From Statistical to Deep Relational Learning

Maria L. Pacheco

Microsoft Research

University of Colorado Boulder

Dan Goldwasser

Purdue University

Oct. 2022

This Section in one Slide

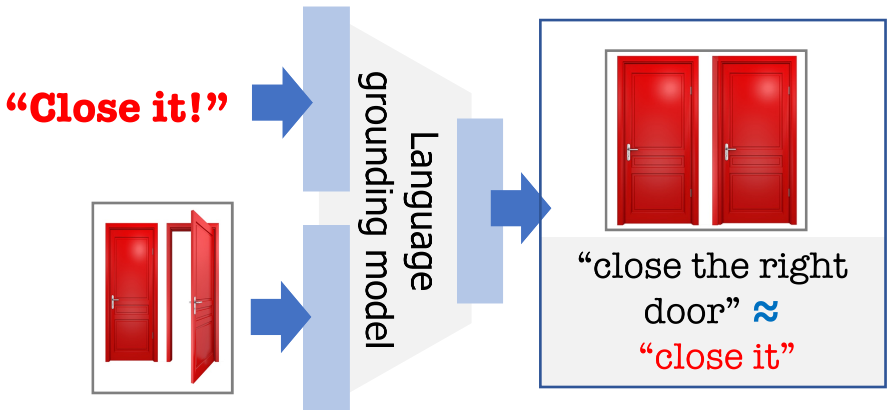

Real world context | **Linguistic and real-world inferences**

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

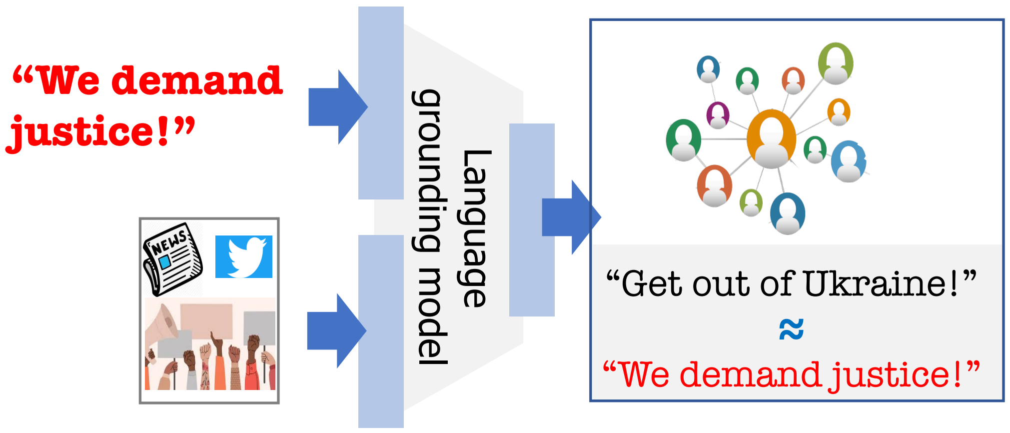

Real world context | **Linguistic and real-world inferences**

- How about now?
- Can we ask <u>*people*</u> to help us?

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

Neuro-Symbolic NLP and CSS

- A case for computational social science
  - **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!

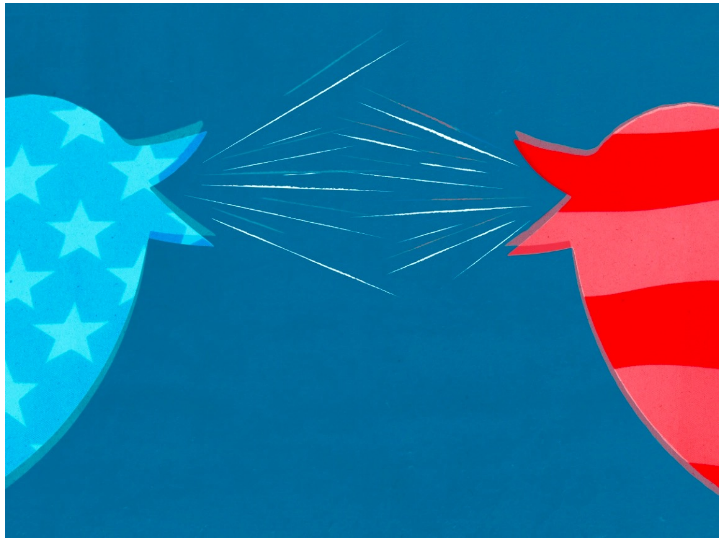

"if you talk about healthcare as a human right then…"

"..probably voted in favor of Obamacare"

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

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"

Characterizing Context through Inference

- Often easier to think about structure in a declarative way
- Define **entities** and **relations** and **probabilistic rules**
- **Replace classification with inference**: many decisions that should agree with each other, to support the decision

MakesClaim(user, claim) $\wedge$ *HasFrame(claim, safety)* $\rightarrow$ *IsPro(user, gun-control)*

"Banning guns will create a ***safer environment****"*

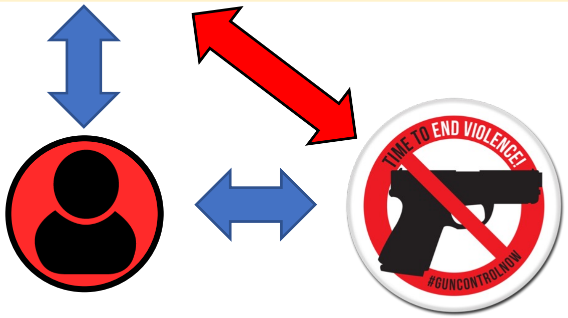

Characterizing Context through <u>Representation</u>

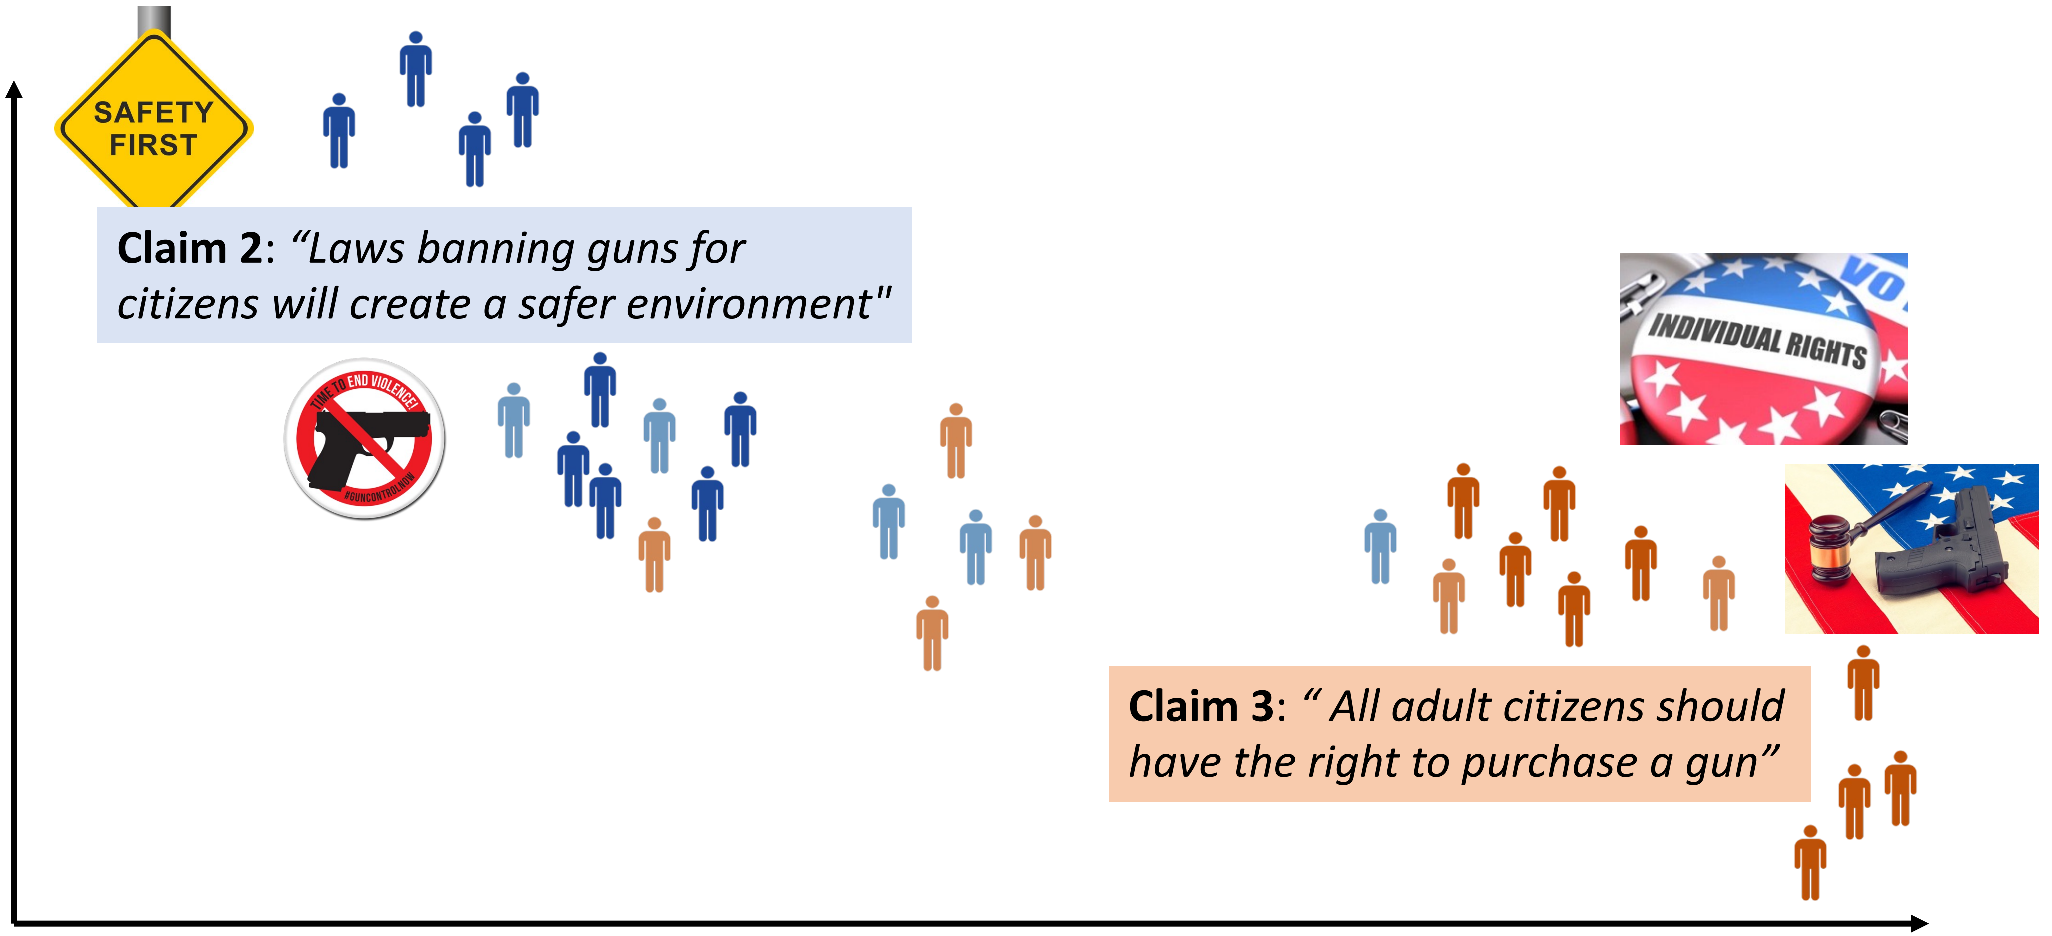

Socially Grounded Language Representation

“ban guns to avoid mass shootings in schools”

We should keep our teachers safe!

$\neq$

We should keep our teachers safe!

“arming schoolteachers stops active shooters”

Generalized View: can we create *<u>context dependent</u>* language representations that will support textual inferences and classification tasks?

ACL’19, EMNLP’21

Capturing Symbolic Dependencies

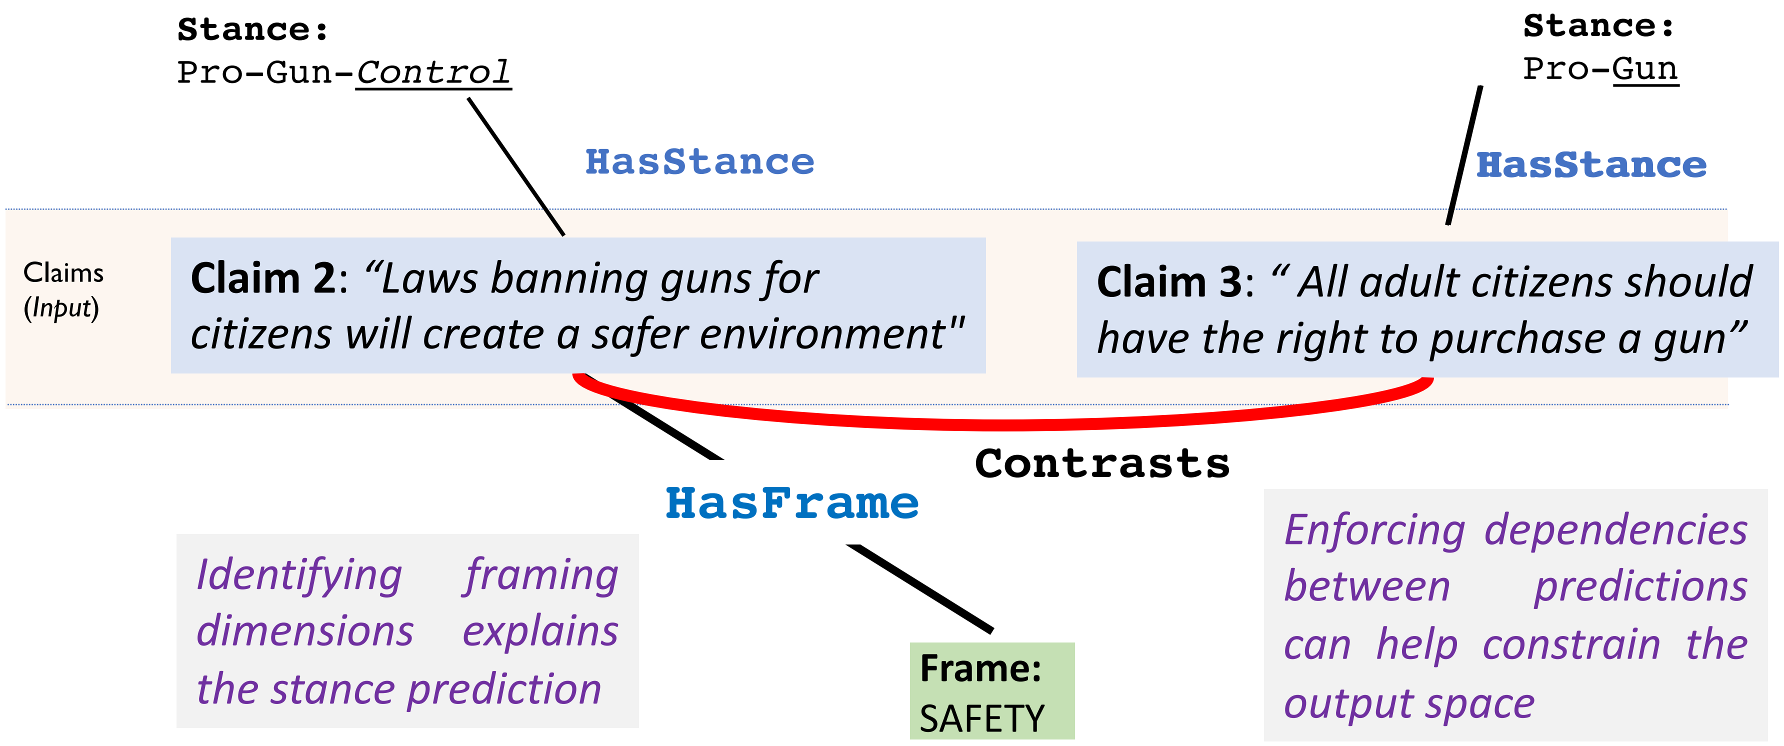

Generalized View: what are the relevant symbols and inferences needed for characterizing opinions? Explaining social group membership? Ideological differences?

Framing theory

- A **lens** through which a topic is perceived, organized and communicated
- Very often, used as a tool to bias the discussion on social issues towards a given stance by creating associations beneficial for holding it
  - Gun regulation as a question of **rights** or **safety**?
- **Challenge for NLP: *what are the relevant framing dimensions?***
  - Policy Frames (Boydstun et-al '14): general policy related framing dimensions (health, safety, crime, economy, etc.), **applicable across different issues**
  - **Domain specific frames**, which can emerge from data directly (Tsur et al., 2015, Demszky et al., 2019, Roy et al. 2020), or developed and coded by humans (Morstatter et al., 2018, Liu et al. 2019, Mendelsohn et al., 2021).

Political Issue Stance and Framing

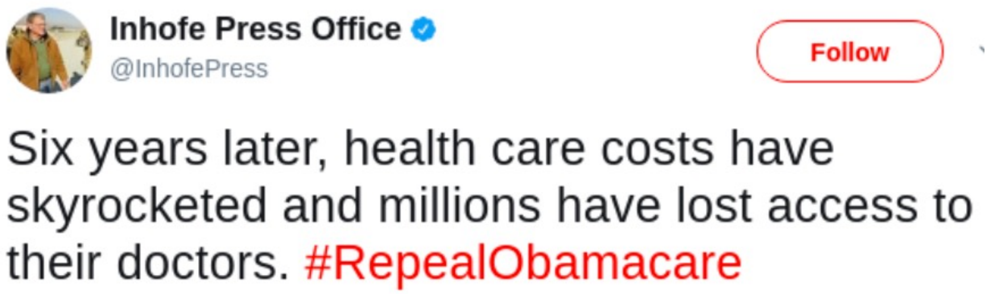

Stance:
Clearly, not a fan.

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*

Policy Frames vs. Domain Specific Frames

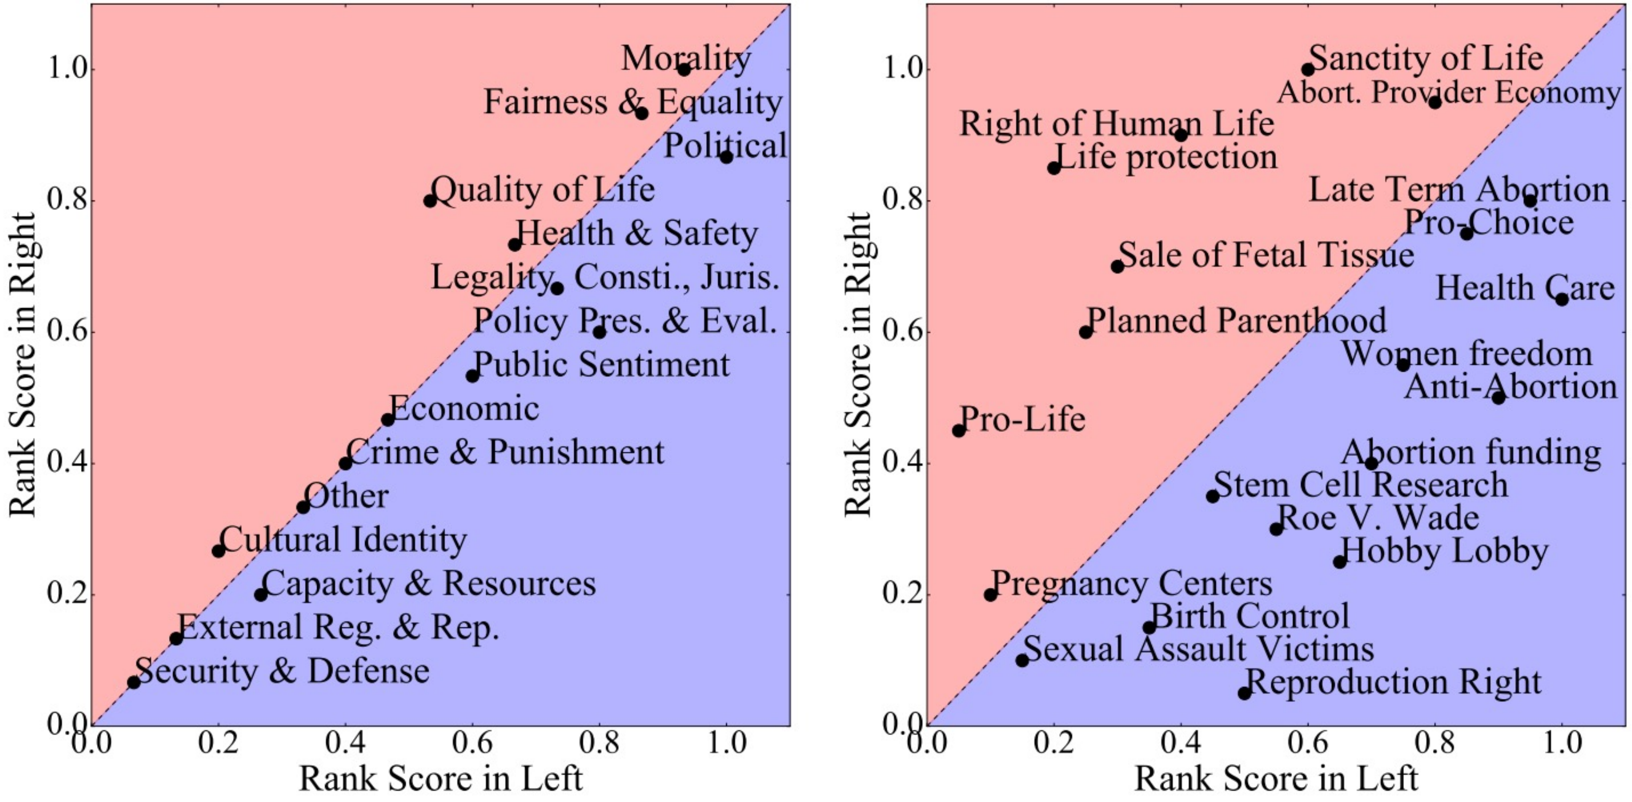

(a) Frame Usage

(b) Subframe Usage

Roy et al. 2020

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, 2004)

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

1. Care/ Harm: care for others, generosity, compassion, sensitivity to suffering of 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, an elevated life.

Rising popularity in the NLP community, used for analyzing news media (Fulgoni et al., 2016 Shahid Et al. 2020), social media (Johnson et al 2018, Hoover et al., 2020), explain moral values (Forbes et al., 2020, Hulpuș et al., 2020)

From Moral Foundations to 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!***

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

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

Morality Frames: *identifies the moral roles of different entities. Distinguishes between agents/targets, as well as positive/negative roles*

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

Analysis: 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
In Left	Woman	Target of fairness/cheating
	Reproduction Right	Target of fairness/cheating
	Planned Parenthood	Target of loyalty/betrayal
	Reproductive Care	Target of fairness/cheating
	SCOTUS	Entity ensuring fairness
In Right	Life	Target of purity/degradation
	Planned Parenthood	Entity doing cheating
	Democrats	Failing authority
	Born Alive	Target of purity/degradation
	Woman	Target of care/harm

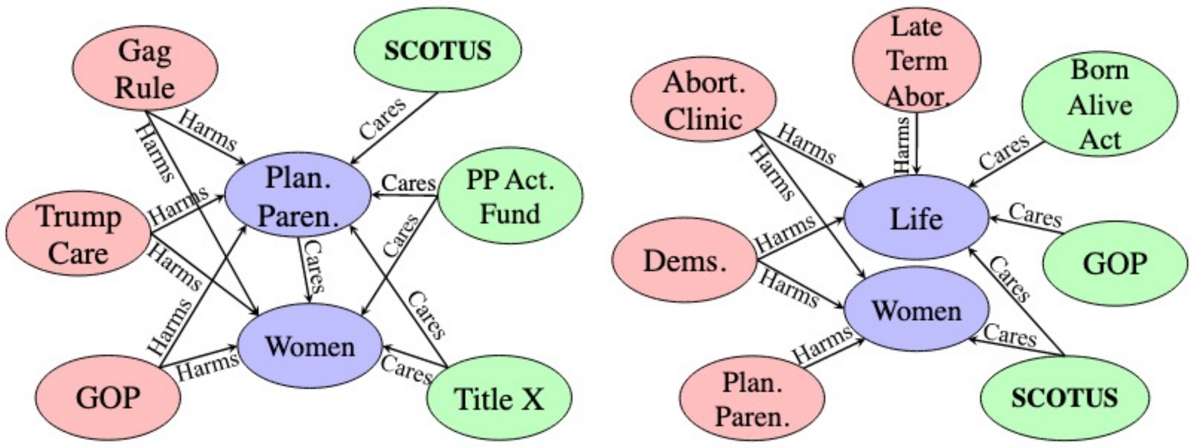

Aggregated results from all Congressional Tweets

Roy et al. 2021

How do we approach this challenge?

MakesClaim(user, claim) ∧
HasFrame(claim, safety) →
HasStance(user, pro-gun-control)

Inference

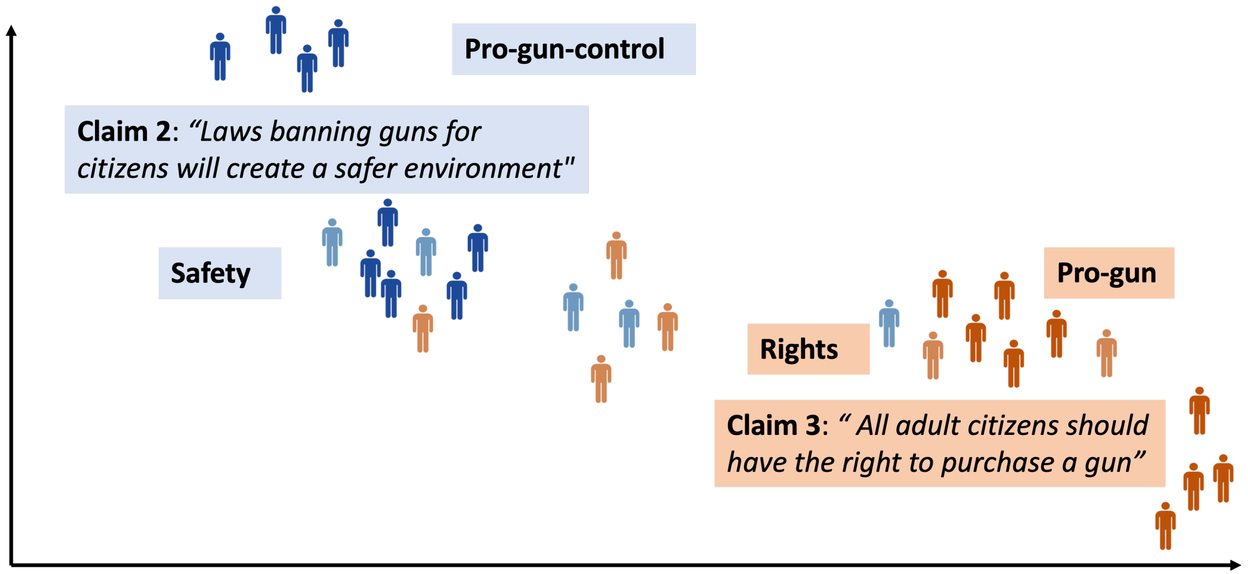

Representation

Lessons from the Past: Statistical Rel. Learning

- Key idea: **add weights to first-order formulae**
  - Expresses the strength of the formula
- Problem as a set of pairs <Formula_i, weight_i>
- Describes an **undirected graphical model**
- Ground it in data and use it for **inference**

Learn them from data!

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

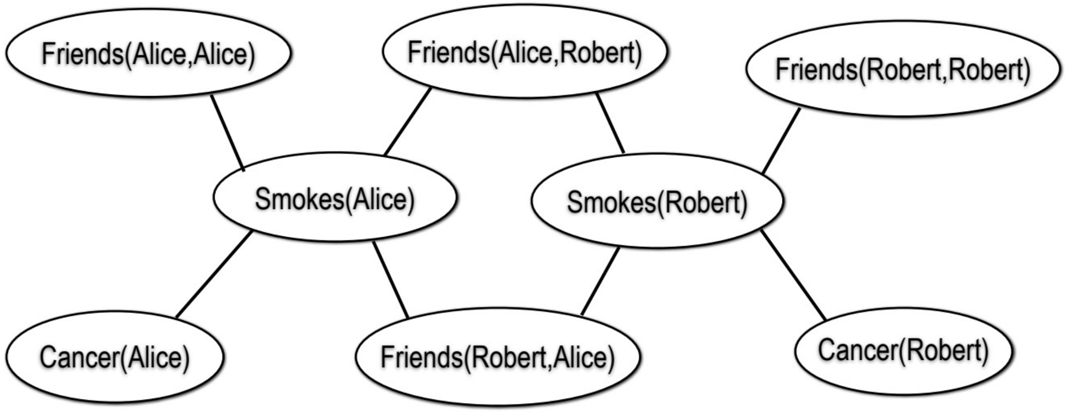

Richardson, M., Domingos, P. "Markov Logic Networks" **Machine Learning 2006**

Bach S. et al. "Hinge-Loss Markov Fields and Probabilistic Soft Logic", **JMLR 2017**

Limitations of SRL in Abstract Settings

- How to model CSS+NLP problems in a framework like PSL?

1. Enumerating relevant surface forms
(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", **ACL 2015**
Johson, K. et al. "Leveraging Behavioral and Social Information for Weakly Supervised Collective Classification of Political Discourse", **ACL 2017**

Declarative Deep Relational Learning

- **Rules as context**, using a graphical model

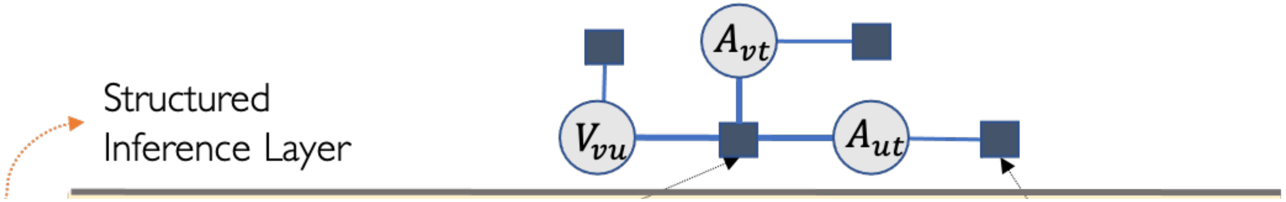


```
rule_def:
 Debates(u,t) -> Agree(u,t)
```

```
rule def:
 Agree(v,t) & VoteFor(v,u)
   -> Agree(u,t)
```

Declarative Deep Relational Learning

- **Rules as context**, using a graphical model *with **neural potentials***

DRaiL Program

```
rule_def:
 Debates(u,t) -> Agree(u,t)
 net def:  network1
```

```
rule def:
 Agree(v,t) & VoteFor(v,u)
   -> Agree(u,t)
 net_def:  network2
```

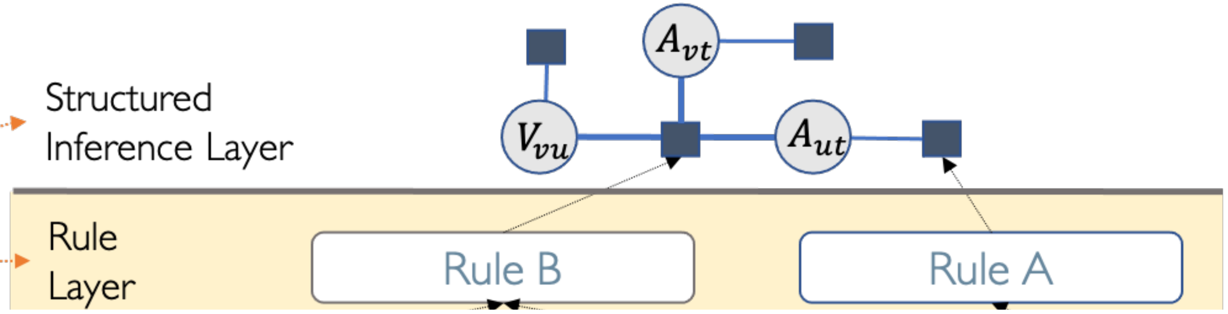

Pacheco and Goldwasser, "Modeling Content and Context with Deep Relational Learning", **TACL 2021**

Declarative Deep Relational Learning

- **Rules as context**, using a graphical model *with **neural potentials***
- **Representation as context**, using neural architectures

DRaiL Program

```
rule_def:
 Debates(u,t) -> Agree(u,t)
 net def:  network1
```

```
rule def:
 Agree(v,t) & VoteFor(v,u)
   -> Agree(u,t)
 net_def:  network2
```

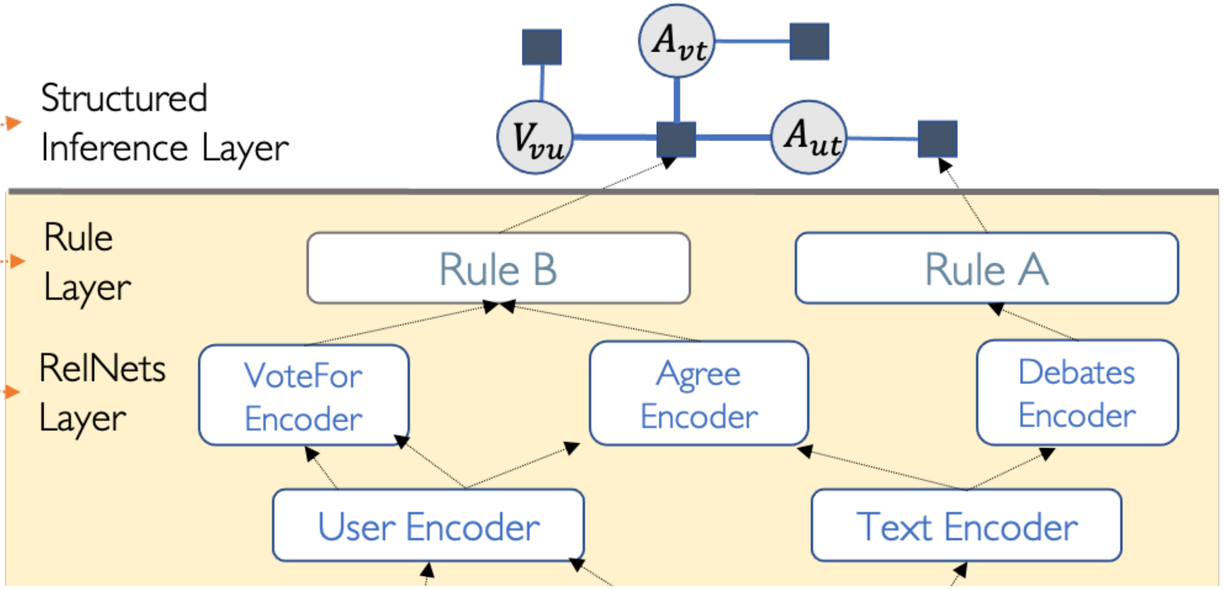

Pacheco and Goldwasser, "Modeling Content and Context with Deep Relational Learning", **TACL 2021**

Declarative Deep Relational Learning

- **Rules as context**, using a graphical model *with **neural potentials***
- **Representation as context**, using neural architectures

DRaiL Program

```
rule_def:
 Debates(u,t) -> Agree(u,t)
 net def:  network1
```

```
rule def:
 Agree(v,t) & VoteFor(v,u)
   -> Agree(u,t)
 net_def:  network2
 feat_def: [t:feat2;
             u:feat1;
             v:feat1]
```

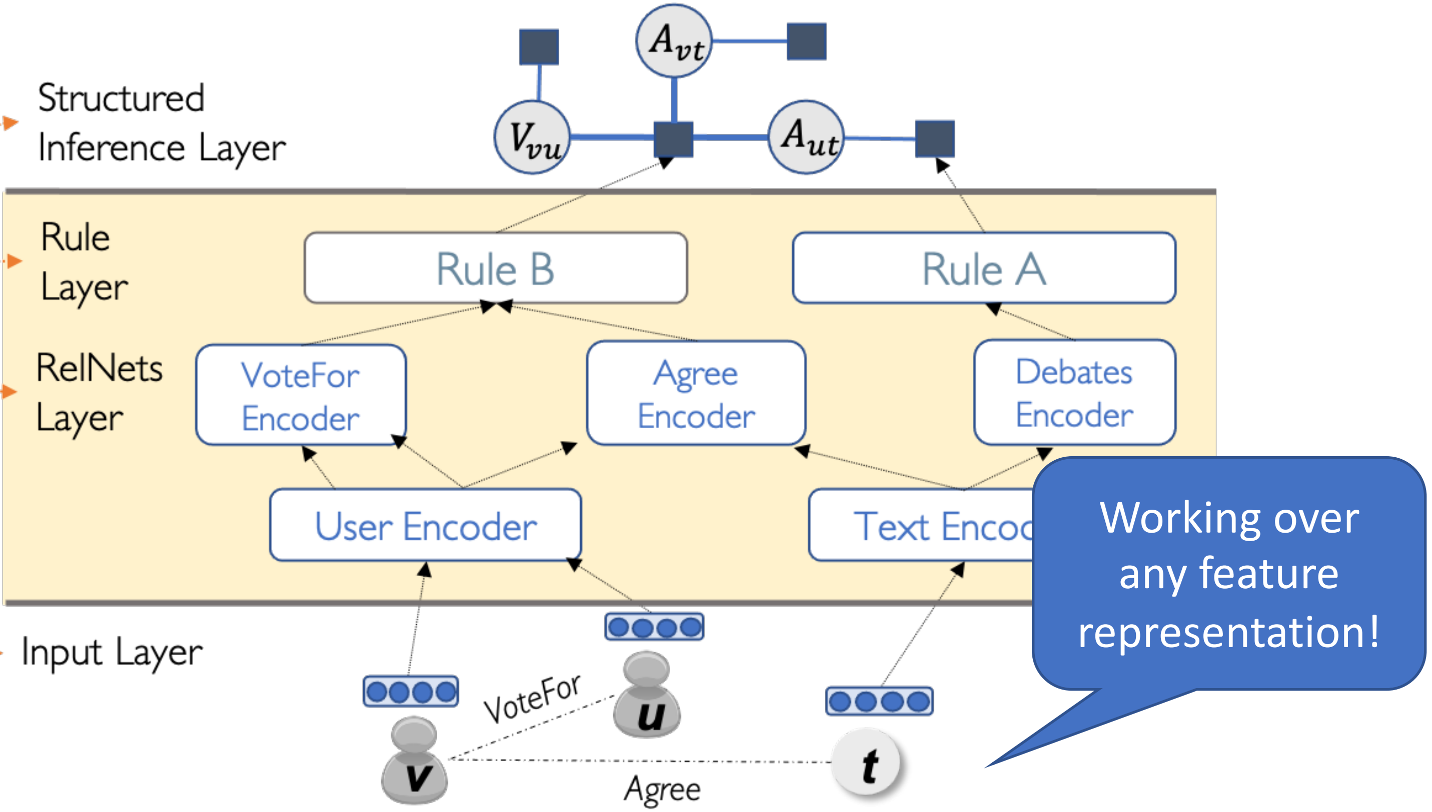

Pacheco and Goldwasser, "Modeling Content and Context with Deep Relational Learning", **TACL 2021**

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

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

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

Embedding Space is Updated!

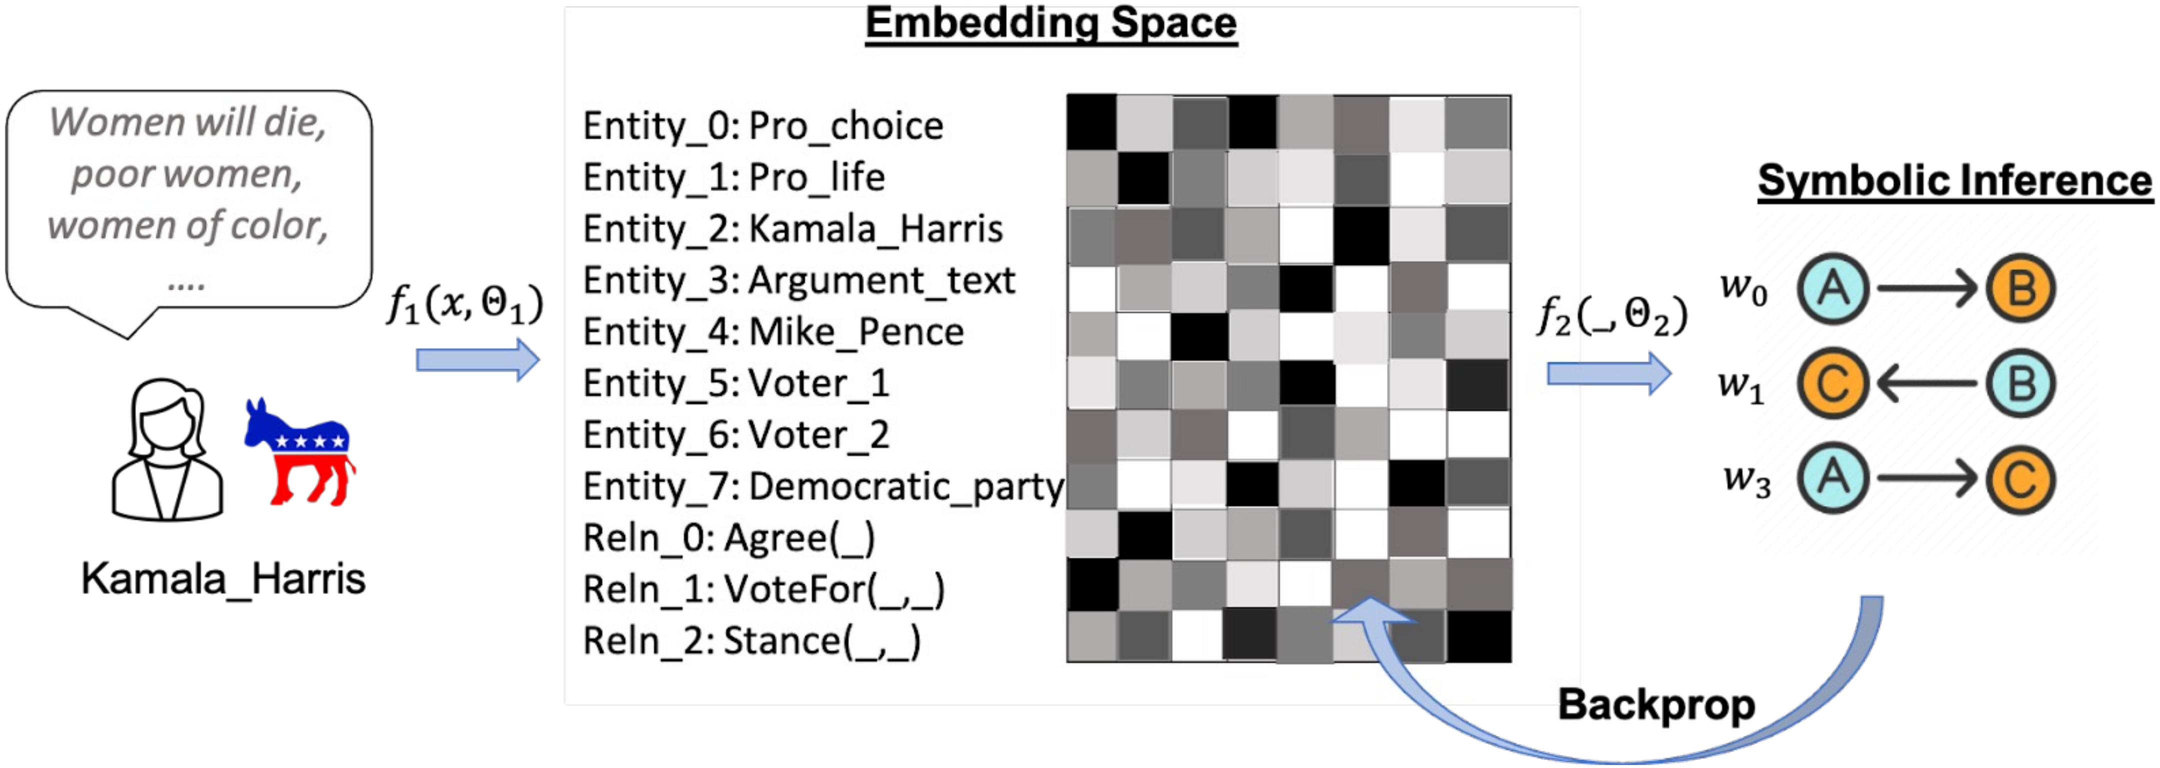

Pacheco and Goldwasser, "Modeling Content and Context with Deep Relational Learning", **TACL 2021**

Scenario: Understanding Debate Networks

Debated Claim :
"Limiting gun sales is unconstitutional"

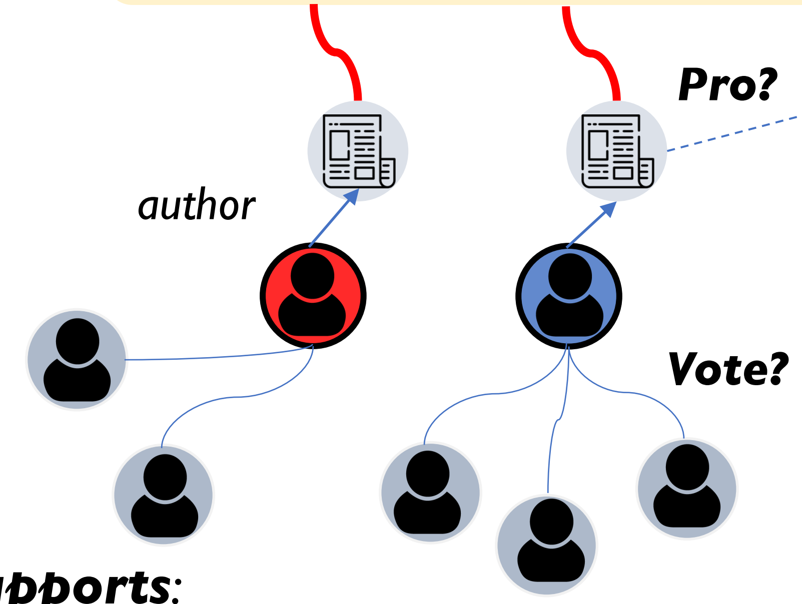

"this is an interpretation that is politically motivated. The right to bear arms does not mean that it cannot be regulated"

<u>Supports</u>:
Death Penalty
Border Fence

<u>Supports</u>:
Same-sex marriage
Gun regulations

This results in several reasoning tasks:

- **Textual Inference**
- **Authors and Text**
- **Authors and other Users**
- **<u>Clearly, there's structure to be exploited!</u>**

Pacheco and Goldwasser, "Modeling Content and Context with Deep Relational Learning", **TACL 2021**

Evaluating Modeling Choices

Exploiting Global Social <u>Embedding</u>

Exploiting Social Context at <u>Inference</u>

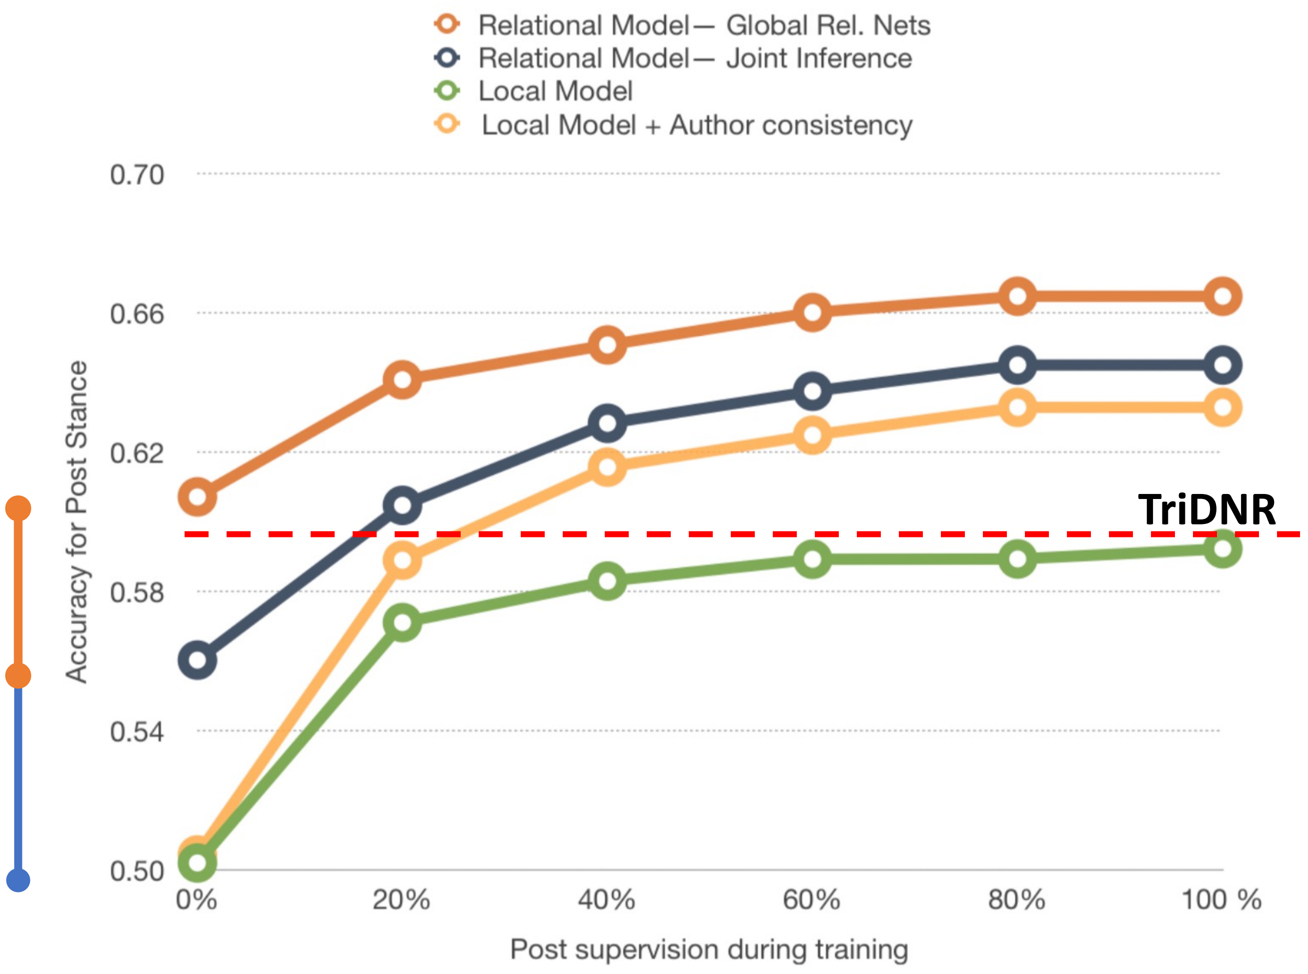

Pacheco and Goldwasser, "Modeling Content and Context with Deep Relational Learning", **TACL 2021**

Other Examples of this Paradigm

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

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

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

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

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

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

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

DeepProbLog: Inference and Learning

- **Inference**: As in ProbLog, forward pass on Neural predicates
- **Jointly learn parameters** for probabilistic facts and Nnets
- Loss based on single query output

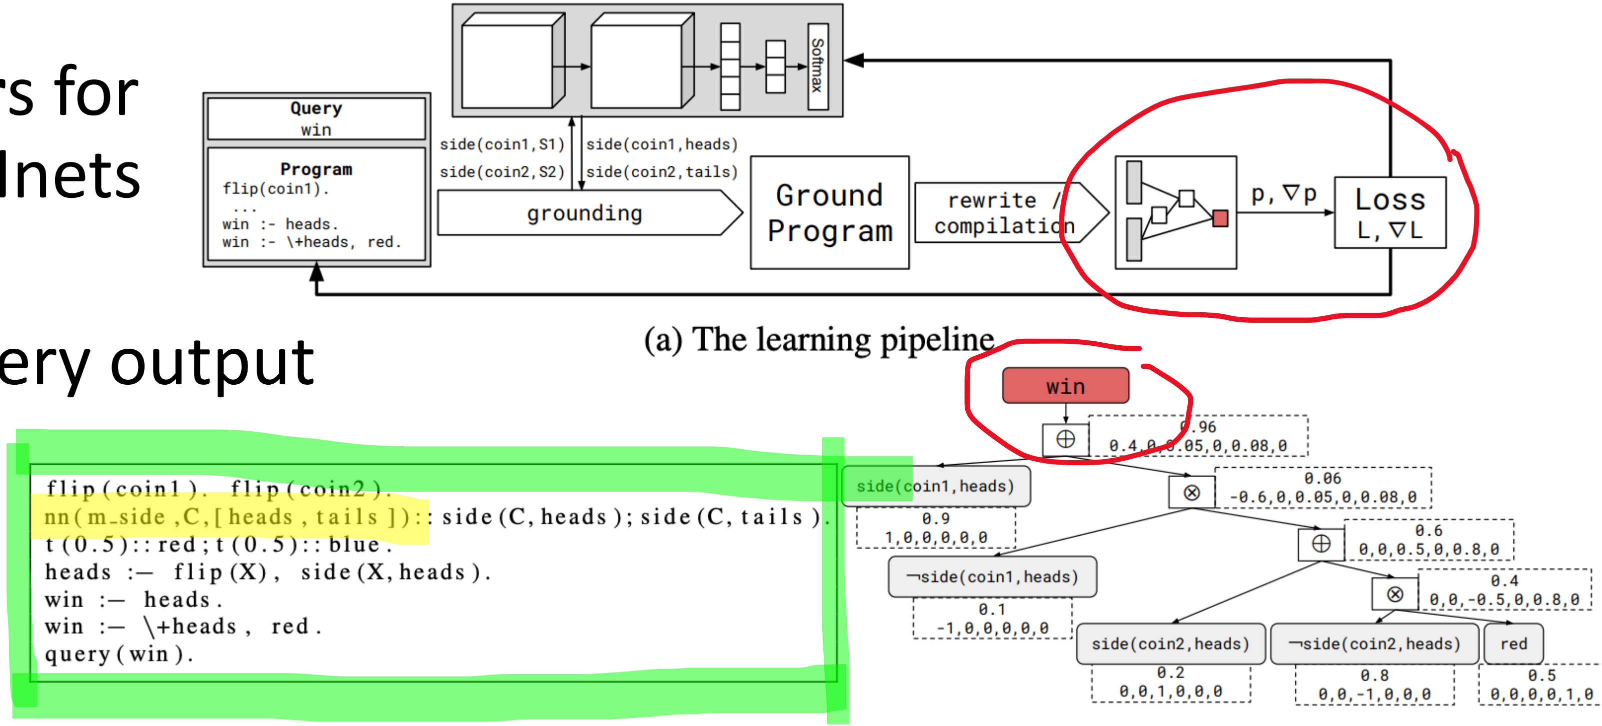

(a) The learning pipeline

(b) The DeepProbLog program.

(c) SDD for query `win`.

Figure 2: Parameter learning in DeepProbLog.

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

DPL: The Case for Indirect Supervision

- We want to learn model $P(y|x)$ using a NNet, but Y is unobserved
- Weak labeling functions $K = (\Phi_1, \ldots, \Phi_V)$
- Dependencies between weak labels and output
- Constraints on instances or model expectations

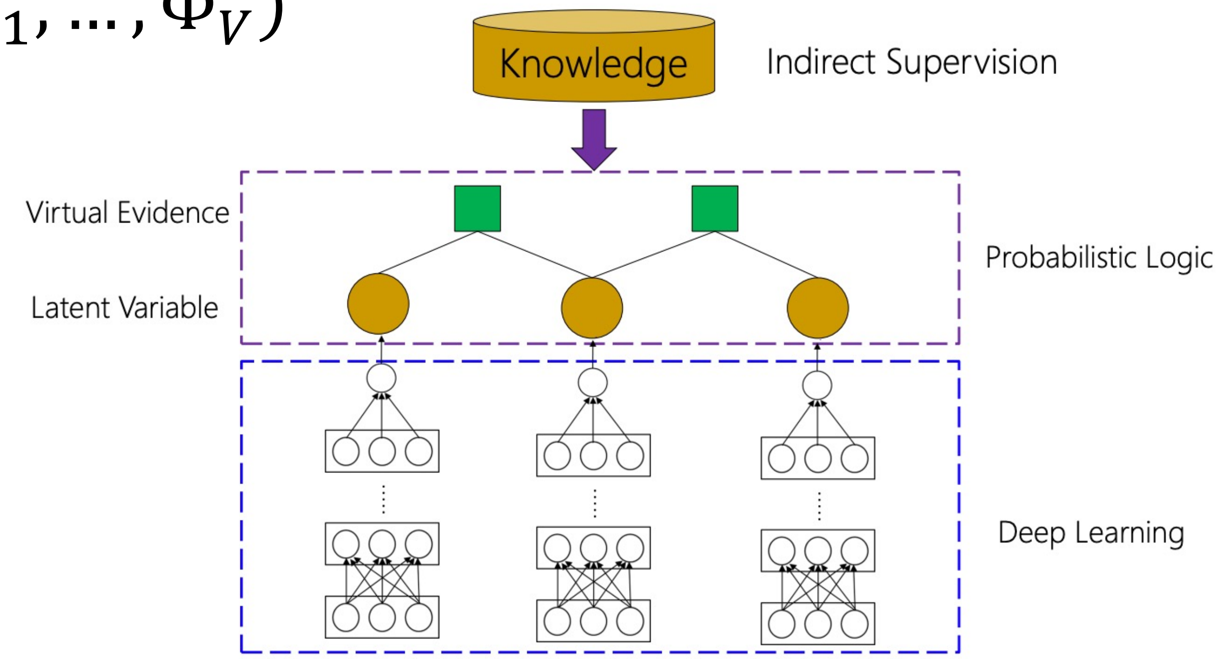

Wang H. and Poon H., "Deep Probabilistic Logic: A Unifying Framework for Indirect Supervision", **EMNLP 2018**

DPL: The Case for Indirect Supervision

0.5 Relation in Toy KB (distant supervision)
3.2 No more than one "et al." (data programming)
10 Relation holds for at least one instance (joint inference)

K

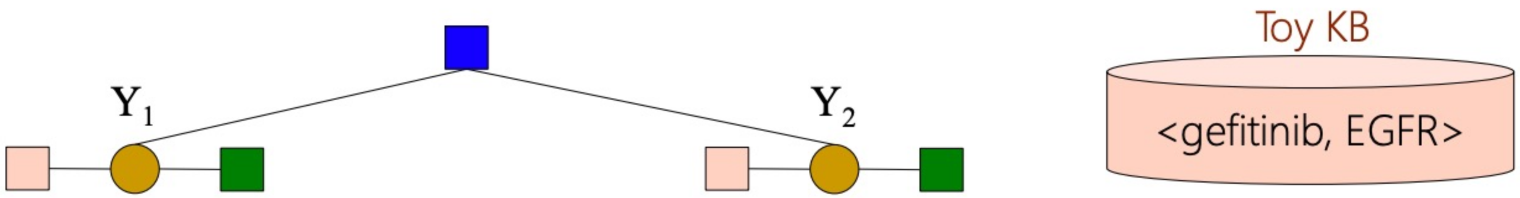

Patients with **EGFR** mutations show partial response to **gefitinib**.

X_1

Horn et al., 2001. Activities of **gefitinib** in NSCLC patients. J Clin Onco.
Zhang et al., 2006. Resistant mechanisms of **EGFR** mutations. J Thorac Onco.

X_2

Y_1	Y_2	$P(K,Y\|X) \propto$	$P(K, Y\|X)$
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 X_1 but not in X_2.

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

Human Interaction as Indirect Supervision

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

- 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

The Role of Symbols in Human Interaction

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

Learning from Explanations

"Since vaccine, the number of patients that have come in for miscarriages has quadrupled"

NL

AntiVax(x) ← harm(x, PREGNANT_WOMEN)

Human-readable Symbolic Representation

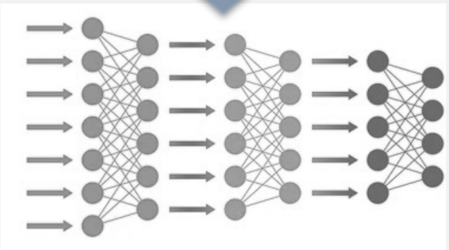

Neural Representation (Machine-representation)

Interaction over the symbolic representation

- Dependencies between concepts.

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?

Putting all Pieces Together: Opinion Analysis

Domain-specific!

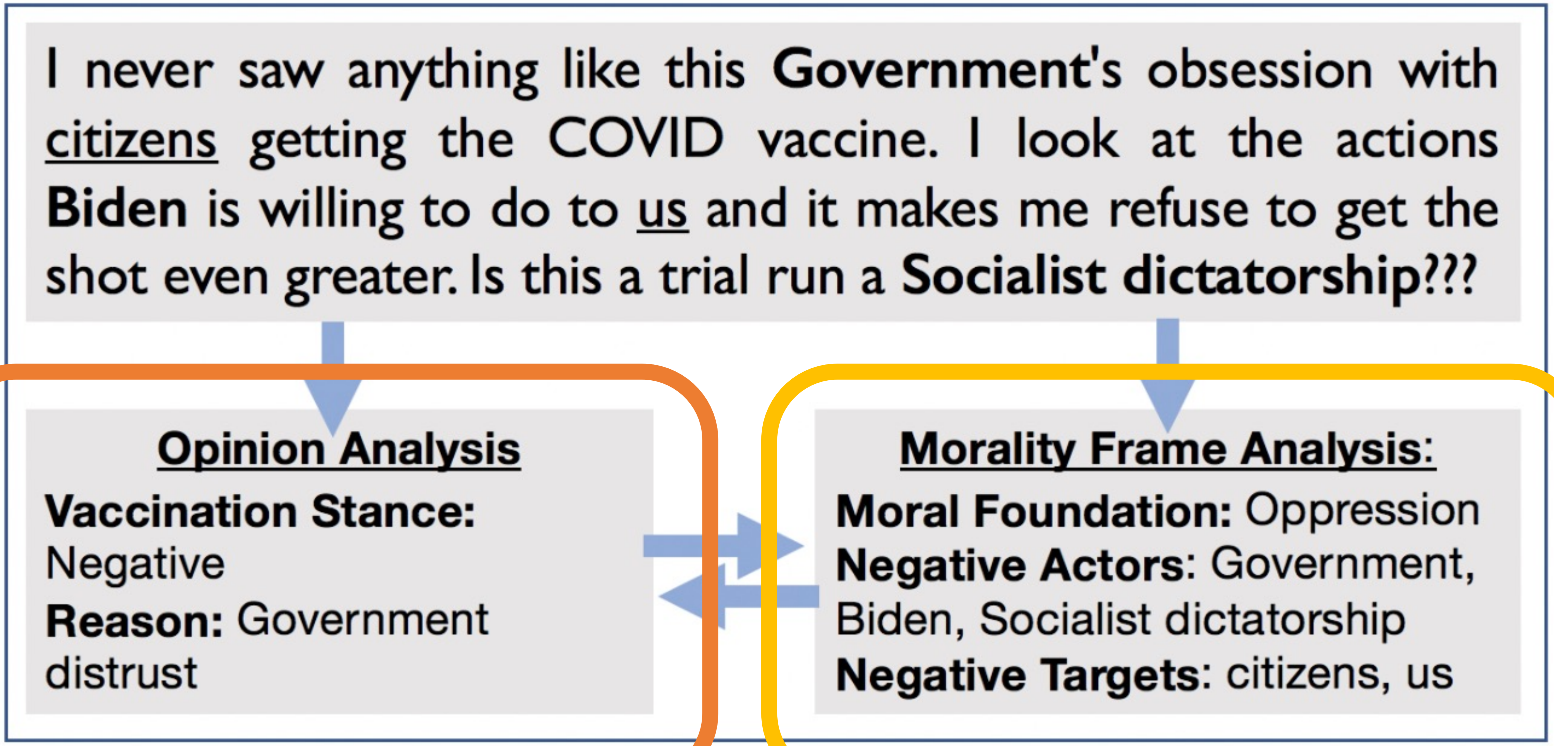

General, cross-domain

Pacheco et al., "A Holistic Framework for Analyzing the COVID-19 Vaccine Debate", **NAACL 2022**

Morality Frames Can Help us Explain Opinions

- **Reminder**: Morality frames capture differences in the **actors / targets** of moral sentiment
- They can help us ***explain opinions***

If the actor of AUTHORITY is "Fauci", the author is more likely to be pro-vaccine, and to express Trust in Science

If the actor of HARM is "Fauci", the author is more likely to be anti-vaccine, and to express Distrust in Government

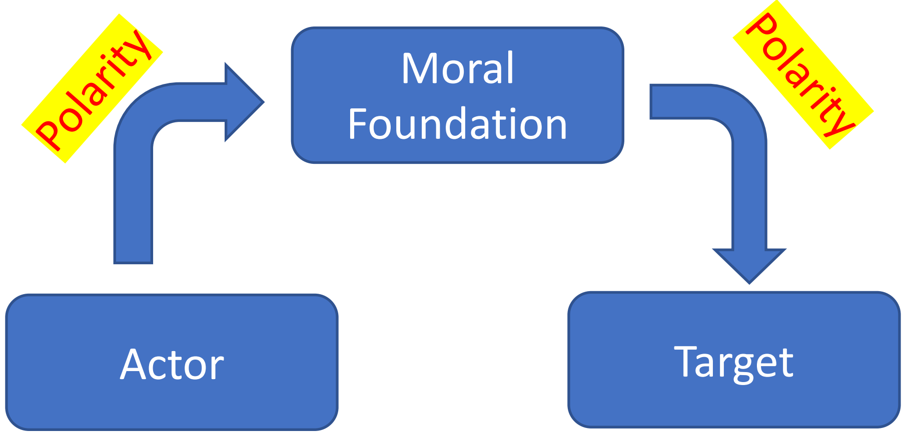

Roy, Pacheco and Goldwasser, "Identifying Morality Frames in Political Tweets Using Relational Learning", **EMNLP 2021**

Humans Can Help tie General Concepts to Domain-Specific Outputs

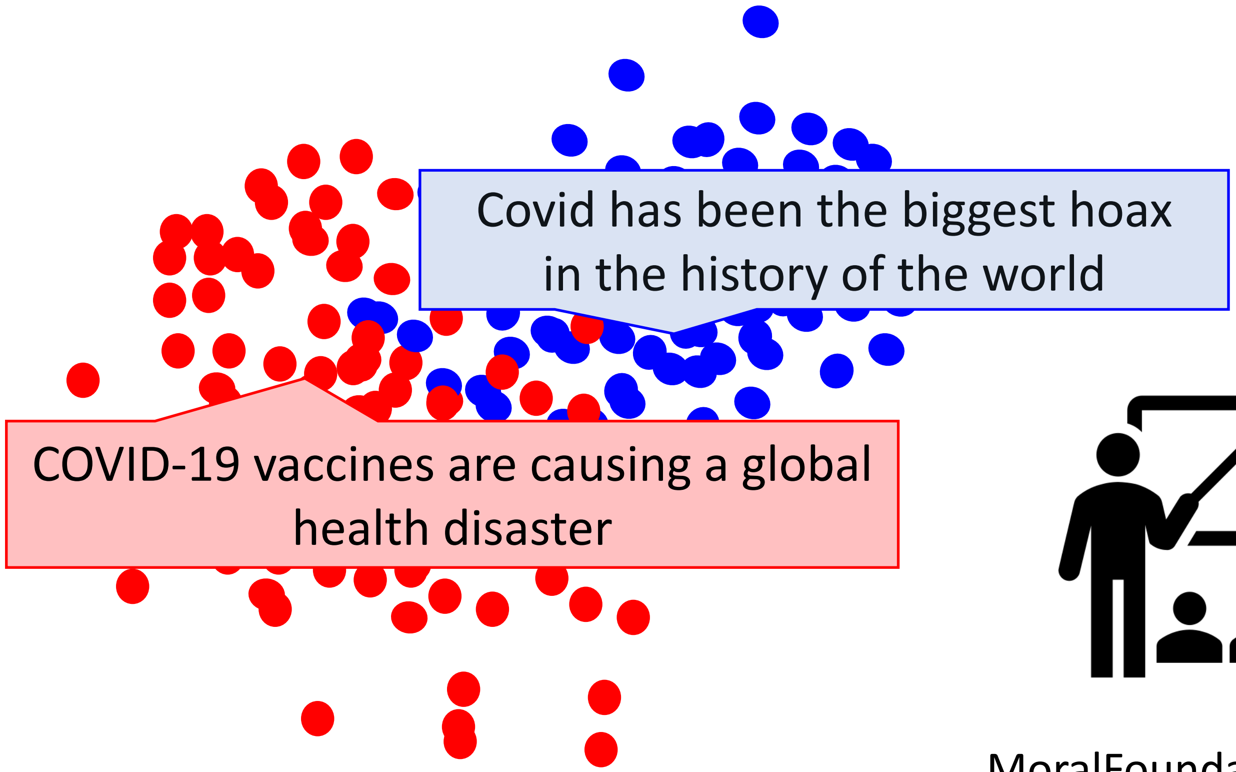

MoralFoundation(Cheat)
& Actor(Authority_Fig, Cheating)
& Target(Public, Cheating)
=> AntiVax & Covid_Fake

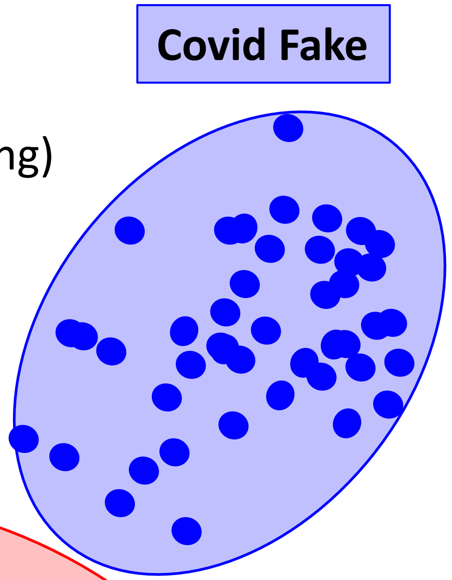

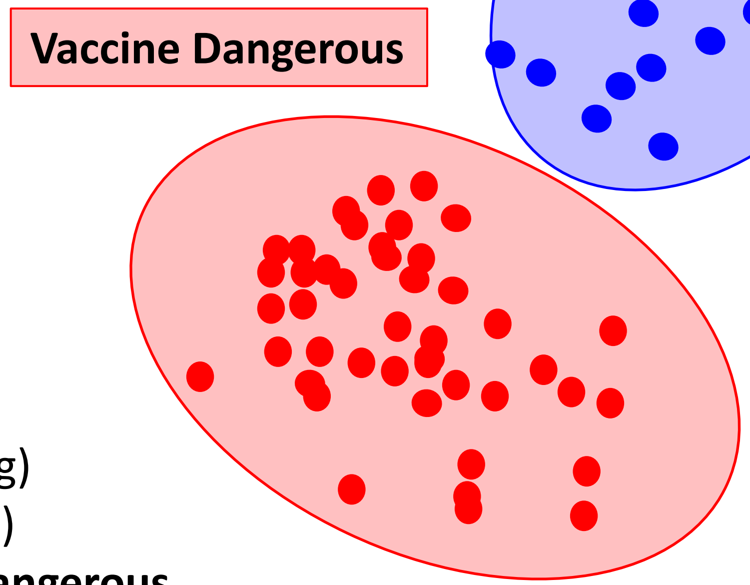

MoralFoundation(Harm)
& Actor(Vaccine, Harming)
& Target(Public, Harming)
=> AntiVax & Vaccine_Dangerous

Pacheco et al., "A Holistic Framework for Analyzing the COVID-19 Vaccine Debate", **NAACL 2022**

A Small Experiment

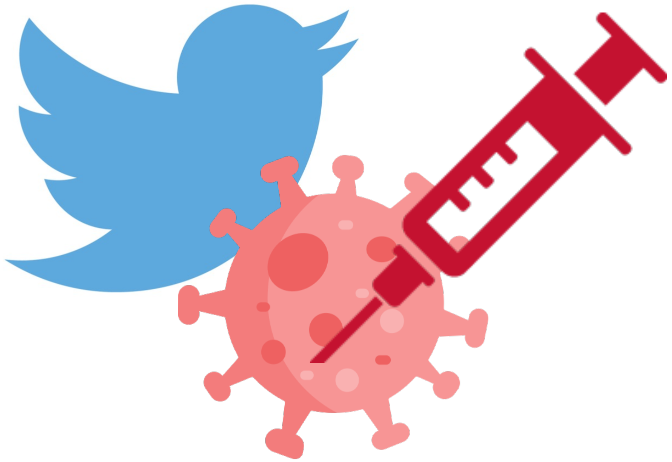

- 3 NLP/CSS Researchers in **two** 1-hour sessions

- 85,000 unlabeled tweets about the covid vaccine

- **Initial set of themes:** main reasons people cite to refuse the vaccine

  e.g: **“The vaccine is dangerous”** (Wawrutza et. al, 2021)

- **Interactive session**: identifying **high-level argumentative patterns** and contributing 2-5 examples

Pacheco et al., “A Holistic Framework for Analyzing the COVID-19 Vaccine Debate”, **NAACL 2022**

Joint Model for Opinions and Morality Frames

Basic Classifiers

- Map tweets to stance/reason/MF
- Map entities to role and polarity

Dependencies between different dimensions

- If Fauci harming, likely anti-vax

Stance consistency preferences

- If two tweets talk about Fauci, and they are both anti-vax, likely same polarity

DRaiL: Deep Relational Learning
[Pacheco and Goldwasser, 2020]

W: Harm(Fauci) -> AntiVax(Tweet)

Pacheco et al., "A Holistic Framework for Analyzing the COVID-19 Vaccine Debate", **NAACL 2022**

Joint Model for Opinions and Morality Frames

Basic Classifiers

- Map tweets to stance/reason/MF
- Map entities to role and polarity

Dependencies between different dimensions

- If Fauci harming, likely anti-vax

Stance consistency preferences

- If two tweets talk about Fauci, and they are both anti-vax, likely same polarity

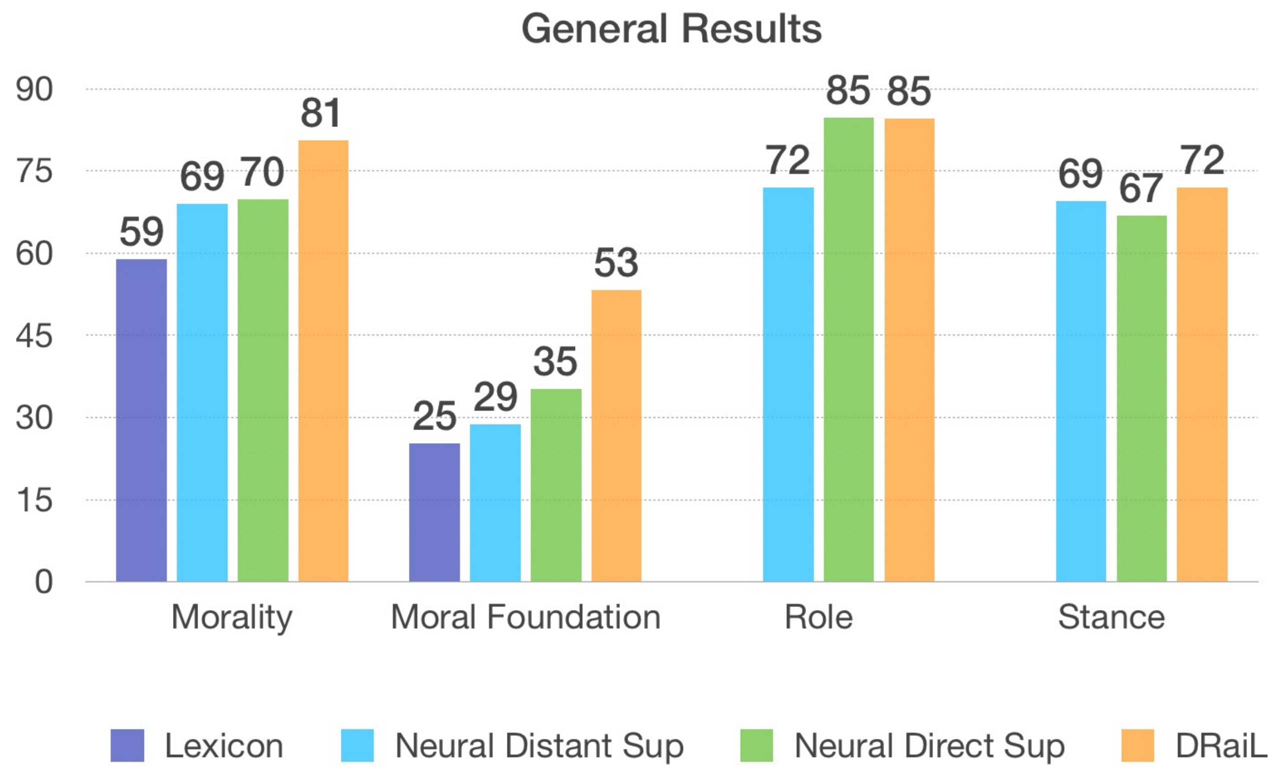

Pacheco et al., "A Holistic Framework for Analyzing the COVID-19 Vaccine Debate", **NAACL 2022**

The Impact of Inference

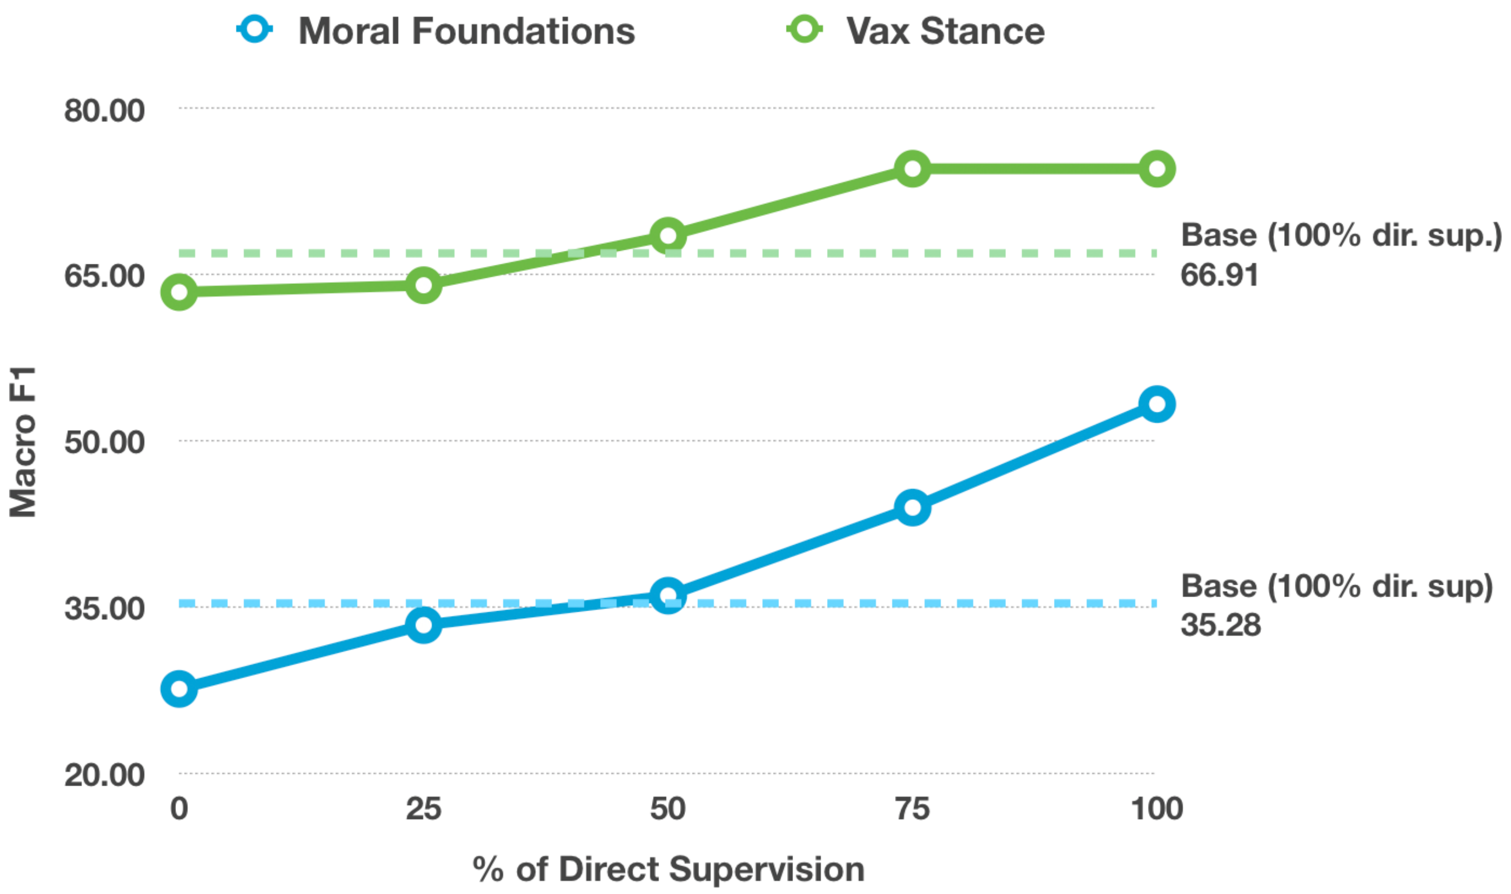

- **Joint inference** makes our model competitive with just 25% of direct supervision
- We beat the fully supervised base model with 50% of direct supervision

The Impact of Interaction

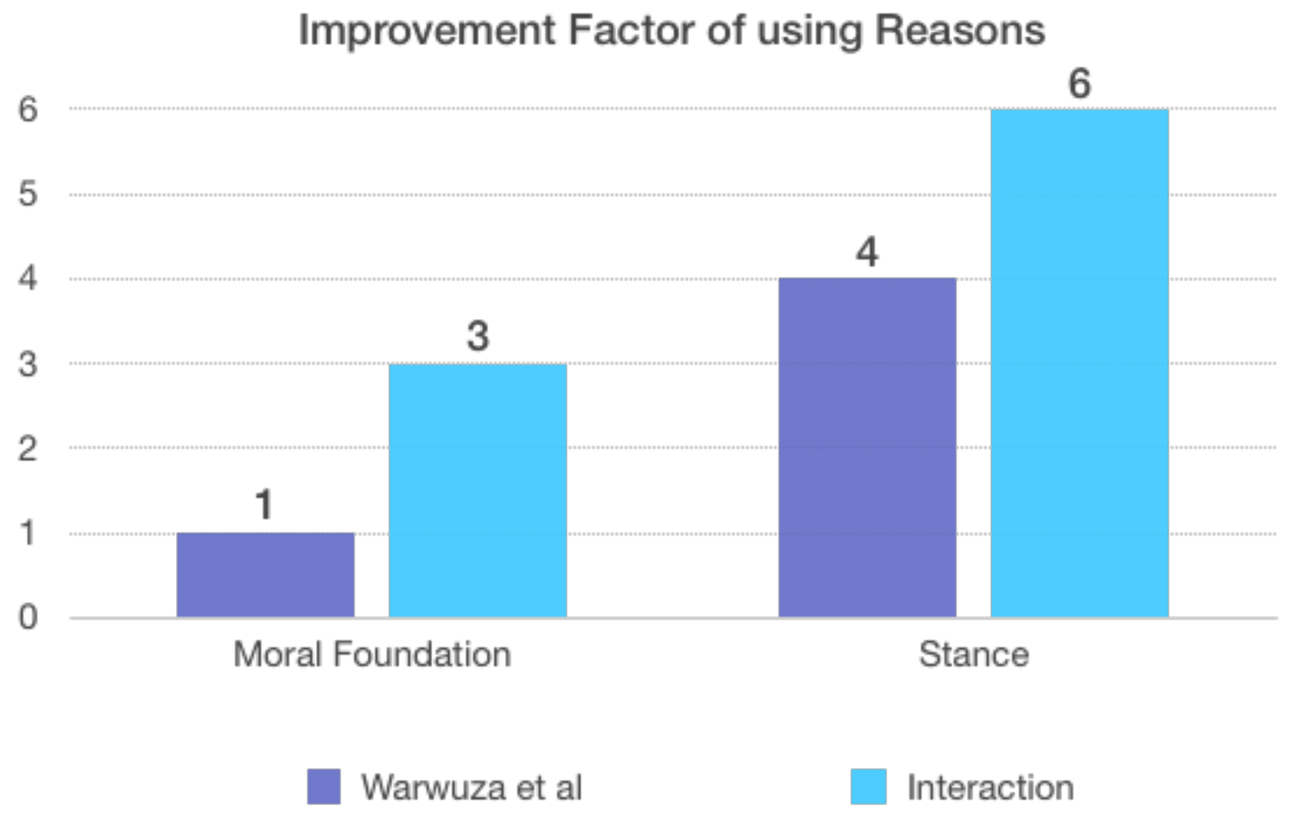

Pacheco et al., "A Holistic Framework for Analyzing the COVID-19 Vaccine Debate", **NAACL 2022**

Summary: Deep Relational Learning

- **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, **amenable to human intervention**