



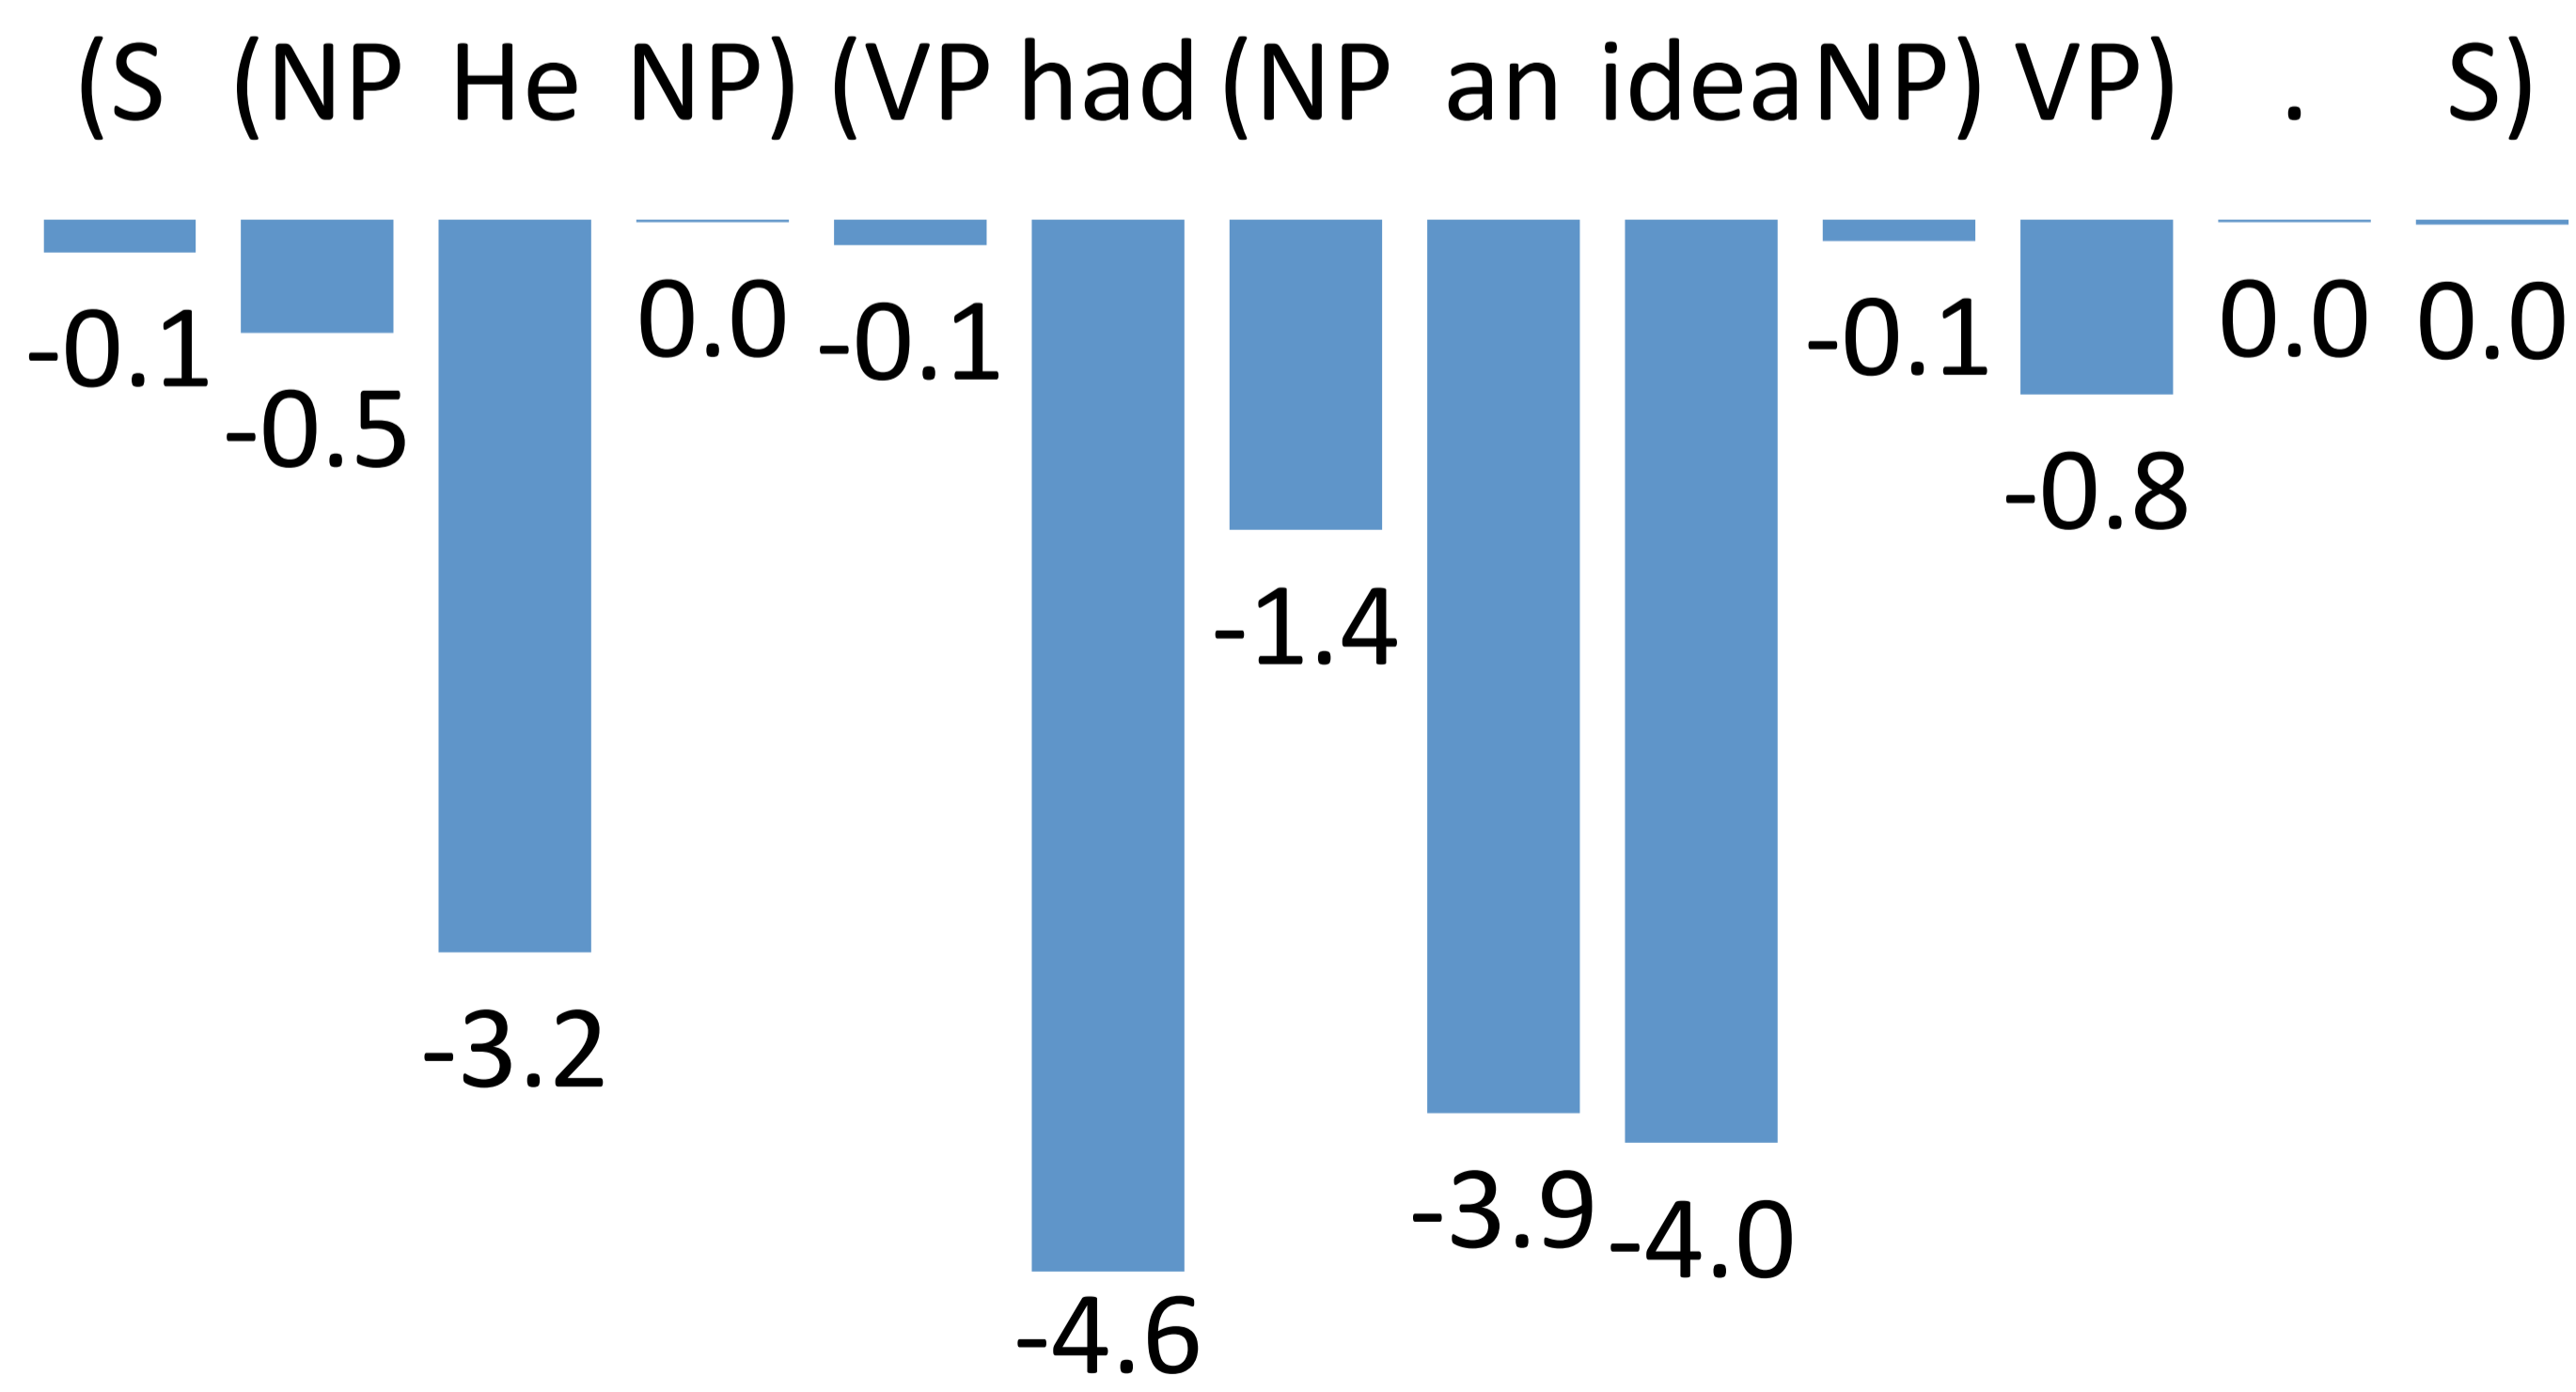

Effective Inference for Generative Neural Parsing

Mitchell Stern, Daniel Fried, Dan Klein · University of California, Berkeley

Overview

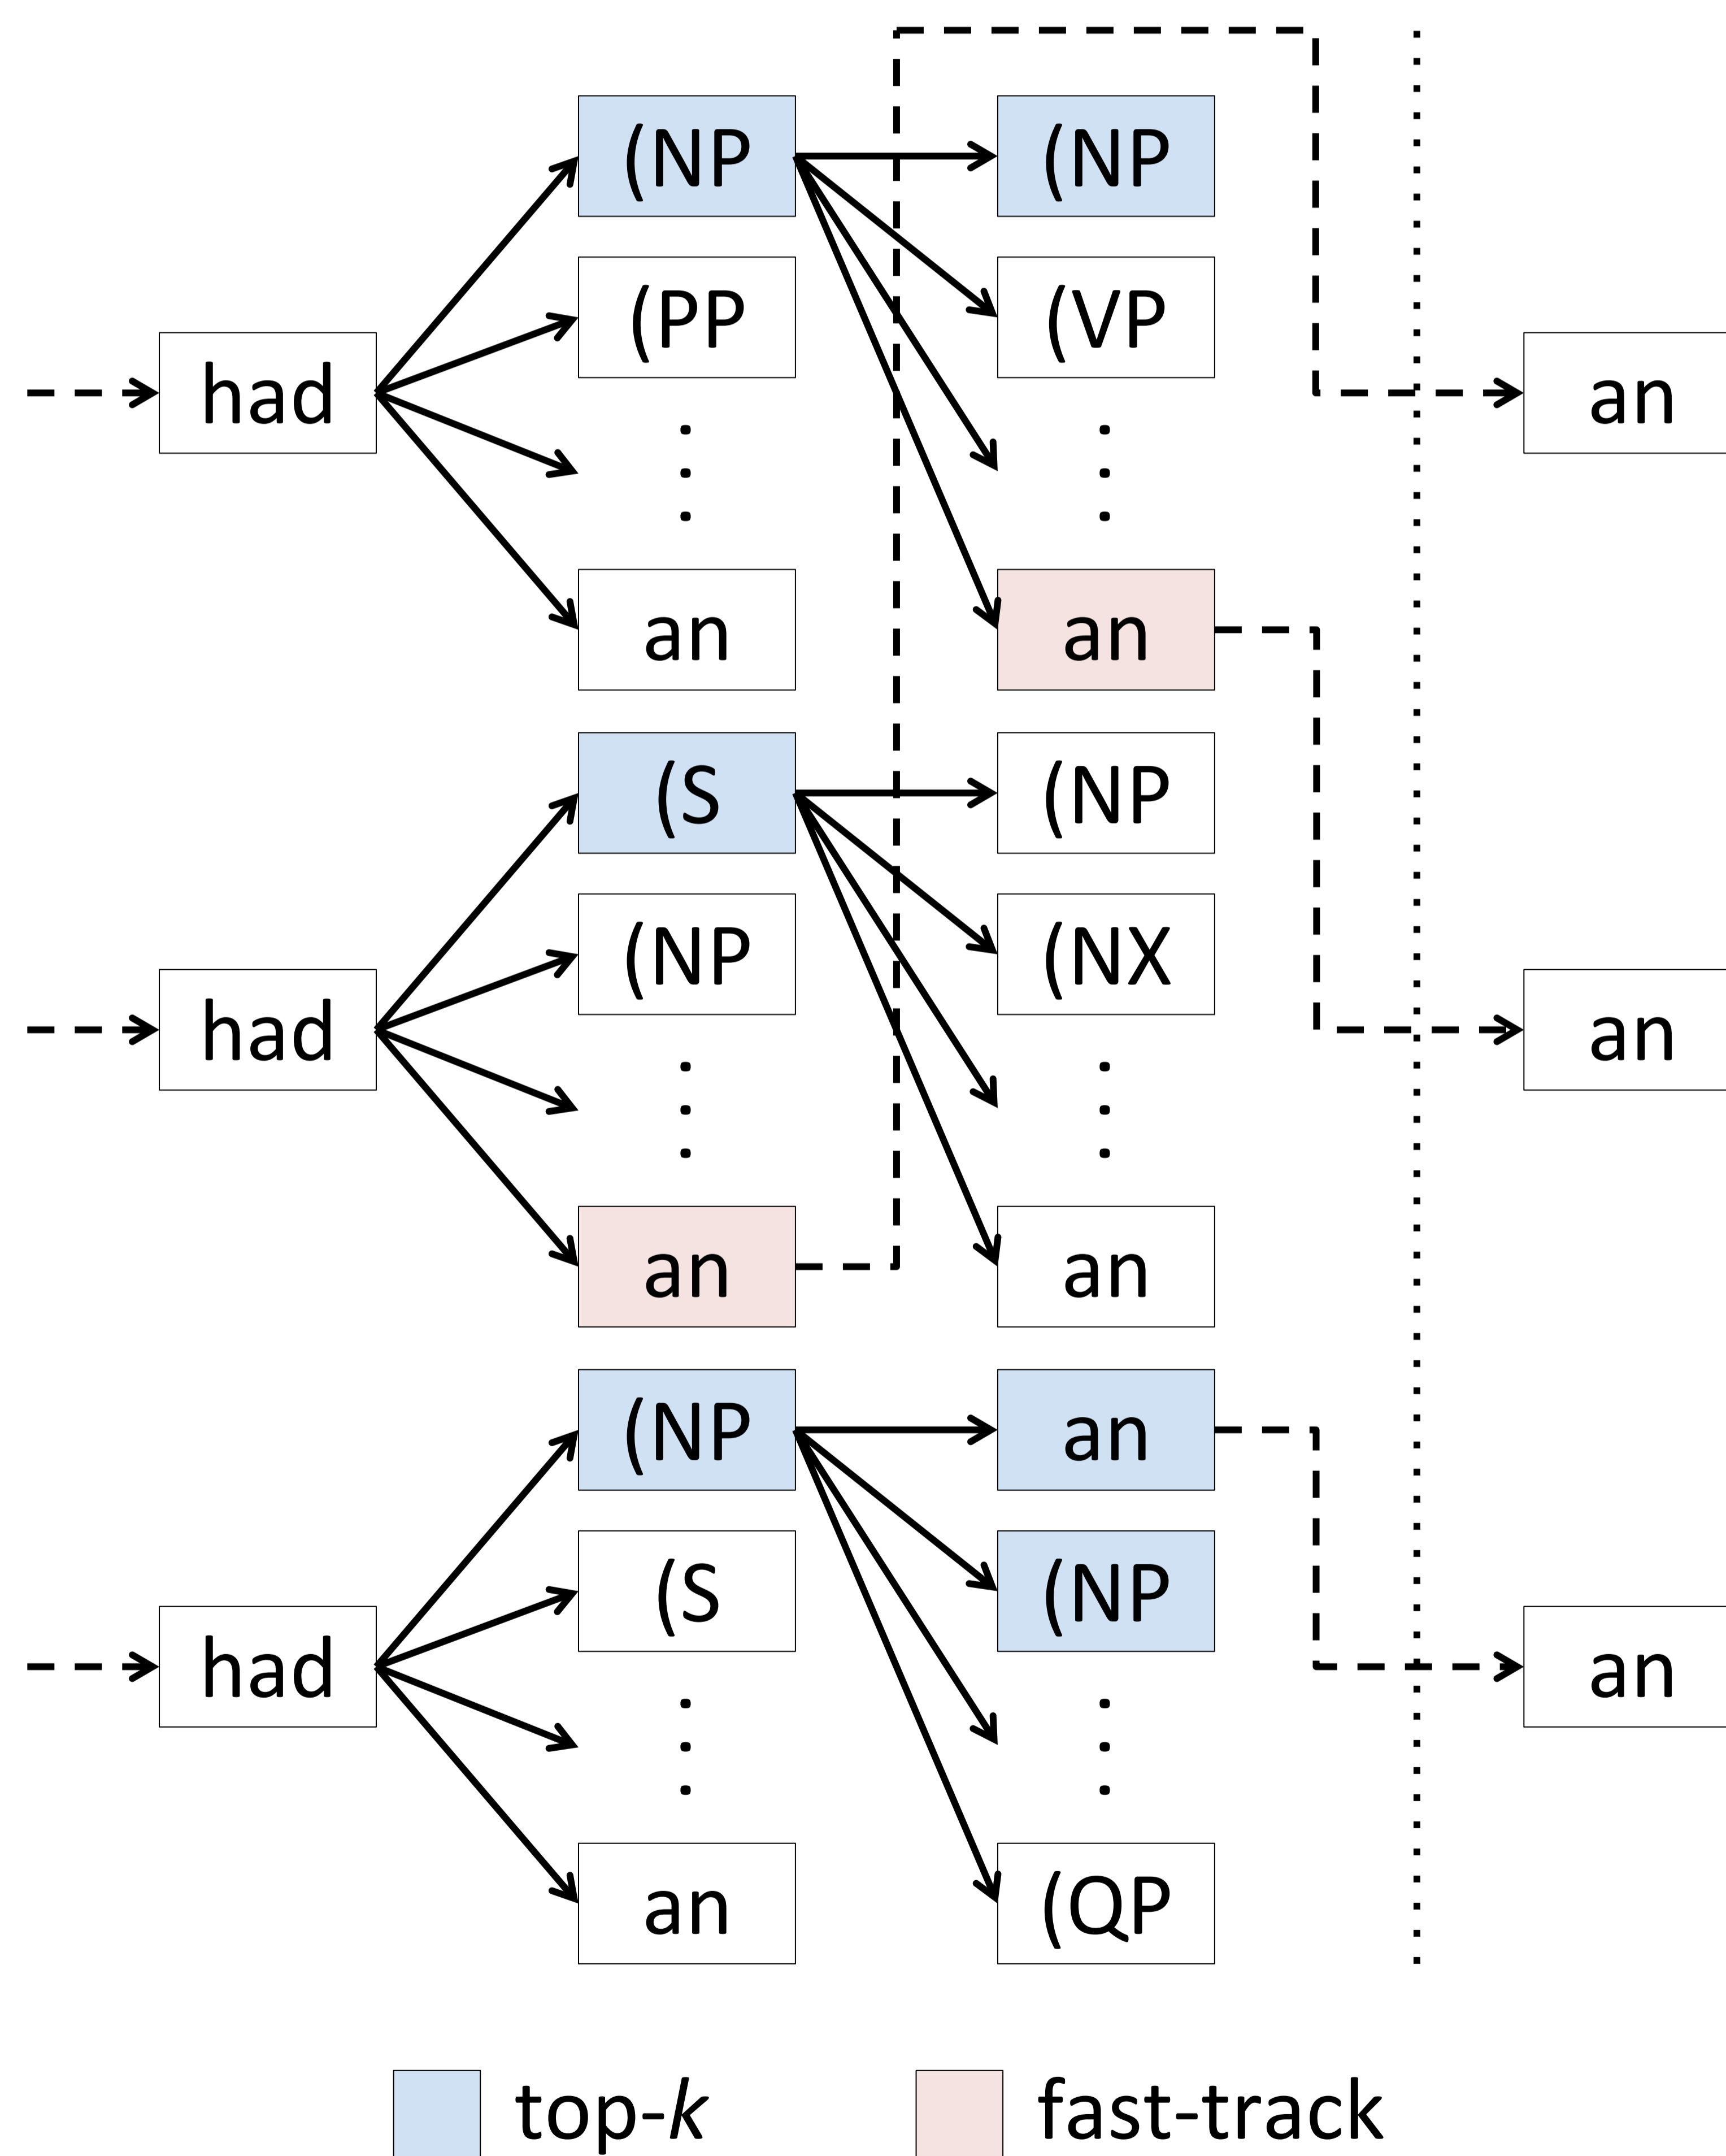
We develop an effective inference procedure for the Choe and Charniak (2016) LSTM language model over linearized parse trees, achieving state-of-the-art single-model performance on the Penn Treebank.

Generative Model



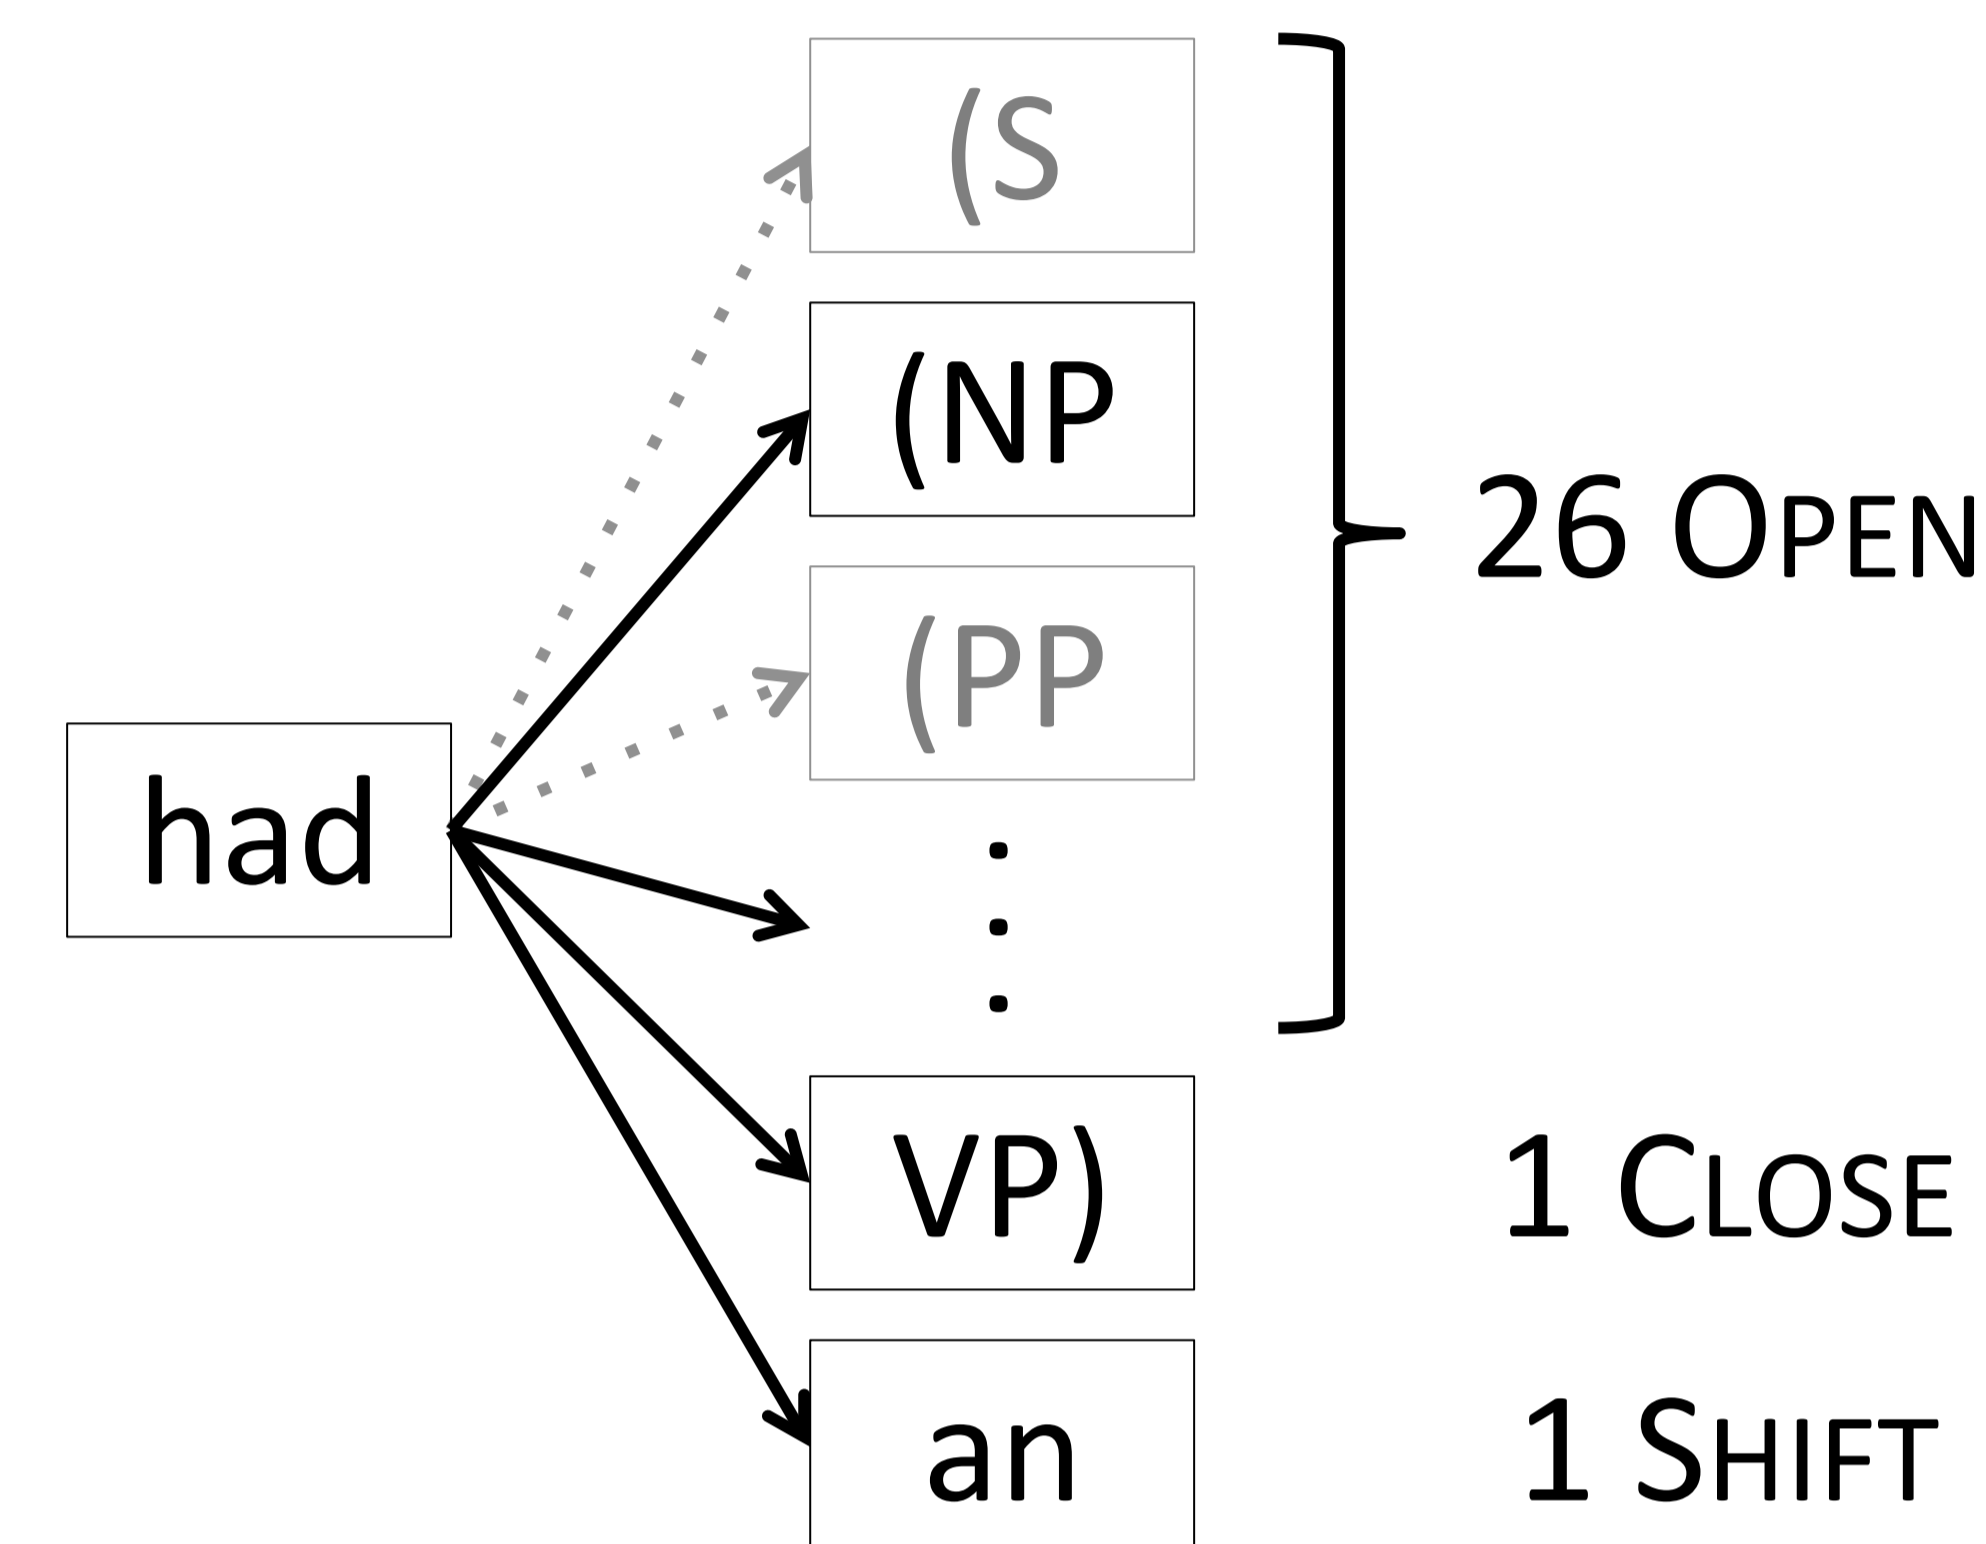
Disparities between the log-probabilities of structural vs. word generation actions cause standard beam search to fail.

Word-Level Search with Fast-Track Candidates



Grouping candidates by the current word addresses the imbalance in probabilities.

OPEN Action Pruning



Using the last 2 actions and the next word, we can prune 70% of OPEN actions with negligible effect on F1.

Final Results

Parser	F1
Vinyals et al. (2015)	88.3
Cross and Huang (2016)	91.3
Dyer et al. (2016)	91.7
Stern et al. (2017)	91.79
Our Best Result	92.56
Our Best Result (with pruning)	92.53
Vinyals et al. (2015) (ensemble)	90.5
Choe and Charniak (2016) (rerank)	92.6
Dyer et al. (2016) (rerank)	93.3
Fried et al. (2017) (ensemble, rerank)	94.25

State-of-the-art single-model performance.