

The Same-head Heuristic for Coreference

Micha Elsner and Eugene Charniak

Do phrases with same head nouns refer to the same thing?

Coreferent

Alice sat on the bank...

same person

Alice thought, "What good is a book without pictures?"

Not coreferent

Her sister was reading a book

not the same!

Most unsupervised work assumes they always do...
True only about half the time!

Results depend on experimental design

- Most unsupervised work:

Gold mentions boundaries:

input: *[Alice] waited a few minutes*

b^3 scoring:

punish errors in large clusters

Maximum loss caused by same-head assumption: 3.3%

- More sensitive to small errors:

Parser mention boundaries:

input: *[Alice] waited a few [minutes]*

CEAF scoring:

all errors equal

Loss caused by same-head assumption is greater: 11.2%

Standard experiments are insensitive to same-head NPs

Modeling improves precision

- IBM model 2:

For **coreferent** phrases: choose antecedent based on syntax

source (possible antecedents):

Alice sat on the bank by her sister.

hidden alignment:

$p(\text{ante}=\text{sister}) \propto \exp(w \cdot f)$

target (context of next NP):

[TARGET] was reading a book...

Distribution over antecedents: log-linear model based on syntactic features
Tree distance, syntactic role, modifiers, etc.

For **non-coreferent** phrases: choose antecedent at random

- Generative mixture model

Unsupervised learning via EM: M-step estimates log-linear model

Model trades recall for precision (marks fewer NPs coreferent)

Standard unsupervised: b^3 prec **+13%**, b^3 rec **-10%**: b^3 F **-13%**

More sensitive to small errors: CEAF **+5%** (reduces loss by half)

Syntactic model for same-head NPs can improve over heuristic