

LUCIE FLEKOVA, JORDAN CARPENTER, SALVATORE GIORGI, LYLE UNGAR, DANIEL PREOTIUC-PIETRO

- How accurate are people in judging traits of other users?
- Are there systematic biases humans are subject to?
- What are the implications of using human perception as a proxy for truth?
- Which textual cues lead to a false perception of the truth?
- Which textual cues make people more or less confident in their ratings?

GENDER

AGE

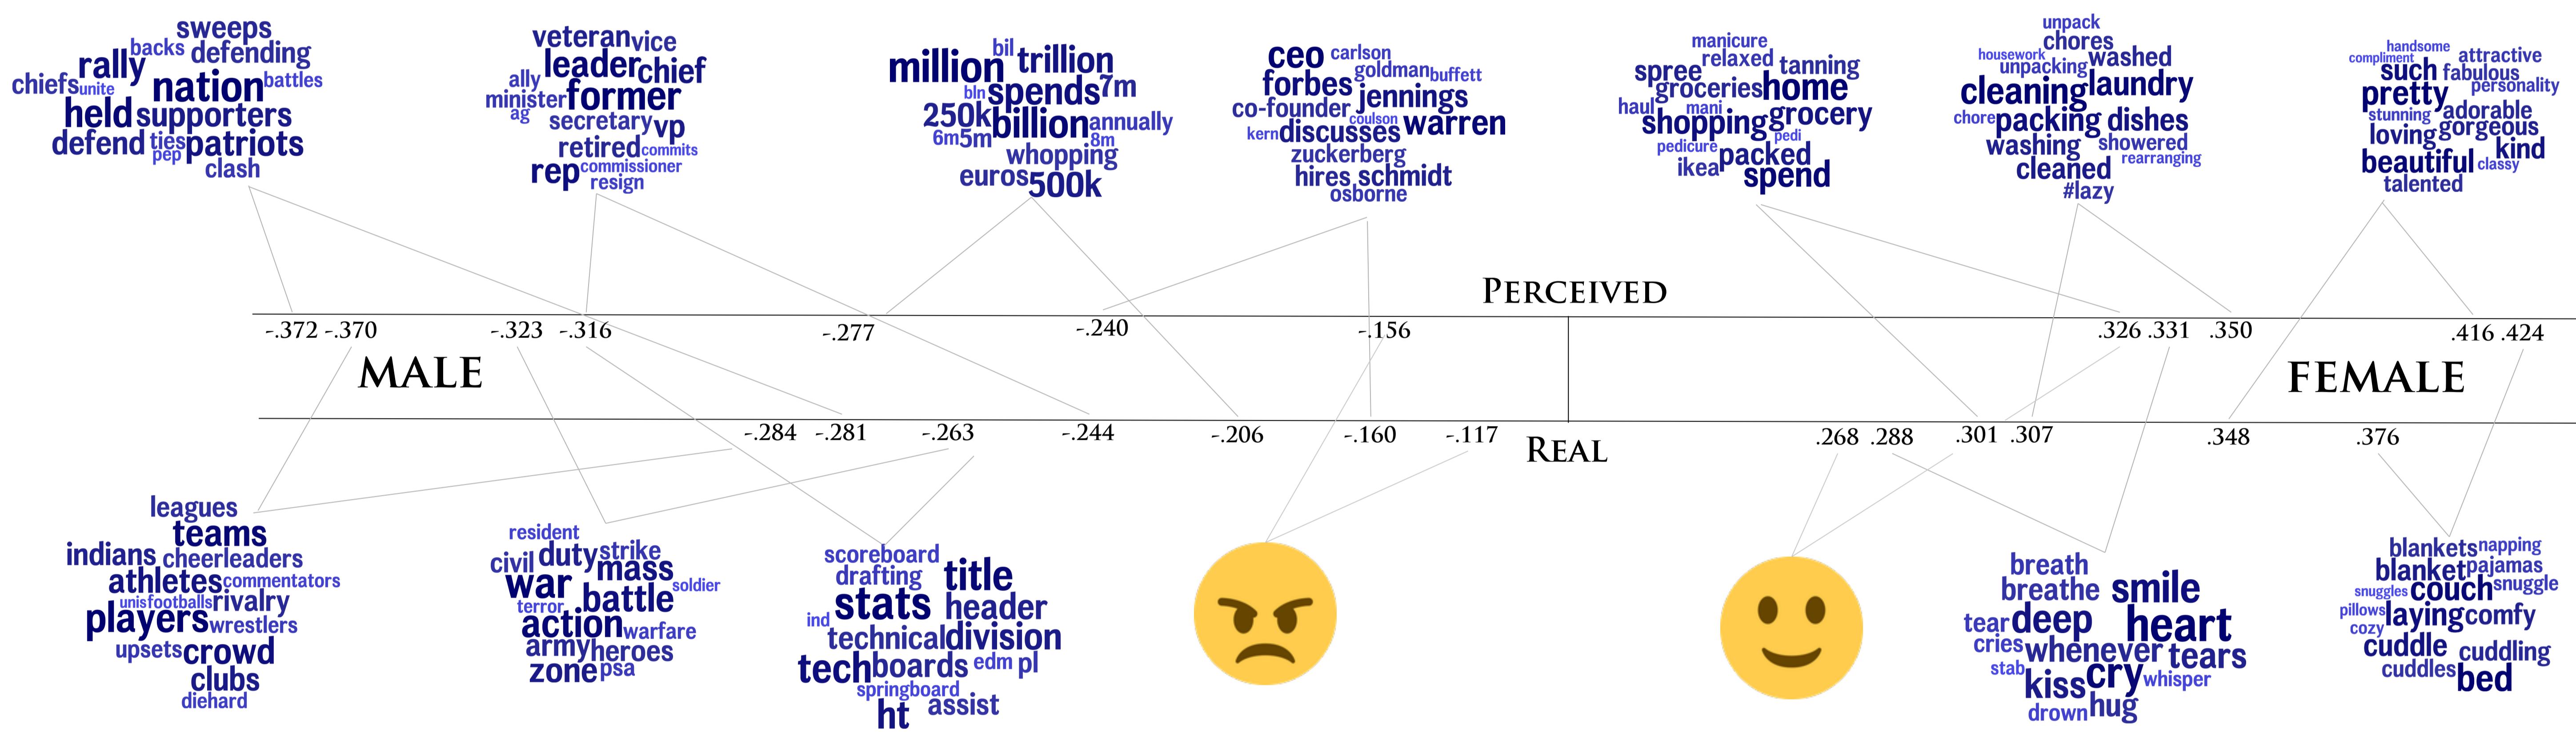
FEMALE MALE CONFIDENCE: 1 2 3 4 5

DATA

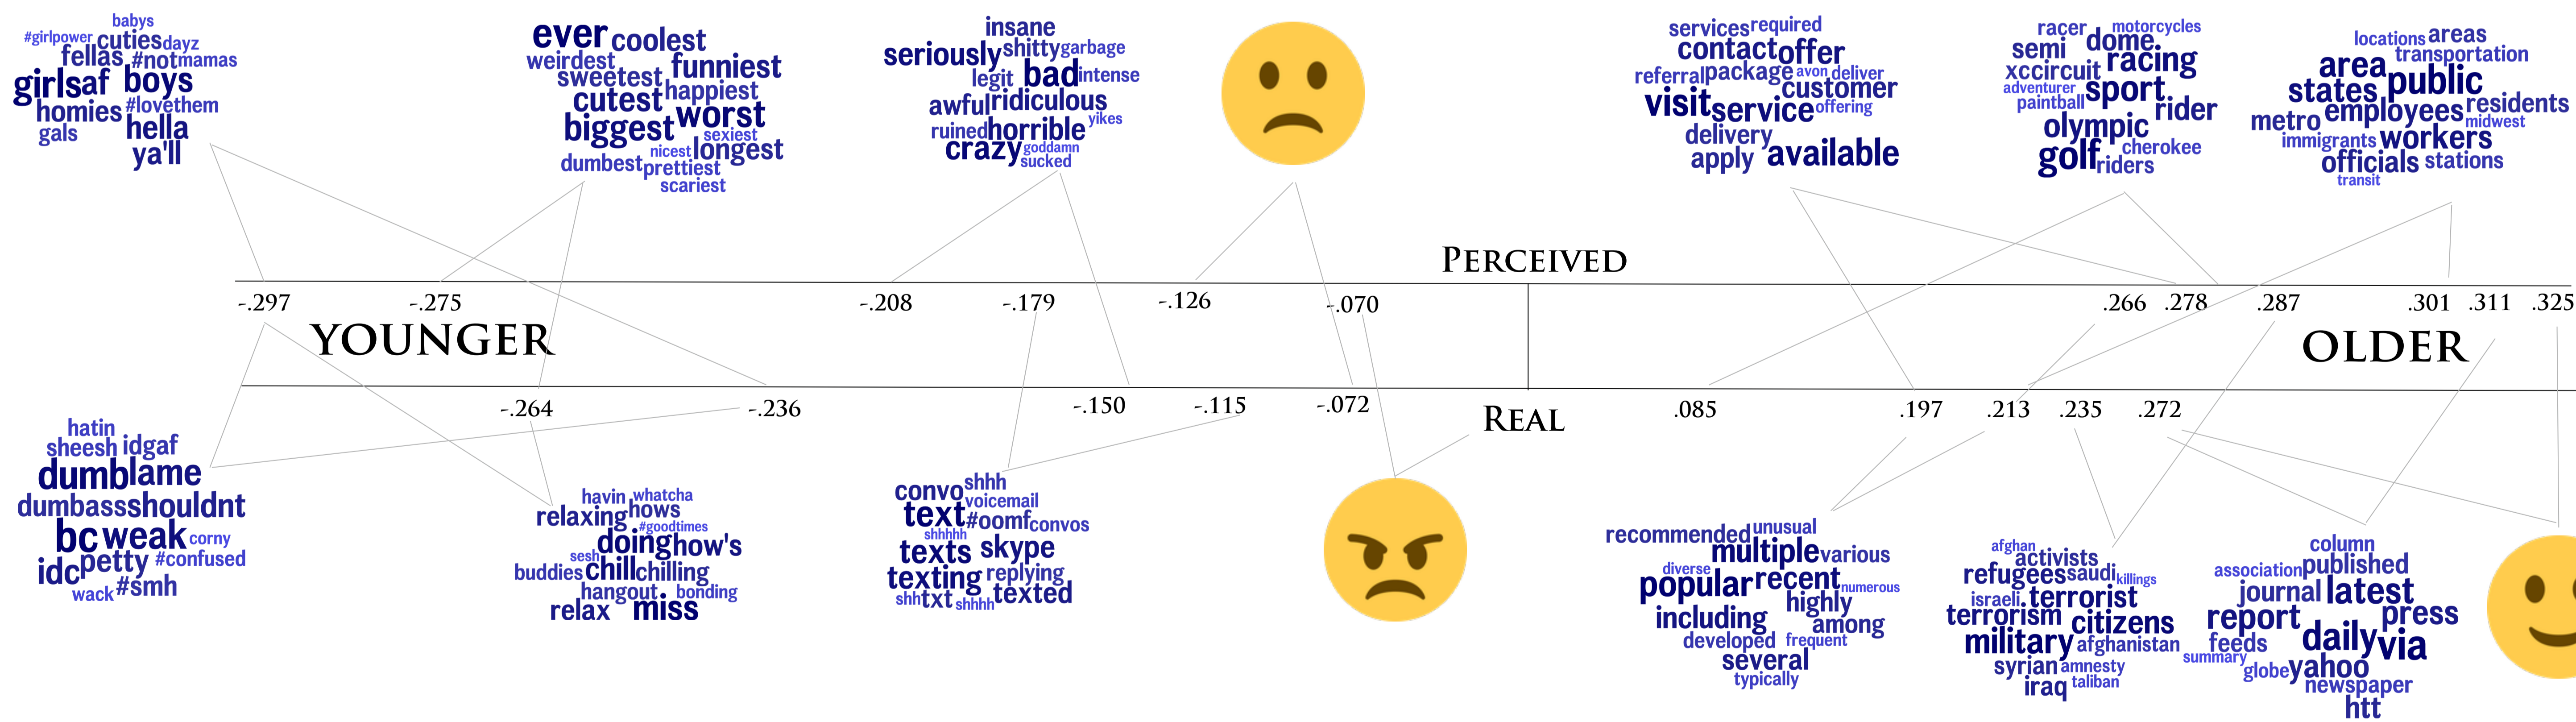


- Gender - 2,607 authors, age - 826 authors
- we use 100 tweets per author, 9 Mturk votes per author
- URLs and mentions anonymized, English only filtered, duplicates eliminated, same 6 month time interval


GENDER PERCEPTION



AGE PERCEPTION



CLASSIFIERS TRAINED ON REAL AND PERCEIVED LABELS PERFORM DIFFERENTLY



- The first study to systematically analyze differences between real and perceived user traits
- People are generally accurate, but overemphasize stereotypes
- Confidence feature analysis in the full paper
- Using human perception, we introduce systematic biases in classifiers