MOKA: Moral Knowledge Augmentation for Moral Event Extraction

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

[1]University of Michigan,
[2]University of Hawaii at Hilo
[3]Northeastern University

Moral Event Extraction (MEE)

Preamble: "Human beings make **moral decisions** based on how they feel when confronted with an **experience.**" -- Jonathan Haidt

Moral Foundation Theory: 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.

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

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

Basic statistics of MORAL EVENTS dataset. For a **balanced view**, each news story consists of 3 articles reported by media of different ideology. MORAL EVENTS also features the **first** dataset annotating **structured** moral events.

Sample moral event extractions (MEEs) for a target sentence from our MORAL EVENTS dataset.

|	Virtue	Vice	Proportion
Care/Harm	1,348	2,060	51.6%
Fairness/Cheating	531	453	14.9%
Loyalty/Betrayal	329	257	8.9%
Authority/Subversion	1,140	418	23.6%
Sanctity/Degradation	19	46	1.0%

Distribution of moralities in MORAL EVENTS dataset.

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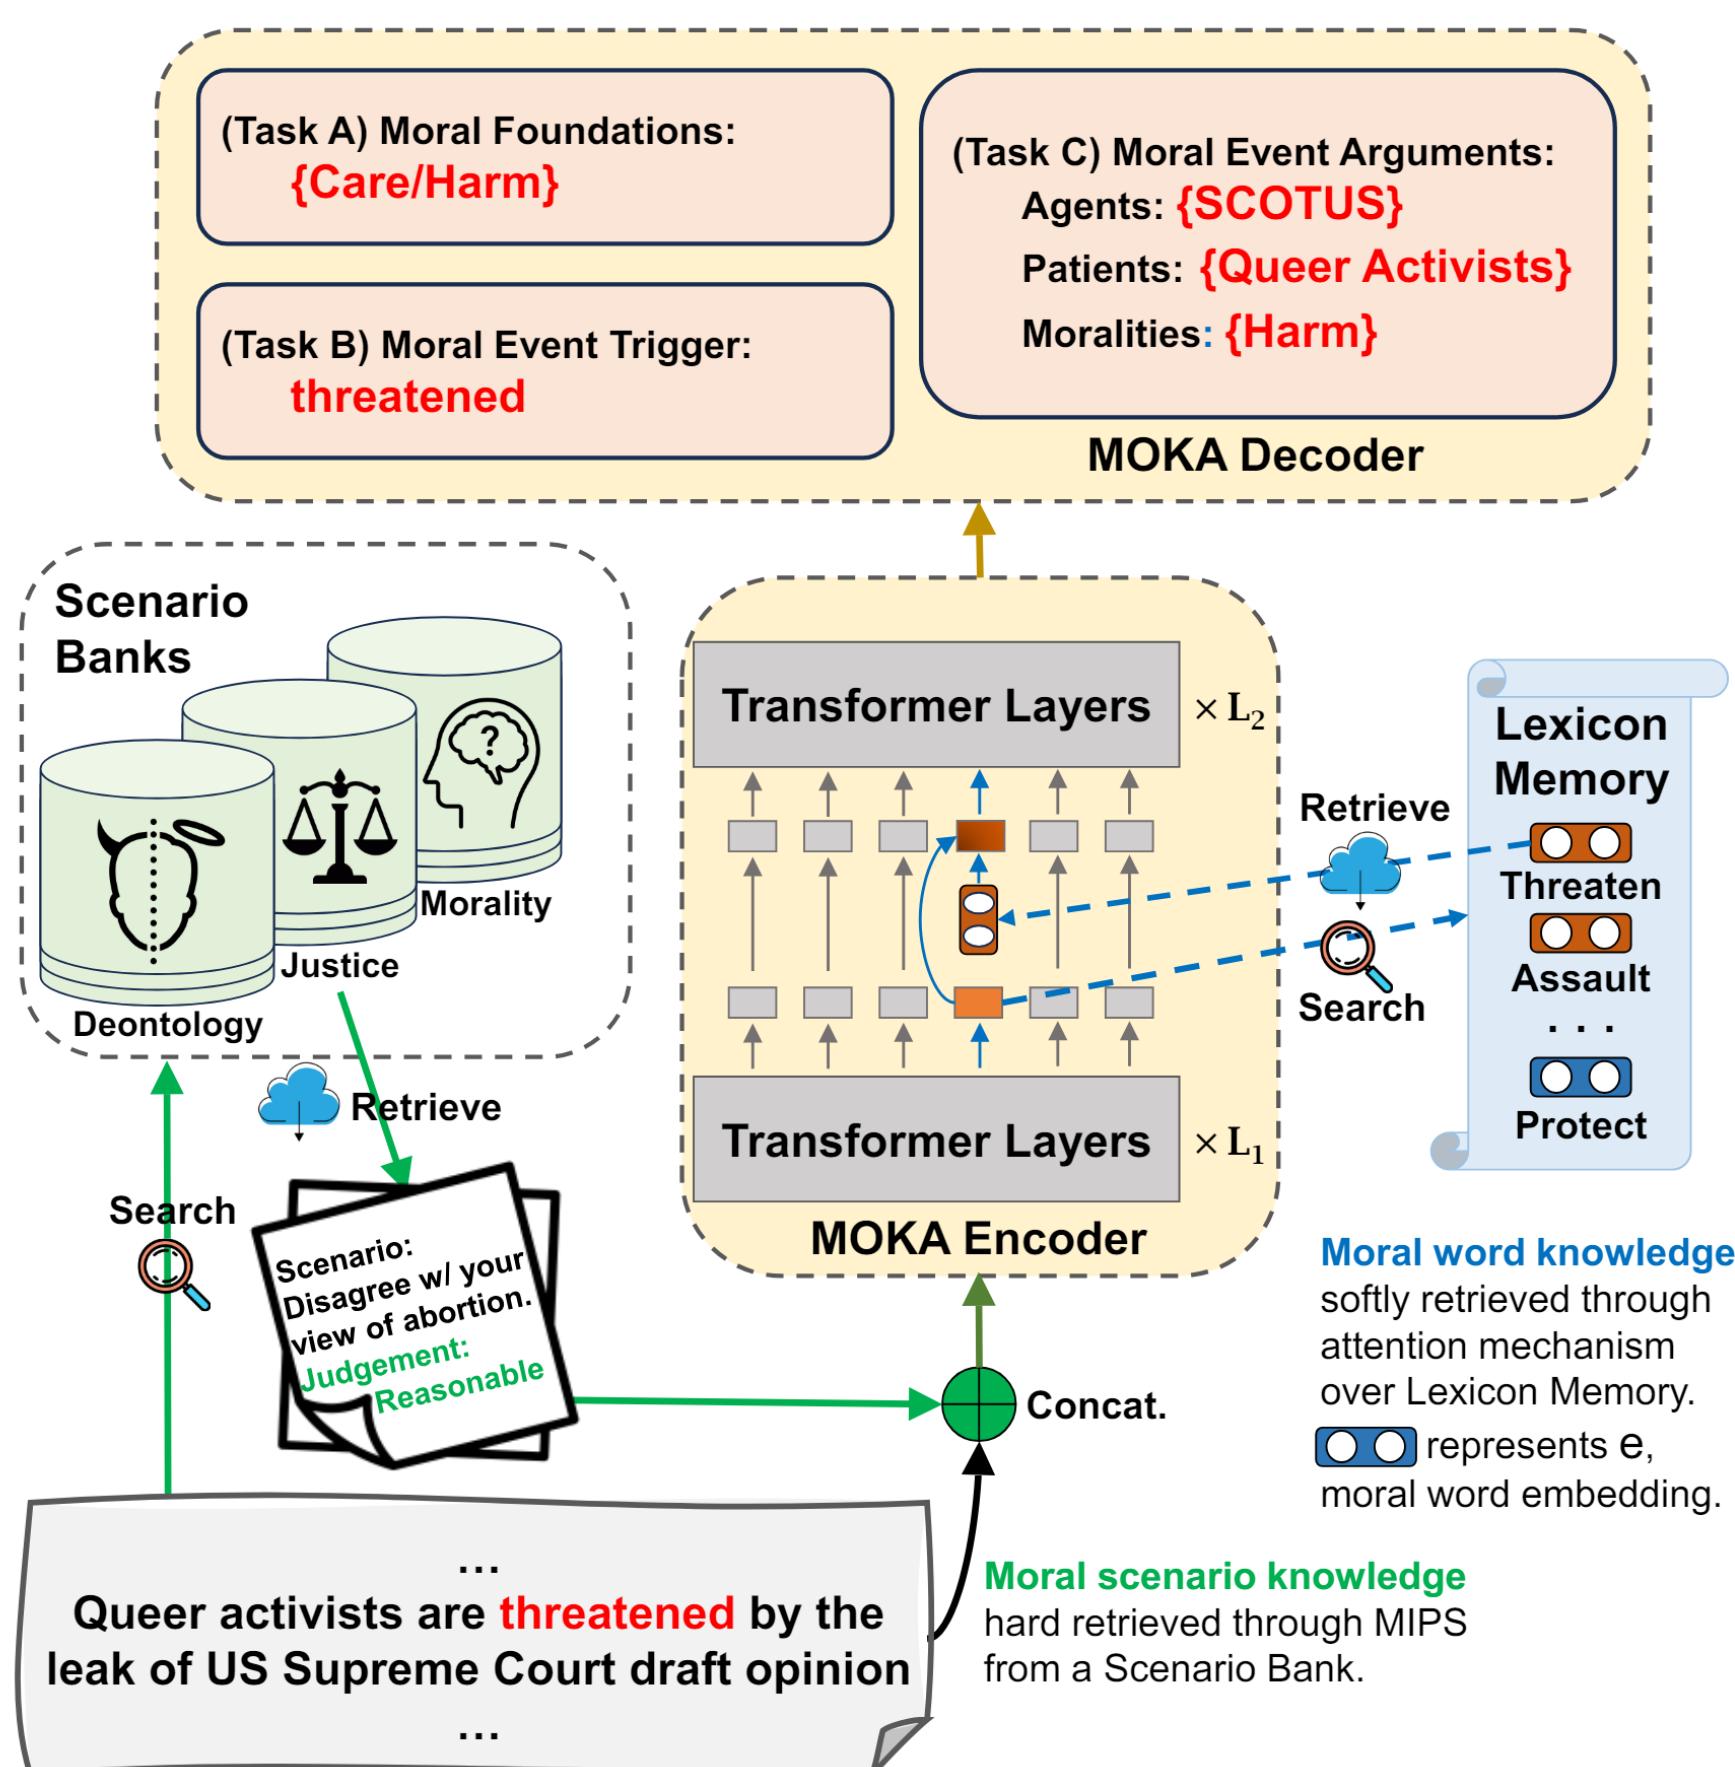
Model Architecture and Information Flow of MOKA. Our MOKA first **retrieves** K **moral scenario** pairs for the input passage. We then **tag** the moral knowledge-enriched input by identifying **moral mentions**. These mentions trigger **Moral Lexicon Memory** access to integrate moral word knowledge.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

Results
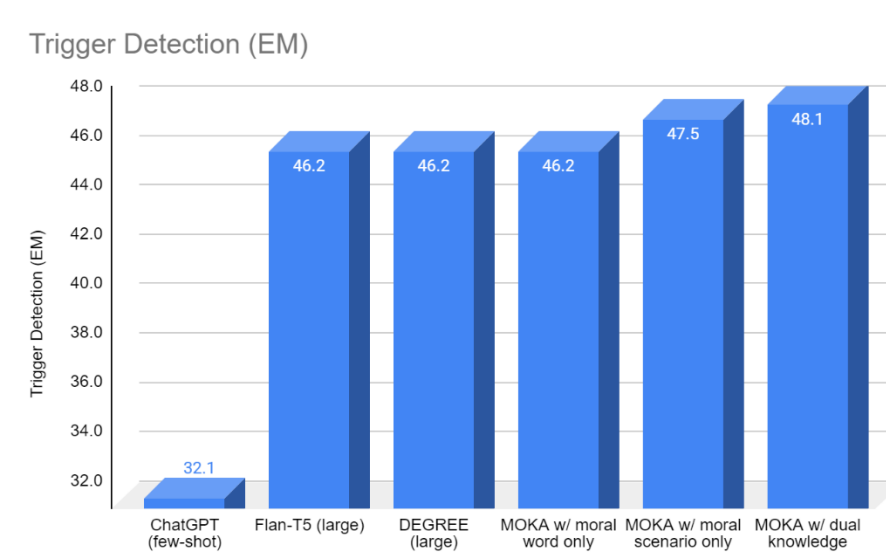
Results on MORAL EVENTS for Moral Event Trigger Detection task.
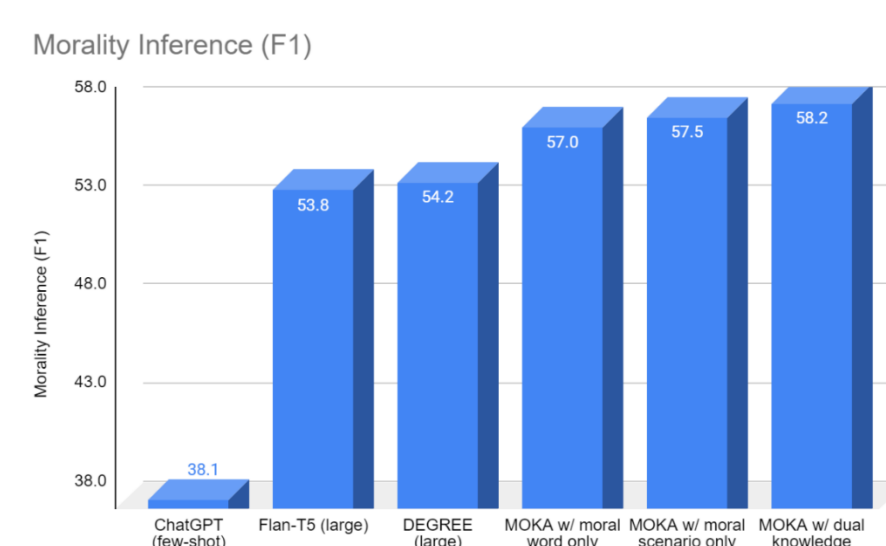
Results on MORAL EVENTS for Morality Identification task (part of Moral Event Argument Extraction task).

Take-home messages

- **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.**

Code: github.com/launchnlp/MOKA
Contact: xlfzhang@umich.edu