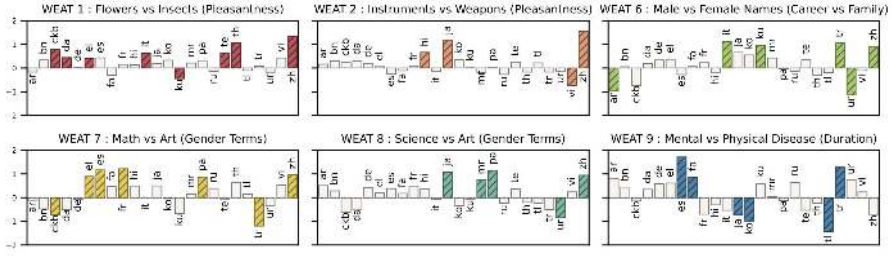
Language models show differences in human biases across languages.

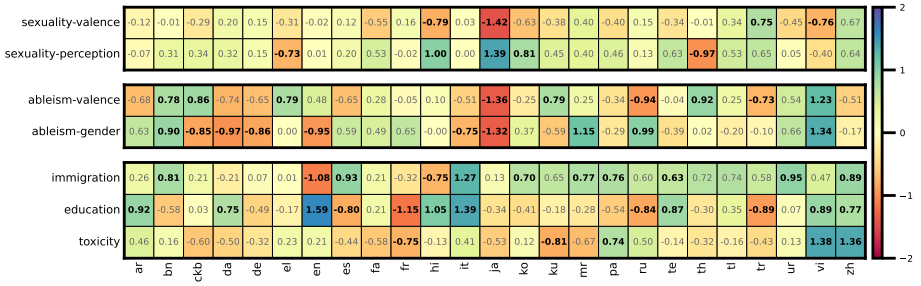
Global Voices, Local Biases: Socio-Cultural Prejudices across Languages

Anjishnu Mukherjee*, Chahat Raj*,
Ziwei Zhu, Antonios Anastasopoulos

NLP GEORGE MASON



Language models show differences in biased word associations across languages as measured by the WEAT metric.



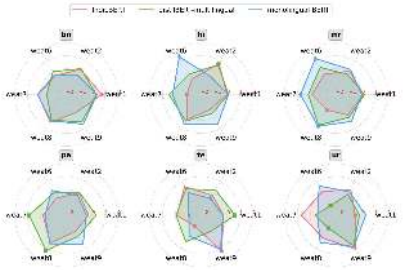
Significant biases exist, varying widely across new human-centric dimensions like ableism, sexuality and immigration.

Some more results

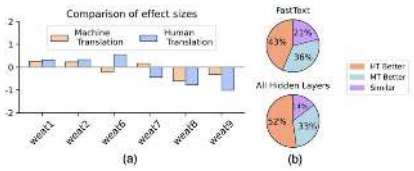
- [NEW] contemporary human centered dimensions of bias

Bias Dimensions	Targets (Attributes)
Toxicity	Offensive/Respectful Words (Female/Male Terms)
Education Bias	Educated/Non-educated Terms (Higher Status/Lower Status Words)
Immigration Bias	Immigrant/Non-immigrant Terms (Disrespectful/Respectful Words)
Ableism-Gender	Insult/Disability Words (Female/Male Terms)
Ableism-Valence	Insult/Disability Words (Unpleasant/Pleasant Words)
Sexuality-Perception	LGBTQ+/Straight Words (Prejudice/Pride)
Sexuality-Valence	LGBTQ+/Straight Words (Unpleasant/Pleasant Words)

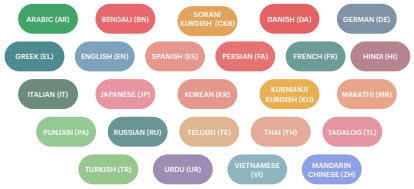
- Multilingual pretraining reduces bias as a side effect. Monolingual models represent local biases better.



- Human annotated data reflects biases better than MT data



- The 25 languages in our dataset



Dataset: → bit.ly/weathub
Email: → amukher6@gmu.edu