Persuasion Under Information Overload

Emaad Manzoor George H. Chen Dokyun Lee Michael D. Smith

Online Deliberation

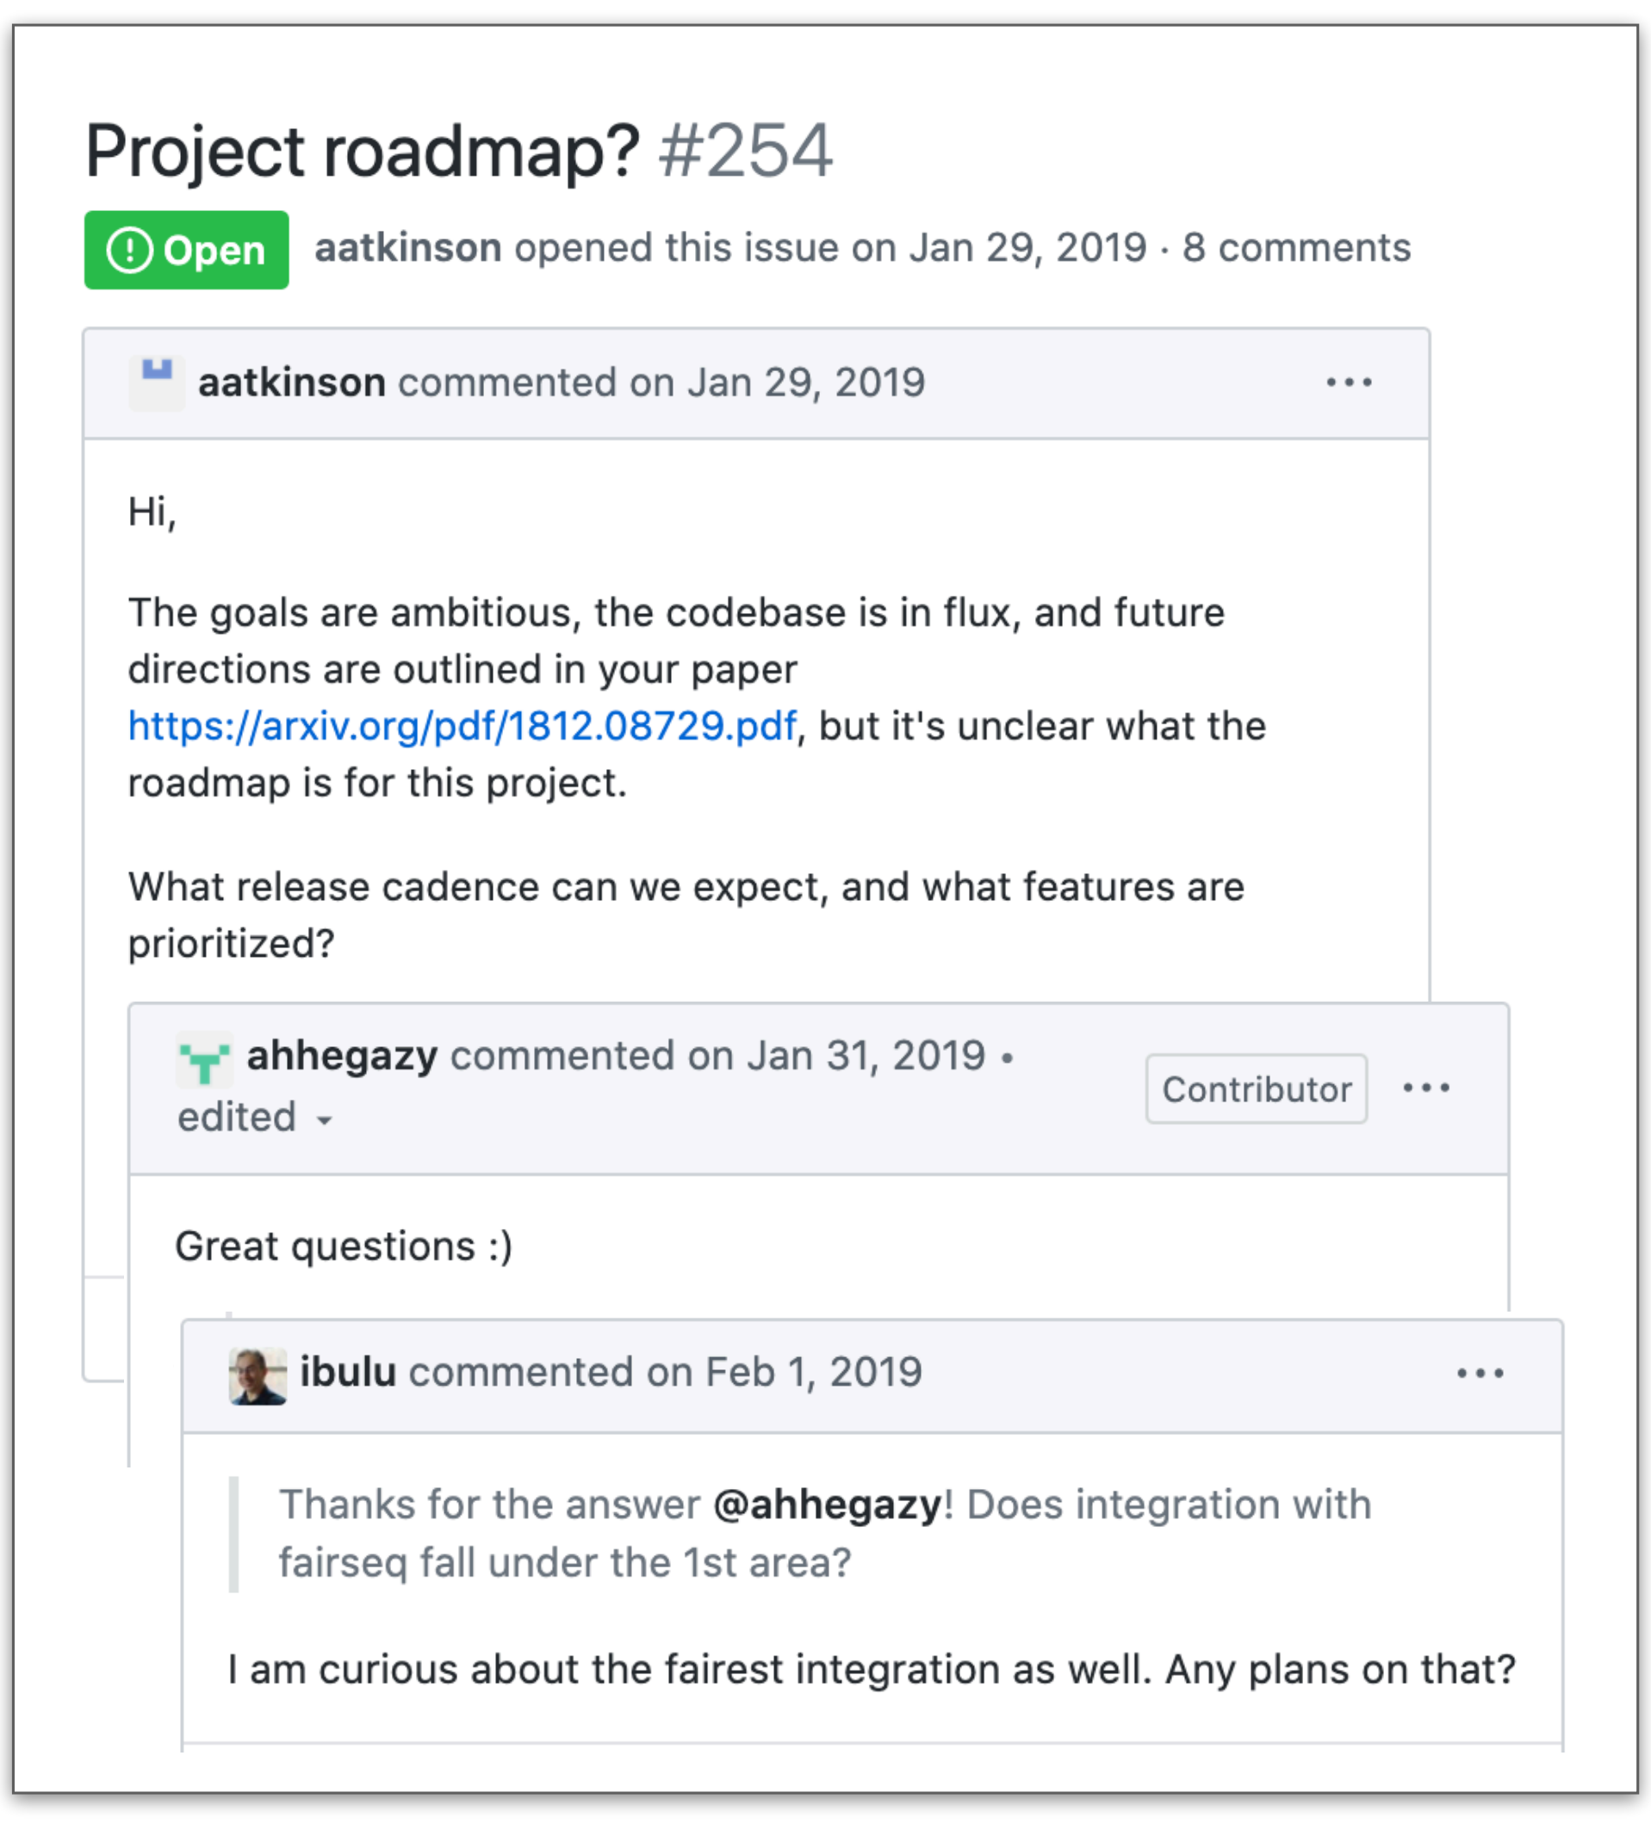

Project roadmap? #254

aatkinson opened this issue on Jan 29, 2019 · 8 comments

aatkinson commented on Jan 29, 2019

Hi,

The goals are ambitious, the codebase is in flux, and future directions are outlined in your paper https://arxiv.org/pdf/1812.08729.pdf, but it's unclear what the roadmap is for this project.

What release cadence can we expect, and what features are prioritized?

ahhegazy commented on Jan 31, 2019 • edited

Great questions :)

ibulu commented on Feb 1, 2019

> Thanks for the answer **@ahhegazy**! Does integration with fairseq fall under the 1st area?

I am curious about the fairest integration as well. Any plans on that?

Verified Verified Verified Verified Verified

Reputation Signals

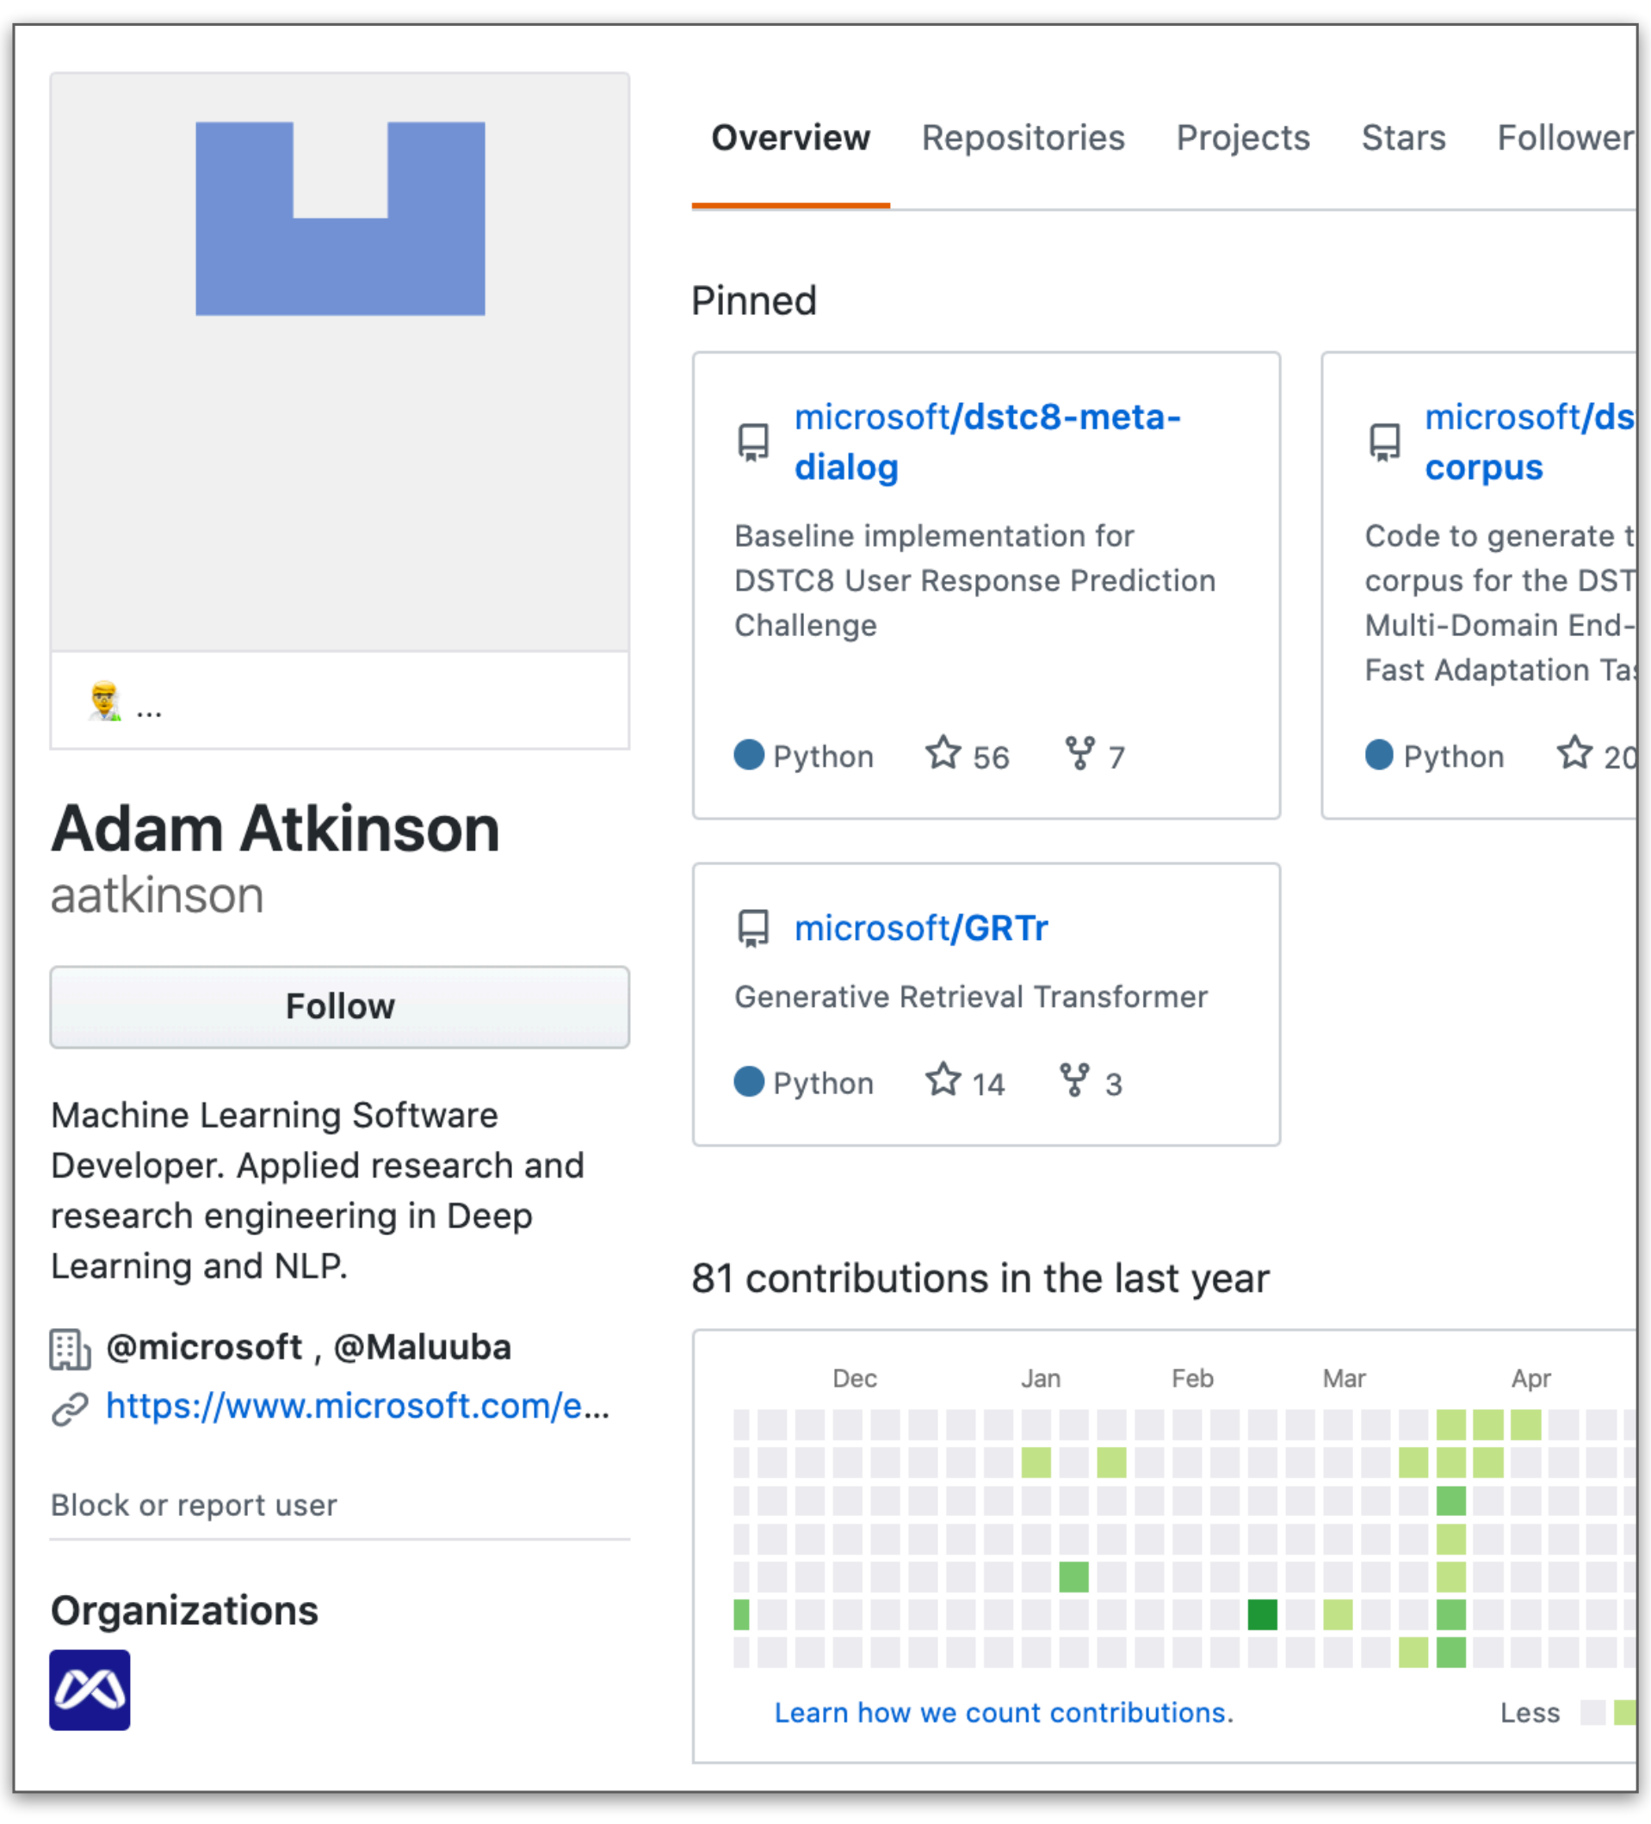

Adam Atkinson
aatkinson

Machine Learning Software Developer. Applied research and research engineering in Deep Learning and NLP.

@microsoft , @Maluuba
https://www.microsoft.com/e...

Block or report user

Organizations

Pinned

microsoft/dstc8-meta-dialog
Baseline implementation for DSTC8 User Response Prediction Challenge
Python ☆ 56 7

microsoft/ds
corpus
Code to generate t
corpus for the DST
Multi-Domain End-
Fast Adaptation Ta
Python ☆ 20

microsoft/GRTr
Generative Retrieval Transformer
Python ☆ 14 3

81 contributions in the last year

(used as heuristics)

Research Question

Q. Does reputation have persuasive power in online deliberation?

A. +10 reputation
→ +26% persuasion probability

(reference cues theory)

Empirical Challenges

I. Identifying Changes in Opinions

Dataset of online deliberation from **ChangeMyView**

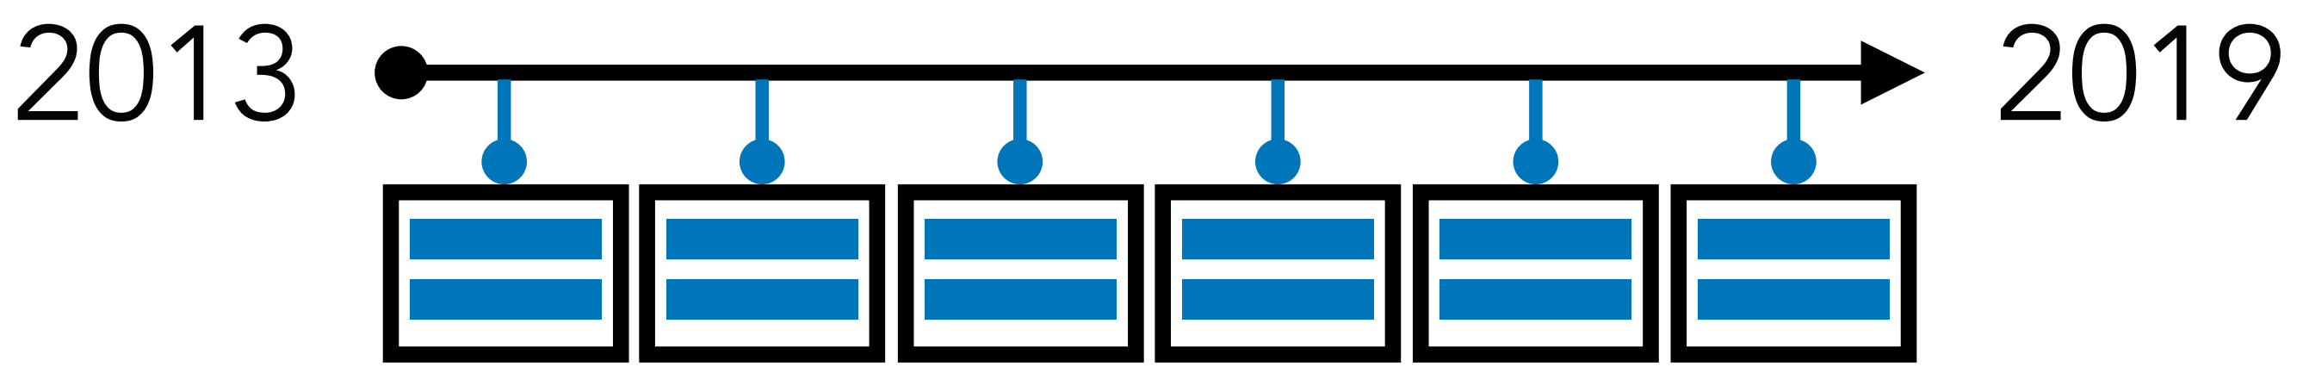

7 years, 1.1 million debates
800,000+ users, 20+ moderators

Δ

Explicit indicator of opinion-change

Reputation ≈ no. persuaded previously

II. Skill & Opinion Selection

A. Exploit multiple debates per user

$$\text{skill} = \frac{\text{no. persuaded previously}}{\text{no. previous debates}}$$

Controls for **time-invariant** challenger characteristics that affect persuasion

B. Exploit multiple debates per opinion

Addresses confounding due to opinion selection

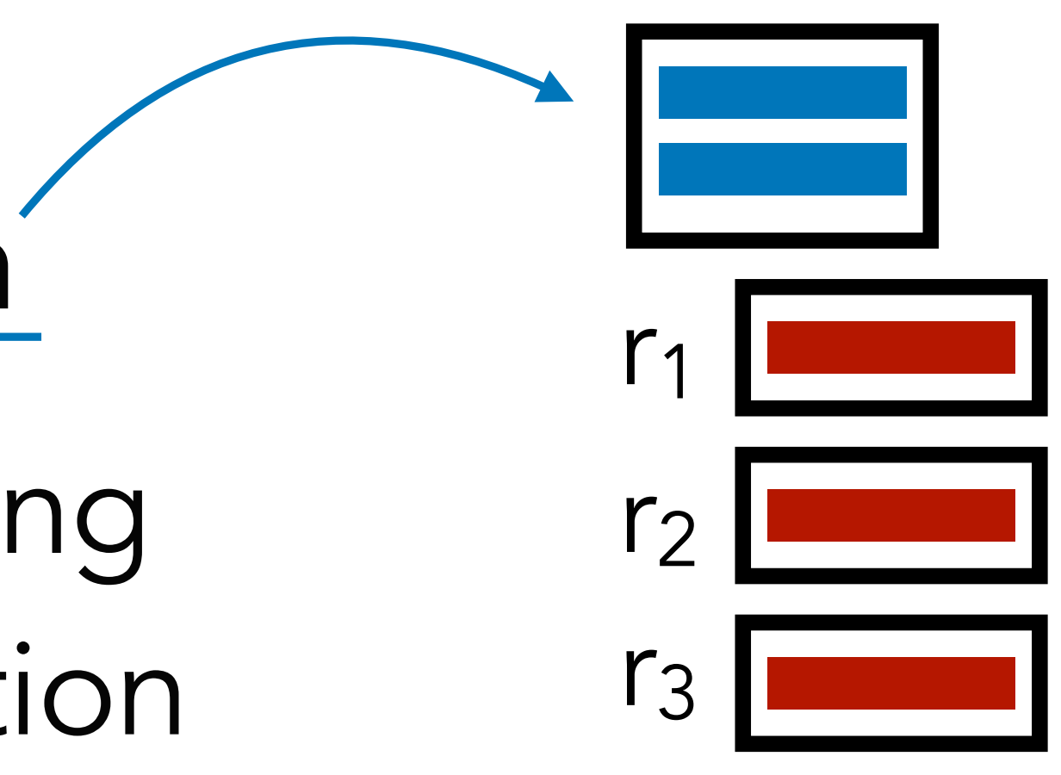

III. Unobserved Confounders

Instrument

Mean past challenger position (F > 1000)

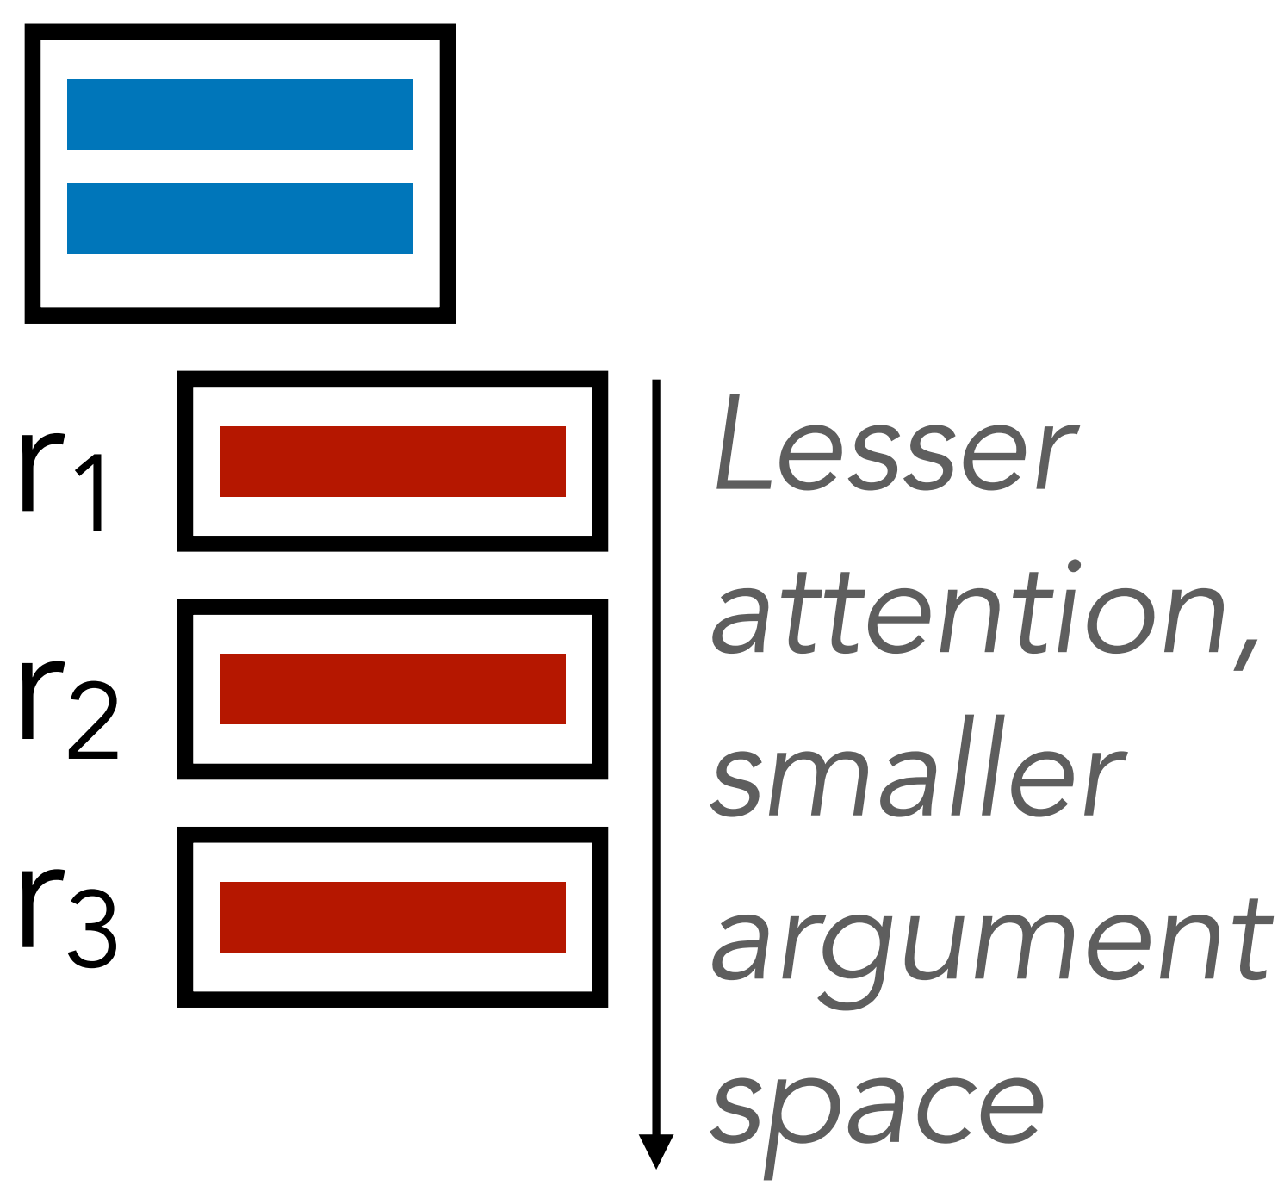

Intuition

- Higher (worse) position → lower persuasion probability
- Reputation ≈ no. of posters persuaded previously

IV. Controlling for Text

Why? Confounders likely to affect the debate outcome through the text of the challenger's response

How? Partially-linear IV model with deep (ReLU) neural networks to model text — estimated via double machine-learning

Consistent estimates + valid inference — in contrast with NLP dimensionality reduction

Main Results

Reputation is persuasive

+10 reputation units → +26% persuasion rate increase over the platform average persuasion rate (≈3.5%)

Heterogeneity patterns consistent with "reference cues" persuasion theory

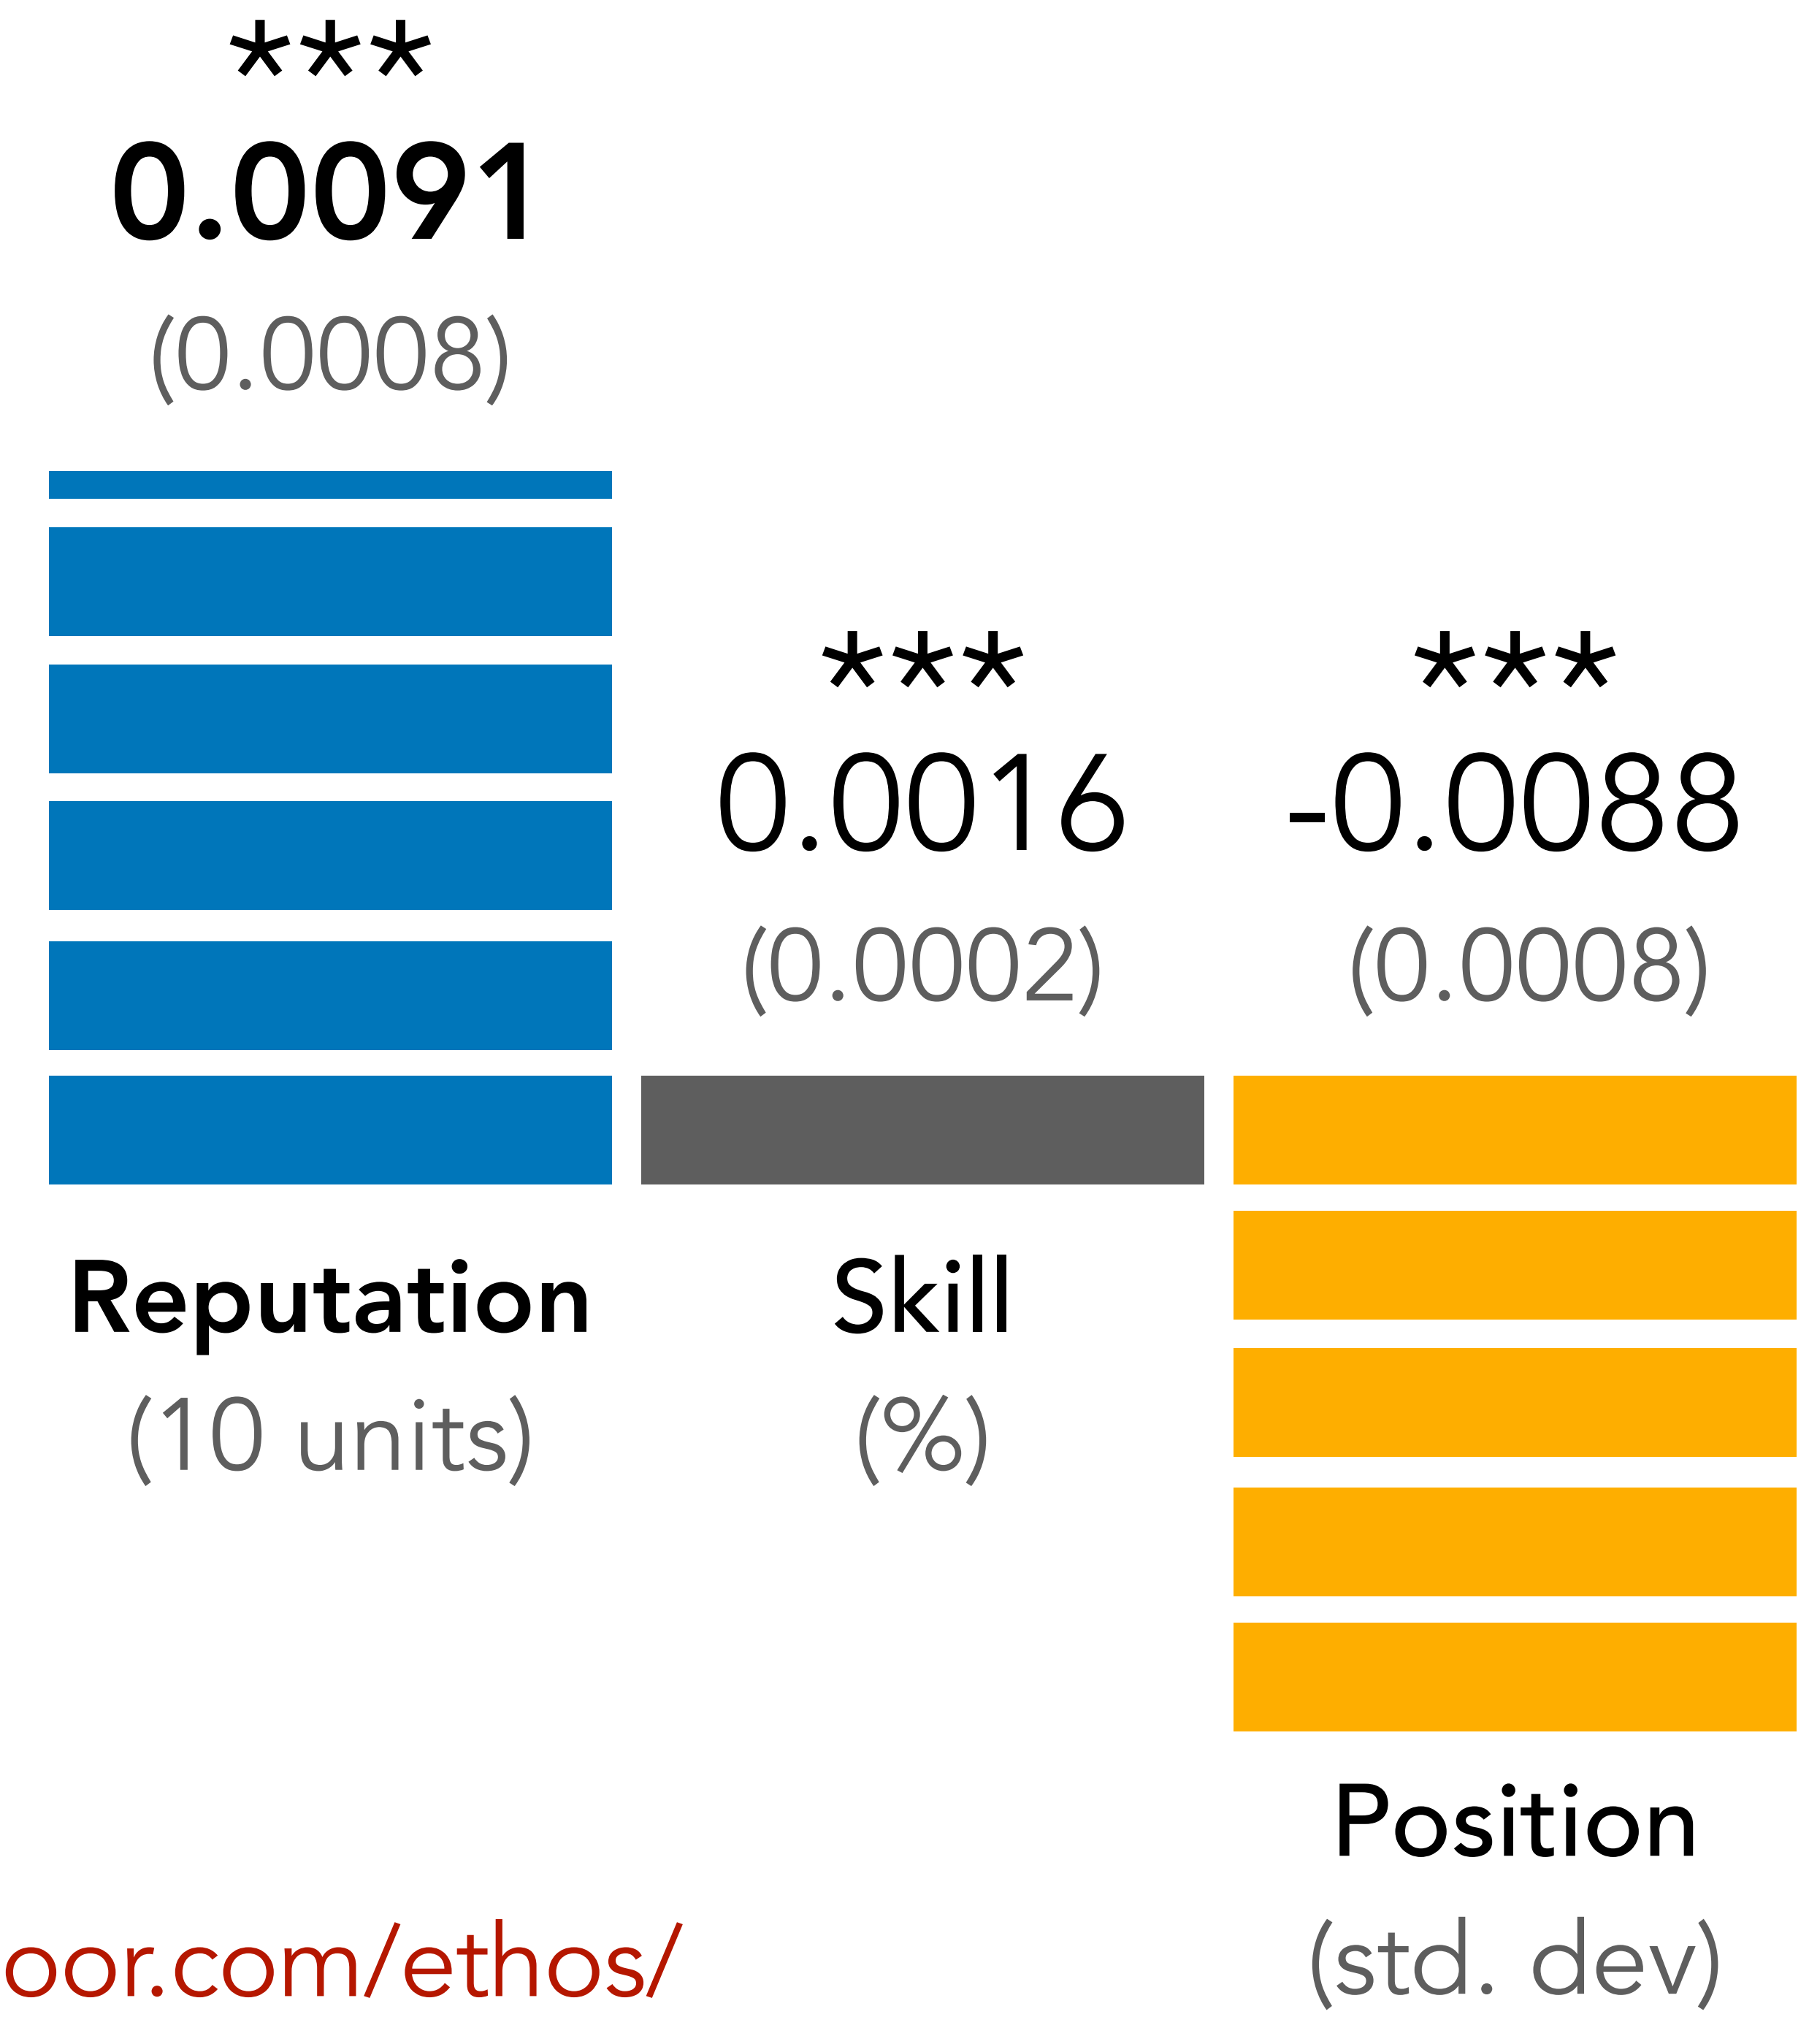

IC²S² Carnegie Mellon University Preprint, slides and video: emaadmanzoor.com/ethos/