

# **NATURAL LANGUAGE INFERENCE: FOR HUMANS AND MACHINES**

*Valeria de Paiva*

One hears much about the incredible results of recent neural nets methods in Natural Language Processing (NLP). In particular much has been made of the results on the Natural Language Inference (NLI) task using the huge new corpora SNLI, MultiNLI, SciTail, etc, constructed since 2015. The main idea of the NLI task is that to test understanding of the language one checks whether the system can detect entailment, contradiction, or neutrality between pairs of sentences. Thus the sentence "John arrived" entails the sentence "A person arrived", contradicts the sentence "No one arrived" and it is neutral with respect to the sentence "John and Mary arrived", while the sentence "John and Mary arrived" entails "John arrived". Wanting to join in the fun of doing logic with sentences instead of formulae, we decided to check the results on the corpus SICK (Sentences Involving Compositional Knowledge), which is two orders of magnitude smaller than SLNI, but curated by linguists and hence presumably easier to deal with.

We discovered that there were many results that did not agree with our logical intuitions. As a result, we have written several papers on the subject of NLI on SICK. I want to show you a potted summary of this work, to explain why we think the work on NLI is not near completion, yet, despite claims that modern neural nets systems achieve superhuman performance on the existing benchmarks. I plan to also describe how we're tackling NLI, why we think this problem is very important and how we envisage the next steps.