“We Demand Justice!”: Towards Social Context Grounding of Political Texts

Rajkumar Pujari, Chengfei Wu, and Dan Goldwasser

PURDUE UNIVERSITY | Department of Computer Science

SNL: Republican or not?

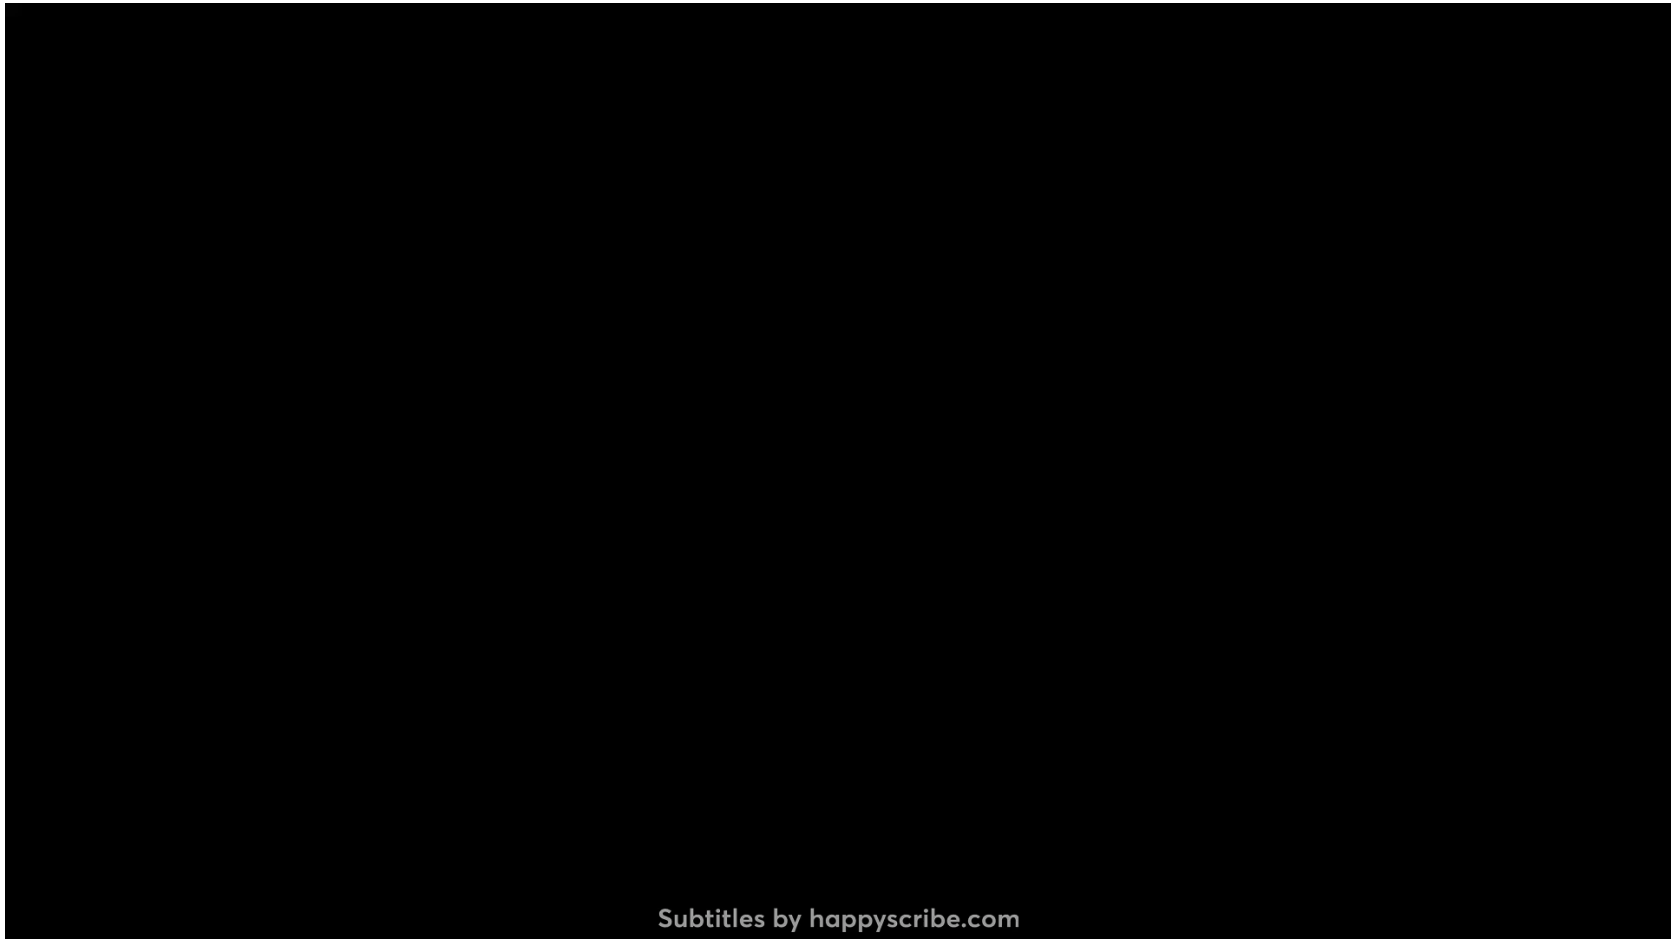

Ambiguity

Can you pass me the salt?

A: Yes, I can!

- Pragmatic vs. Literal meaning

Social Ambiguity

- Commonsense, factual knowledge, physical cues, cultural understanding, etc. help us disambiguate "conveyed meaning"

- **Social Ambiguity:** Pragmatic ambiguity that requires social understanding to be resolved

E.g.: 'How are you doing?' in the US vs. other places

rhetorical questions or a sincere inquiry about well-being?

Shared Understanding of Reality

- Reflects an understanding of the events and political debate surrounding George Floyd's death caused by a police officer.

Social Commonsense

- It could be *dynamically drawn* or *inferred* from various sources *in response* to a presented situation

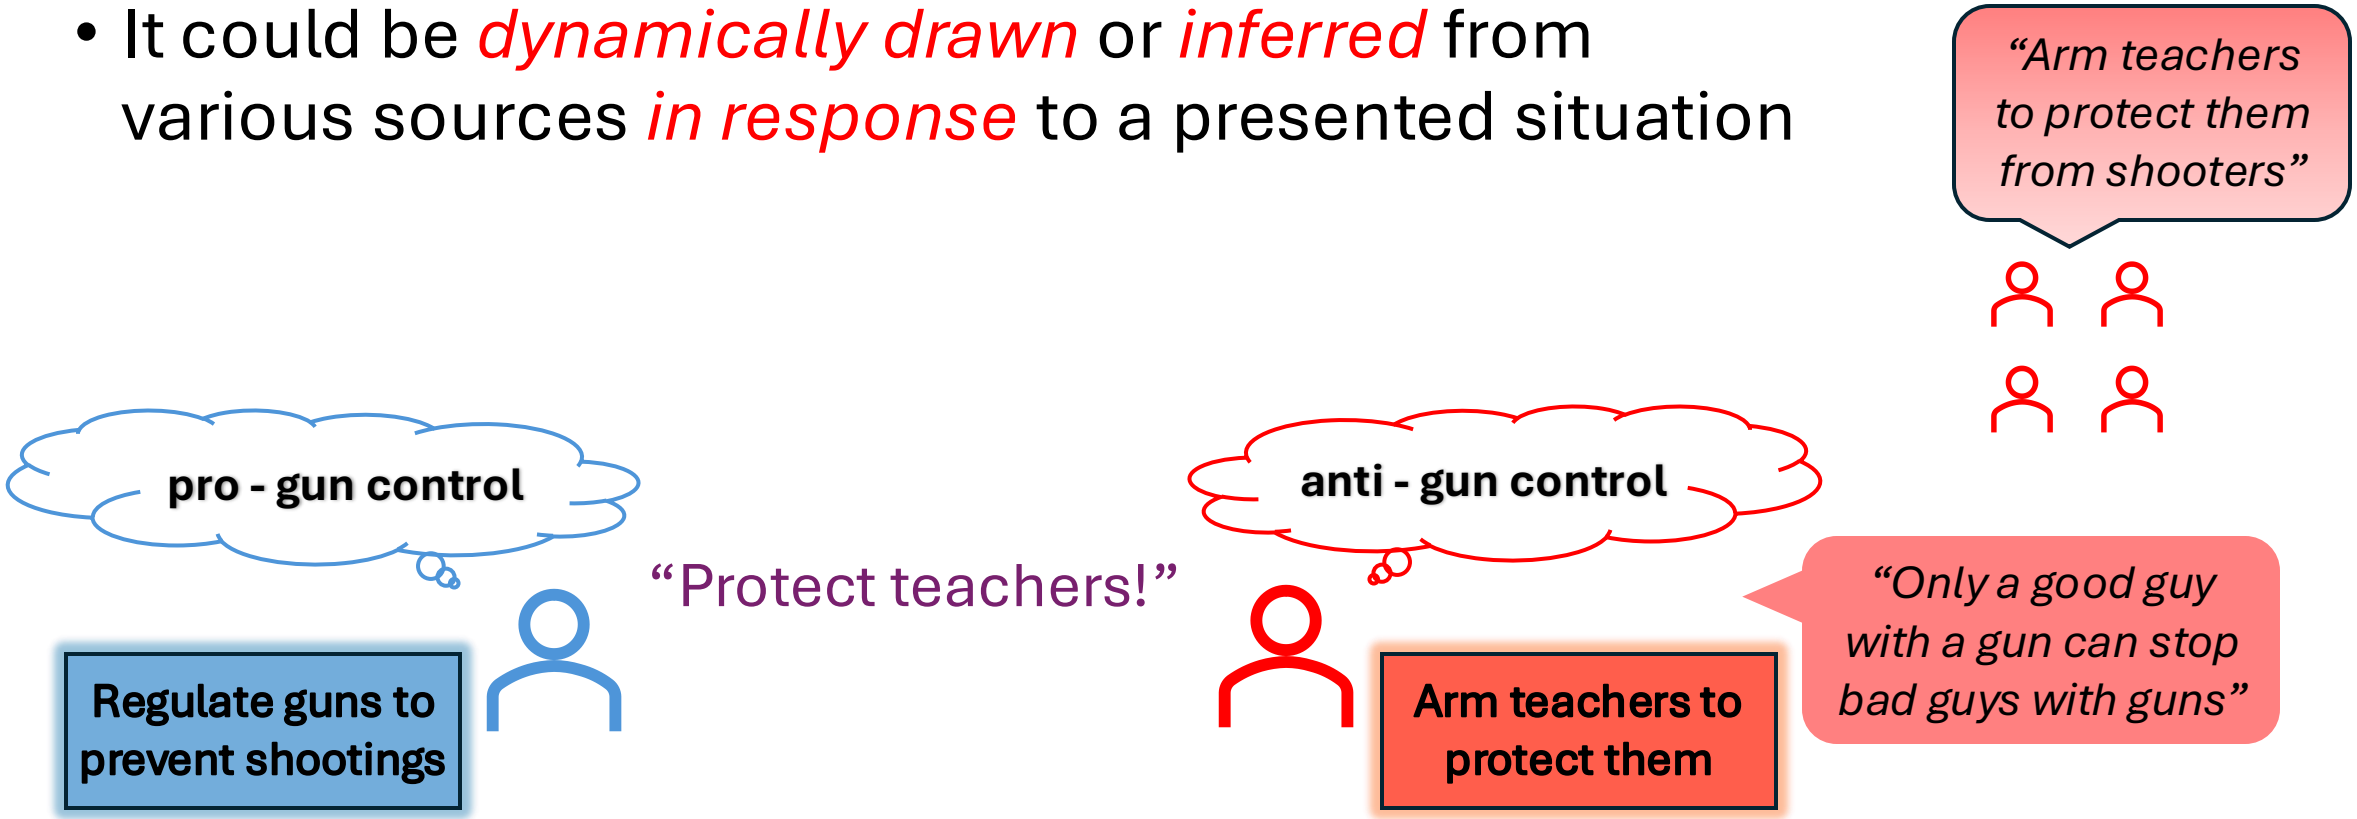

Research Questions

- How do we create *tasks* that require *holistic social context understanding*?
- How do *humans* perform on these tasks?
- How do *existing models* perform? Can we *use* them better? How can we *improve* them?

Grounding in NLP

- **Grounding:** Connecting language to real-world objects/concepts.

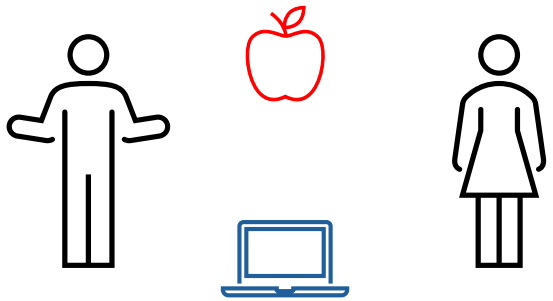

Jon bought Jane an ***apple***

How do we evaluate if an NLP model "***understands***" which type of ***apple*** Jane received?

Option 1: Classification task
- food / electronics

Option 2: QA task
- Did Jane find it filling?
- Can Jane finish her presentation?

Evaluating Social Grounding

Grounding language in the social settings

- *entities, attitudes, events, preferences*

1) Does the model identify ***who*** the speaker is talking about?
2) Can an NLP model identify the ***attitude*** of the speaker?
3) Is the model able to ***discriminate*** between ***plausible*** and ***improbable*** explanations of the situation?

Approach

- Operationalize two social grounding tasks on US political data:
  1. Target Entity & Sentiment Detection
  2. Vague Text Disambiguation

- Evaluate state-of-the-art NLP models on these tasks

- Analyze ***human performance***, ***model performance***, and ***challenges*** posed by the tasks

Target Entity-Sentiment Task

Task: Given an ***opinionated tweet*** from a politician, identify ***intended target entities*** and ***sentiment*** towards them

Brett Kavanaugh Supreme Court Nomination

Targets: Brett Kavanaugh (negative), Julie Swetnick (positive), Christine Ford (positive), Deborah Ramirez (positive)

Vague Text Disambiguation Task

Task: Given ***party affiliation*** and a ***vague statement*** in context of an ***event***, identify a ***plausible interpretation*** of the text

Vague statement: First, but not the last.

Author affiliation: Republican

Event: US withdraws from Paris climate agreement

The **withdrawal** from the **Paris climate agreement** is the first of many **positive** actions for **American economy** to come for the Trump administration

Trump's **inauguration** marks the **first day** of a **new era** of progress and prosperity **(Incorrect Event)**

It's time for America to **move forward & make progress** without being held back by a **global agreement** that doesn't serve our interests **(Doesn't match vague text)**

The **Paris Climate Agreement withdrawal** is the **first of many** backward steps Trump administration is sure to take in **destroying our environment (Improbable Stance)**

How do Humans Perform?

- Three annotators answered 97 questions from Vague Text Dataset
- Easy task for humans, challenging for NLP models

Model	Accuracy
Best Model	64.79
Humans	**94.85**

Benchmark Models

We evaluate ***four*** genres of models:

1. No-context baselines
2. Textual Context baselines (Wikipedia descriptions)
3. LLM (GPT-3 in-context learning)
4. Discourse contextualization models

Benchmarking Summary

Model	Target-Entity (Macro-F1)	Target-Sentiment (Macro-F1)	Vague Text (Macro-F1)
No-Context	68.83	61.36	54.53
Text-Context	69.34	60.13	66.87
GPT-3 (no-context, few-shot)	69.77	55.00	62.58
Contextualized Models	**73.56**	**65.34**	**71.71**

Summary

- We conceptualize and operationalize two holistic social context grounding tasks in English on the US political domain

- We evaluate existing state-of-the-art models and humans on these tasks and present interesting observations

- **What are other challenges in social grounding?**

Future Work

- Building explicit and interpretable models for Social Context Grounding
- Expanding datasets for these tasks and combining diverse set of tasks which also require social understanding
- Other flavors of social context - cross-cultural understanding, emergency response, etc.

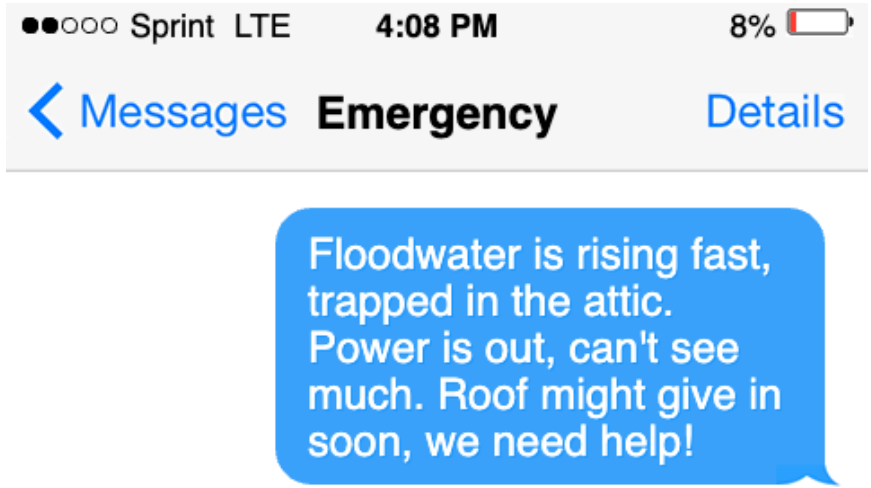

Thank you!

Questions?