


Why isn't linguistics more useful in NLP?

- Be clear about your goals
 - ▶ science \neq engineering
- More linguistic theory is not necessarily better
 - ▶ features should be *relevant* to task
 - ▶ true $\not\Rightarrow$ useful for some specific task
- Commonplace linguistic insights are often most useful
 - ▶ especially if we want good *average-case* accuracy
- No good models of “world knowledge” or “common-sense reasoning”
 - ▶ but they are necessary for understanding language
 - ▶ formalizing world knowledge is not a goal of linguistics (nor should it be)
 - ▶ lexical statistics (e.g., head-to-head dependencies) are a crude approximation

Lessons from parsing and related applications

- Banal linguistic insights are sometimes very useful
 - ▶ words group together to form phrases
 - ▶ phrases have head words
 - ▶ relative locations within a phrase matter
- Linguists' insights are sometimes useful, but their representations and formalisms are not
 - ▶ machine-learning based approaches map parses to feature vectors
 - ⇒ details of parse representations don't matter (as long as the features can be read off somehow)
- Corpus annotation seems a more economical way of getting linguistic information into a system than writing grammars
 - ▶ linguistic grammars are *closed world* (everything not explicitly permitted is ungrammatical)
 - ▶ stochastic models are typically *open world* (everything is possible)

Accuracy reduction removing a feature class



Doing feature selection well is hard!

