


A Holistic Analysis Framework for Analyzing the Covid-19 Vaccine Debate

Maria L. Pacheco*, **Tunazzina Islam***, Monal Mahajan,
Andrey Shor, Ming Yin, Lyle Ungar and Dan Goldwasser



The Covid-19 *Infodemic*

- Low quality and polarizing information




The Covid-19 *Infodemic*

- Low quality and polarizing information
- Real-world implications
 - Acceptance of treatments, vaccines and other prevention measures



The Covid-19 *Infodemic*

- Low quality and polarizing information
- Real-world implications
 - Acceptance of treatments, vaccines and other prevention measures
- Combating the *infodemic*:



The Covid-19 *Infodemic*

- Low quality and polarizing information
- Real-world implications
 - Acceptance of treatments, vaccines and other prevention measures
- Combating the *infodemic*:



Most NLP Work:
Detecting Misinformation

The Covid-19 *Infodemic*

- Low quality and polarizing information
- Real-world implications
 - Acceptance of treatments, vaccines and other prevention measures
- Combating the *infodemic*:



Most NLP Work:
Detecting Misinformation

Our Work:
Characterizing the *reasons*
and attitudes that guide
people's choices

Explaining Opinions about the Covid Vaccine



...

There is no way someone can tell me that the COVID vaccine does not cause harm to pregnant women. The amount of patients that have come in for miscarriages here lately has almost quadrupled and 9 out of 10 of them have gotten the vaccine.

8:38 AM · Aug 16, 2021 · Twitter for iPhone

Explaining Opinions about the Covid Vaccine



...

There is no way someone can tell me that the COVID vaccine does not cause harm to pregnant women. The amount of patients that have come in for miscarriages here lately has almost quadrupled and 9 out of 10 of them have gotten the vaccine.

8:38 AM · Aug 16, 2021 · Twitter for iPhone



Stance: Not a fan

Explaining Opinions about the Covid Vaccine



...

There is no way someone can tell me that the COVID vaccine does not cause harm to pregnant women. The amount of patients that have come in for miscarriages here lately has almost quadrupled and 9 out of 10 of them have gotten the vaccine.

Stance: *Not a fan*

Why? *The vaccine is dangerous*

Explaining Opinions about the Covid Vaccine



...

There is no way someone can tell me that the COVID vaccine does not cause harm to pregnant women. The amount of patients that have come in for miscarriages here lately has almost quadrupled and 9 out of 10 of them have gotten the vaccine.

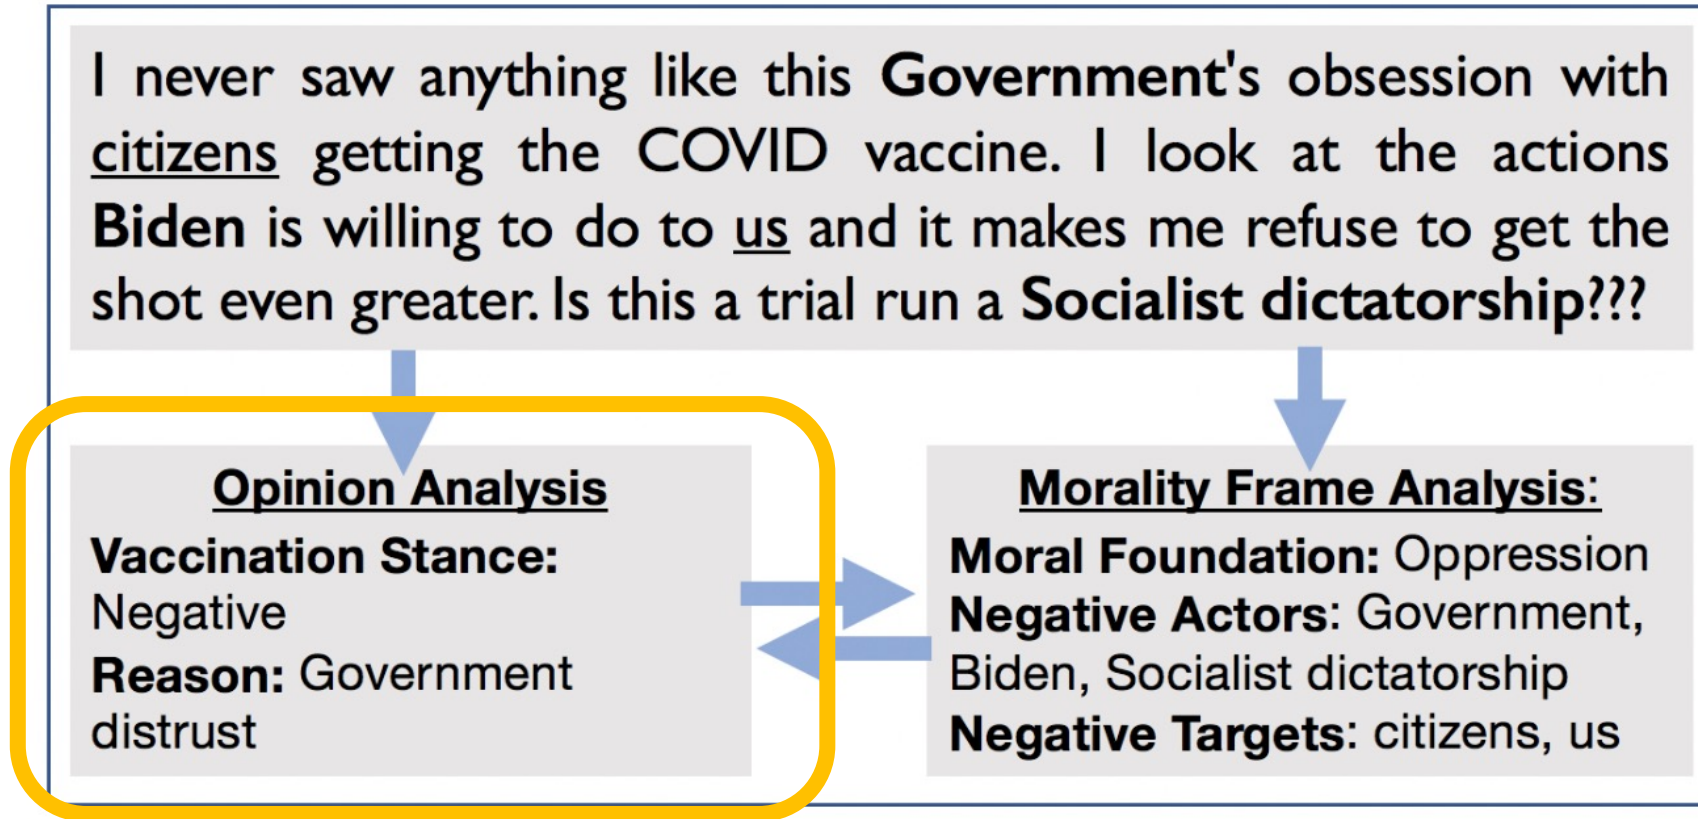
Stance: *Not a fan*

Why? *The vaccine is dangerous*

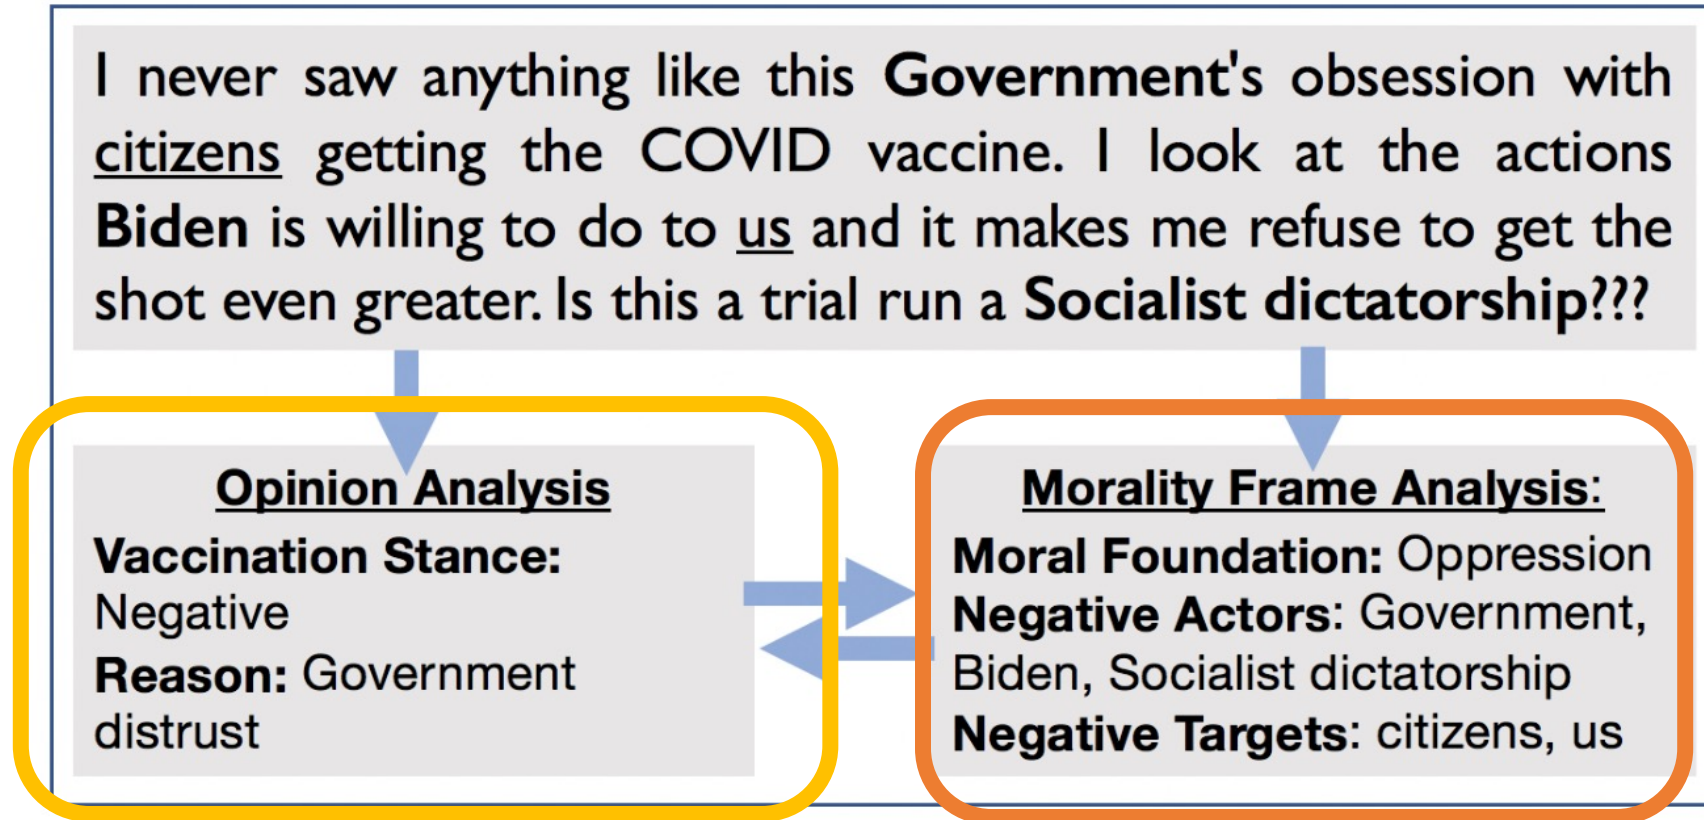


The vaccine framed as
harmful to pregnant women

An **Interconnected** View of Opinions



An **Interconnected** View of Opinions



Opinion Analysis

- Pro or anti vaccine, and the *reasons why*



Opinion Analysis

- Pro or anti vaccine, and the *reasons why*
- Three main challenges:



Opinion Analysis

- Pro or anti vaccine, and the *reasons why*
- Three main challenges:
 1. Predicting the stance



Opinion Analysis

- Pro or anti vaccine, and the *reasons why*
- Three main challenges:
 1. Predicting the stance
 2. **Constructing the space of possible reasons**



Opinion Analysis

- Pro or anti vaccine, and the *reasons why*
- Three main challenges:
 1. Predicting the stance
 2. **Constructing the space of possible reasons**
 3. **Mapping text to reasons**



Opinion Analysis

- Pro or anti vaccine, and the reasons why
- Three main challenges:
 1. Predicting the stance
 2. **Constructing the space of possible reasons**
 3. **Mapping text to reasons**



Established NLP classification task

Opinion Analysis

- Pro or anti vaccine, and the reasons why
- Three main challenges:



1. Predicting the stance

Established NLP classification task

2. **Constructing the space of possible reasons**

Big challenge

3. **Mapping text to reasons**

Interactive, humans-in-the-loop protocol

Interactive HIL Protocol


- **Initial set of themes:** main reasons people cite to refuse the vaccine
e.g: **“The vaccine is dangerous”** (Wawrutza et. al, 2021)

Interactive HIL Protocol

- **Initial set of themes:** main reasons people cite to refuse the vaccine
e.g: **“The vaccine is dangerous”** (Wawrutza et. al, 2021)
- **How can we ground them?**

Interactive HIL Protocol


- **Initial set of themes:** main reasons people cite to refuse the vaccine
e.g: “**The vaccine is dangerous**” (Wawrutza et. al, 2021)
- **How can we ground them?**



Experts can **explain themes** in **natural language**

Interactive HIL Protocol

- **Initial set of themes:** main reasons people cite to refuse the vaccine
e.g: “The vaccine is dangerous” (Wawrutza et. al, 2021)
- **How can we ground them?**




Experts can **explain themes** in **natural language**

Interactive HIL Protocol

- **Initial set of themes:** main reasons people cite to refuse the vaccine
e.g: “The vaccine is dangerous” (Wawrutza et. al, 2021)
- **How can we ground them?**



Model




Experts can **explain themes** in
natural language

Experts can **examine groundings** using
exploratory operations

Interactive HIL Protocol

- **Initial set of themes:** main reasons people cite to refuse the vaccine
e.g: “The vaccine is dangerous” (Wawrutza et. al, 2021)
- **How can we ground them?**

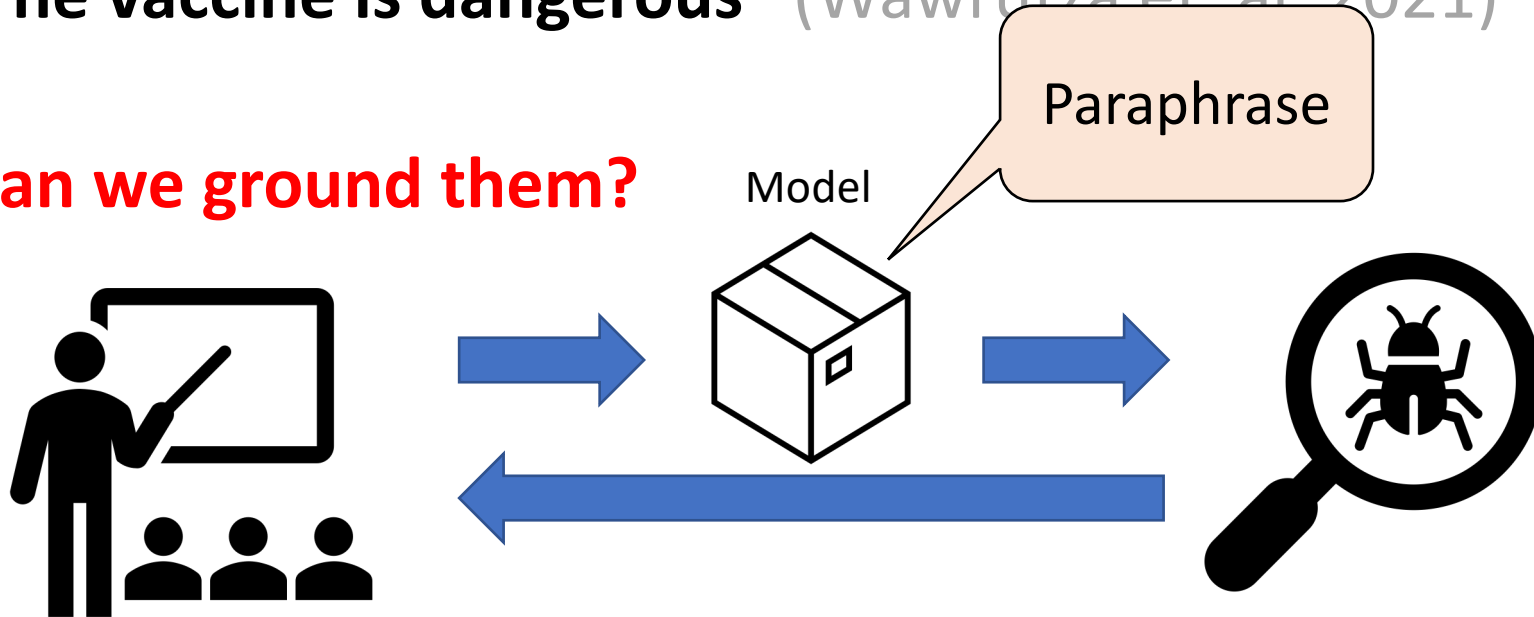


Experts can **explain themes** in natural language

Experts can **examine groundings** using exploratory operations

Interactive HIL Protocol


- **Initial set of themes:** main reasons people cite to refuse the vaccine
e.g: “The vaccine is dangerous” (Wawrutza et al 2021)
- **How can we ground them?**




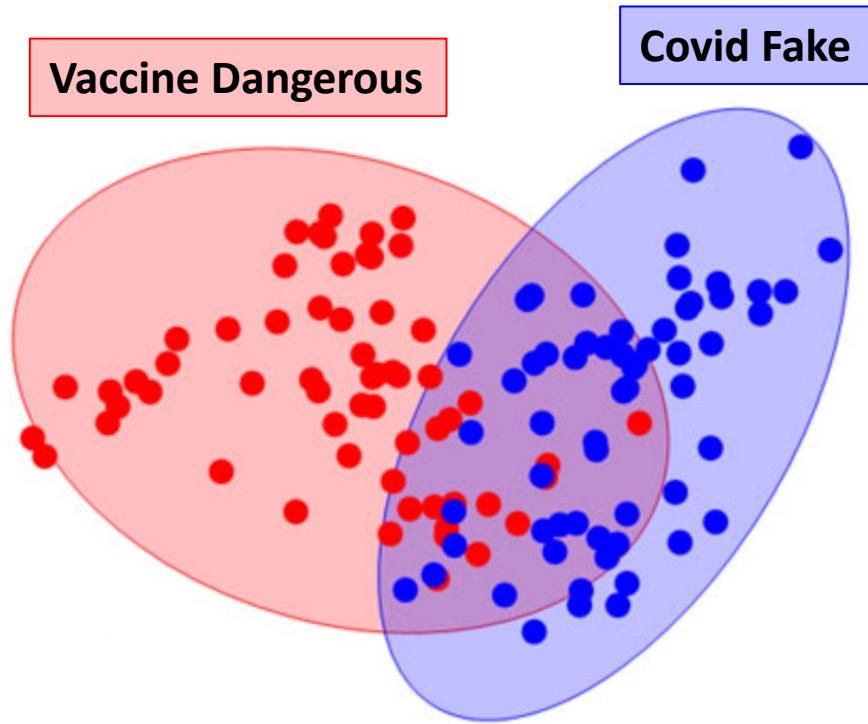
Experts can **explain themes** in natural language

Experts can **examine groundings** using exploratory operations


Experts Can Help us Refine the Themes



Experts Can Help us Refine the Themes



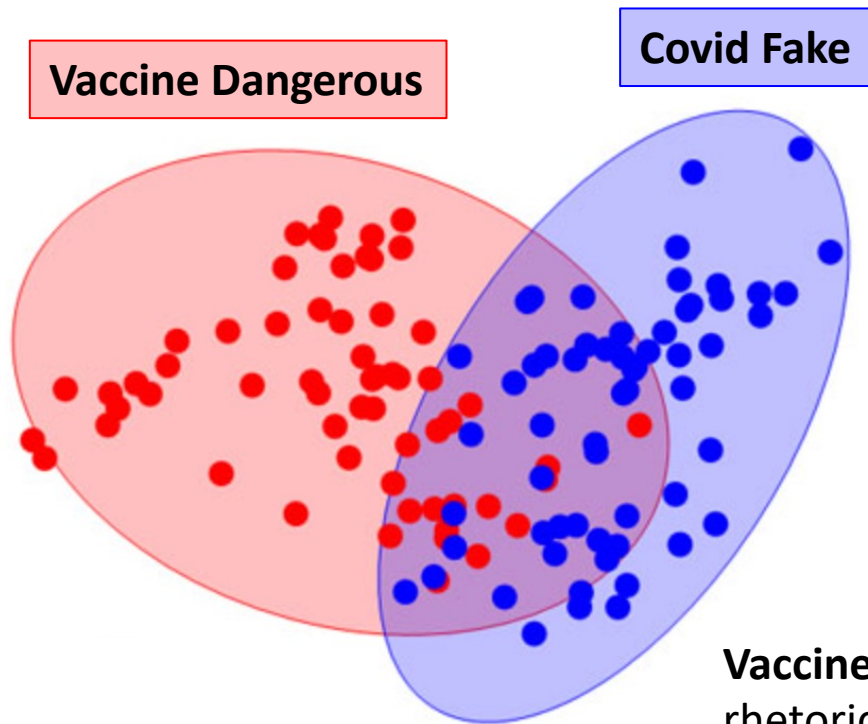
Experts Can Help us Refine the Themes



Covid Fake: wording related to be cheated, scammed or hoaxed.



Experts Can Help us Refine the Themes



Vaccine Dangerous

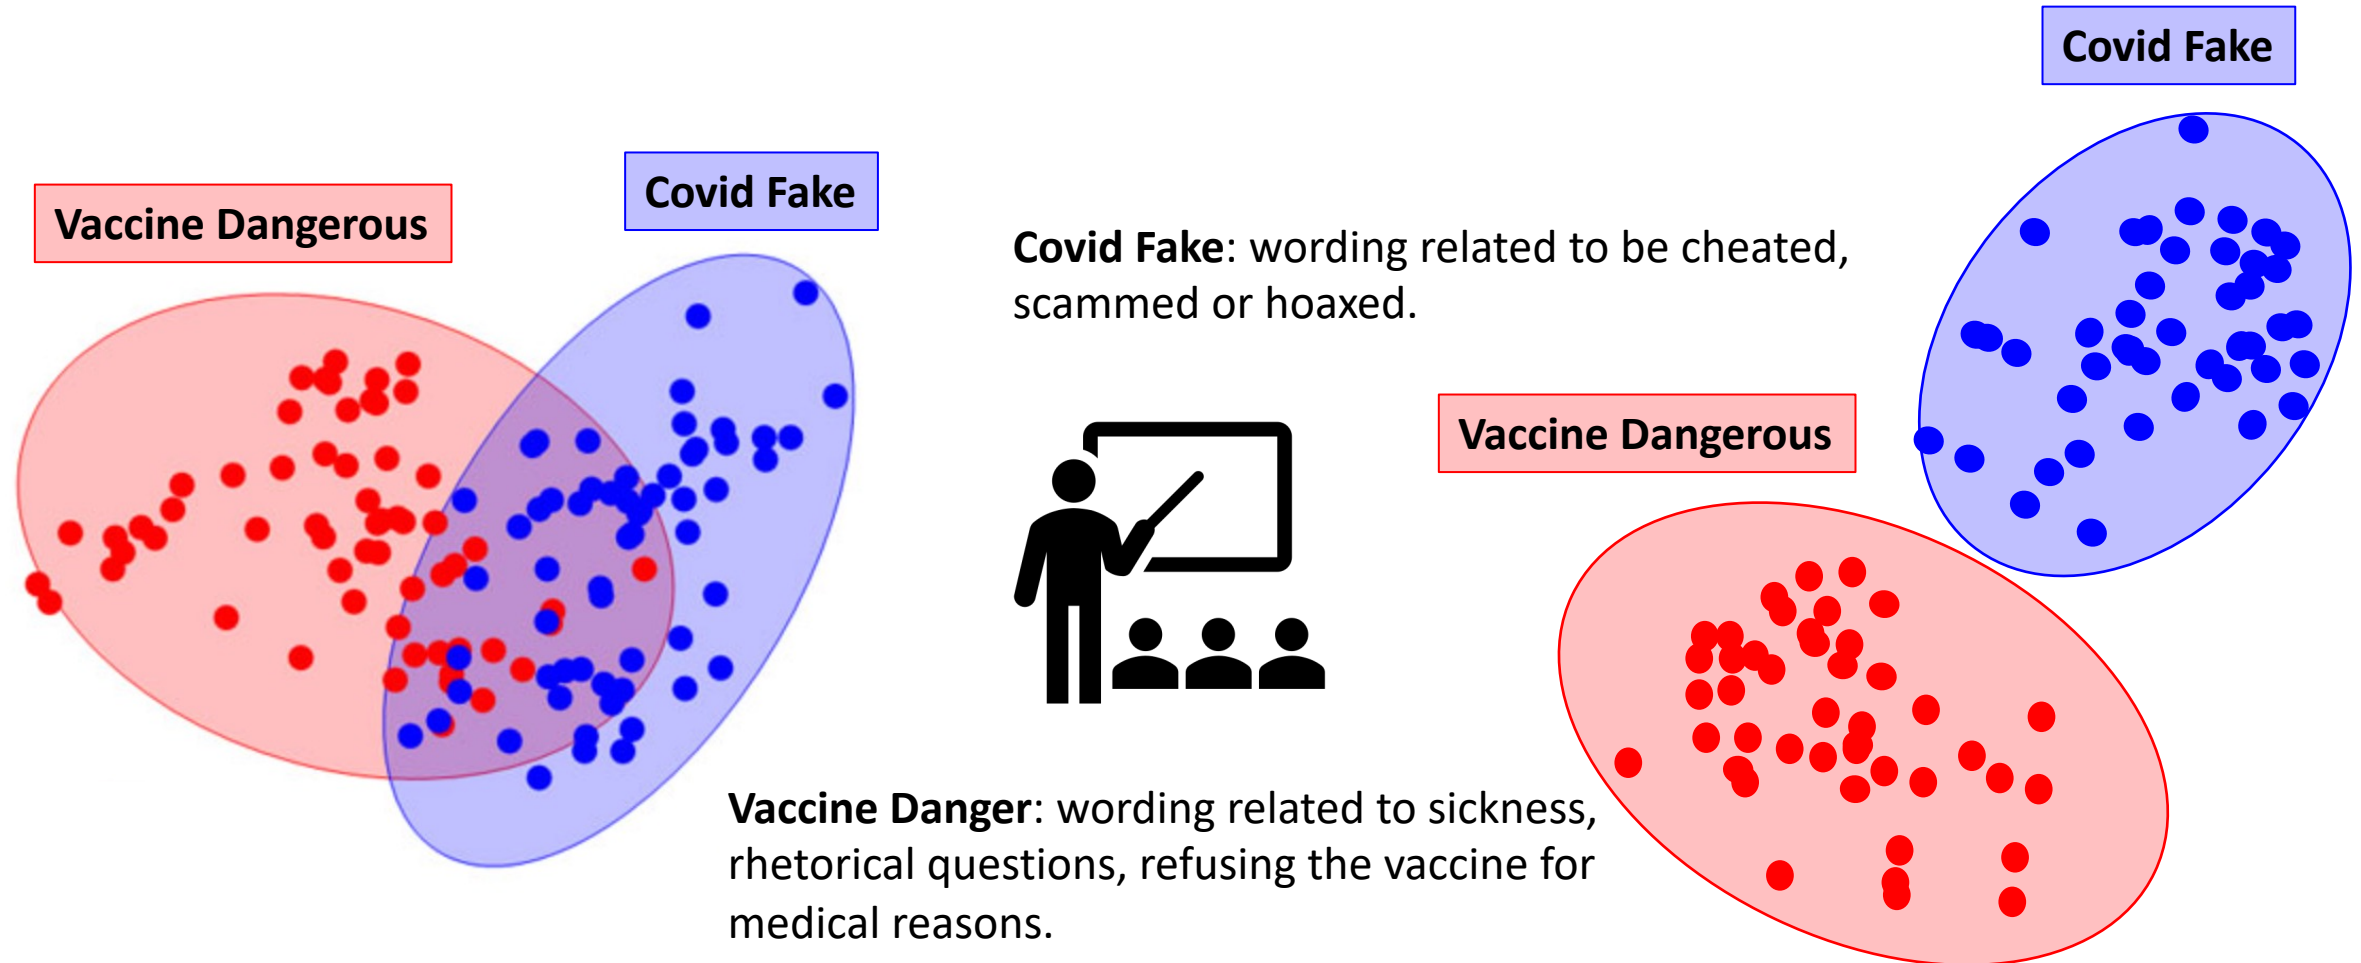
Covid Fake

Covid Fake: wording related to be cheated, scammed or hoaxed.




Vaccine Danger: wording related to sickness, rhetorical questions, refusing the vaccine for medical reasons.

Experts Can Help us Refine the Themes




Interactive Sessions

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




Interactive Sessions

- 3 NLP/CSS Researchers in **two** 1-hour sessions
- 85,000 unlabeled tweets about the covid vaccine




Interactive Sessions

- 3 NLP/CSS Researchers in **two** 1-hour sessions
- 85,000 unlabeled tweets about the covid vaccine
- **First session:** **adding new themes**, **removing themes** that were not prevalent in the data




Interactive Sessions

- 3 NLP/CSS Researchers in **two** 1-hour sessions
- 85,000 unlabeled tweets about the covid vaccine
- **First session:** **adding new themes**, **removing themes** that were not prevalent in the data
- **Second session:** identifying **high-level argumentative patterns** that were not being captured, **and contributing 2-5 examples**



Morality Frames Can Help us Explain Opinions

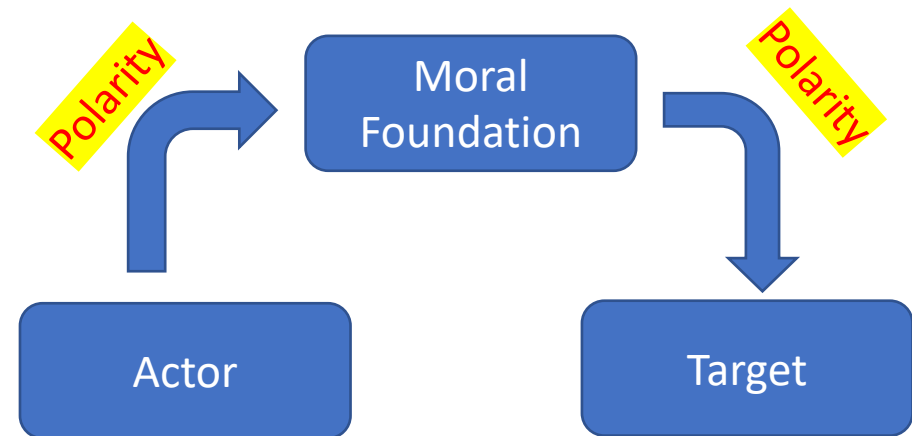
- **Reminder:** Morality frames capture differences in the **actors / targets** of moral sentiment
- They can help us *explain opinions*



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**




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**



Joint Model for Opinions and Morality Frames

Basic Classifiers

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

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

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

Joint Model for Opinions and Morality Frames

Basic Classifiers

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

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

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

Joint Model for Opinions and Morality Frames

Basic Classifiers

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

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

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

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

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)

- **Entities relations and prob. rules**
- Learn **embeddings** for ent/rels
- Learn **weights** for rules
- **Constrained optimization**

Joint Model for Opinions and Morality Frames

Basic Classifiers

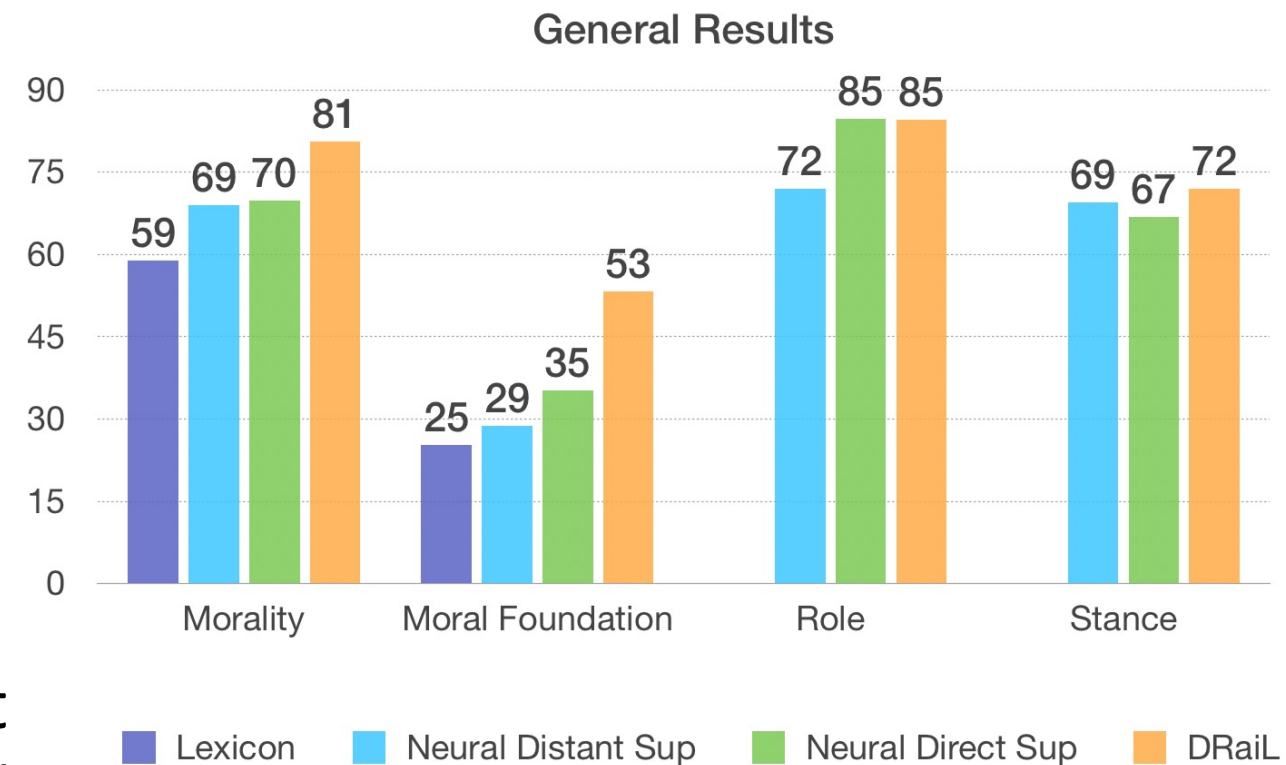
- Map tweets to stance/reason/MF
- Map entities to role and polarity

Dependencies between different dimensions


- If Fauci harming, likely anti-vax

Stance consistency preferences

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




Joint Inference in Low Supervision Settings



- **Joint inference** makes our model competitive with just 25% of direct supervision

Joint Inference in Low Supervision Settings



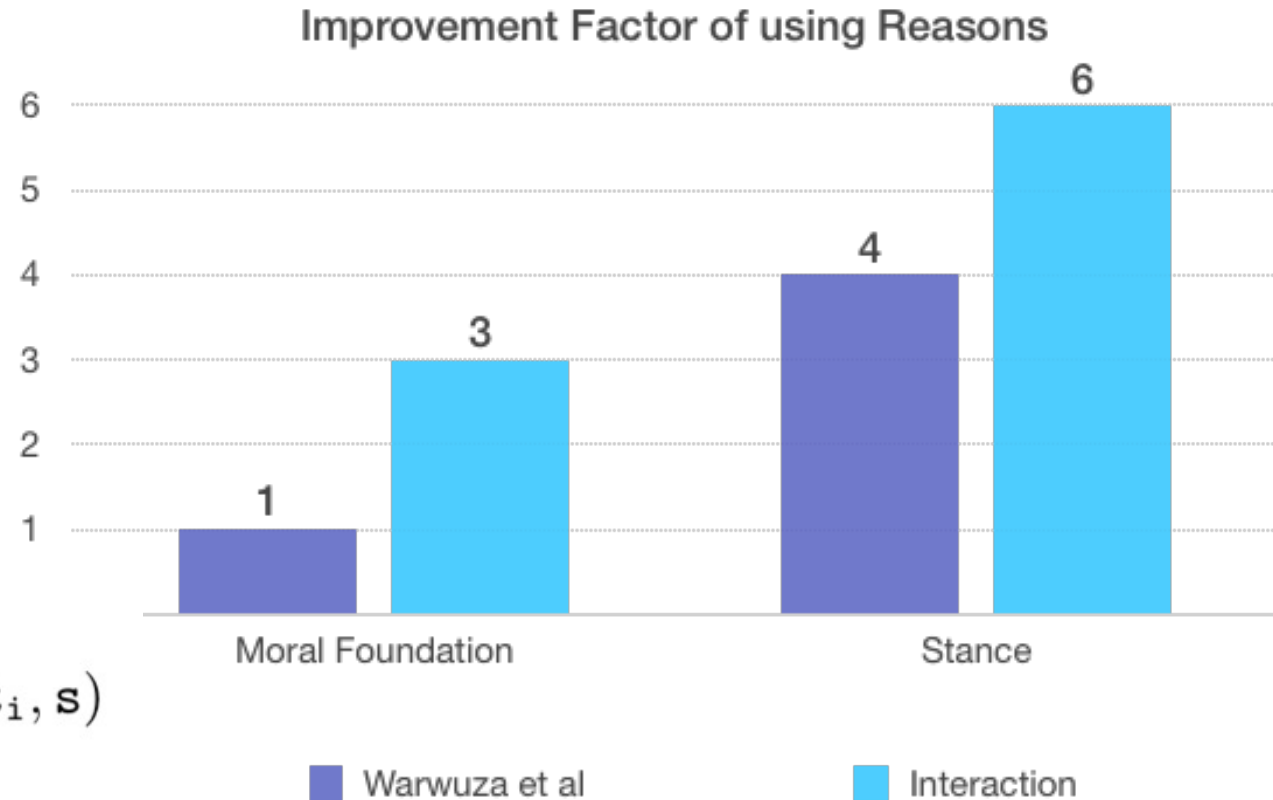
- **Joint inference** makes our model competitive with just 25% of direct supervision
- We beat the fully supervised base model with 50% of direct supervision

The Impact of Reasons **and Interaction**

Adding rules that
condition
stance/MF on
themes

$r_7 : \text{Mentions}(t_i, r) \Rightarrow \text{HasMf}(t_i, m)$
 $r_8 : \text{Mentions}(t_i, r) \Rightarrow \text{VaxStance}(t_i, s)$


+ all other rules



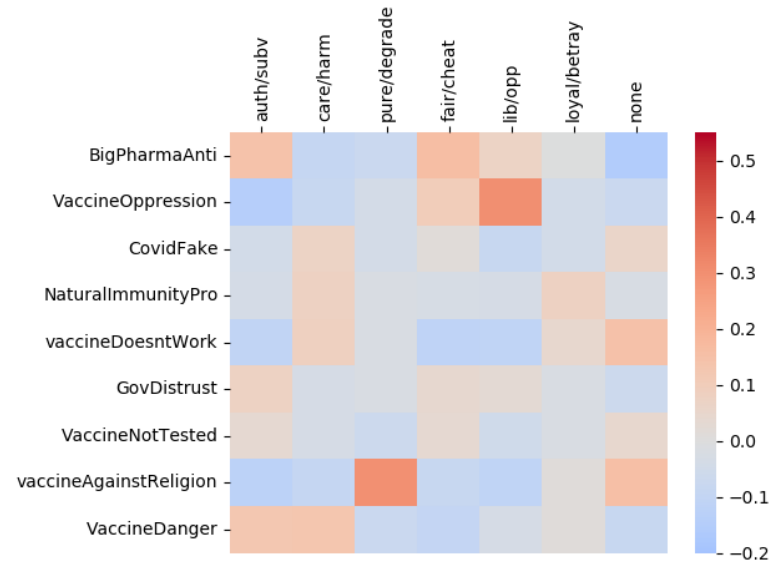
Is Human Intervention Helpful?

- Stronger correlations!

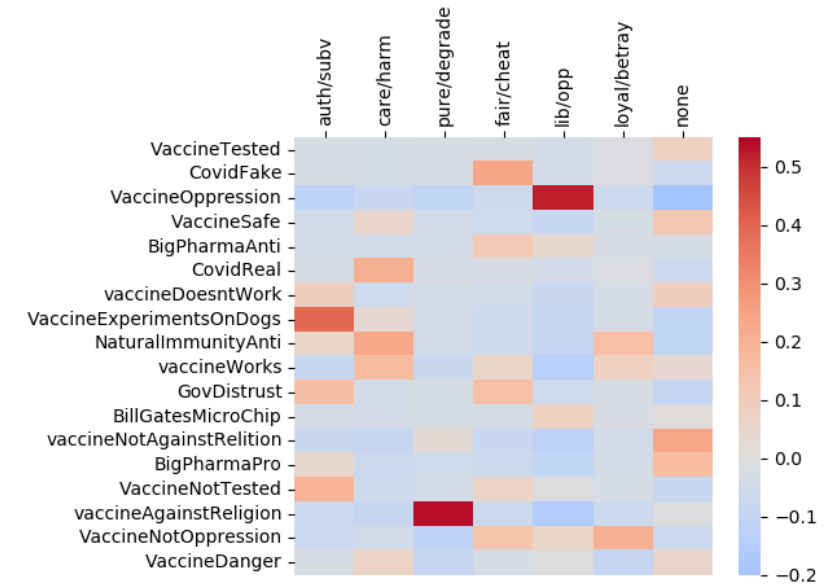
Correlation heatmaps between **Themes** and **Moral Foundations**



LDA Topics



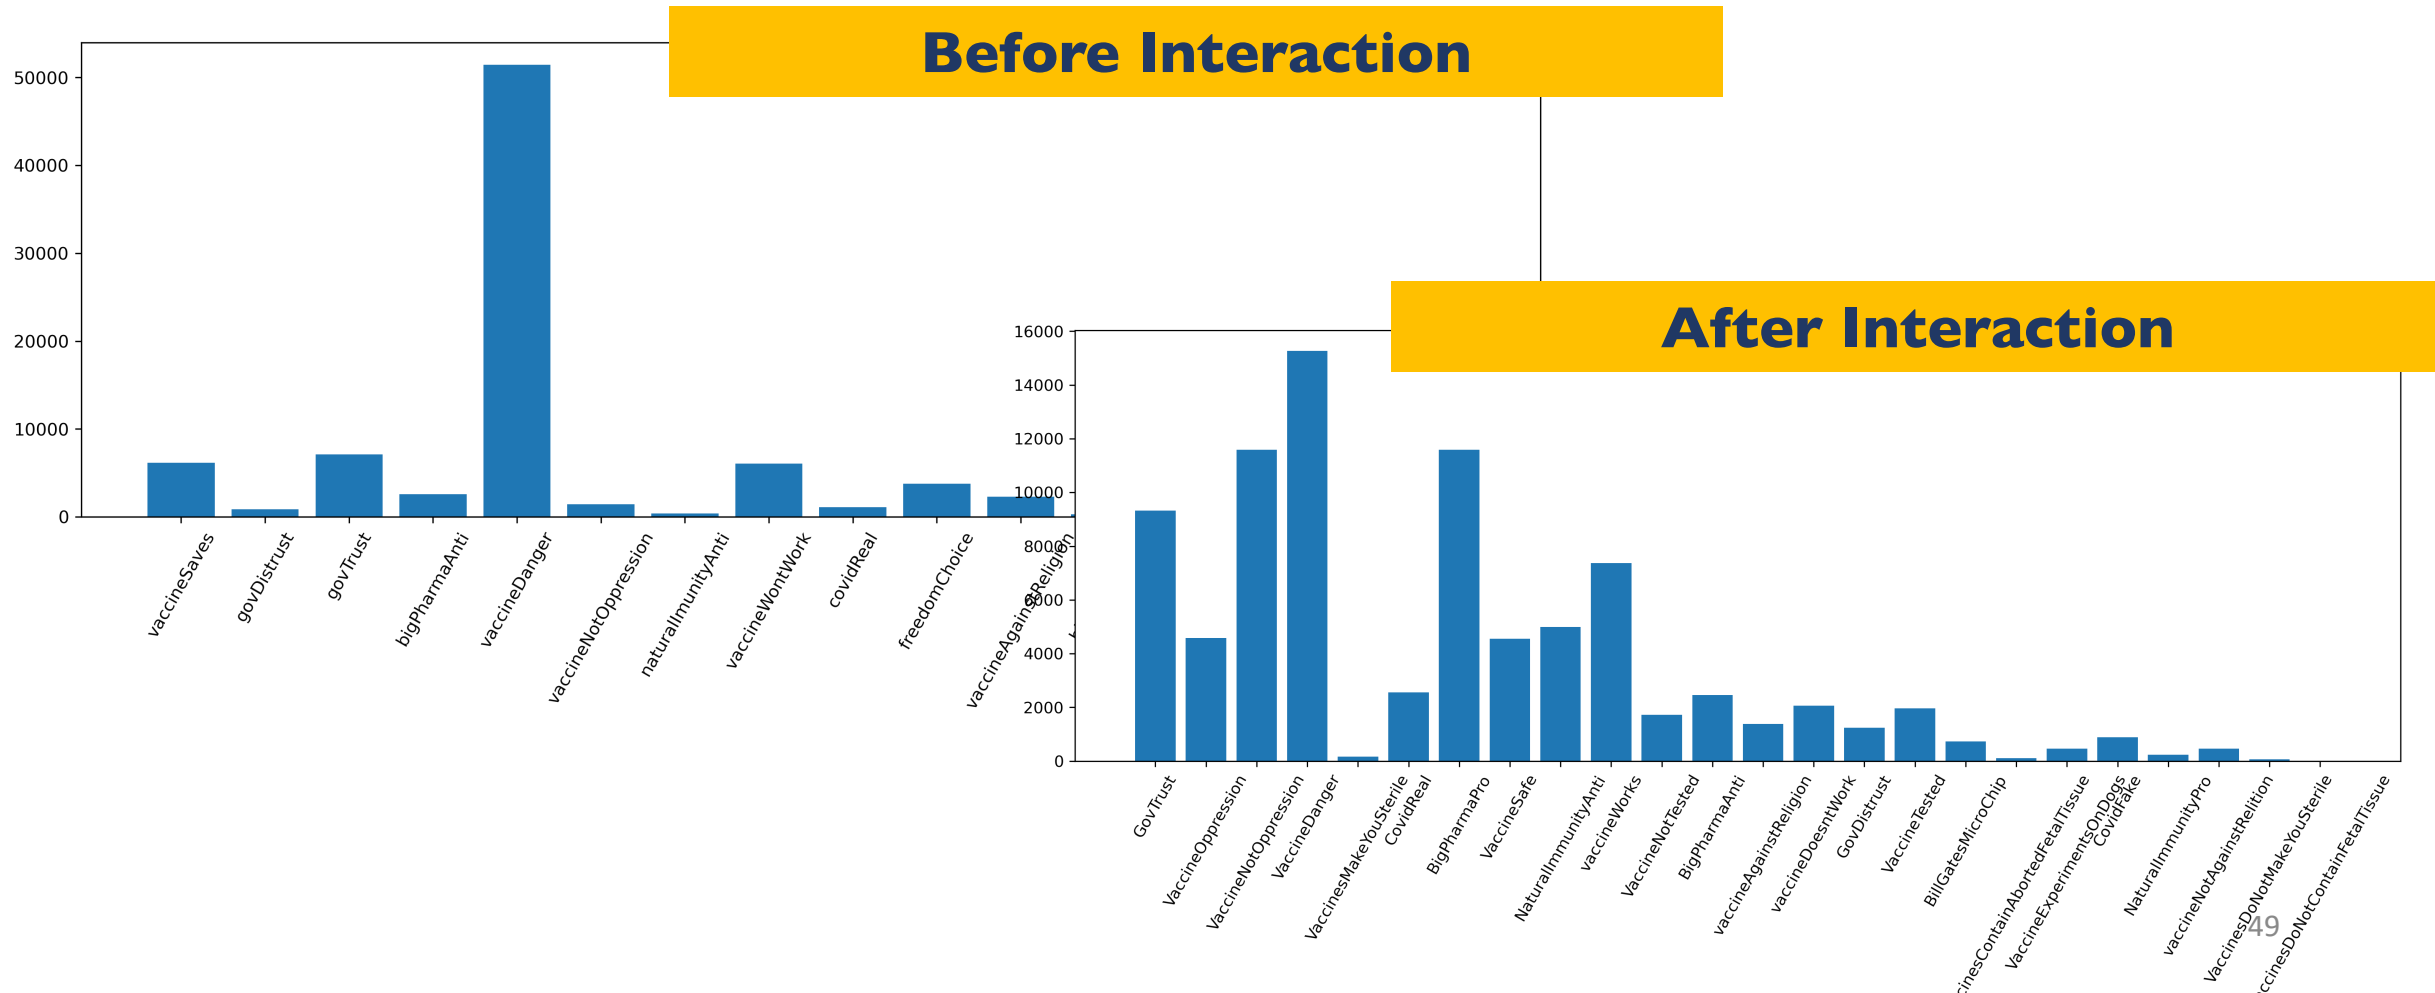
Wawrutza et. al, 2021 Themes



Interactive Themes

Is Human Intervention Helpful?

- More diverse distribution of reasons!



Discussion

- It ***makes sense*** to decompose the decision into multiple views
 - Better predictions
 - Interpretability

Discussion

- It *makes sense* to decompose the decision into multiple views
 - Better predictions
 - Interpretability
- Limitations: size of annotated dataset studied

Discussion

- It ***makes sense*** to decompose the decision into multiple views
 - Better predictions
 - Interpretability
- Limitations: size of annotated dataset studied
- Annotating for morality frames is expensive: **a case for semi-supervised methods.**

Discussion

- It **makes sense** to decompose the decision into multiple views
 - Better predictions
 - Interpretability
- Limitations: size of annotated dataset studied
- Annotating for morality frames is expensive: **a case for semi-supervised methods.**
- **Structured inference** and **human interaction** are good alternatives