

Introducing the SPMRL 2014 Shared Task on Parsing Morphologically-Rich Languages

Djamé Seddah
INRIA & Univ. Paris Sorbonne
Paris, France
djame.seddah@paris-sorbonne.fr

Sandra Kübler
Indiana University
Bloomington, IN, USA
skuebler@indiana.edu

Reut Tsarfaty
Weizman Institute
Rehovot, Israel
reut.tsarfaty@weizmann.ac.il

1 Introduction

This first joint meeting on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical English (SPMRL-SANCL) featured a shared task on statistical parsing of morphologically rich languages (SPMRL). The goal of the shared task is to allow to train and test different participating systems on comparable data sets, thus providing an objective measure of comparison between state-of-the-art parsing systems on data sets from a range of different languages. This 2014 SPMRL shared task is a continuation and extension of the SPMRL shared task, which was co-located with the SPMRL meeting at EMNLP 2013 (Seddah et al., 2013).

This paper provides a short overview of the 2014 SPMRL shared task goals, data sets, and evaluation setup. Since the SPMRL 2014 largely builds on the infrastructure established for the SPMRL 2013 shared task, we start by reviewing the previous shared task (§2) and then proceed to the 2014 SPMRL evaluation settings (§3), data sets (§4), and a task summary (§5). Due to organizational constraints, this overview is published prior to the submission of all system test runs, and a more detailed overview including the description of participating systems and the analysis of their results will follow as part of (Seddah et al., 2014), once the shared task is completed.

2 The SPMRL Shared Task 2013

The SPMRL Shared Task 2013 (Seddah et al., 2013) was organized with the goal of providing standard data sets, streamlined evaluation metrics, and a set of strong baselines for parsing morphologically rich languages (MRLs). The goals were both to provide a focal point for researchers interested in parsing MRLs and consequently to advance the state of the art in this area of research.

The shared task focused on parsing nine morphologically rich languages, from different typological language families, in both a constituent-based and a dependency-based format. The set of nine typologically diverse languages comprised data sets for Arabic, Basque, French, German, Hebrew, Hungarian, Korean, Polish, and Swedish. Compared to previous multilingual shared tasks (Buchholz and Marsi, 2006; Nivre et al., 2007), the SPMRL shared task targeted parsing in realistic evaluation scenarios, in which the analysis of morphologically ambiguous input tokens is not known in advance. An additional novelty of the SPMRL shared task is that it allowed for both a dependency-based and a constituent-based parse representation. This setting relied on an intricate and careful data preparation process which ensured consistency between the constituent and the dependency version by aligning the two representation types at the token level and at the level of part-of-speech tags. For all languages, we provided two versions of the data sets: an *all* data set, identical in size to the one made available by the individual treebank providers, and a *small* data set, with a training set of 5,000 sentences, and a test set of about 500 sentences. Controlling the set sizes across languages allows us to level the playing field across languages and treebanks.

This work is licenced under a Creative Commons Attribution 4.0 International License. Page numbers and proceedings footer are added by the organizers. License details: <http://creativecommons.org/licenses/by/4.0/>

The shared task also advanced the state of the art by introducing different levels of complexity in parsing. In general, parsing is reduced to the parsing proper step, assuming gold segmentation of the text into sentences and words as well as gold POS tags and morphological analyses. This is a serious simplification of the task since especially in Semitic languages, the segmentation into input tokens is a task that is best performed in combination with parsing because of the ambiguities involved.

The shared task deviated from this standard configuration by adding conditions in which more realistic settings were given: In the gold setting, unambiguous gold morphological segmentation, POS tags, and morphological features for each input token were given. In the predicted setting, disambiguated morphological segmentation was provided, but the POS tags and morphological features for each input segment were not. In the raw setting, there was no gold information, i.e., morphological segmentation, POS tags and morphological features for each input token had to be predicted as part of the parsing task. To lower the entry cost, participants were provided with reasonable baseline (if not state-of-the-art) morphological predictions (either disambiguated – in most cases – or ambiguous prediction in lattice forms).

As a consequence of the raw scenario, it was not possible to (only) rely on the accepted parsing metrics, labeled bracket evaluation via EVALB¹ (Black et al., 1991), Leaf-Ancessor (Sampson and Babarczy, 2003) for constituents and CONLL X’s Labeled/Unlabeled Attachment Score for dependencies (Buchholz and Marsi, 2006). When the segmentation of words into input tokens is not given, there may be discrepancies on the lexical levels, which neither EVALB and LEAF-ANCESTOR nor LAS/UAS are prepared to handle. Thus, we also used TedEval, a distance-based metric that evaluates a morphosyntactic structure as a complete whole (Tsarfaty et al., 2012b). Note that given the workload brought to the participants, we did not try to enforce function label evaluation for constituent parsing. We hope that further shared tasks will try to generalize such an evaluation. Indeed, having predicted function labels would ease labeled TEDEVAl evaluation and favor a full parsing chain evaluation. Nevertheless, the choice of TEDEVAl allowed us to go beyond the standard cross-parser evaluation within one setting and approach cross-framework (constituent vs. dependency (Tsarfaty et al., 2012a)) and cross-language evaluation, thus pushing the envelope on parsing evaluation. Additionally, we performed a specialized evaluation of multi-word expressions in the French treebank.

The SPMRL Shared Task 2013 featured seven teams who approached the dependency parsing task and one team that approached constituent parsing. The best performing system (Björkelund et al., 2013) in either framework consisted of an ensemble system, combining several dependency parsers or several instantiations of a PCFG-LA parser by a (re-)ranker, both on top of state-of-the-art morphological analyses. The results show that parser combination helps to reach a robust performance across languages. However, the integration of morphological analysis into the parsing needs to be investigated thoroughly, and new, morphologically aware approaches are needed. The cross-parser, cross-scenario, and cross-framework evaluation protocols show that performance on gold morphological input is significantly higher than that in more realistic scenarios, and more training data is beneficial. Additionally, differences between dependency and constituents are smaller than previously assumed, and languages which are typologically farthest from English, such as Semitic and Asian languages, are still amongst the hardest to parse, regardless of the parsing method used.

3 SPMRL 2014 Parsing Scenarios

As in the previous edition, this year, we consider three parsing scenarios, depending on how much of the morphological information is provided. The scenarios are listed below, in increasing order of difficulty.

- **Gold:** In this scenario, the parser is provided with unambiguous gold morphological segmentation, POS tags, and morphological features for each input token.
- **Predicted:** In this scenario, the parser is provided with disambiguated morphological segmentation. However, the POS tags and morphological features for each input segment are unknown.
- **Raw:** In this scenario, the parser is provided with morphologically ambiguous input. The morphological segmentation, POS tags, and morphological features for each input token are unknown.

¹We extended the usual EVALB to penalize unparsed sentences.

Scenario	Segmentation	PoS+Feat.	Tree
Gold	✓	✓	–
Predicted	✓	1-best	–
Raw (1-best)	1-best	1-best	–
Raw (all)	–	–	–

Table 1: A summary of the parsing and evaluation scenarios. ✓ depicts gold information, – depicts unknown information, to be predicted by the system.

The **Predicted** and **Raw** scenarios require predicting morphological analyses. This may be done using a language-specific morphological analyzer, or it may be done jointly with parsing. We provide inputs that support these different scenarios:

- **Predicted**: Gold treebank segmentation is given to the parser. The POS tags assignment and morphological features are automatically predicted by the parser or by an external resource.
- **Raw (1-best)**: The 1-best segmentation and POS tags assignment is predicted by an external resource and given to the parser.
- **Raw (all)**: All possible segmentations and POS tags are specified by an external resource. The parser selects jointly a segmentation and a tree.

An overview of all scenarios is shown in table 1. For languages in which terminals equal tokens, only **Gold** and **Predicted** scenarios are considered. For the Semitic languages, we further provide input for both **Raw (1-best)** and **Raw (all)** scenarios.²

4 SPMRL 2014 Data Sets

The main innovation of the SPMRL 2014 shared task with respect to the previous edition is the availability of additional, unannotated data, for the purpose of semi-supervised training. This section provides a description of the unlabeled-data preparation that is required in the context of parsing MRLs, and the core labeled data that is used in conjunction with it.

4.1 SPMRL Unlabeled Data Set

One of the common problems when dealing with morphologically rich languages (MRLs) is lexical data sparseness due to the high level of variation in word forms (Tsarfaty et al., 2010; Tsarfaty et al., 2012c). The use of large, unlabeled corpora in a semi-supervised setting, in addition to the relatively small MRL data sets, can become a valid option to overcome such issues. For instance, using Brown clusters (Brown et al., 1992) has been shown to boost the performance of a PCFG-LA based parser for French (Candito and Crabbé, 2009; Candito and Seddah, 2010). External lexical acquisition was successfully used for Arabic (Habash, 2008) and Hebrew (Goldberg et al., 2009), self-training increased accuracy for parsing German (Rehbein, 2011), and more recently, the use of word embeddings led to some promising results for some MRLs (Cirik and Şensoy, 2013).

By releasing large, unlabeled data sets and by providing accurate pre-annotation in a format directly compatible with models trained on the SPMRL Shared Task treebanks, we hope to foster the development of interesting and feature-rich parsing models that build on larger, morphologically rich, lexicons. Table 2 presents basic facts about the data sets. Details on the unlabeled data and their pre-annotations will be provided in (Seddah et al., 2014). Note that we could not ensure the same volume of data for all languages, nor we could run the same parser, or morphology prediction, on all data. Potential future work could focus on ensuring a stricter level of comparability of these data or on investigating the feasibility of such a normalization of procedures.

²The raw Arabic lattices were made available later than the other data. They are now included in the shared task release.

Language	Source (main)	type	size (tree tokens)	morph	parsed
Arabic	news domain	news	120M	✓*	✓*
Basque	web	balanced	150M	✓	✓
French	news domain	newswire	120M	✓+mwe	✓*
German	Wikipedia	wiki (edited)	205M	✓	✓
Hebrew	Wikipedia	wiki (edited)	160M	✓	✓
Hungarian	news domain	newswire	100M	✓	✓
Korean	news domain	newswire	40M	✓	✓*
Polish	Wikipedia	wiki (edited)	100M	✓	✓
Swedish	PAROLE	balanced	24M	✓	✓

Table 2: Unlabeled data set properties. *: *made available mid-july*

4.2 SPMRL Core Labeled Data Set

In order to provide a faithful evaluation of the impact of these additional sets of unlabeled data, we used the exact same data sets for training and testing as in the previous edition. Specifically, we used an Arabic data set, originally provided by the LDC (Maamouri et al., 2004), in a dependency form, derived from the Columbia Catib Treebank (Habash and Roth, 2009; Habash et al., 2009) and in a constituency instance, following the Stanford pre-processing scheme (Green and Manning, 2010) and extended according to the SPMRL 2013 extension scheme (Seddah et al., 2013). For Basque, the data was provided by Aduriz et al. (2003) in both dependency and constituency, we removed sentences with non-projective trees so both instances could be aligned at the token level. Regarding French, we used a new instance of the French Treebank (Abeillé et al., 2003) that includes multi-word expression (MWE) annotations, annotated at the morpho-syntactic level in both instances. Predicted MWEs were added this year, using the same tools as Constant et al. (2013). The German data are based on the Tiger corpus (Brants et al., 2002), and converted to constituent and dependency following (Seeker and Kuhn, 2012). The Hebrew data set is based on the Modern Hebrew Treebank (Sima’an et al., 2001), with the Goldberg (2011) dependency version, in turn aligned with the phrase structure instance described in (Tsarfaty, 2010; Tsarfaty, 2013). Note that in order to match the Hebrew unlabeled data encoding, the Hebrew treebank was converted back to UTF-8. The Hungarian data are derived from the Szeged treebank (Csendes et al., 2005; Vincze et al., 2010), while the Korean data originate from the Kaist Treebank (Choi et al., 1994) which was converted to dependency for the SPMRL shared task by Choi (2013). The Polish treebank we used is described in (Woliński et al., 2011; Świdziński and Woliński, 2010; Wróblewska, 2012). Compared to the last year’s edition, we added explicit feature names in the relevant data fields. The Swedish data originate from (Nivre et al., 2006), we added function labels extracted from the original Swedish XML data. Note that in addition to constituency and dependency versions, the Polish, German and Swedish data sets are also available in the Tiger XML format (Mengel and Lezius, 2000), allowing a direct representation of discontinuous structures in their phrase-based structures.

5 Conclusion

At the time of writing this short introduction, the shared task is ongoing, and neither results nor the final submitting teams are known. At this point, we can say that 15 teams registered for the 2014 shared task edition, indicating an increased awareness of and continued interest in the topic of the shared task. Results, cross-parser and cross-data analysis, and shared task description papers will be made available at <http://www.spmrl.org/spmrl2014-sharedtask.html>.

Acknowledgments

We would like to express our gratitude to the original treebank labeled and unlabeled data contributors for the considerable time they devoted to our shared task. Namely, **Arabic**: Nizar Habash, Ryan Roth (Columbia University); Spence Green (Stanford University), Ann Bies, Seth Kulick, Mohamed Maamouri (the Linguistic Data Consortium); **Basque**: Koldo Gojenola, Iakes Goenaga (University of the Basque Country); **French**: Marie Candito (Univ. Paris 7 & Inria), Djamé Seddah (Univ. Paris Sorbonne & Inria), Matthieu Constant (Univ. Marne la Vallée); **German**: Wolfgang Seeker (IMS

Stuttgart), Wolfgang Maier (Univ. of Dusseldorf), Yannick Versley (Univ. of Tuebingen) ; **Hebrew:** Yoav Goldberg (Bar Ilan Univ.), Reut Tsarfaty (Weizmann Institute of Science) ; **Hungarian:** Richárd Farkas, Veronika Vincze (Univ. of Szeged) ; **Korean:** Jinho D. Choi (Univ. of Massachusetts Amherst), Jungyeul Park (Kaist); **Polish:** Adam Przepiórkowski, Marcin Woliński, Alina Wróblewska (Institute of Computer Science, Polish Academy of Sciences) ; **Swedish:** Joakim Nivre (Uppsala Univ.), Marco Kuhlmann (Linköping University).

We gratefully acknowledge the contribution of Språkbanken and the University of Gothenburg for providing the PAROLE corpus. We are also very grateful to the Philosophical Faculty of the Heinrich-Heine Universität Düsseldorf for hosting the shared task data via their dokuwiki.

References

- Anne Abeillé, Lionel Clément, and François Toussenet. 2003. Building a treebank for French. In Anne Abeillé, editor, *Treebanks*. Kluwer, Dordrecht.
- I. Aduriz, M. J. Aranzabe, J. M. Arriola, A. Atutxa, A. Díaz de Ilarraza, A. Garmendia, and M. Oronoz. 2003. Construction of a Basque dependency treebank. In *Proceedings of the Second Workshop on Treebanks and Linguistic Theories*, pages 201–204, Växjö, Sweden.
- Anders Björkelund, Ozlem Cetinoglu, Richárd Farkas, Thomas Mueller, and Wolfgang Seeker. 2013. (Re)ranking meets morphosyntax: State-of-the-art results from the SPMRL 2013 shared task. In *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 134–144, Seattle, WA.
- Ezra Black, Steven Abney, Dan Flickinger, Claudia Gdaniec, Ralph Grishman, Philip Harrison, Donald Hindle, Robert Ingria, Frederick Jelinek, Judith Klavans, Mark Liberman, Mitchell Marcus, Salim Roukos, Beatrice Santorini, and Tomek Strzalkowski. 1991. A procedure for quantitatively comparing the syntactic coverage of English grammars. In *Proceedings of the DARPA Speech and Natural Language Workshop 1991*, pages 306–311, Pacific Grove, CA.
- Sabine Brants, Stefanie Dipper, Silvia Hansen, Wolfgang Lezius, and George Smith. 2002. The TIGER treebank. In *Proceedings of the First Workshop on Treebanks and Linguistic Theories (TLT)*, pages 24–41, Sozopol, Bulgaria.
- Peter F. Brown, Vincent J. Della, Peter V. Desouza, Jennifer C. Lai, and Robert L. Mercer. 1992. Class-based n-gram models of natural language. *Computational Linguistics*, 18(4):467–479.
- Sabine Buchholz and Erwin Marsi. 2006. CoNLL-X shared task on multilingual dependency parsing. In *Proceedings of CoNLL*, pages 149–164, New York, NY.
- Marie Candito and Benoît Crabbé. 2009. Improving generative statistical parsing with semi-supervised word clustering. In *Proceedings of the 11th International Conference on Parsing Technologies (IWPT’09)*, pages 138–141, Paris, France.
- Marie Candito and Djamé Seddah. 2010. Parsing word clusters. In *Proceedings of the NAACL/HLT Workshop on Statistical Parsing of Morphologically Rich Languages (SPMRL 2010)*, Los Angeles, CA.
- Key-sun Choi, Young S. Han, Young G. Han, and Oh W. Kwon. 1994. KAIST Tree Bank Project for Korean: Present and Future Development. In *Proceedings of the International Workshop on Sharable Natural Language Resources*, pages 7–14, Nara, Japan.
- Jinho D. Choi. 2013. Preparing Korean data for the shared task on parsing morphologically rich languages. arXiv:1309.1649.
- Volkan Cirik and Hüsnü Şensoy. 2013. The AI-KU system at the SPMRL 2013 shared task: Unsupervised features for dependency parsing. In *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 68–75, Seattle, WA.
- Matthieu Constant, Marie Candito, and Djamé Seddah. 2013. The LIGM-Alpage architecture for the SPMRL 2013 shared task: Multiword expression analysis and dependency parsing. In *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 46–52, Seattle, WA.
- Dóra Csendes, János Csirik, Tibor Gyimóthy, and András Kocsor. 2005. The Szeged treebank. In *Proceedings of the 8th International Conference on Text, Speech and Dialogue (TSD)*, Lecture Notes in Computer Science, pages 123–132, Berlin / Heidelberg. Springer.

- Yoav Goldberg, Reut Tsarfaty, Meni Adler, and Michael Elhadad. 2009. Enhancing unlexicalized parsing performance using a wide coverage lexicon, fuzzy tag-set mapping, and EM-HMM-based lexical probabilities. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL)*, pages 327–335, Athens, Greece.
- Yoav Goldberg. 2011. *Automatic syntactic processing of Modern Hebrew*. Ph.D. thesis, Ben Gurion University of the Negev.
- Spence Green and Christopher D. Manning. 2010. Better Arabic parsing: Baselines, evaluations, and analysis. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 394–402, Beijing, China.
- Nizar Habash and Ryan Roth. 2009. CATiB: The Columbia Arabic Treebank. In *Proceedings of ACL-IJCNLP*, pages 221–224, Suntec, Singapore.
- Nizar Habash, Reem Faraj, and Ryan Roth. 2009. Syntactic Annotation in the Columbia Arabic Treebank. In *Proceedings of MEDAR International Conference on Arabic Language Resources and Tools*, Cairo, Egypt.
- Nizar Habash. 2008. Four techniques for online handling of out-of-vocabulary words in Arabic-English statistical machine translation. In *Proceedings of ACL-08: HLT, Short Papers*, pages 57–60, Columbus, OH.
- Mohamed Maamouri, Ann Bies, Tim Buckwalter, and Wigdan Mekki. 2004. The Penn Arabic treebank: Building a large-scale annotated Arabic corpus. In *Proceedings of NEMLAR International Conference on Arabic Language Resources and Tools*, pages 102–109, Cairo, Egypt.
- Andreas Mengel and Wolfgang Lezius. 2000. An XML-based encoding format for syntactically annotated corpora. In *Proceedings of the Second International Conference on Language Resources and Engineering (LREC 2000)*, pages 121–126, Athens, Greece.
- Joakim Nivre, Jens Nilsson, and Johan Hall. 2006. Talbanken05: A Swedish treebank with phrase structure and dependency annotation. In *Proceedings of LREC*, pages 1392–1395, Genoa, Italy.
- Joakim Nivre, Johan Hall, Sandra Kübler, Ryan McDonald, Jens Nilsson, Sebastian Riedel, and Deniz Yuret. 2007. The CoNLL 2007 shared task on dependency parsing. In *Proceedings of the CoNLL Shared Task Session of EMNLP-CoNLL 2007*, pages 915–932, Prague, Czech Republic.
- Ines Rehbein. 2011. Data point selection for self-training. In *Proceedings of the Second Workshop on Statistical Parsing of Morphologically Rich Languages*, pages 62–67, Dublin, Ireland.
- Geoffrey Sampson and Anna Babarczy. 2003. A test of the leaf-ancestor metric for parse accuracy. *Natural Language Engineering*, 9(04):365–380.
- Djamé Seddah, Reut Tsarfaty, Sandra Kübler, Marie Candito, Jinho D. Choi, Richárd Farkas, Jennifer Foster, Iakes Goenaga, Koldo Gojenola Gallettebeitia, Yoav Goldberg, Spence Green, Nizar Habash, Marco Kuhlmann, Wolfgang Maier, Joakim Nivre, Adam Przepiórkowski, Ryan Roth, Wolfgang Seeker, Yannick Versley, Veronika Vincze, Marcin Woliński, Alina Wróblewska, and Eric Villemonte de la Clergerie. 2013. Overview of the SPMRL 2013 shared task: A cross-framework evaluation of parsing morphologically rich languages. In *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 146–182, Seattle, WA.
- Djamé Seddah, Reut Tsarfaty, Sandra Kübler, Marie Candito, Jinho Choi, Matthieu Constant, Richárd Farkas, Iakes Goenaga, Koldo Gojenola, Yoav Goldberg, Spence Green, Nizar Habash, Marco Kuhlmann, Wolfgang Maier, Joakim Nivre, Adam Przepiórkowski, Ryan Roth, Wolfgang Seeker, Yannick Versley, Veronika Vincze, Marcin Woliński, Alina Wróblewska, and Eric Villemonte de la Clérgerie. 2014. Overview of the spmrl 2014 shared task on parsing morphologically rich languages. In *Notes of the SPMRL 2014 Shared Task on Parsing Morphologically-Rich Languages*, Dublin, Ireland.
- Wolfgang Seeker and Jonas Kuhn. 2012. Making Ellipses Explicit in Dependency Conversion for a German Treebank. In *Proceedings of the 8th International Conference on Language Resources and Evaluation*, pages 3132–3139, Istanbul, Turkey.
- Khalil Sima’an, Alon Itai, Yoad Winter, Alon Altmann, and Noa Nativ. 2001. Building a tree-bank of Modern Hebrew text. *Traitement Automatique des Langues*, 42:347–380.
- Marek Świdziński and Marcin Woliński. 2010. Towards a bank of constituent parse trees for Polish. In *Proceedings of Text, Speech and Dialogue*, pages 197–204, Brno, Czech Republic.

- Reut Tsarfaty, Djame Seddah, Yoav Goldberg, Sandra Kübler, Marie Candito, Jennifer Foster, Yannick Versley, Ines Rehbein, and Lamia Tounsi. 2010. Statistical parsing for morphologically rich language (SPMRL): What, how and whither. In *Proceedings of the First workshop on Statistical Parsing of Morphologically Rich Languages (SPMRL)*, Los Angeles, CA.
- Reut Tsarfaty, Joakim Nivre, and Evelina Andersson. 2012a. Cross-framework evaluation for statistical parsing. In *Proceeding of EACL*, Avignon, France.
- Reut Tsarfaty, Joakim Nivre, and Evelina Andersson. 2012b. Joint evaluation for segmentation and parsing. In *Proceedings of ACL*, Jeju, Korea.
- Reut Tsarfaty, Djame Seddah, Sandra Kübler, and Joakim Nivre. 2012c. Parsing morphologically rich languages: Introduction to the special issue. *Computational Linguistics*, 39(1):15–22.
- Reut Tsarfaty. 2010. *Relational-Realizational Parsing*. Ph.D. thesis, University of Amsterdam.
- Reut Tsarfaty. 2013. A unified morpho-syntactic scheme of Stanford dependencies. In *Proceedings of ACL*, Sofia, Bulgaria.
- Veronika Vincze, Dóra Szauter, Attila Almási, György Móra, Zoltán Alexin, and János Csirik. 2010. Hungarian Dependency Treebank. In *Proceedings of LREC*, Valletta, Malta.
- Marcin Woliński, Katarzyna Głowińska, and Marek Świdziński. 2011. A preliminary version of Składnica—a treebank of Polish. In *Proceedings of the 5th Language & Technology Conference*, pages 299–303, Poznań, Poland.
- Alina Wróblewska. 2012. Polish Dependency Bank. *Linguistic Issues in Language Technology*, 7(1):1–15.