MOKA: Moral Knowledge Augmentation for Moral Event Extraction

Xinliang Frederick Zhang[1], Winston Wu[2], Nick Beauchamp[3], and Lu Wang[1]

[1]Computer Science and Engineering, University of Michigan
[2]Department of Computer Science, University of Hawaii at Hilo
[3]Political Science and Network Science Institute, Northeastern University

NAACL 2024

Introduction: Morality

Preamble: “Human beings make **moral decisions** based on how they feel when confronted with an **experience**” discovered by Jonathan Haidt.

Introduction: Morality

Preamble: “Human beings make **moral decisions** based on how they feel when confronted with an **experience**” discovered by Jonathan Haidt.

Background: News media frame their stories to further a particular viewpoint, often employing moral values rather than explicitly biased language to **subtly affect readers** (Scheufele, 1999; Haidt and Graham, 2007; Haidt et al., 2009; Lakoff, 2010; Feinberg and Willer, 2015).

Introduction: Morality

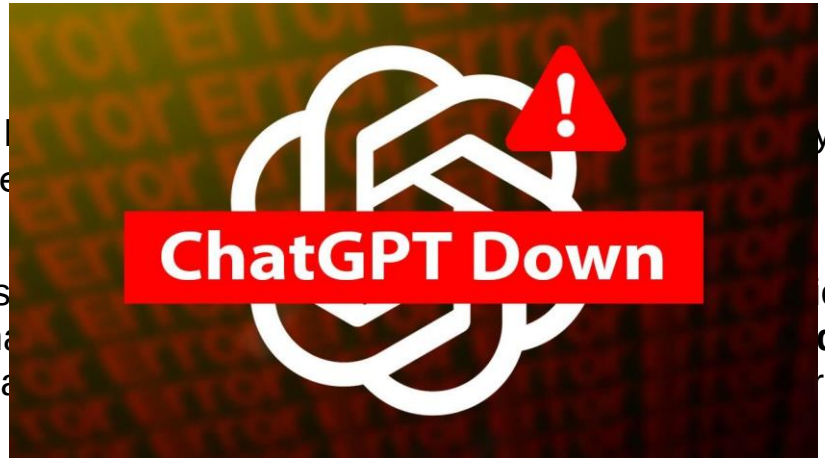

Preamble: “Human l y feel when confronted with an **experience**” discovere

Background: News ewpoint, often employing moral values rather tha **ders** (Scheufele, 1999; Haidt and Graham, 2007; Ha r, 2015).

Status Quo: Existing NLP methods, including **LLMs** like ChatGPT, face significant challenges in discerning moral values, duo to the subtlety challenge posed by morally-laden events.

Introduction: Morality

Preamble: “Human beings make **moral decisions** based on how they feel when confronted with an **experience**” discovered by Jonathan Haidt.

Background: News media frame their stories to further a particular viewpoint, often employing moral values rather than explicitly biased language to **subtly affect readers** (Scheufele, 1999; Haidt and Graham, 2007; Haidt et al., 2009; Lakoff, 2010; Feinberg and Willer, 2015).

Status Quo: Existing NLP methods, including **LLMs** like ChatGPT, face significant challenges in discerning moral values, duo to the subtlety challenge posed by morally-laden events.

Moral Foundation Theory (MFT): there is an innate psychological systems at the core of our “intuitive ethics”. MFT posits five moral foundations, each containing two polarities of virtue and vice.

Care/Harm
Fairness/Cheating
Loyalty/Betrayal
Authority/Subversion
Sanctity/Degradation

Task: Moral Event Extraction (MEE)

NEWS TITLE

Holder Orders Equal Treatment For Married Same-Sex Couples

NEWS BODY

· · · <Target> The new policy follows similar moves by the Department of Homeland Security and the IRS after the U.S. Supreme Court ① last year invalidated a key part of the Defense of Marriage Act ② which had defined the institution of marriage for federal purposes as limited to heterosexual couples. </Target> · · ·

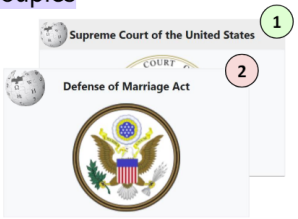

Wiki pages for linked names

Task formulation:

Given unstructured text, produce structural representations for morally-laden events:

- Detect morality-bearing event triggers
- Extract participants (agent & patient)
- Infer embodied moralities
- Determine event status

Task: Moral Event Extraction (MEE)

NEWS TITLE

Holder Orders Equal Treatment For Married Same-Sex Couples

NEWS BODY

· · · <Target> The new policy follows similar moves by the Department of Homeland Security and the IRS after the U.S. Supreme Court ① last year invalidated a key part of the Defense of Marriage Act ② which had defined the institution of marriage for federal purposes as limited to heterosexual couples. </Target> · · ·

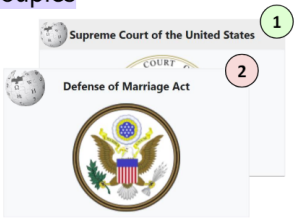

Wiki pages for linked names

EVENT RECORD 1

invalidated a key part of the Defense of Marriage Act

AGENT ①
Supreme Court of the United States

PATIENT
Same-Sex Couples

MORALITY

Care	Fairness	Loyalty	Authority	Sanctity
Harm	Cheating	Betrayal	Subversion	Degradation

EVENT STATUS

Actual	Intentional	Speculation

EVENT RECORD 2

defined the institution of marriage for federal purposes as **limited** to heterosexual couples

AGENT ②
Defense of Marriage Act

PATIENT
Same-Sex Couples

MORALITY

Care	Fairness	Loyalty	Authority	Sanctity
Harm	Cheating	Betrayal	Subversion	Degradation

EVENT STATUS

Actual	Intentional	Speculation

Task formulation:

Given unstructured text, produce structural representations for morally-laden events:

- Detect morality-bearing event triggers
- Extract participants (agent & patient)
- Infer embodied moralities
- Determine event status

Dataset: MORAL EVENTS

Stories count:	158
Articles count:	474
Moral Events count	5,494
Outlets coverage:	63 (26 Left, 18 Center, 9 Right)
Distinct entities count:	1,952 (E.g., Americans, Donald Trump)
Time range:	2012 – 2022
Quality:	Each article annotated by 2+ people; krippendorff's α > 0.9 for all attributes

- Balanced view
- Diverse
- Large
- High quality
- FIRST dataset annotating moral events

Model

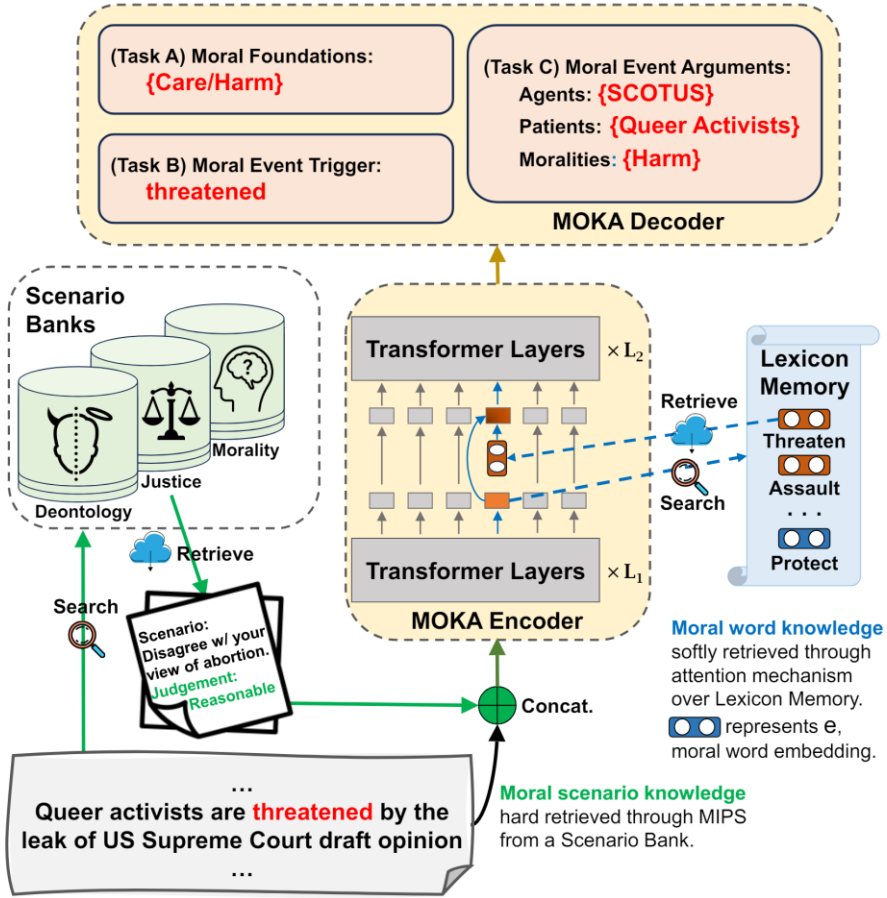

Our MOKA first retrieves K moral scenario pairs for the input passage. We then tag the moral knowledge-enriched input by identifying moral mentions. These mentions trigger Moral Lexicon Memory access to integrate moral word knowledge. Lastly, the dual knowledge-augmented representation is decoded to produce task-specific outputs.

Experiments Results (Selected)

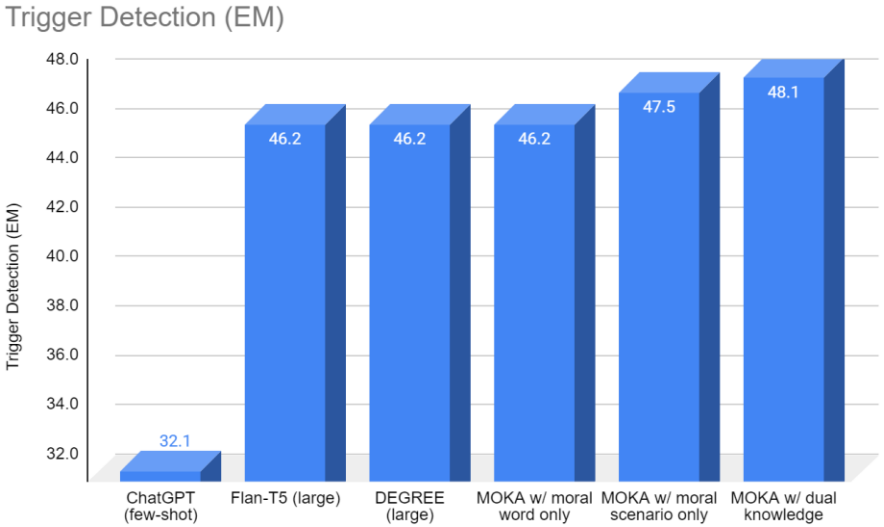

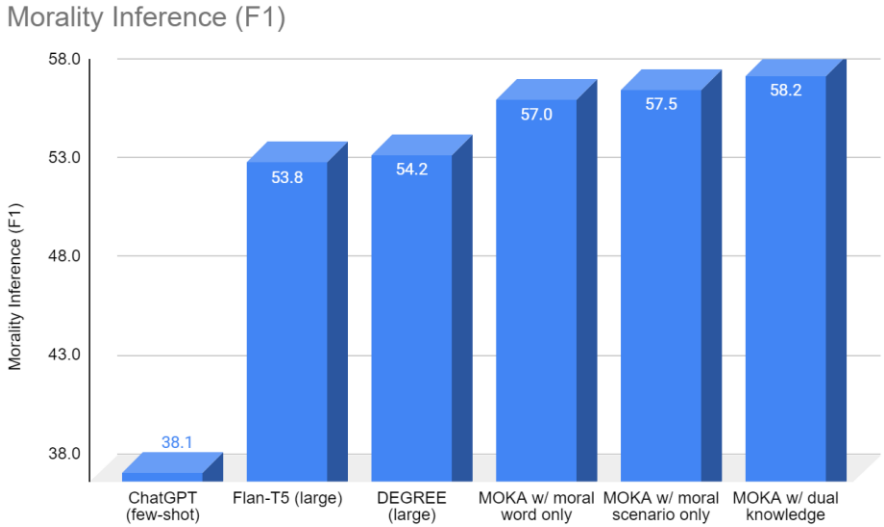

ChatGPT is significantly lagging due to insufficient moral knowledge seen in training.
Knowledge augmentation is essential for NLP models to excel at moral event extraction.
Our dual-knowledge augmented MOKA leads to the best performance across all tasks.

Further Analyses

RQ1: how do media of different ideologies focus on different moralities?

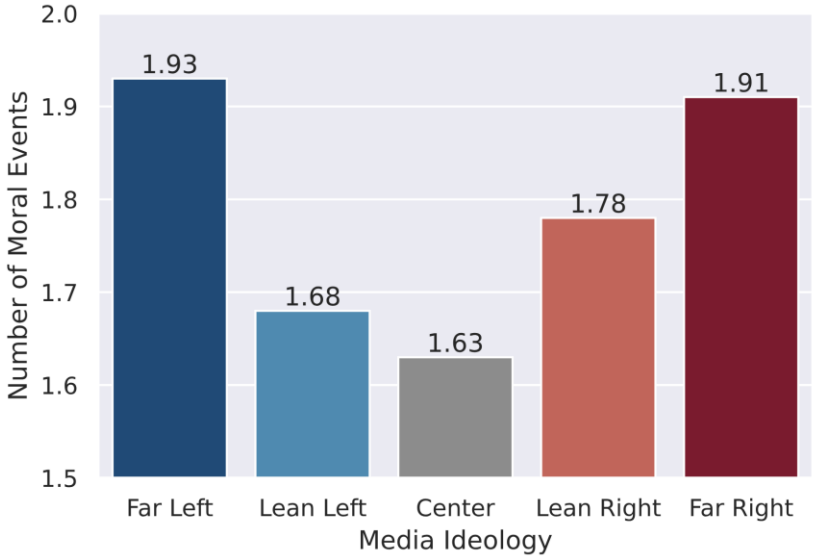

RQ2: how does selective reporting of moral narratives reveals more subtle and asymmetrical forms of bias?

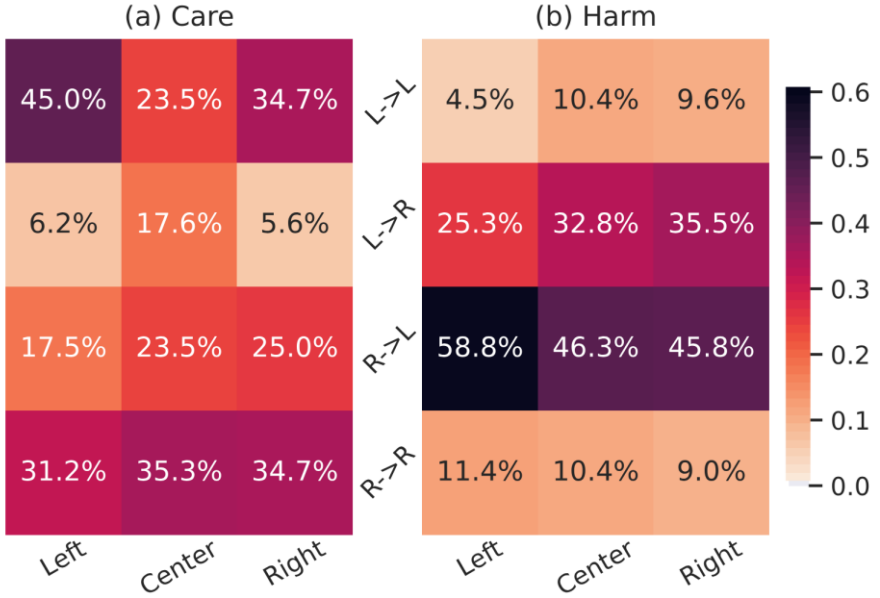

MOKA Conclusion

- We define a new **structured schema** for a **moral event**, rooted in a well-regarded psychosocial theory: **Moral Foundation Theory**.
- We present and investigate a novel task: **Moral Event Extraction.**
- To support this study, we curate a large dataset, **MORAL EVENTS**, with 5k+ fine-grained structural event annotations.
- We propose **MOKA**, a moral event extraction framework with **MO**ral **K**nowledge **A**ugmentation, leveraging knowledge derived from moral words and moral scenarios to comprehend morally-laden events.
- Further analyses of moral event reporting reveal the **left-right asymmetries** and **the distinctive behavior of centrist media.**

Thanks!

Codebase and dataset are available at https://github.com/launchnlp/MOKA.

Paper

Codebase & Data

Contact: xlfzhang@umich.edu

This work supported by NSF, AFOSR, and UM Advanced Research Computing.