

Precision-biased Parsing and High-Quality Parse Selection

Yoav Goldberg

Ben Gurion University of the Negev
Department of Computer Science
POB 653 Be'er Sheva, 84105, Israel
yoav.goldberg@gmail.com

Michael Elhadad

Ben Gurion University of the Negev
Department of Computer Science
POB 653 Be'er Sheva, 84105, Israel
elhadad@cs.bgu.ac.il

Abstract

We introduce precision-biased parsing: a parsing task which favors precision over recall by allowing the parser to abstain from decisions deemed uncertain. We focus on dependency-parsing and present an ensemble method which is capable of assigning parents to 84% of the text tokens while being over 96% accurate on these tokens. We use the precision-biased parsing task to solve the related high-quality parse-selection task: finding a subset of high-quality (accurate) trees in a large collection of parsed text. We present a method for choosing over a third of the input trees while keeping unlabeled dependency parsing accuracy of 97% on these trees. We also present a method which is not based on an ensemble but rather on directly predicting the risk associated with individual parser decisions. In addition to its efficiency, this method demonstrates that a parsing system can provide reasonable estimates of confidence in its predictions without relying on ensembles or aggregate corpus counts.

1 Introduction and Methodology

Parsing technology has made great progress over the last decade, and current state-of-the-art parsers for English have reported accuracies in the low 90%'s. Current parsing systems are designed to provide a *complete* parse to *every* sentence, and are evaluated based on their average number of correct decisions over a test corpus. Such evaluation, however, does not tell a complete story, as the mistakes are never uniformly distributed across sentences. In

practice, parsers usually perform very well on some sentences, but also perform very poorly on others. In addition, parsers do not provide confidence estimates in their predictions, making it hard for downstream applications to rely on parsers' output, even if the average parsing quality is high.

We advocate a modification to the parsing task, which we term *precision-biased* parsing. Rather than providing a complete parse to every sentence, we advocate providing *partial analyses to some* sentences, while trying to guarantee that the structures that are provided are of high quality. In other words, we advocate parsing systems which are able to trade recall for precision. The trade-off can occur either at the sentence level (abstaining from providing a parse to some sentences) or at the individual attachment level (abstaining from attaching some words or phrases to the rest of the sentence in case the attachment is uncertain).

Such trade-off is useful for tasks which rely on precise structures.¹ These include information extraction and question answering systems, sentence simplification and summarization systems and syntax-based translation, as well as linguistically-oriented tasks such as learning selectional preferences, case frames or lexical ontologies. In addition, partial but precise output may prove useful for self-training (McClosky et al., 2006), uptraining (Petrov et al., 2010) and active-learning setups.

Some previous efforts (Alexander Yates and Etzioni, 2006; Reichart and Rappoport, 2007) attempt

¹The other direction, of trading precision for recall, is also of interest. A solution to this inverse problem is already available to some extent in the form of k-best parsing and packed forests.

to identify high-quality parse trees among parser’s output. This is an instance of trading recall for precision at the sentence level. Here, we focus on trading recall for precision at the individual attachment level. As discussed in section 3, solving the problem at the individual attachment level entails a natural solution to the problem at the sentence level as well. Moreover, we believe there is benefit in solving the problem at the attachment level – useful information can be extracted from a partial parse tree, even when some attachment decisions are missing or marked as unreliable.

The precision-biased task was explored in the past in the context of parsers based on manually-developed grammars (Carroll and Briscoe, 2002; Watson et al., 2005). However, state-of-the-art data-driven statistical parsers do not allow trading recall for precision. In this paper, we focus on data-driven dependency parsing.

We begin by defining the precision-biased parsing task and its evaluation measures (Section 2), and provide a strong baseline based on parse-ensembles (Section 5). While effective, the ensemble system requires substantial computational effort, and is of little theoretical interest as it is well known that committee (dis)agreement is a good indicator of confidence. We propose another method based on parser-modeling in Section 6. We train a probabilistic classifier to try and predict the risk associated with attachment decisions in the parser’s output. The classifier learns the error patterns of the parser, and assigns a reliability score to parse edges. By thresholding these reliability scores, we can effectively trade recall for precision. The method comes close to the ensemble-based baseline in terms of precision and coverage, while running faster and providing a straightforward way to control the recall/precision tradeoff. More importantly, it demonstrates that reasonable confidence estimates of the correctness of parser predictions can be attained without relying on committee, diversity, or aggregate corpus-based counts of recurring structures. Inspecting the behavior of the learned model on PP-attachment, a classic case of syntactic ambiguity, reveal that it is not judged by the model as being categorically hard. Rather, some PP-attachment cases are marked as unreliable, while others are not.

2 Precision-biased Dependency Parsing

In the traditional dependency parsing task, the input is a set of sentences (the test corpus), and we are interested in the most precise analysis for each sentence. The task performance is measured based on the average number of tokens that got assigned a correct parent over the entire test corpus. Crucially, the parsing process in this task must assign a parent to each of the tokens in the test corpus.

In contrast, in the *precision-biased* task we are concerned with precision more than recall. We allow the parsing system to *abstain* from providing an analysis to some of the input by *skipping* some decisions. That is, we require the parser to assign parents to as many of the input tokens as possible, but allow it to leave the parents of some tokens unassigned.

Metrics Let T be the set of input tokens, A be the set of tokens that got assigned a parent, and S be the set of tokens that were not assigned a parent ($T = A \cup S$, $A \cap S = \emptyset$). Let $C \subset A \subset T$ be the set of tokens that got assigned a correct parent. Then:

$$precision = \frac{|C|}{|A|} \quad recall = \frac{|C|}{|T|} \quad coverage = \frac{|A|}{|T|}$$

By requiring complete coverage of the input tokens ($coverage = \frac{|A|}{|T|} = 1$) we get $A = T$, and then $precision = recall = accuracy$, where *accuracy* is the traditional dependency parsing accuracy.

The precision-biased setting allows $coverage < 1$. The aim is to maximize *precision*, while still retaining sufficient *coverage* of the input tokens.²

Parse-selection Previous work addresses the *parse-selection* task: selecting a subset of the input sentences for which we have high-accuracy parses. This is an instance of precision-biased parsing in which abstaining on a token requires abstaining also on all the other tokens in the same sentence. When discussing the parse selection task we distinguish between *token coverage* which is identical to coverage as defined above, and *sentence coverage* which is the number of selected sentences divided by the total number of input sentences $\frac{|\{sent | sent \subset A\}|}{|\{sent | sent \subset T\}|}$. The definitions of precision and recall remain as above.

²We chose to balance *precision* against *coverage* rather than against *recall* because *coverage* is an upper-bound on *recall*: if $precision = 1$ then $coverage = recall$.

3 Our Approach

We tackle precision-biased parsing by defining a *riskiness* function on individual attachment decisions. The riskiness $R(tok, par)$ of a token/parent pair is the inverse of our confidence in the attachment decision. A high riskiness indicates uncertainty in the decision, and low riskiness indicates that we believe the decision to be correct. We then set a riskiness-threshold, and abstain from any parent assignment for which the riskiness is above the threshold. Setting a low riskiness-threshold results in higher precision (considering even low-risk decisions as too risky), and setting a high riskiness-threshold results in higher coverage.

Parse-selection Having defined a risk function and a risk-threshold, we get a natural selection criteria for the parse-selection task: high-quality parses are those for which at most K attachments are above the riskiness threshold. The precision/coverage balance can be controlled either by changing the risk threshold, or by changing K .

3.1 Related Work

While little research attention was dedicated to the precision-biased task (Carroll and Briscoe, 2002; Watson et al., 2005), several studies address the parse-selection task. Yates et.al. (2006) perform parse-selection by filtering out parses containing “semantically implausible” relations, where semantic-plausibility is estimated by high co-occurrence of the words in relation in a large corpora (*i.e.*, the web).

Reichart and Rappoport (2007) perform committee based selection of high-quality constituency-parses by calculating an agreement measure between 20 copies of a lexicalized parser trained on different subsets of a training corpus.

Sagae and Tsujii (2007) select high-quality dependency parses by using two dependency parsers and selecting only sentences on which both parses agree on the entire parse.

Kawahara and Uchimoto (2008) identify high-quality dependency parses by training a classifier based on sentence level features: sentence length, average dependency length, number of unknown words, number of commas and conjunctions, and corpus frequencies of sentence words.

Reichart and Rappoport (2009) identify high-quality parses of an unsupervised parser by looking for parses with many reliable constituents, where reliability of a constituent is calculated based on the number of times its POS-sequence appears in the automatically parsed text. Finally, Dell’Orletta *et.al.*, (2011) assign quality-scores to dependency-parses using a metric which measures various syntactic properties of the parse tree and compares them the aggregate measurements over the entire parsed corpora.

To summarize, there are three lines of work addressing the parse-selection task: selecting parses based upon agreement between a committee of parsers, selecting parses based on agreement between the parses and aggregate counts over a large corpora (either lexicalized “semantic” agreement or syntactic agreement), and selecting parses based on sentence-specific features (length, vocabulary, number of commas and so on).

In contrast, we are primarily interested in selecting high-quality edges rather than complete parses. We view parse-selection as an extension of the precision-biased parsing task, and perform parse-selection based on the number of risky attachment decisions. Our assessment of the riskiness or reliability of a particular decision is not based on aggregate corpus counts nor on global features of the input sentence (though such kinds of information may be integrated in the future). In our first method, we adopt a committee-based approach, but apply it primarily for edge-selection. In the second method we present below, we investigate features which may help the parser assess edge riskiness. We note that the marginal edge probabilities obtained from a log-linear parsing model as in (Smith and Smith, 2007) are not reliable predictors of edge riskiness: indeed, the pruning procedure used in (Carreras et al., 2008) consider edges with marginal scores of up to 10^{-6} of the highest scoring edge as possible candidates in order to ensure sufficient coverage, indicating that such models may greatly overestimate the marginals of incorrect edges, while underestimating the marginal values of correct edges.

Kawahara (2001) present an automatic method for Case-Frame dictionary construction for Japanese. Their method identify verb case-frames by identifying reliable syntactic constructions, where the re-

liability is learned using a hand-crafted heuristic and aggregate corpus counts. This demonstrates the usefulness of identifying reliable instances of specific constructions. Our proposal is to try and identify reliable instances of many different constructions, without relying on hand-crafted heuristics.

4 Data

Our experiments are based on the dependency-version of the Penn WSJ corpus, as converted using the Penn2Malt³ software with Collins’ head-rules. The data is POS-tagged using the HMM-based Hunpos tagger⁴.

While our work is not directly comparable to any previous work⁵, this presents an opportunity to stop following the standard train/test/dev splits, and in particular to stop testing only on section 23. Instead, we adopt a setup in which we use sections 2-11 (about 18k sentences) for training the parser(s), sections 12-15 (8900 sentences) for training the riskiness-estimator (where appropriate), section 16 (2780 sentences) for development and sections 17-21 (9500 sentences) for testing. This setup leaves a reasonable amount of training data for the statistical models (the parser training set is roughly the same size as the one used in the CoNLL shared task), while retaining a much larger test set than the standard setting.

5 Parser-ensemble Riskiness Estimation

Our first method of estimating the riskiness function is using an ensemble method. Ensemble methods have been shown to provide good results for dependency parsing (Sagae and Lavie, 2006; Hall et al., 2007), as well as for parse-selection as discussed above. Here, we use ensembles to estimate the riskiness of individual edges in a dependency tree. To estimate the riskiness, we parse the input sentence using k different parsers and take the intersection of their predictions. The riskiness of a token/parent pair is 0 if all k parsers agree on that prediction, and 1 otherwise. In the final output, we take only edges with 0 riskiness.

³<http://http://w3.msi.vxu.se/~nivre/research/Penn2Malt.html>

⁴<http://code.google.com/p/hunpos/>

⁵We are not aware of previous work on the precision-biased task, while previous work for the parse-selection task either focus on constituency-structures, or use non-standard datasets.

We use an ensemble of 3 parsers: a linear-time shift-reduce parser as described in (Huang et al., 2009) (SHIFTREDUCE), the globally optimized first-order projective dependency parser of (McDonald et al., 2005) (MST1), and the easy-first parser of (Goldberg and Elhadad, 2010) (EASYFIRST). Such ensemble was shown in (Goldberg and Elhadad, 2010) to provide good oracle accuracies, as well as state-of-the-art accuracies in a non-oracle setting due to the diversity among its parsers. The runtime of this ensemble is dominated by the $O(n^2)$ feature extraction stage and the $O(n^3)$ inference of the globally optimized MST1 parser.

5.1 Results and Discussion

The individual parser’s scores on the test set are 87.4 (SHIFTREDUCE) 88.6 (MST1) and 88.4 (EASYFIRST).

Precision-biased Scores The precision-biased scores of the ensemble system on the test-set are 96% precision with a coverage of 84.2% (recall of 80.8%). By not providing an analysis for about 15% of the input tokens, we get an impressive gain in precision.

Parse-selection Scores As discussed above, we reduce parse-selection to risk-based precision-biased parsing by selecting parses with at most K risky attachments. Table 5.1 shows the precision and sentence-coverage on the development and test set for various values of K . With a K value of 0 (forcing the parsers to agree on all edges) achieves a precision of 97.5 while covering about a quarter of the sentences in the test set. By allowing one disagreement between the parsers the precision drops to 95.0, but we gain a better sentence-coverage – about 36%. Increasing the value of K decreases the selected parses accuracy while increasing their quantity.

K	Precision (%)	Sentence-Coverage (%)
	dev / test	dev / test
0	97.5 / 97.8	23.2 / 24.6
1	95.0 / 95.6	36.2 / 36.8
2	93.0 / 93.4	47.2 / 47.3
3	91.0 / 91.3	56.9 / 57.1
4	89.4 / 89.5	64.5 / 66.0

Table 1: Ensemble-based Parse-Selection Precision and Coverage for various risk-cutoffs (K)

6 Single-parser Riskiness Estimation

While the ensemble method is effective at the precision-biased task, it has two shortcomings: (1) it takes a long time to run due to the runtime complexity of the MST1 parser, and (2) it does not provide a way of tuning the precision/coverage balance.

Here we take a different route. We use a single parser (we use the EASYFIRST parser for its balance between speed and accuracy and its incremental parse construction), and train a discriminative probabilistic classifier to predict the risk associated with its predictions.

EASYFIRST is a greedy parser that work by incrementally adding dependency edges in a bottom-up fashion (see (Goldberg and Elhadad, 2010) for the details). It is trained to take easy decisions before harder ones, but does not provide confidence in its predictions: at a given step, the highest scoring action can still be very ambiguous, yet easier than the alternatives. The two measures of the best possible (“easiest”) action and the riskiness of an action are interrelated, but not identical. The easiest action, at a given stage, may still be risky. For example, consider the case in which the parser sees a configuration consisting of [Verb Noun Prep]. This is a PP attachment ambiguity, where the Prep should be the child of either the Noun or the Verb. Concretely, the parser should choose to either attach Noun under Verb and then Prep under Verb+Noun, or to first attach Prep under Noun and then Noun+Prep under Verb. At this stage, the two possible attachments (Verb+Noun and Noun+Prep) are risky (even though the Verb+Noun edge will turn out in the final parse in any case), but the parser should nevertheless choose one of them. It will choose, based on its training experience and on the specific properties of the VP, NP and PP at hand, the action which it finds is most correct. This would be the least-risky attachment *at a given stage*, but it does not reflect directly on the objective riskiness of the decision at large.

In the precision-biased setting, we are interested in assessing the objective riskiness of various parser decisions.

Riskiness Predictor We train a separate classifier to assess the riskiness involved in each prediction. We interpret the riskiness as a probability function: $risk(context) = Pr(decision\ is\ wrong|context)$.

In words: the riskiness is the probability of the parser making a wrong choice in a given situation (context). The riskiness function does not necessarily depend on the actual decision (*i.e.*, it should be interpreted as “when faced with situation X you are likely to make a mistake” rather than as “attaching t_i below t_j in situation X is likely to be wrong”), but the decision can be encoded in the context if desired.

We treat riskiness prediction as a binary classification task, and fit a Maximum Entropy model⁶ based on training data as described in Section 7.1 below.

We experimented with several alternative interpretations of the riskiness function, capturing different kinds of information (features).

Riskiness of parser actions An interesting question that arises is whether the information available during parsing is sufficient for determining the risk associated with a parsing decision, and which kinds of information are most useful.

The first set of experiments attaches risk to *parser actions*. These aim to answer the question “can the parser assess the quality of its own actions”. Note that parser actions are not equivalent to attachment decisions: the easy-first parser may choose to attach a token to its correct parent and still be wrong, because it is not yet the correct time to do so (because the child node is not yet saturated), and this action, while resulting in one correct edge, prevents future correct edges from being added (consider the Verb+Noun edge in the PP-attachment example above). Thus, this set of experiments can be used only for the parse-selection task (selecting as good parses those for which there were less than K risky actions), and not for the precision-biased parsing task.

We experimented with the following feature sets: *Process-based features*: `action_process` is a minimal set of 5 numerical features which relate only to the parsing process itself. These include: sentence-length, current number of parent-less tokens, score of the best action (to be applied), score of the second best action, and the difference between the best and second-best actions.

State-based features: `action_state` is a set of features which relate only to what the parser sees.

⁶We use the Megam optimizer (Daumé III, 2004)⁷

Here, we use the exact same feature set which is used by the parser for predicting the scores of the various actions.

Riskiness of predicted edges The second set of experiments associates riskiness with edge predictions. That is, riskiness is interpreted as “what is the probability of this particular predicted edge to be wrong”. In contrast to the previous experiments, this definition of riskiness addresses the full precision-biased parsing problem, by abstaining from providing (or ruling-out) attachments for edges that are considered too risky.

We experimented with the following feature sets: *State-based features*: as above, the `edge_state` feature set encodes exactly what the parser sees when making an attachment decision, *i.e.*, the feature set used by the parser when scoring actions. However, here wrong parser actions which result in a correct edge are considered as correct (non-risky) examples.

Edge-factored features: the `edge_factored` feature set is not related to the parsing process, and can be extracted from the parse tree in a post-processing step. Here, we use the same features as used in Ryan McDonald’s first-order edge-factored MST parser (McDonald et al., 2005).

Higher-order features: the `edge_higher` feature-set does not depend on the parsing process, and uses more information than the edge-factored one: the features of a (token,parent) pair include information on the token and the parent, as well as on the siblings of the token, siblings of the parent, children of the token, and parent of the parent.

`edge_state` has only a negligible effect on the parsing time (as above, the features are already extracted by the parser), while `edge_factored` and `edge_higher` have a noticeable (though still small) effect on the parsing time by adding n feature extraction and scoring operations.

7 Experiments and Results

7.1 Training

We followed the following procedure:

1. Train the easy-first parser on the parser-training set, and use it to parse the rest of the data (riskiness-training, test and dev, see Sec. 4) while keeping

track of the parser’s predictions.⁸

2. Extract correct and incorrect decisions and their corresponding features (according to the definitions above) from the automatically parsed data.

3. Train a MaxEnt binary classifier on the riskiness-training set.

7.2 Evaluation

ROC We begin by plotting the ROC curves for identifying risky decisions using the different MaxEnt risk predictors with varying risk-thresholds. Figure 2 presents the results. The curves are not entirely comparable: the two lower curves (`action_process` and `action_state`) identify risky parser actions, while the higher curves identify risky edge attachments. There is a clear hierarchy between the different predictors, but even the simplest ones are quite effective at identifying risky decisions. The predictors that attach riskiness to edges are more effective than those that attach riskiness to parsing actions, even when the same feature-set is used (`edge_state` vs. `action_state`), and the two predictors that use information external to the parser (`edge_factored` and `edge_higher`) are better than those using information internal to the parser. Still, it is interesting to note that the exact same feature-set which is available to the parser during parsing is sufficient to assess the riskiness of many of the decisions which are based on the same feature-set.

There is only a small difference between the `edge_factored` feature-set and the one including higher-order features (`edge_higher`): while the extra contextual information does help, most of the riskiness associated with an edge can already be determined based on the edge itself and sentence-level properties (without considering proposed surrounding edges).

Precision-biased Scores We now turn to evaluating the precision-biased results for the various predictors. As discussed above, these can be calculated only for the three edge-based riskiness predictors. For the precision-biased results, the parser abstains from predicting edges with an associated risk-

⁸In case of a smaller treebank, a k-fold jackknifing scheme should be used.




Figure 1: Precision-biased results on the dev-set.

ness above a certain riskiness-threshold. Figure 1 plots the precision and coverage of the parser for varying riskiness-thresholds, using the three different riskiness-predictors. The third plot in the figure plots precision against coverage for the same predictors. The overall trends are similar to those observed in the ROC curves, though here the difference between `edge_factored` and `edge_higher` is somewhat more pronounced.

With appropriate riskiness thresholds we could achieve a coverage as high as 95%, or precision of above 97%. Unfortunately, we cannot get both: higher precisions mean lower coverages and vice-versa. Compared to the ensemble-based riskiness estimation (96% precision with 84% coverage) the single-parser results are not as strong. The same level of coverage (84.6%) results in a precision of around 94.2%, and a precision of 96.5% leads to a coverage of 70%. A riskiness-threshold of 0.15 (`edge_higher` predictor) strikes a nice balance of just over 95% precision with 80% coverage. A coverage of 90% gets us a precision of above 92%, still substantially higher than the 88.4% of the full-coverage baseline parser. The numbers are practically the same for the development and test sets.

Parse-selection scores The parse-selection task has two tunable parameters: the riskiness-threshold R , and the number of risky decisions (K) above which we regard the complete parse-tree as unreliable. For each predictor we performed a grid-search over these parameters with R ranging in value from 0 to 0.5 with increments of 0.01, and K ranging from 0 to 4. For each point we recorded the precision and the sentence-coverage over the development set. We then chose the parameters yielding the




Figure 2: ROC Curves for the various riskiness predictors on the dev set. The True Positive Rate is the percentage of incorrect decisions which were identified as risky, and the False Positive Rate is the percentage of correct decisions which were identified as risky.

best sentence-coverage for a given precision level (we varied the precision levels from 89 to 99 with increments of 0.5). We processed the test-set with the selected parameter-values. Figure 3 plots precision against sentence-coverage on the test-set using the (K, R) obtained on the development set.

8 Riskiness of PP-attachments

Having developed a model of assessing the riskiness associated with parsing decisions (i.e., the chances of a certain decision being wrong), what does it find




Figure 3: Parse-selection (precision vs. sentence-coverage) results on the test sets based on best (K,R) values obtained on dev set.

to be risky? A complete analysis of the behaviour is well beyond the scope of this paper, but we provide a glimpse into what is possible by inspecting the model’s behaviour on a particular case of syntactic ambiguity: PP attachment. Not surprisingly, attaching a preposition to its parent is judged very risky. When sorting POS-tags by the number of times their parent-attachment is judged to be risky, prepositions are at the top of the list. But do the models learn that all PP attachment decisions are risky and abstain from attaching any PP to their parent? Or maybe some kinds of prepositions riskier than others? Table 2 shows the confusion matrix for prepositions as judged by the `edge_higher` model with a risk-threshold of 0.15. Table 3 breaks down the numbers by preposition type.

reality / model	Risky	Safe
Incorrect	TP: 961	FN: 326
Correct	FP: 1353	FP: 4302

Table 2: Preposition’s riskiness confusion matrix over dev-set. `edge_higher` features, riskiness-threshold of 0.15.

Interestingly, while PP attachment is the most risky phenomena, most PP attachment cases are correctly judged by the model to be non-risky (4302 cases). 326 other PP attachment cases are judged by the model to be safe, but are incorrect. Finally, the model marks 2314 PP-attachment cases as risky, and 961 of these are indeed parsing mistakes.

When breaking down by preposition type, we can observe that “of” is by far the least ambiguous (1457 cases, or 93%, are correctly marked as safe), “in”, “for”, “on”, “as”, “at” are the most am-

Preposition	TP	FP	TN	FN	Total
as	41	80	123	19	263
with	34	76	168	12	290
at	55	68	177	21	321
on	77	90	179	16	362
from	32	59	222	17	330
that	31	62	223	16	332
by	35	49	261	9	354
for	120	186	317	28	651
in	245	319	575	75	1214
of	18	40	1457	39	1554

Table 3: Preposition’s riskiness by type over dev-set. `edge_higher` features, riskiness-threshold of 0.15. TP: risky/incorrect, FP:risky/correct, TN:safe/correct, FN:safe/incorrect

biguous (marked as risky about half of the time) while “by”, “that”, “from” are in between (25-35% of the cases are judged to be risky).

9 Discussion

We advocate a modified version of the parsing task – precision-biased parsing – which favors precision over recall by allowing the parser to abstain from decisions about which it is uncertain. In our view, partial but highly accurate structural information is in many cases more valuable than complete but less accurate structural information. The precision-biased parsing problem is related to confidence estimation, that is, attaching reliability scores to model predictions.

In order to address the precision-biased parsing task we introduce the notion of riskiness of parser decisions. On the basis of riskiness assessment, the parser can abstain from risky predictions. This gives rise to a natural solution to the parse-selection task: reliable parse-trees are those associated with few risky actions.


After verifying that disagreement in a parser-ensemble is a good indicator of risky edges, we presented a novel approach that does not rely on a parser-ensemble, but instead learns to predict the riskiness involved with individual actions of a single parser. While the method sacrifices more coverage than the parser-ensemble in order to achieve the same level of accuracy, the results are encouraging and demonstrate that a single parsing system can monitor the confidence of its own predictions.

Single parser riskiness assessment turns out to be a good indicator of confidence on aggregate: the single-parser system is as capable as the ensemble-based one at selecting high-quality complete parses.


References

- Stefan Schoenmackers Alexander Yates and Oren Etzioni. 2006. Detecting parser errors using web-based semantic filters. In *Proc. of EMNLP 2006*.
- Xavier Carreras, Michael Collins, and Terry Koo. 2008. Tag, dynamic programming, and the perceptron for efficient, feature-rich parsing. In *Proc. of CoNLL*.
- John Carroll and Ted Briscoe. 2002. High precision extraction of grammatical relations. In *Proc. of COLING*.
- Hal Daumé III. 2004. Notes on CG and LM-BFGS optimization of logistic regression. Paper available at <http://pub.hal3.name#daume04cg-bfgs>, implementation available at <http://hal3.name/megam/>, August.
- Felice Dell’Orletta, Giulia Venturi, and Simonetta Montemagni. 2011. Ulisse: an unsupervised algorithm for detecting reliable dependency parses. In *In Proc. of CoNLL*.
- Yoav Goldberg and Michael Elhadad. 2010. An efficient algorithm for easy-first non-directional dependency parsing. In *Proc. of NAACL*.
- Johan Hall, Jens Nilsson, Joakim Nivre, Gülsen Eryigit, Beáta Megyesi, Mattias Nilsson, and Markus Saers. 2007. Single malt or blended? a study in multilingual parser optimization. In *Proc. of the CoNLL 2007 Shared Task*.
- Liang Huang, Wenbin Jiang, and Qun Liu. 2009. Bilingually-constrained (monolingual) shift-reduce parsing. In *Proc. of EMNLP*.
- Daisuke Kawahara and Sadao Kurohashi. 2001. Japanese case frame construction by coupling the verb and its closest case component. In *Proc. of HLT-2001, HLT ’01*.
- Daisuke Kawahara and Kiyotaka Uchimoto. 2008. Learning reliability of parses for domain adaptation of dependency parsing. In *Proc. of IJCNLP 2008*.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. Effective Self-Training for Parsing. In *Proc. of HLT-NAACL 2006*.
- Ryan McDonald, Koby Crammer, and Fernando Pereira. 2005. Online large-margin training of dependency parsers. In *Proc. of ACL*.
- Slav Petrov, Pi-Chuan Chang, Michael Ringgaard, and Hiyan Alshawi. 2010. Uptraining for accurate deterministic question parsing. In *Proc. of EMNLP-2010*, October.
- Roi Reichart and Ari Rappoport. 2007. An ensemble method for selection of high quality parses. In *Proc. of ACL-2007*.
- Roi Reichart and Ari Rappoport. 2009. Automatic selection of high quality parses created by a fully unsupervised parser. In *Proc. of CoNLL-2009*.
- Kenji Sagae and Alon Lavie. 2006. Parser combination by reparsing. In *Proc. of HLT-NAACL*.
- K. Sagae and J. Tsujii. 2007. Dependency parsing and domain adaptation with LR models and parser ensembles. In *Proc. of CoNLL-2007 Shared Task*.
- David A. Smith and Noah A. Smith. 2007. Probabilistic models of nonprojective dependency trees. In *Proc. of EMNLP*.
- Rebecca Watson, John Carroll, and Ted Briscoe. 2005. Efficient extraction of grammatical relations. In *Proc. of IWPT*.

Supplementary Material
edge_state, risk threshold 0.15:



edge_factored, risk threshold 0.15:



edge_higher, risk threshold 0.15:




Figure 4: Precision-biased parse examples of the single-parser systems' predictions on the dev set.




Figure 5: Parse-selection based on single-parser risk predictors results (precision vs. sentence-coverage) on dev set.