


Distributed representations of **words** have boomed in the past years improving many NLP tasks:



What could we do with a more refined semantic representation that takes distinct **word senses** into account?



Our method does just so building on

- word2vec,
- SALDO,
- and a bit of statistics.

It's simple, efficient, and shows some neat results.

Drop by to check how it works and see some examples!

Or to complain about how this has nothing to do with word senses.

Do word senses exist at all?

All discussion welcome!