Sharad Goel @5harad

In our new @PNASNews paper by @allisonkoe et al., we show that leading speech recognition tools by @Apple, @amazon, @Google, @IBM & @Microsoft misunderstand black speakers twice as often as whites, likely due to lack of diversity in training data. (thread)

fairspeech.stanford.edu
The Race Gap in Speech Recognition
The leading speech recognition tools misunderstand black speakers twice as often as whites.

3:07 PM · Mar 23, 2020

Sharad Goel @5harad · Mar 23, 2020

Automated speech recognition is everywhere, from our pockets to our homes, from the doctor's office to the courtroom. It powers tools that help people with physical impairments use their digital devices. But not all of us can take advantage of this new technology.

Sharad Goel @5harad · Mar 23, 2020

Last spring we ran thousands of short audio snippets of white and black speakers through five leading speech-to-text services. For every 100 words, the systems made on average 19 errors for white speakers compared to 35 errors for black speakers.

Sharad Goel @5harad · Mar 23, 2020

We found similar racial disparities for all 5 of the systems we audited. In every case the speech recognition tools made about twice as many errors for black speakers than for whites. They performed particularly poorly for black men, with more than 40 errors per 100 words.

Sharad Goel @5harad · Mar 23, 2020

These systems also performed far worse for speakers who used more linguistic features characteristic of African American Vernacular English, likely because the systems were trained primarily on audio samples of white Americans.

Sharad Goel @5harad · Mar 23, 2020

To close these performance gaps, we need to create and use datasets that reflect a diversity of accents and dialects. We hope tech companies, academic researchers, and gov't agencies invest the necessary resources to ensure speech recognition tools are broadly inclusive.

Sharad Goel @5harad · Mar 23, 2020

Listen to audio samples, and compare human and machine transcriptions at: fairspeech.stanford.edu.

fairspeech.stanford.edu
The Race Gap in Speech Recognition
The leading speech recognition tools misunderstand black speakers twice as often as whites.

Sharad Goel @5harad · Mar 23, 2020

This work was done by an amazing, interdisciplinary team of linguists, psychologists, computer scientists, and statisticians at @CompPolicyLab: @allisonkoe, Andrew Nam, @emilyrlake, @JoeNudell, Minnie Quartey, Zion Mengesha, @ctoups123, @ayeshazarah_, @jrickford, and @jurafsky.