldsymbol{x_r}, \boldsymbol{y_r})}$$

$$s.t.\ c(\boldsymbol{x_c}, \boldsymbol{y_c}) \leq 0; \quad \forall c \in C$$

Characterizing Social Context through Representations

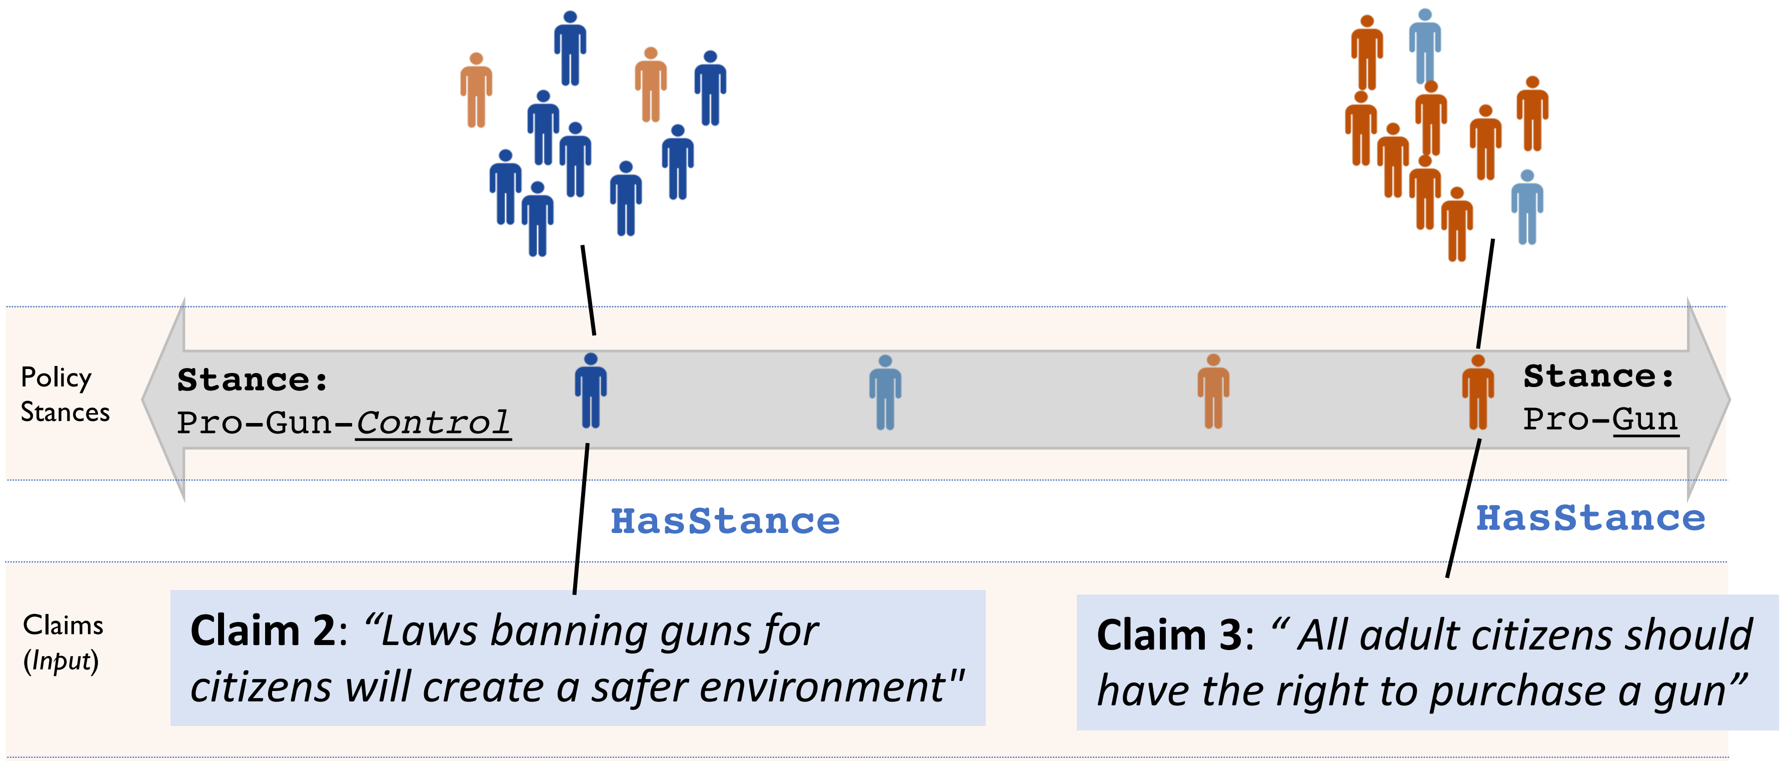

Deep Relational Learning

- **DRaiL**, Declarative framework for Deep Relational Learning
  - <u>Rules as context</u>, using a graphical model
  - <u>Representation as context</u>, using neural architectures

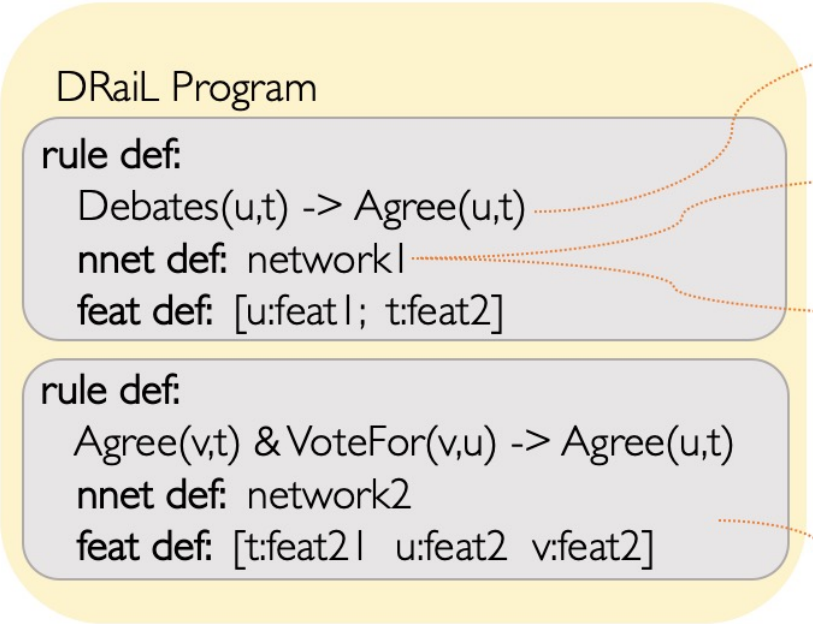

Each rule:

(1) Defines a factor template to capture dependencies between decisions contained in it

(2) Defines a sub-graph for compositional representation learning

Pacheco and Goldwasser, “Modeling Content and Context with Deep Relational Learning”, 2021

Deep Relational Learning

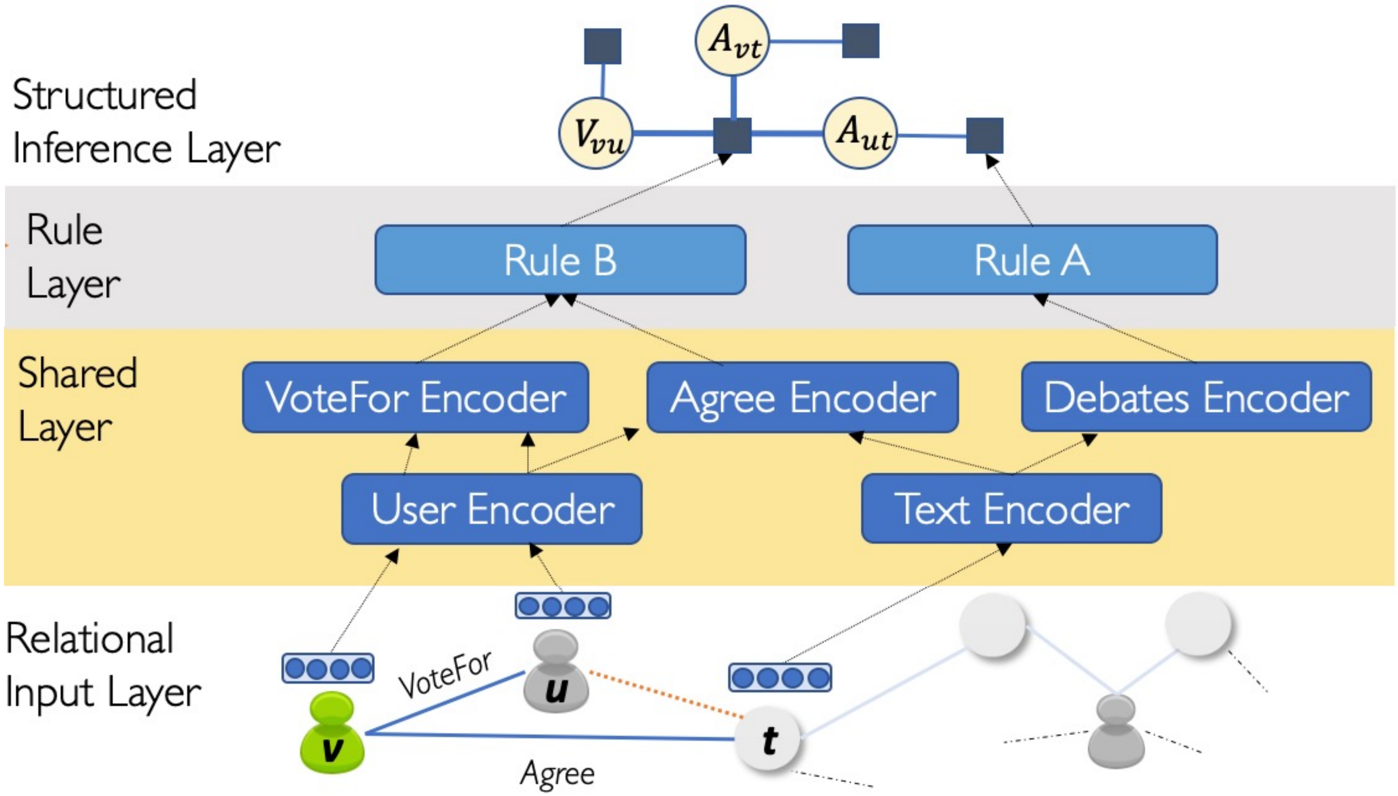

DRaiL compiles logic programs into a neural factor graph and learns its parameters given data

Deep Relational Learning

Neural potentials over a relational embedding space

Neural part

Symbolic part

$$\arg\max_{\boldsymbol{y}\in\{0,1\}^n} P(\boldsymbol{y}|\boldsymbol{x}) \equiv \arg\max_{\boldsymbol{y}\in\{0,1\}^n} \sum_{\psi_{r,t}\in\Psi} w_r\, \psi_r(\boldsymbol{x_r}, \boldsymbol{y_r})$$

$$s.t.\; c(\boldsymbol{x_c}, \boldsymbol{y_c}) \leq 0; \quad \forall c \in C$$

Learning: We use the structured hinge-loss over the neural representation

$$\max(0, \max_{\mathbf{y}\in Y}(\Delta(\mathbf{y}, \hat{\mathbf{y}}) + \sum_{\psi_r\in\Psi} \Phi_t(\boldsymbol{x_r}, \boldsymbol{y_r}; \theta^t)) - \sum_{\psi_r\in\Psi} \Phi_t(\boldsymbol{x_r}, \boldsymbol{y_r}; \theta^t))$$

PURDUE UNIVERSITY

Understanding Debates Networks

Debated Claim :
"Limiting gun sales is unconstitutional"

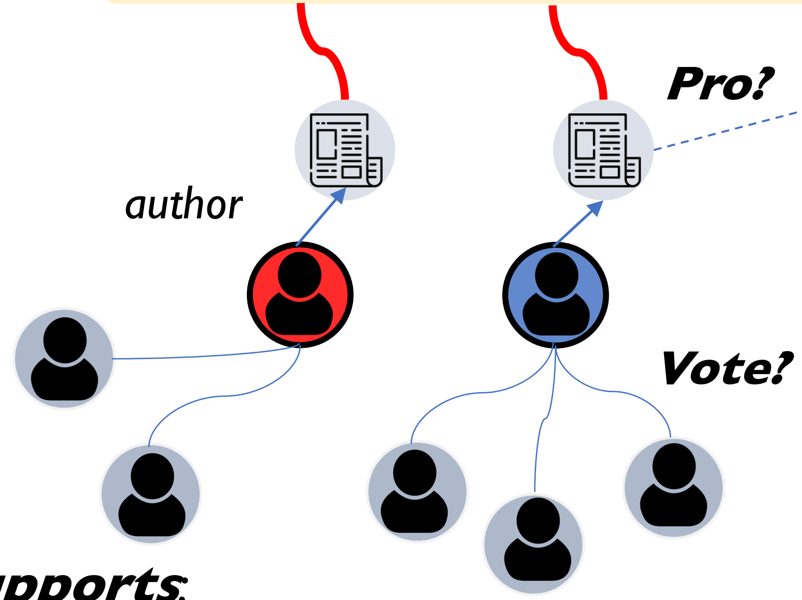

"this is an interpretation that is politically motivated. The right to bear arms does not mean that it cannot be regulated"

Supports:
Death Penalty
Border Fence

Supports:
Same-sex marriage
Gun regulations

Some Stats:

Affiliation	# Debating Users	# Voting Users	# Initiated Debates	# Authored Posts
Liberal	4,876	1,105	8,323	64,111
Conservative	3,009	726	5,391	42,753
Unknown	7,894	1,554	8,080	55,173
Total	15,779	3,385	21,794	162,037

Understanding Debates Networks

This results in several reasoning tasks:

- **Textual Inference:** does the post support the debate claim?
- **Authors and Text**: *who is likely to write/support such claims?*
- **Authors and other Users**: who would they vote for?

- **Clearly, there's structure to be exploited!**

"Limiting gun sales is unconstitutional"

"this is an interpretation that is politically motivated. The right to bear arms does not mean that it cannot be regulated"

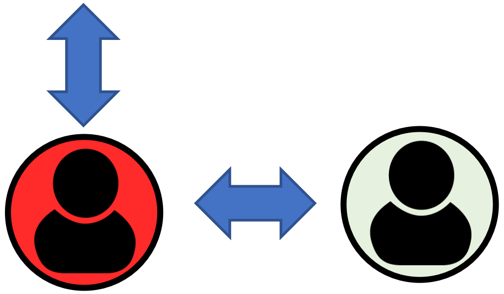

Evaluating Modeling Choices

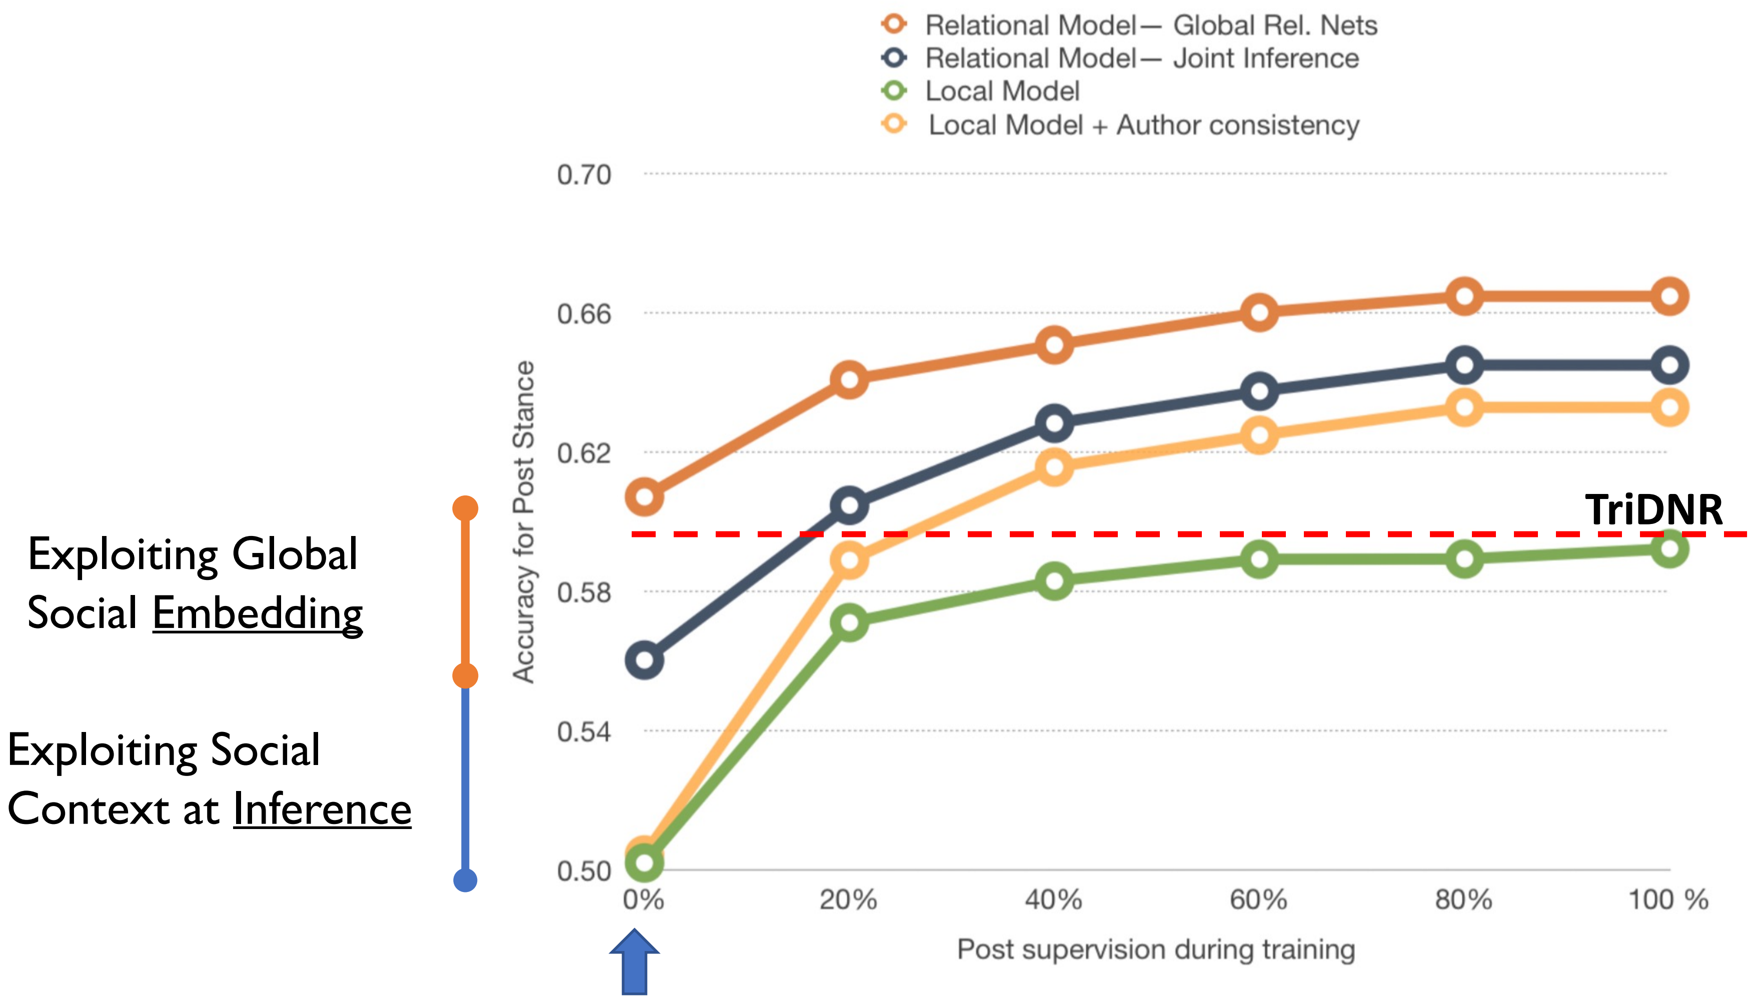

Distant Supervision Using Explanation Inference

Summary: DRaiL

- **A general framework for combining symbolic and neural representations**

- **Neural**: captures "implicit" interactions between entities in the embedding space.
- **Symbolic:** explicit interactions between entities, forced to provide a consistent view
- **Neuro-Symbolic:** consistency constraints are propagated to the embedding space

- Provides a convenient way to compile symbolic explanations into neural classifiers

Symbolic Explanations
for
Discourse Analysis

Capturing Symbolic Dependencies

Stance:
Pro-Gun-*Control*

HasStance

Stance:
Pro-Gun

HasStance

Claims
(*Input*)

Claim 2: *"Laws banning guns for citizens will create a safer environment"*

Claim 3: *" All adult citizens should have the right to purchase a gun"*

Contrasts

HasFrame

Identifying framing dimensions explains the stance prediction

Frame:
SAFETY

Enforcing dependencies between predictions can help constrain the output space

Symbolic Explanations for …?

- In "tightly" closed worlds, relevant symbolic vocabulary emerges naturally
  - Colors, shapes, spatial relations, etc.
- **When looking to explain peoples' intent, motivations and opinions, this is not always so clear!**
- We often need to look outside of NLP to reason about these concepts
  - *Morality, ideology, psychological state, etc.*
- **Open problems:**
  - What are the relevant concepts to look at?
  - Can they be operationalized? Predicted reliably?
  - Do they have enough "explanatory power" to be used for inference?

Modeling Mental States

- A lot of effort in modeling emotions, intent and motivation
  - **Personality traits** (McCrae and Costa 1997, Myers-Briggs), in **NLP** (Resnik'13, Neuman Tutorial EMNLP'15, Plank and Hovy'15, Pizzolli and Strapparava'19, Lynn et-al'20)
  - **Agency and Power** (Danescu-Niculescu-Mizil'12, Mohammad '18, Sap et-al '17, Jaidka et-al'18, Sap et-al'20)
  - **Psychological state**

    (Rashkin et-al'18)

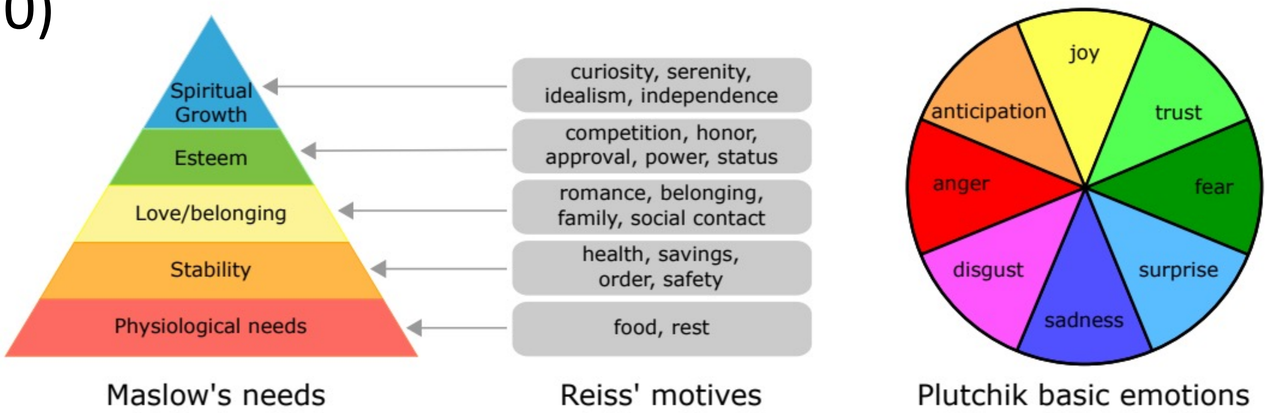

Obtaining Reliable Human Ratings of Valence, Arousal, and Dominance for 20,000 English Words. Mohammad, ACL'18

Modeling and Visualizing Locus of Control with Facebook Language. Jaidka et-al, ICWSM'18

SOCIAL BIAS FRAMES: Reasoning about Social and Power Implications of Language. Sap et-al '20

Gifts differing: Understanding personality type. Isabel Briggs Myers and Peter Myers. 2010.

Personality trait structure as a human universal. McCrae and Costa 1997

Echoes of power: Language effects and power differences in social interaction. Danescu-Niculescu-Mizil et-al '12

SOCIAL BIAS FRAMES: Reasoning about Social and Power Implications of Language. Sap et-al '20

Modeling Naive Psychology of Characters in Simple Commonsense Stories. Rashkin et-al ACL'18

Modeling Psychological States: a N-S view

Cindy really likes apples.

She wanted to try something new with them.

She decided to try to make baked apples for the first time.

She gathered everything she needed and began cooking.

It's now her favorite apple dish!

Desire Expression: try something new with them
Motivation (Reiss): Curiosity
Emotion (Plutchik): Joy, Anticipation

Desire Fulfilled!
Motivation (Reiss): Independence
Emotion (Plutchik): Joy

Break down the text into character-driven storylines, represented as a graph

➔ Modeling the characters' mental states helps in understanding the narrative flow

Modeling Human Mental States with an Entity-based Narrative Graph. Lee, Pacheco et-al NAACL'21

Modeling Psychological States: a N-S view

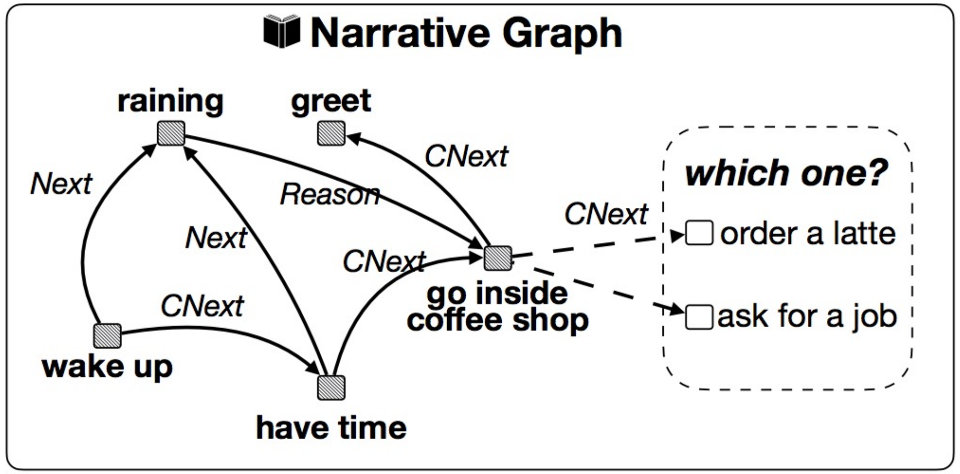

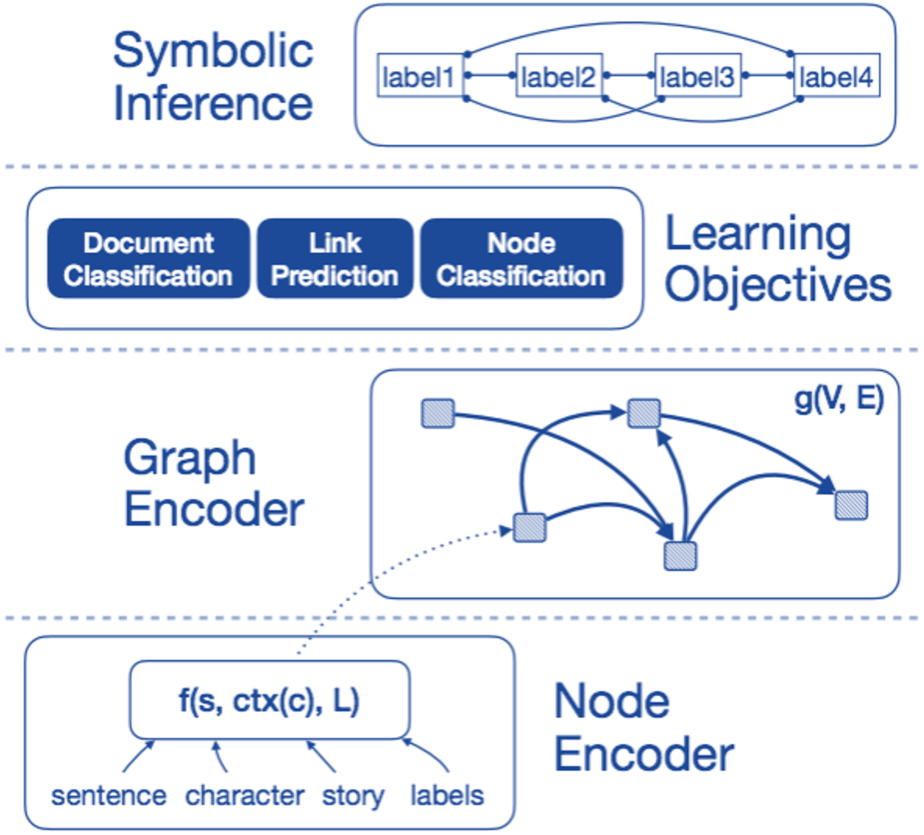

<u>Entity Narrative Graph represents the text</u>

***Nodes**: events associated with characters' roles*

***Edges:** represents the narrative flow using*

Next/Cnext edges ,(*event order***)** *and* **Discourse relations** *(cause, contrast, temporal, etc.)*

- Train a relational NN to recover the graph edges

ates with an Entity-based Narrative Graph. Lee, Pacheco et-al NAACL'21

Modeling Psychological States: a N-S view

Shapes code event type (verb), **colors** code character's emotional states based on Maslow's pyramid

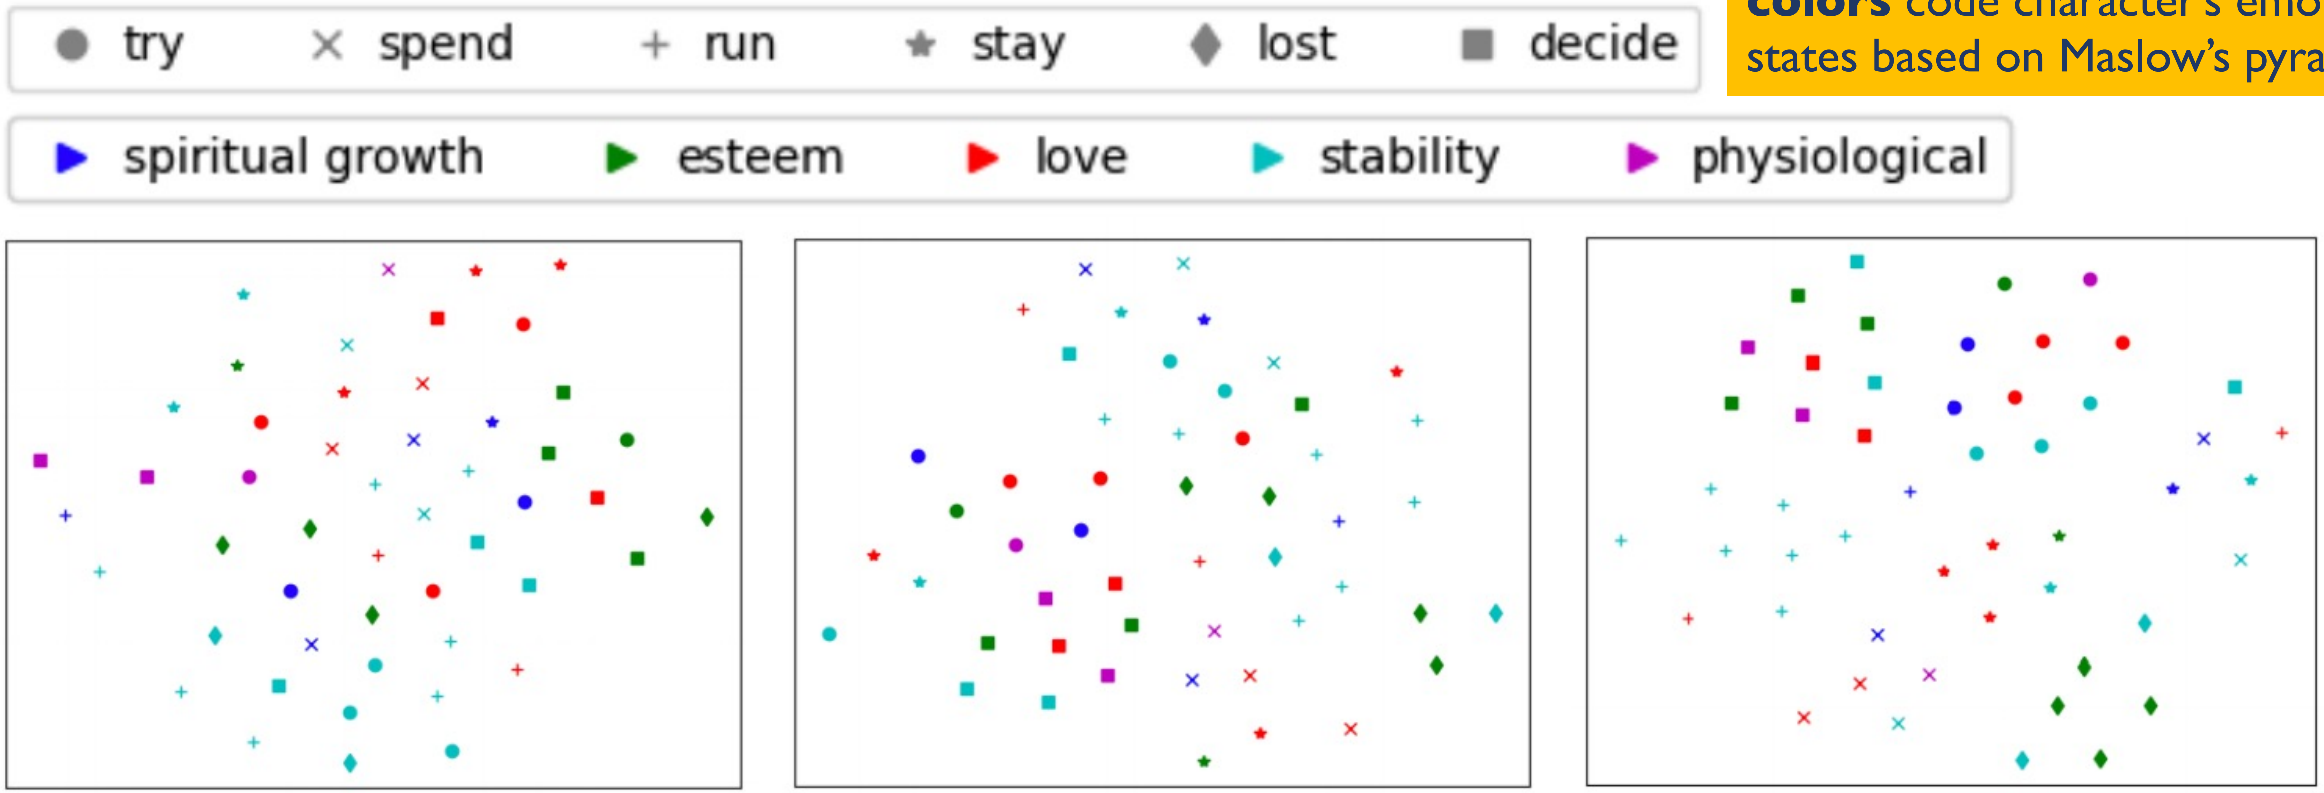

(a) ENG-CTX

(b) W-CTX-STORY

(c) W-CTX-SENT

Modeling Human Mental States with an Entity-based Narrative Graph. Lee, Pacheco et-al NAACL'21

Modeling Psychological States: a N-S view

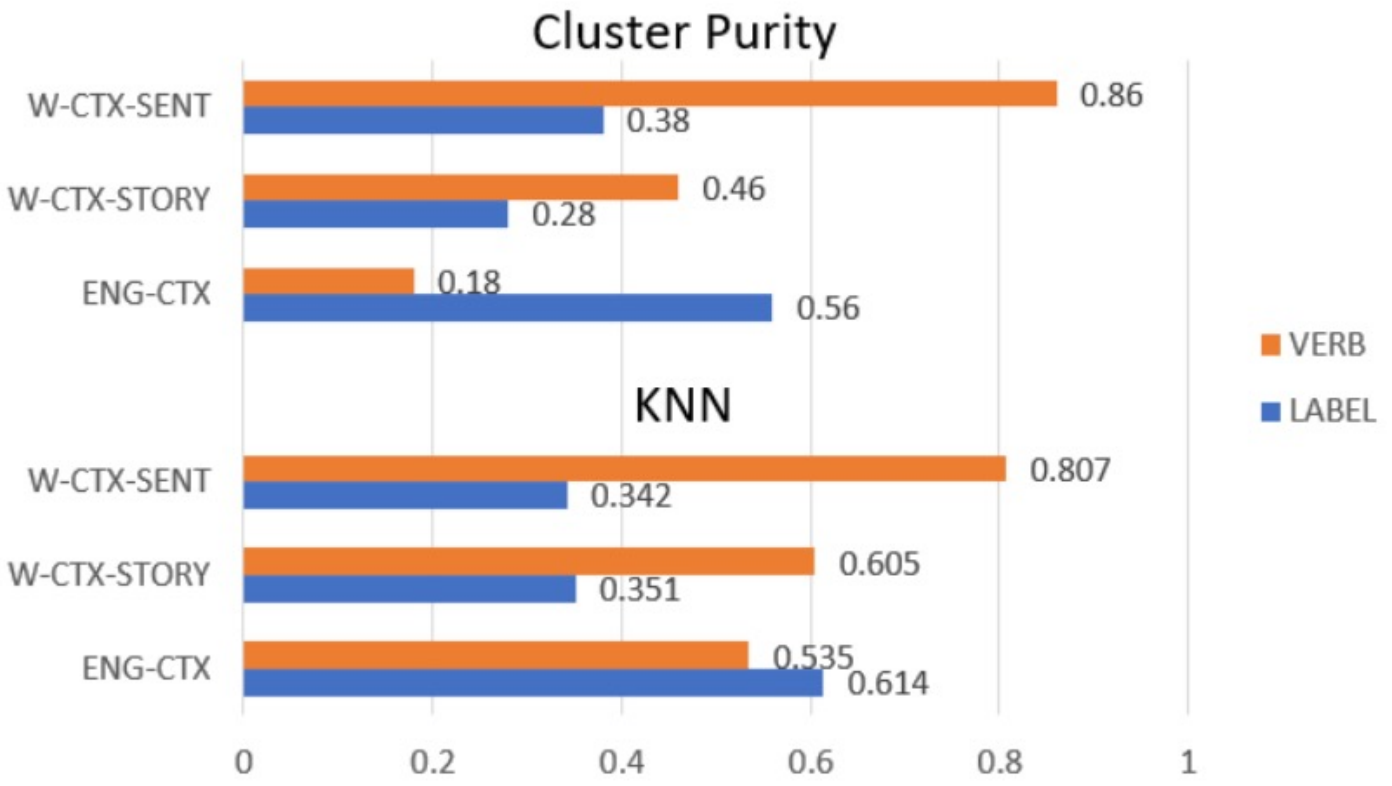

Modeling Human Mental States with an Entity-based Narrative Graph. Lee, Pacheco et-al NAACL'21

Decoding Political Messaging

- Social media is the primary tool for political discourse
- Unlike traditional outlets, **this is a conversation!**
  - a way to test ideas and get feedback, actively build support
- Highly dynamic and diverse text
  - Ungrammatical, short, coded language
  - Constant adaptation to new issues and styles

Can we automatically decode this information?

Johnson and Goldwasser (Coling ’16, ACL ‘17, ACL ’18) Roy and Goldwasser (EMNLP ’20)

Political Issue Stance and Framing

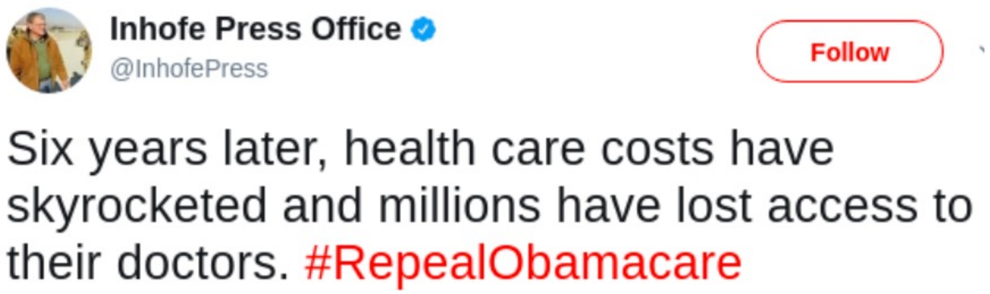

Stance:
Clearly, not a fan.

Framing: *what are the right abstractions of the tweet, capturing the arguments supporting the stance?*

"Obamacare should be repealed since it is too expensive"

Healthcare is framed as an ***economic*** issue

Framing

- ***Perspective through which information is presented*** (Denis Druckman'07)
- Often *an intentional decision*, priming the discussion towards a stance
  - US presence in Afghanistan: **cost** or **commitment to democracy**?
  - Abortion: **life of the unborn** or **women's rights?**
- **Where do framing dimensions come from?**
  - Developed for a specific issue (Choi et-al'12), or general policy frames across multiple issues (Boydstun et-al'14, Card et-al'15, Johnson et-al'16'17,Field'18)
  - Emerging from data (Tsur et-al'15, Demszky et-al'19)

The Media Frames Corpus: Annotations of Frames Across Issues. Card et-al, ACL'15

Framing Theory. Denis and Druckman'07 *Annual Review of Political Science*

Hedge detection as a lens on framing in the GMO debates: A position paper. Choi'12

Leveraging behavioral and social information for weakly supervised collective classification of political discourse on Twitter Johnson et-al ACL'17

A Frame of Mind: Using Statistical Models for Detection of Framing and Agenda Setting Campaigns Tsur et-al ACL'15

Analyzing Polarization in Social Media: Method and Application to Tweets on 21 Mass Shootings. Demszky et-al, 2019

Framing and Agenda-setting in Russian News: a Computational Analysis of Intricate Political Strategies Field et-al, EMNLP'18

Moral Foundations in Tweets

America woke up to heartbreaking news from Las Vegas. We stand united in our shock, our condolences, & our prayers.

Another horrific shooting. Another unspeakable horror. My thoughts are with everyone at Marjory Stoneman Douglas High School after this terrible day.

.. Stance can be harder to determine..

Moral Foundations

Human morality organized around 5 foundations, emerging from evolutional, cultural and social origins (Haidt, 2007)

- ***Each foundation has a positive and negative aspect (praise/judgement)***

1. Care/ Harm: care for others, generosity, compassion, sensitivity to suffering of others, prohibits actions that harm others

2. Fairness/ Cheating: Fairness, justice, reciprocity, rights, autonomy, prohibits cheating

3. Loyalty/ Betrayal: Group affiliation and solidarity, patriotism, self-sacrifice

4. Authority/ Subversion: Fulfilling social roles, authority, hierarchy, tradition.

5. Purity/ Degradation: association with sacred and holy, disgust contamination, striving to live in an elevated way.

When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. Haidt and Graham'07

Context for Twitter Analysis

- Model factorizes according to different aspects: social, behavioral and linguistic heuristic indicators
- Used graphical model inference to combine these aspects (PSL)
- Weak Supervision - Use EM to learn how to combine simple heuristic models

Modeling Social and Behavioral Information

```
SameTempAct(P1,P2) ∧
SameSentiment(P2,P1) →
SameStance(P1,P2)
```

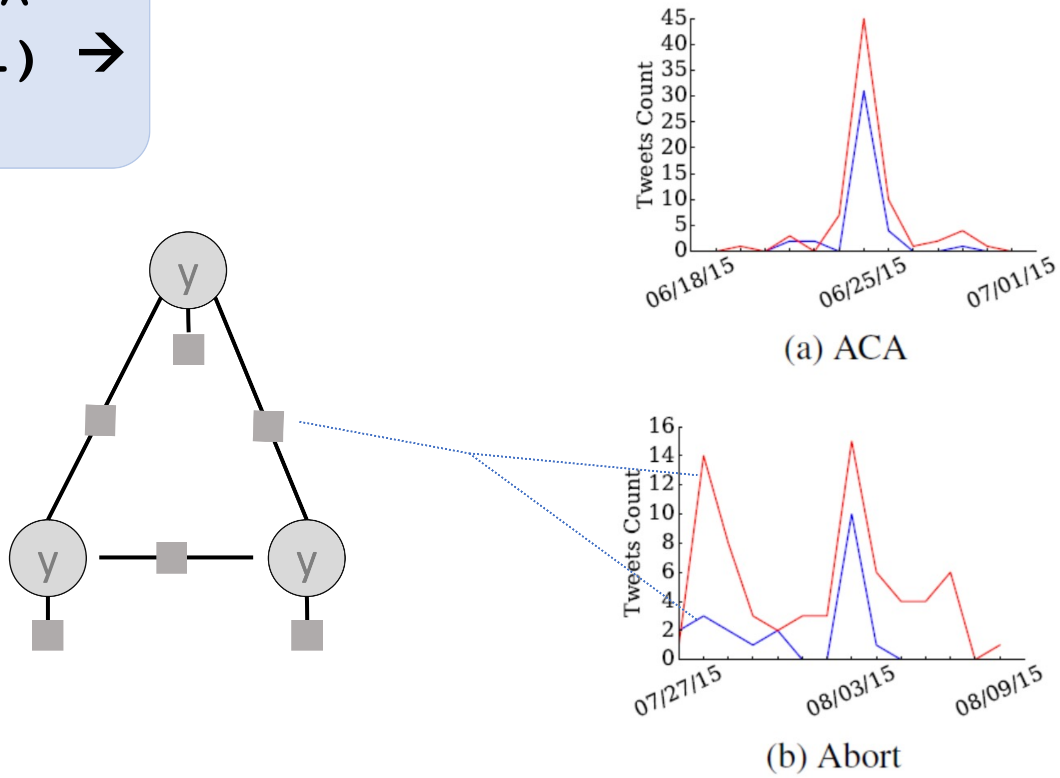

Align temporal Twitter activity as an indication of agreement

Modeling Social and Behavioral Information

```
Retweets(T1,T2) ∧ Frame(T1,F) → Frame(T2,F)
Follows(T1,T2) ∧ Frame(T1,F) → Frame(T2,F)
```

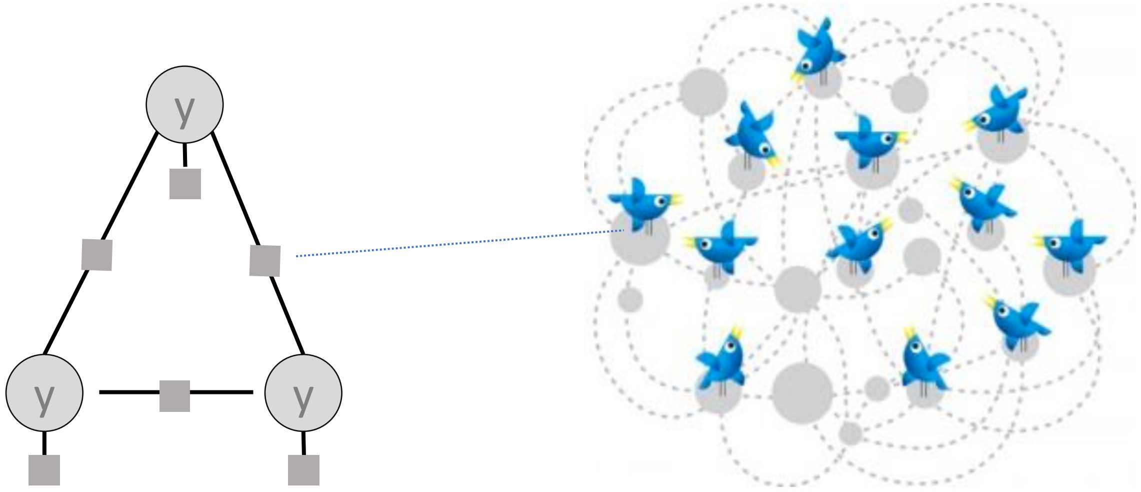

Use twitter social activity (e.g., `follows`, `retweet`) as an indication of agreement

Results: *Framing*

Frame Number	Frame	Results of Unsupervised PSL Model Frame Predictions						
		Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 6
1	Economic	31.82	31.52	69.57	72.22	72.22	**73.23**	**89.88**
2	Capacity & Resources	23.38	28.51	40.00	**41.18**	**41.18**	**41.18**	**79.55**
3	Morality	28.63	29.41	47.67	53.98	43.06	**53.99**	**87.43**
4	Fairness	33.49	47.19	59.15	63.50	63.50	**64.74**	**82.35**
5	Legality	44.58	46.93	58.02	60.64	60.63	**64.54**	82.16
6	Crime	7.89	7.62	73.33	75.00	75.00	**76.92**	**76.92**
7	Security	42.50	40.24	51.83	62.09	61.68	**64.09**	**88.48**
8	Health	48.36	48.79	79.43	86.49	86.49	**86.67**	**79.71**
9	Quality of Life	17.82	21.99	48.89	52.63	52.63	**54.35**	**82.93**
10	Cultural	15.38	15.67	51.22	52.63	52.63	**55.56**	**85.71**
11	Public Sentiment	15.22	15.72	50.79	53.97	41.03	**54.69**	**29.41**
12	Political	49.06	48.20	50.29	46.99	46.99	**47.23**	**74.52**
13	Policy	39.88	39.39	37.02	42.77	42.77	**43.79**	**65.06**
14	External Regulation	12.66	14.22	44.44	66.67	66.67	**71.43**	**85.71**
15	Factual	24.64	19.21	70.95	70.37	70.41	**78.95**	**82.85**
16	(Self) Promotion	40.11	46.41	48.16	50.96	50.96	**52.89**	**91.76**
17	Personal	45.36	46.15	59.66	62.99	62.13	**71.20**	**77.55**
	Weighted Average	37.14	38.79	53.13	56.49	55.54	**58.66**	**77.79**

Only text heuristics

No Label
(full model)

Supervised
Skyline
(full model)

Healthcare Framing Analysis

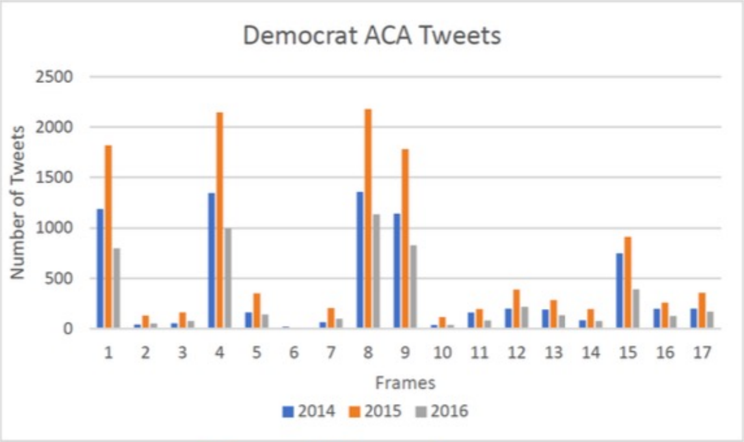

Frame 1: Economic
Frame 4: Fairness & Equality
Frame 8: Healthy & Safety
Frame 9: Quality of Life
Frame 15: Factual

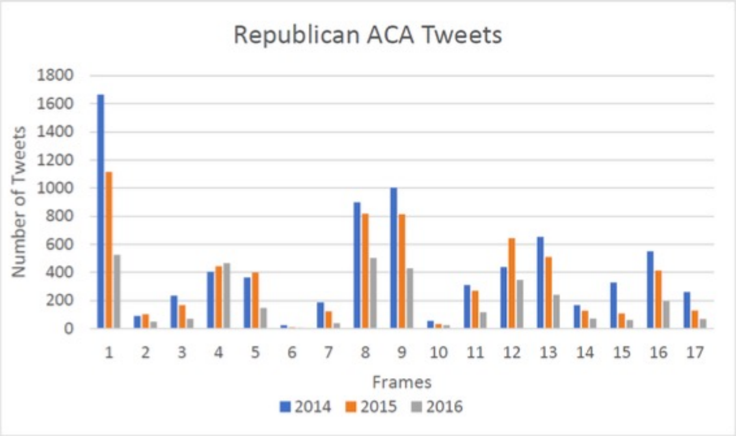

Frame 1: Economic
Frame 8: Healthy & Safety
Frame 9: Quality of Life

Our analysis also showed that aisle-crossing Republicans used similar frames as Democrats!

Question: *when Republicans and Democrats talk about health case as an Economic issue, do they make the same point?*

Weakly Supervised Nuanced Frame Extraction

- Generalized issue frames might not capture ideological lines

Adapted from **Alternet** (*Left*)	*Adapted from* **Breitbart** (*Right*)
Employees-many of whom are undocumented immigrants from Mexico, Ecuador and elsewhere-toil seven days a week for less than minimum wage, with no overtime pay.	Mass immigration has come at the expense of America's working and middle class, which suffered from poor job growth, stagnant wages, and increased public costs.

- Adapt general frames to topic-specific **sub frames**
  - Construct a **lexicon** of repeating talking points
  - Collapse similar talking points to **sub-frames**
  - **Contextualize** by embedding frames, subframes and text in a shared space
  - **Evaluate** by ability to explain ideological viewpoints

Weakly Supervised Nuanced Frame Extraction

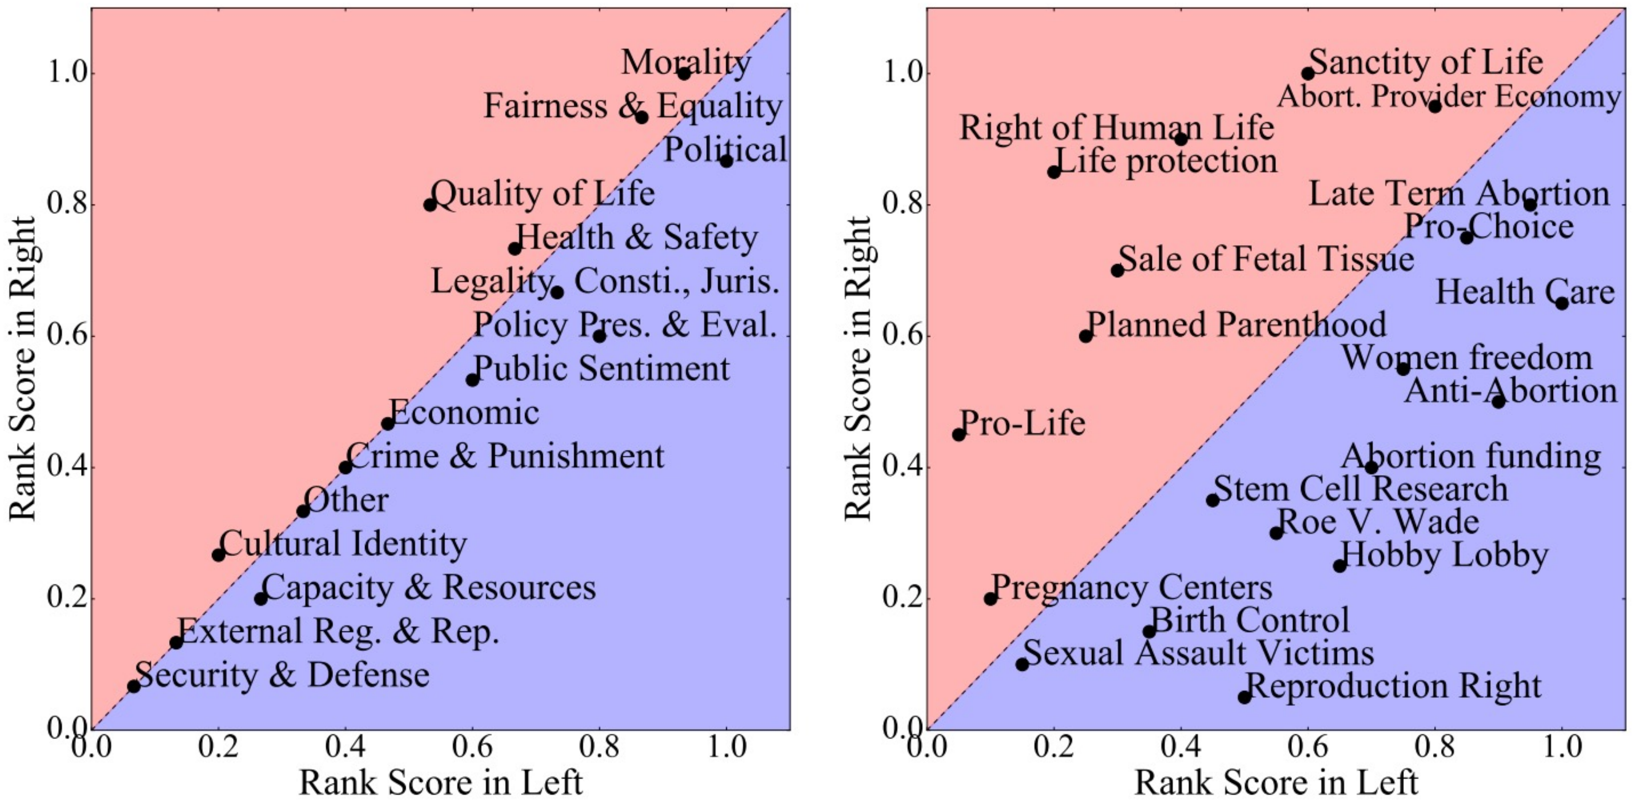

(a) Frame Usage (b) Subframe Usage

Morality Frames

- Moral Foundation Theory was repeatedly used to explain behaviors.
  - Liberals emphasize Fairness, Conservatives emphasize Loyalty and Authority
- ***But.. Everybody CAREs*** ***… but not about the same things!***

[@SenThadCochran and I]$_{CARING}$ are working to protect [MS small businesses]$_{CARE-FOR}$ from more expensive [#Obamacare mandates]$_{HARMING}$.

[The ACA]$_{CARING}$ was a life saver for the more than [130 million Americans]$_{CARE-FOR}$ with a preexisting condition – including covid now. [Republicans]$_{HARMING}$ want to take us back to coverage denials.

Morality Frames

- We define a morality frame structure to capture differences in the **targets** of moral sentiment

If the target of CARE is "illegal immigrants", author more likely to be a…

If the causer of HARM is "illegal immigrants", author more likely to be a…

MORAL FOUNDATIONS	MORAL ROLES
CARE/HARM: Care for others, generosity, compassion, ability to feel pain of others, sensitivity to suffering of others, prohibiting actions that harm others.	1. Target of care/harm 2. Entity causing harm 3. Entity providing care
FAIRNESS/CHEATING: Fairness, justice, reciprocity, reciprocal altruism, rights, autonomy, equality, proportionality, prohibiting cheating.	1. Target of fairness/cheating 2. Entity ensuring fairness 3. Entity doing cheating
LOYALTY/BETRAYAL: Group affiliation and solidarity, virtues of patriotism, self-sacrifice for the group, prohibiting betrayal of one's group.	1. Target of loyalty/betrayal 2. Entity being loyal 3. Entity doing betrayal
AUTHORITY/SUBVERSION: Fulfilling social roles, submitting to authority, respect for social hierarchy/traditions, leadership, prohibiting rebellion against authority.	1. Justified authority 2. Justified authority over 3. Failing authority 4. Failing authority over
PURITY/DEGRADATION: Associations with the sacred and holy, disgust, contamination, religious notions which guide how to live, prohibiting violating the sacred.	1. Target of purity/degradation 2. Entity preserving purity 3. Entity causing degradation

Morality Frames

- DRaiL Model

Basic Classifiers

$$r_1 : \texttt{Tweet}(\texttt{t}) \Rightarrow \texttt{MF}(\texttt{t}, \texttt{m})$$
$$r_2 : \texttt{Tweet}(\texttt{t}) \wedge \texttt{Ent}(\texttt{t}, \texttt{e}) \Rightarrow \texttt{Role}(\texttt{t}, \texttt{e}, \texttt{r})$$

Party's messaging preferences

$$r_3 : \texttt{Tweet}(\texttt{t}) \wedge \texttt{Ideo}(\texttt{t}, \texttt{i}) \wedge \texttt{Topic}(\texttt{t}, \texttt{k}) \Rightarrow \texttt{MF}(\texttt{t}, \texttt{m})$$
$$r_4 : \texttt{Tweet}(\texttt{t}) \wedge \texttt{Ideo}(\texttt{t}, \texttt{i}) \wedge \texttt{Topic}(\texttt{t}, \texttt{k}) \wedge \texttt{Ent}(\texttt{t}, \texttt{e}) \Rightarrow \texttt{Role}(\texttt{t}, \texttt{e}, \texttt{r})$$

Party's consistency preferences

$$c_3 : \texttt{SameIdeo}(\texttt{t}_1, \texttt{t}_2) \wedge \texttt{SameTopic}(\texttt{t}_1, \texttt{t}_2) \wedge \texttt{Ent}(\texttt{t}_1, \texttt{e}) \wedge \texttt{Ent}(\texttt{t}_2, \texttt{e}) \wedge \texttt{Role}(\texttt{t}_1, \texttt{e}, \texttt{r}_1) \wedge \texttt{Role}(\texttt{t}_2, \texttt{e}, \texttt{r}_2) \Rightarrow \texttt{SamePolarity}(\texttt{r}_1, \texttt{r}_2)$$

GROUP	MODEL	MACRO ROLE	MACRO MF	WEIGHTED ROLE	WEIGHTED MF
Lexicon Matching	MF Dictionary	-	30.37	-	37.32
	PMI Lexicon	-	36.44	-	35.94
	MFD + PMI	-	39.78	-	42.12
Seq-Tagging	BiLSTM-CRF	35.18	-	45.91	-
End-to-end Classifiers	BiLSTM	39.75	58.61	45.61	59.90
	BERT-base	49.32	59.99	57.37	62.17
	BERT-tapt	54.73	66.44	62.18	68.29
	+ Ideo + Issue	54.81	66.13	62.83	68.34
Multi-task	BERT-base	44.37	61.63	57.74	67.71
	BERT-tapt	52.08	63.46	61.96	69.20
	+ Ideo + Issue	52.11	63.44	63.61	68.61
Relational Learning	PSL	56.51	68.98	64.02	71.85
	DRaiL Local	58.07	71.20	64.38	73.85
	DRaiL Global	**59.23**	**72.34**	**64.98**	**74.39**
Skyline	DRaiL Global (Fixed MF)	79.35	-	84.52	-

Morality Frames as Explanations

- On the topic of Abortion Rights

If the text describes X as Y then it reflects a Right/Left perspective

	Most Frequent Entities	Most Associated Moral Roles
	Woman	Target of fairness/cheating
	Reproduction Right	Target of fairness/cheating
In Left	Planned Parenthood	Target of loyalty/betrayal
	Reproductive Care	Target of fairness/cheating
	SCOTUS	Entity ensuring fairness
	Life	Target of purity/degradation
In Right	Planned Parenthood	Entity doing cheating
	Democrats	Failing authority
	Born Alive	Target of purity/degradation
	Woman	Target of care/harm

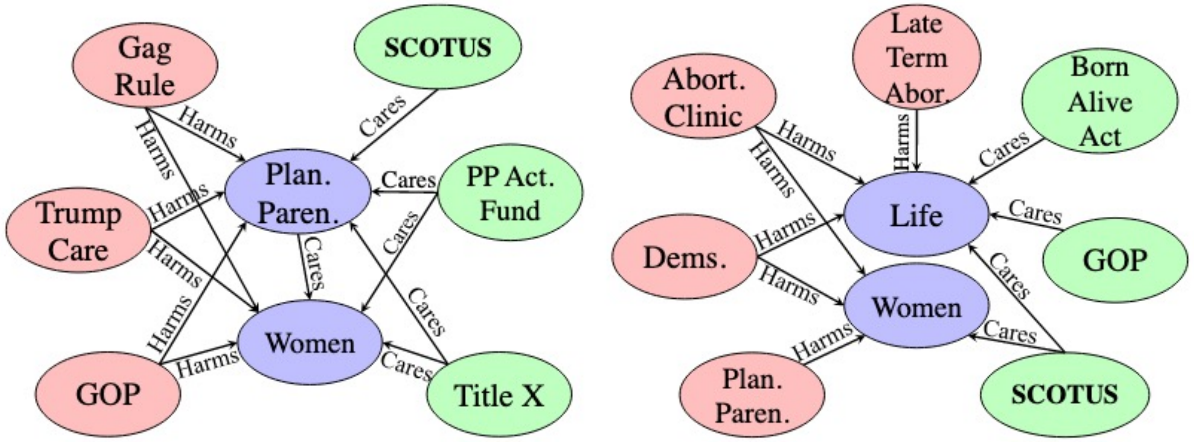

Aggregated results from all Congressional Tweets

Break!

Demo

Analyzing Political Discourse

- Analyze tweets written by US congress members on the abortion issue
- Use morality frames (moral foundations, roles)
- ~1k tweets by democrats and republicans
- Polarizing issue in the US political discourse
- Opinions range from "pro-life above all", to "women's choices above all"
- We will look at the entities at the center of this debate
  - Women, babies, life, US government institutions, legislative bills

Morality Frames to Analyze The Abortion Debate

- ***Entities***

Event -> *One Global Instance*

Tweet

Topic -> *Abortion*

Ideology -> *Left, Right*

Entity-Mention -> *"The ACA"*

Entity-Group -> *ACA*

Moral Foundation -> *Fairness, Care, etc.*

Role -> *target-care*, *provide-care, etc.*

Polarity -> *positive, negative*

- ***Relations***

InInstance(Tweet, Event)

HasEntity(Tweet, Entity)

HasTopic(Tweet, Topic)

HasIdeology(Tweet, Ideology)

HasRole(Entity, Role)

RoleHasMF(Role, MF)

RoleHasPolarity(Role, Polarity)

Using DRaiL to Analyze the Abortion Debate

Base Classifiers:

- InInstance(T, Z) & HasEntity(T, E) => HasRole(T, E, R^RoleLabel?)
- InInstance(T, Z) => HasMf(T, M^MfLabel?)

Party Messaging Preferences:

- InInstance(T, Z) & HasIdeology(T, I) => HasMf(T, M^MfLabel?)
- InInstance(T, Z) & HasEntity(T, E) & HasIdeology(T, I) => HasRole(T, E, R^RoleLabel?)

Joint Inference:

- InInstance(T, Z) & HasEntity(T, E) & RoleHasMf(R, M) & HasRole(T, E, R)^? => HasMf(T, M)^?

Live Demo