

Twitter Thread by Michael Auli



Michael Auli

@MichaelAuli



Today we are announcing our work on building speech recognition models without any labeled data! wav2vec-U rivals some of the best supervised systems from only two years ago.

Paper: <https://t.co/cYzF9MGU56>


Blog: <https://t.co/iiGmgdnCiV>

Code:

This shows how completely unsupervised speech recognition with wav2vec-U compares to the best supervised systems on the Librispeech benchmark over the past few years.

Amount of labeled data used

960 hrs.+ ■=1hr.




0 hrs.

Word error rate



Here is how it works:



It also works in languages other than English, see the Swahili demo below. So far we tried it on Kyrgyz, Tatar, German, Dutch, French, Spanish, Portuguese, Italian.

