Feature-Rich Part-of-speech Tagging
for Morphologically Complex Languages: Application to Bulgarian

Georgi Georgiev and Valentin Zhikov
Ontotext AD
135 Tsarigradsko Sh., Sofia, Bulgaria
{georgi.georgiev,valentin.zhikov}@ontotext.com

Petya Osenova and Kiril Simov
IICT, Bulgarian Academy of Sciences
25A Acad. G. Bonchev, Sofia, Bulgaria
{petya,kivs}@bultreebank.org

Preslav Nakov
Qatar Computing Research Institute, Qatar Foundation
Tornado Tower, floor 10, P.O. Box 5825, Doha, Qatar
pnakov@qf.org.qa

Abstract

We present experiments with part-of-speech tagging for Bulgarian, a Slavic language with rich inflectional and derivational morphology. Unlike most previous work, which has used a small number of grammatical categories, we work with 680 morpho-syntactic tags. We combine a large morphological lexicon with prior linguistic knowledge and guided learning from a POS-annotated corpus, achieving accuracy of 97.98%, which is a significant improvement over the state-of-the-art for Bulgarian.

1 Introduction

Part-of-speech (POS) tagging is the task of assigning each of the words in a given piece of text a contextually suitable grammatical category. This is not trivial since words can play different syntactic roles in different contexts, e.g., *can* is a noun in "*I opened a can of coke.*" but a verb in "*I can write.*" Traditionally, linguists have classified English words into the following eight basic POS categories: noun, pronoun, adjective, verb, adverb, preposition, conjunction, and interjection; this list is often extended a bit, e.g., with determiners, particles, participles, etc., but the number of categories considered is rarely more than 15.

Computational linguistics works with a larger inventory of POS tags, e.g., the Penn Treebank (Marcus et al., 1993) uses 48 tags: 36 for part-of-speech, and 12 for punctuation and currency symbols. This increase in the number of tags is partially due to finer granularity, e.g., there are special tags for determiners, particles, modal verbs, cardinal numbers, foreign words, existential *there*, etc., but also to the desire to encode morphological information as part of the tags.

For example, there are six tags for verbs in the Penn Treebank: VB (verb, base form; e.g., *sing*), VBD (verb, past tense; e.g., *sang*), VBG (verb, gerund or present participle; e.g., *singing*), VBN (verb, past participle; e.g., *sung*) VBP (verb, non-3rd person singular present; e.g., *sing*), and VBZ (verb, 3rd person singular present; e.g., *sings*); these tags are morpho-syntactic in nature. Other corpora have used even larger tagsets, e.g., the Brown corpus (Kučera and Francis, 1967) and the Lancaster-Oslo/Bergen (LOB) corpus (Johansson et al., 1986) use 87 and 135 tags, respectively.

POS tagging poses major challenges for morphologically complex languages, whose tagsets encode a lot of additional morpho-syntactic features (for most of the basic POS categories), e.g., gender, number, person, etc. For example, the BulTreeBank (Simov et al., 2004) for Bulgarian uses 680 tags, while the Prague Dependency Treebank (Hajič, 1998) for Czech has over 1,400 tags.

Below we present experiments with POS tagging for Bulgarian, which is an inflectional language with rich morphology. Unlike most previous work, which has used a reduced set of POS tags, we use all 680 tags in the BulTreeBank. We combine prior linguistic knowledge and statistical learning, achieving accuracy comparable to that reported for state-of-the-art systems for English.

The remainder of the paper is organized as follows: Section 2 provides an overview of related work, Section 3 describes Bulgarian morphology, Section 4 introduces our approach, Section 5 describes the datasets, Section 6 presents our experiments in detail, Section 7 discusses the results, Section 8 offers application-specific error analysis, and Section 9 concludes and points to some promising directions for future work.

2 Related Work

Most research on part-of-speech tagging has focused on English, and has relied on the Penn Treebank (Marcus et al., 1993) and its tagset for training and evaluation. The task is typically addressed as a sequential tagging problem; one notable exception is the work of Brill (1995), who proposed non-sequential transformation-based learning.

A number of different sequential learning frameworks have been tried, yielding 96-97% accuracy: Lafferty et al. (2001) experimented with conditional random fields (CRFs) (95.7% accuracy), Ratnaparkhi (1996) used a maximum entropy sequence classifier (96.6% accuracy), Brants (2000) employed a hidden Markov model (96.6% accuracy), Collins (2002) adopted an averaged perception discriminative sequence model (97.1% accuracy). All these models fix the order of inference from left to right.

Toutanova et al. (2003) introduced a cyclic dependency network (97.2% accuracy), where the search is bi-directional. Shen et al. (2007) have further shown that better results (97.3% accuracy) can be obtained using guided learning, a framework for bidirectional sequence classification, which integrates token classification and inference order selection into a single learning task and uses a perceptron-like (Collins and Roark, 2004) passive-aggressive classifier to make the easiest decisions first. Recently, Tsuruoka et al. (2011), proposed a simple perceptron-based classifier applied from left to right but augmented with a lookahead mechanism that searches the space of future actions, yielding 97.3% accuracy.

For morphologically complex languages, the problem of POS tagging typically includes morphological disambiguation, which yields a much larger number of tags. For example, for Arabic, Habash and Rambow (2005) used support vector machines (SVM), achieving 97.6% accuracy with 139 tags from the Arabic Treebank (Maamouri et al., 2003). For Czech, Hajič et al. (2001) combined a hidden Markov model (HMM) with linguistic rules, which yielded 95.2% accuracy using an inventory of over 1,400 tags from the Prague Dependency Treebank (Hajič, 1998). For Icelandic, Dredze and Wallenberg (2008) reported 92.1% accuracy with 639 tags developed for the Icelandic frequency lexicon (Pind et al., 1991), they used guided learning and tag decomposition:

First, a coarse POS class is assigned (e.g., noun, verb, adjective), then, additional fine-grained morphological features like case, number and gender are added, and finally, the proposed tags are further reconsidered using non-local features. Similarly, Smith et al. (2005) decomposed the complex tags into factors, where models for predicting part-of-speech, gender, number, case, and lemma are estimated separately, and then composed into a single CRF model; this yielded competitive results for Arabic, Korean, and Czech.

Most previous work on Bulgarian POS tagging has started with large tagsets, which were then reduced. For example, Dojchinova and Mihov (2004) mapped their initial tagset of 946 tags to just 40, which allowed them to achieve 95.5% accuracy using the transformation-based learning of Brill (1995), and 98.4% accuracy using manually crafted linguistic rules. Similarly, Georgiev et al. (2009), who used maximum entropy and the BulTreeBank (Simov et al., 2004), grouped its 680 fine-grained POS tags into 95 coarse-grained ones, and thus improved their accuracy from 90.34% to 94.4%. Simov and Osenova (2001) used a recurrent neural network to predict (a) 160 morpho-syntactic tags (92.9% accuracy) and (b) 15 POS tags (95.2% accuracy).

Some researchers did not reduce the tagset: Savkov et al. (2011) used 680 tags (94.7% accuracy), and Tanev and Mitkov (2002) used 303 tags and the BULMORPH morphological analyzer (Krushkov, 1997), achieving P=R=95%.

3 Bulgarian Morphology

Bulgarian is an Indo-European language from the Slavic language group, written with the Cyrillic alphabet and spoken by about 9-12 million people. It is also a member of the Balkan Sprachbund and thus differs from most other Slavic languages: it has no case declensions, uses a suffixed definite article (which has a short and a long form for singular masculine), and lacks verb infinitive forms. It further uses special evidential verb forms to express unwitnessed, retold, and doubtful activities.

Bulgarian is an inflective language with very rich morphology. For example, Bulgarian verbs have 52 synthetic wordforms on average, while pronouns have altogether more than ten grammatical features (not necessarily shared by all pronouns), including case, gender, person, number, definiteness, etc.

This rich morphology inevitably leads to ambiguity proliferation; our analysis of BulTreeBank shows four major types of ambiguity:

1. Between the wordforms of the same lexeme, i.e., in the paradigm. For example, дивана, an inflected form of диван ('sofa', masculine), can mean (a) 'the sofa' (definite, singular, short definite article) or (b) a count form, e.g., as in два дивана ('two sofas').

2. Between two or more lexemes, i.e., conversion. For example, като can be (a) a subordinator meaning 'as, when', or (b) a preposition meaning 'like, such as'.

3. Between a lexeme and an inflected wordform of another lexeme, i.e., across-paradigms. For example, политика can mean (a) 'the politician' (masculine, singular, definite, short definite article) or (b) 'politics' (feminine, singular, indefinite).

4. Between the wordforms of two or more lexemes, i.e., across-paradigms and quasi-conversion. For example, върви can mean (a) 'walks' (verb, 2nd or 3rd person, present tense) or (b) 'strings, laces' (feminine, plural, indefinite).

Some morpho-syntactic ambiguities in Bulgarian are occasional, but many are systematic, e.g., neuter singular adjectives have the same forms as adverbs. Overall, most ambiguities are local, and thus arguably resolvable using n-grams, e.g., compare хубаво дете ('beautiful child'), where хубаво is a neuter adjective, and "Пея хубаво." ('I sing beautifully.'), where it is an adverb of manner. Other ambiguities, however, are non-local and may require discourse-level analysis, e.g., "Видях го." can mean 'I saw him.', where го is a masculine pronoun, or 'I saw it.', where it is a neuter pronoun. Finally, there are ambiguities that are very hard or even impossible[1] to resolve, e.g., "Детето влезе весело." can mean both 'The child came in happy.' (весело is an adjective) and 'The child came in happily.' (it is an adverb); however, the latter is much more likely.

---

[1] The problem also exists for English, e.g., the annotators of the Penn Treebank were allowed to use tag combinations for inherently ambiguous cases: JJ|NN (adjective or noun as prenominal modifier), JJ|VBG (adjective or gerund/present participle), JJ|VBN (adjective or past participle), NN|VBG (noun or gerund), and RB|RP (adverb or particle).

In many cases, strong domain preferences exist about how various systematic ambiguities should be resolved. We made a study for the newswire domain, analyzing a corpus of 546,029 words, and we found that ambiguity type 2 (lexeme-lexeme) prevailed for functional parts-of-speech, while the other types were more frequent for inflecting parts-of-speech. Below we show the most frequent types of morpho-syntactic ambiguities and their frequency in our corpus:

- на: preposition ('of') vs. emphatic particle, with a ratio of 28,554 to 38;
- да: auxiliary particle ('to') vs. affirmative particle, with a ratio of 12,035 to 543;
- е: 3rd person present auxiliary verb ('to be') vs. particle ('well') vs. interjection ('wow'), with a ratio of 9,136 to 21 to 5;
- singular masculine noun with a short definite article vs. count form of a masculine noun, with a ratio of 6,437 to 1,592;
- adverb vs. neuter singular adjective, with a ratio of 3,858 to 1,753.

Overall, the following factors should be taken into account when modeling Bulgarian morpho-syntax: (1) locality vs. non-locality of grammatical features, (2) interdependence of grammatical features, and (3) domain-specific preferences.

4 Method

We used the guided learning framework described in (Shen et al., 2007), which has yielded state-of-the-art results for English and has been successfully applied to other morphologically complex languages such as Icelandic (Dredze and Wallenberg, 2008); we found it quite suitable for Bulgarian as well. We used the feature set defined in (Shen et al., 2007), which includes the following:

1. The feature set of Ratnaparkhi (1996), including prefix, suffix and lexical, as well as some bigram and trigram context features;

2. Feature templates as in (Ratnaparkhi, 1996), which have been shown helpful in bidirectional search;

3. More bigram and trigram features and bi-lexical features as in (Shen et al., 2007).

Note that we allowed prefixes and suffixes of length up to 9, as in (Toutanova et al., 2003) and (Tsuruoka and Tsujii, 2005).

We further extended the set of features with the tags proposed for the current word token by a morphological lexicon, which maps words to possible tags; it is exhaustive, i.e., the correct tag is always among the suggested ones for each token.

We also used 70 linguistically-motivated, high-precision rules in order to further reduce the number of possible tags suggested by the lexicon. The rules are similar to those proposed by Hinrichs and Trushkina (2004) for German; we implemented them as constraints in the CLaRK system (Simov et al., 2003).

Here is an example of a rule: If a wordform is ambiguous between a masculine count noun (Ncmt) and a singular short definite masculine noun (Ncmsh), the Ncmt tag should be chosen if the previous token is a numeral or a number.

The 70 rules were developed by linguists based on observations over the training dataset only. They target primarily the most frequent cases of ambiguity, and to a lesser extent some infrequent but very problematic cases. Some rules operate over classes of words, while other refer to particular wordforms. The rules were designed to be 100% accurate on our training dataset; our experiments show that they are also 100% accurate on the test and on the development dataset.

Note that some of the rules are dependent on others, and thus the order of their cascaded application is important. For example, the wordform я is ambiguous between an accusative feminine singular short form of a personal pronoun ('her') and an interjection ('wow'). To handle this properly, the rule for interjection, which targets sentence initial positions, followed by a comma, needs to be executed first. The rule for personal pronouns is only applied afterwards.

Word	Tags
Той	**Ppe-os3m**
обаче	*Cc*; **Dd**
няма	Afsi; Vnitf-o3s; Vnitf-r3s; Vpitf-o2s; Vpitf-o3s; **Vpitf-r3s**
възможност	**Ncfsi**
да	*Ta*;**Tx**
следи	*Ncfpi*; *Vpitf-o2s*; *Vpitf-o3s*; **Vpitf-r3s**; *Vpitz–2s*
. . .	. . .

Table 1: **Sample fragment showing the possible tags suggested by the lexicon.** The tags that are further filtered by the rules are in italic; the correct tag is bold.

The rules are quite efficient at reducing the POS ambiguity. On the test dataset, before the rule application, 34.2% of the tokens (excluding punctuation) had more than one tag in our morphological lexicon. This number is reduced to 18.5% after the cascaded application of the 70 linguistic rules. Table 1 illustrates the effect of the rules on a small sentence fragment. In this example, the rules have left only one tag (the correct one) for three of the ambiguous words. Since the rules in essence decrease the average number of tags per token, we calculated that the lexicon suggests 1.6 tags per token on average, and after the application of the rules this number decreases to 1.44 per token.

5 Datasets

5.1 BulTreeBank

We used the latest version of the BulTree-Bank (Simov and Osenova, 2004), which contains 20,556 sentences and 321,542 word tokens (four times less than the English Penn Treebank), annotated using a total of 680 unique morpho-syntactic tags. See (Simov et al., 2004) for a detailed description of the BulTreeBank tagset.

We split the data into training/development/test as shown in Table 2. Note that only 552 of all 680 tag types were used in the training dataset, and the development and the test datasets combined contain a total of 128 new tag types that were not seen in the training dataset. Moreover, 32% of the word types in the development dataset and 31% of those in the testing dataset do not occur in the training dataset. Thus, data sparseness is an issue at two levels: word-level and tag-level.

Dataset	Sentences	Tokens	Types	Tags
Train	16,532	253,526	38,659	552
Dev	2,007	32,995	9,635	425
Test	2,017	35,021	9,627	435

Table 2: **Statistics about our datasets.**

5.2 Morphological Lexicon

In order to alleviate the data sparseness issues, we further used a large morphological lexicon for Bulgarian, which is an extended version of the dictionary described in (Popov et al., 1998) and (Popov et al., 2003). It contains over 1.5M inflected wordforms (for 110K lemmata and 40K proper names), each mapped to a set of possible morpho-syntactic tags.

6 Experiments and Evaluation

State-of-the-art POS taggers for English typically build a lexicon containing all tags a word type has taken in the training dataset; this lexicon is then used to limit the set of possible tags that an input token can be assigned, i.e., it imposes a hard constraint on the possibilities explored by the POS tagger. For example, if *can* has only been tagged as a verb and as a noun in the training dataset, it will be only assigned those two tags at test time; other tags such as adjective, adverb and pronoun will not be considered. Out-of-vocabulary words, i.e., those that were not seen in the training dataset, are constrained as well, e.g., to a small set of frequent open-class tags.

In our experiments, we used a morphological lexicon that is much larger than what could be built from the training corpus only: building a lexicon from the training corpus only is of limited utility since one can hardly expect to see in the training corpus all 52 synthetic forms a verb can possibly have. Moreover, we did not use the tags listed in the lexicon as hard constraints (except in one of our baselines); instead, we experimented with a different, non-restrictive approach: we used the lexicon's predictions as features or soft constraints, i.e., as suggestions only, thus allowing each token to take any possible tag. Note that for both known and out-of-vocabulary words we used all 680 tags rather than the 552 tags observed in the training dataset; we could afford to explore this huge search space thanks to the efficiency of the guided learning framework. Allowing all 680 tags on training helped the model by exposing it to a larger set of negative examples.

We combined these lexicon features with standard features extracted from the training corpus. We further experimented with the 70 contextual linguistic rules, using them (a) as soft and (b) as hard constraints. Finally, we set four baselines: three that do not use the lexicon and one that does.

#	Baselines	Accuracy (%) (token-level)
1	MFT + unknowns are wrong	78.10
2	MFT + unknowns are Ncmsi	78.52
3	MFT + guesser for unknowns	79.49
4	MFT + lexicon tag-classes	94.40

Table 3: **Most-frequent-tag (MFT) baselines.**

6.1 Baselines

First, we experimented with the most-frequent-tag baseline, which is standard for POS tagging. This baseline ignores context altogether and assigns each word type the POS tag it was most frequently seen with in the training dataset; ties are broken randomly. We coped with word types not seen in the training dataset using three simple strategies: (a) we considered them all wrong, (b) we assigned them Ncmsi, which is the most frequent open-class tag in the training dataset, or (c) we used a very simple guesser, which assigned Ncfsi, Ncnsi, Ncfsi, and Ncmsf, if the target word ended by -a, -o, -и, and -ът, respectively, otherwise, it assigned Ncmsi. The results are shown in lines 1-3 of Table 3: we can see that the token-level accuracy ranges in 78-80% for (a)-(c), which is relatively high, given that we use a large inventory of 680 morpho-syntactic tags.

We further tried a baseline that uses the above-described morphological lexicon, in addition to the training dataset. We first built two frequency lists, containing respectively (1) the most frequent tag in the training dataset for each word type, as before, and (2) the most frequent tag in the training dataset for each class of tags that can be assigned to some word type, according to the lexicon. For example, the most frequent tag for политика is Ncfsi, and the most frequent tag for the tag-class {Ncmt;Ncmsi} is Ncmt.

Given a target word type, this new baseline first tries to assign it the most frequent tag from the first list. If this is not possible, which happens (i) in case of ties or (ii) when the word type was not seen on training, it extracts the tag-class from the lexicon and consults the second list. If there is a single most frequent tag in the corpus for this tag-class, it is assigned; otherwise a random tag from this tag-class is selected.

Line 4 of Table 3 shows that this latter baseline achieves a very high accuracy of 94.40%. Note, however, that this is over-optimistic: the lexicon contains a tag-class for each word type in our testing dataset, i.e., while there can be word types not seen in the training dataset, there are no word types that are not listed in the lexicon. Thus, this high accuracy is probably due to a large extent to the scale and quality of our morphological lexicon, and it might not be as strong with smaller lexicons; we plan to investigate this in future work.

6.2 Lexicon Tags as Soft Constraints

We experimented with three types of features:

1. Word-related features only;

2. Word-related features + the tags suggested by the lexicon;

3. Word-related features + the tags suggested by the lexicon but then further filtered using the 70 contextual linguistic rules.

Table 4 shows the sentence-level and the token-level accuracy on the test dataset for the three kinds of features: shown on lines 1, 3 and 4, respectively. We can see that using the tags proposed by the lexicon as features (lines 3 and 4) has a major positive impact, yielding up to 49% error reduction at the token-level and up to 37% at the sentence-level, as compared to using word-related features alone (line 1).

Interestingly, filtering the tags proposed by the lexicon using the 70 contextual linguistic rules yields a minor decrease in accuracy both at the word token-level and at the sentence-level (compare line 4 to line 2). This is surprising since the linguistic rules are extremely reliable: they were designed to be 100% accurate on the training dataset, and we found them experimentally to be 100% correct on the development and on the testing dataset as well.

One possible explanation is that by limiting the set of available tags for a given token at training time, we prevent the model from observing some potentially useful negative examples. We tested this hypothesis by using the unfiltered lexicon predictions at training time but then making use of the filtered ones at testing time; the results are shown on line 5. We can observe a small increase in accuracy compared to line 4: from 97.80% to 97.84% at the token-level, and from 70.30% to 70.40% at the sentence-level. Although these differences are tiny, they suggest that having more negative examples at training is helpful.

We can conclude that using the lexicon as a source of soft constraints has a major positive impact, e.g., because it provides access to important external knowledge that is complementary to what can be learned from the training corpus alone; the improvements when using linguistic rules as soft constraints are more limited.

6.3 Linguistic Rules as Hard Constraints

Next, we experimented with using the suggestions of the linguistic rules as hard constraints. Table 4 shows that this is a very good idea. Comparing line 1 to line 2, which do not use the morphological lexicon, we can see very significant improvements: from 95.72% to 97.20% at the token-level and from 52.95% to 64.50% at the sentence-level. The improvements are smaller but still consistent when the morphological lexicon is used: comparing lines 3 and 4 to lines 6 and 7, respectively, we see an improvement from 97.83% to 97.91% and from 97.80% to 97.93% at the token-level, and about 1% absolute at the sentence-level.

6.4 Increasing the Beam Size

Finally, we increased the beam size of guided learning from 1 to 3 as in (Shen et al., 2007). Comparing line 7 to line 8 in Table 4, we can see that this yields further token-level improvement: from 97.93% to 97.98%.

7 Discussion

Table 5 compares our results to previously reported evaluation results for Bulgarian. The first four lines show the token-level accuracy for standard POS tagging tools trained and evaluated on the BulTreeBank:[2] TreeTagger (Schmid, 1994), which uses decision trees, TnT (Brants, 2000), which uses a hidden Markov model, SVMtool (Giménez and Màrquez, 2004), which is based on support vector machines, and ACOPOST (Schröder, 2002), implementing the memory-based model of Daelemans et al. (1996). The following lines report the token-level accuracy reported in previous work, as compared to our own experiments using guided learning.

We can see that we outperform by a very large margin (92.53% vs. 97.98%, which represents 73% error reduction) the systems from the first four lines, which are directly comparable to our experiments: they are trained and evaluated on the BulTreeBank using the full inventory of 680 tags.

We further achieved statistically significant improvement ($p < 0.0001$; Pearson's chi-squared test (Plackett, 1983)) over the best pervious result on 680 tags: from 94.65% to 97.98%, which represents 62.24% error reduction at the token-level.

---

[2]We used the pre-trained TreeTagger; for the rest, we report the accuracy given on the Webpage of the BulTreeBank: www.bultreebank.org/taggers/taggers.html

#	Lexicon (source of)	Linguistic Rules (applied to filter): (a) the lexicon features	(b) the output tags	Beam size	Accuracy (%) Sentence-level	Token-level
1	–	–	–	1	52.95	95.72
2	–	–	yes	1	64.50	97.20
3	features	–	–	1	70.40	97.83
4	features	yes	–	1	70.30	97.80
5	features	yes, *for test only*	–	1	70.40	97.84
6	features	–	yes	1	71.34	97.91
7	features	yes	yes	1	71.69	97.93
8	features	yes	yes	3	71.94	97.98

Table 4: **Evaluation results on the test dataset.** Line 1 shows the evaluation results when using features derived from the text corpus only; these features are used by all systems in the table. Line 2 further uses the contextual linguistic rules to limit the set of possible POS tags that can be predicted. Note that these rules (1) consult the lexicon, and (2) always predict a single POS tag. Line 3 uses the POS tags listed in the lexicon as features, i.e., as soft suggestions only. Line 4 is like line 3, but the list of feature-tags proposed by the lexicon is filtered by the contextual linguistic rules. Line 5 is like line 4, but the linguistic rules filtering is only applied at test time; it is not done on training. Lines 6 and 7 are similar to lines 3 and 4, respectively, but here the linguistic rules are further applied to limit the set of possible POS tags that can be predicted, i.e., the rules are used as hard constraints. Finally, line 8 is like line 7, but here the beam size is increased to 3.

Overall, we improved over almost all previously published results. Our accuracy is second only to the manual rules approach of Dojchinova and Mihov (2004). Note, however, that they used 40 tags only, i.e., their inventory is 17 times smaller than ours. Moreover, they have optimized their tagset specifically to achieve very high POS tagging accuracy by choosing not to attempt to resolve some inherently hard systematic ambiguities, e.g., they do not try to choose between second and third person past singular verbs, whose inflected forms are identical in Bulgarian and hard to distinguish when the subject is not present (Bulgarian is a pro-drop language).

In order to compare our results more closely to the smaller tagsets in Table 5, we evaluated our best model with respect to (a) the first letter of the tag only (which is part-of-speech only, no morphological information; 13 tags), e.g., Ncmsf becomes N, and (b) the first two letters of the tag (POS + limited morphological information; 49 tags), e.g., Ncmsf becomes Nc. This yielded 99.30% accuracy for (a) and 98.85% for (b). The latter improves over (Dojchinova and Mihov, 2004), while using a bit larger number of tags.

Our best token-level accuracy of 97.98% is comparable and even slightly better than the state-of-the-art results for English: 97.33% when using Penn Treebank data only (Shen et al., 2007), and 97.50% for Penn Treebank plus some additional unlabeled data (Søgaard, 2011). Of course, our results are only indirectly comparable to English.

Still, our performance is impressive because (1) our model is trained on 253,526 tokens only while the standard training sections 0-18 of the Penn Treebank contain a total of 912,344 tokens, i.e., almost four times more, and (2) we predict 680 rather than just 48 tags as for the Penn Treebank, which is 14 times more.

Note, however, that (1) we used a large external morphological lexicon for Bulgarian, which yielded about 50% error reduction (without it, our accuracy was 95.72% only), and (2) our train/dev/test sentences are generally shorter, and thus arguably simpler for a POS tagger to analyze: we have 17.4 words per test sentence in the Bul-TreeBank vs. 23.7 in the Penn Treebank.

Our results also compare favorably to the state-of-the-art results for other morphologically complex languages that use large tagsets, e.g., 95.2% for Czech with 1,400+ tags (Hajič et al., 2001), 92.1% for Icelandic with 639 tags (Dredze and Wallenberg, 2008), 97.6% for Arabic with 139 tags (Habash and Rambow, 2005).

8 Error Analysis

In this section, we present error analysis with respect to the impact of the POS tagger's performance on other processing steps in a natural language processing pipeline, such as lemmatization and syntactic dependency parsing.

First, we explore the most frequently confused pairs of tags for our best-performing POS tagging system; these are shown in Table 6.

Tool/Authors	Method	# Tags	Accuracy (token-level, %)
*TreeTagger	Decision Trees	680	89.21
*ACOPOST	Memory-based Learning	680	89.91
*SVMtool	Support Vector Machines	680	92.22
*TnT	Hidden Markov Model	680	92.53
(Georgiev et al., 2009)	Maximum Entropy	680	90.34
(Simov and Osenova, 2001)	Recurrent Neural Network	160	92.87
(Georgiev et al., 2009)	Maximum Entropy	95	94.43
(Savkov et al., 2011)	SVM + Lexicon + Rules	680	94.65
(Tanev and Mitkov, 2002)	Manual Rules	303	95.00(=P=R)
(Simov and Osenova, 2001)	Recurrent Neural Network	15	95.17
(Dojchinova and Mihov, 2004)	Transformation-based Learning	40	95.50
(Dojchinova and Mihov, 2004)	Manual Rules + Lexicon	40	98.40
	Guided Learning	680	95.72
	Guided Learning + Lexicon	680	97.83
This work	Guided Learning + Lexicon + Rules	680	97.98
	Guided Learning + Lexicon + Rules	*49*	*98.85*
	Guided Learning + Lexicon + Rules	*13*	*99.30*

Table 5: **Comparison to previous work for Bulgarian.** The first four lines report evaluation results for various standard POS tagging tools, which were retrained and evaluated on the BulTreeBank. The following lines report token-level accuracy for previously published work, as compared to our own experiments using guided learning.

We can see that most of the wrong tags share the same part-of-speech (indicated by the initial uppercase letter), such as V for verb, N for noun, etc. This means that most errors refer to the morphosyntactic features. For example, personal or impersonal verb; definite or indefinite feminine noun; singular or plural masculine adjective, etc. At the same time, there are also cases, where the error has to do with the part-of-speech label itself. For example, between an adjective and an adverb, or between a numeral and an indefinite pronoun.

We want to use the above tagger to develop (1) a rule-based lemmatizer, using the morphological lexicon, e.g., as in (Plisson et al., 2004), and (2) a dependency parser like MaltParser (Nivre et al., 2007), trained on the dependency part of the BulTreeBank. We thus study the potential impact of wrong tags on the performance of these tools.

The lemmatizer relies on the lexicon and uses string transformation functions defined via two operations – *remove* and *concatenate*:

```
if tag = Tag then
      {remove OldEnd; concatenate NewEnd}
```

where Tag is the tag of the wordform, OldEnd is the string that has to be removed from the end of the wordform, and NewEnd is the string that has to be concatenated to the beginning of the wordform in order to produce the lemma.

Here is an example of such a rule:

```
if tag = Vpitf-o1s then
      {remove ох; concatenate a}
```

The application of the above rule to the past simple verb form четох ('I read') would remove ох, and then concatenate а. The result would be the correct lemma чета ('to read').

Such rules are generated for each wordform in the morphological lexicon; the above functional representation allows for compact representation in a finite state automaton. Similar rules are applied to the unknown words, where the lemmatizer tries to guess the correct lemma.

Obviously, the applicability of each rule crucially depends on the output of the POS tagger. If the tagger suggests the correct tag, then the wordform would be lemmatized correctly. Note that, in some cases of wrongly assigned POS tags in a given context, we might still get the correct lemma. This is possible in the majority of the erroneous cases in which the part-of-speech has been assigned correctly, but the wrong grammatical alternative has been selected. In such cases, the error does not influence lemmatization.

In order to calculate the proportion of such cases, we divided each tag into two parts: (a) grammatical features that are common for all wordforms of a given lemma, and (b) features that are specific to the wordform.

Freq.	Gold Tag	Proposed Tag
43	Ansi	Dm
23	Vpitf-r3s	Vnitf-r3s
16	Npmsh	Npmsi
14	Vpiif-r3s	Vniif-r3s
13	Npfsd	Npfsi
12	Dm	Ansi
12	Vpitcam-smi	Vpitcao-smi
12	Vpptf-r3p	Vpitf-r3p
11	Vpptf-r3s	Vpptf-o3s
10	Mcmsi	Pfe-os-mi
10	Ppetas3n	Ppetas3m
10	Ppetds3f	Psot–3–f
9	Npnsi	Npnsd
9	Vpptf-o3s	Vpptf-r3s
8	Dm	A-pi
8	Ppxts	Ppxtd
7	Mcfsi	Pfe-os-fi
7	Npfsi	Npfsd
7	Ppetas3m	Ppetas3n
7	Vnitf-r3s	Vpitf-r3s
7	Vpitcam-p-i	Vpitcao-p-i

Table 6: **Most frequently confused pairs of tags.**

The part-of-speech features are always determined by the lemma. For example, Bulgarian verbs have the lemma features *aspect* and *transitivity*. If they are correct, then the lemma is predicted also correctly, regardless of whether correct or wrong on the grammatical features. For example, if the verb participle form (aorist or imperfect) has its correct aspect and transitivity, then it is lemmatized also correctly, regardless of whether the imperfect or aorist features were guessed correctly; similarly, for other error types. We evaluated these cases for the 711 errors in our experiment, and we found that 206 of them (about 29%) were non-problematic for lemmatization.

For the MaltParser, we encode most of the grammatical features of the wordforms as specific features for the parser. Hence, it is much harder to evaluate the problematic cases due to the tagger. Still, we were able to make an estimation of some cases. Our strategy was to ignore the grammatical features that do not always contribute to the syntactic behavior of the wordforms. Such grammatical features for the verbs are *aspect* and *tense*. Thus, proposing perfective instead of imperfective for a verb or present instead of past tense would not cause problems for the MaltParser. Among our 711 errors, 190 cases (or about 27%) were not problematic for parsing.

Finally, we should note that there are two special classes of tokens for which it is generally hard to predict some of the grammatical features: (1) abbreviations and (2) numerals written with digits. In sentences, they participate in agreement relations only if they are pronounced as whole phrases; unfortunately, it is very hard for the tagger to guess such relations since it does not have at its disposal enough features, such as the inflection of the numeral form, that might help detect and use the agreement pattern.

9 Conclusion and Future Work

We have presented experiments with part-of-speech tagging for Bulgarian, a Slavic language with rich inflectional and derivational morphology. Unlike most previous work for this language, which has limited the number of possible tags, we used a very rich tagset of 680 morpho-syntactic tags as defined in the BulTreeBank. By combining a large morphological lexicon with prior linguistic knowledge and guided learning from a POS-annotated corpus, we achieved accuracy of 97.98%, which is a significant improvement over the state-of-the-art for Bulgarian. Our token-level accuracy is also comparable to the best results reported for English.

In future work, we want to experiment with a richer set of features, e.g., derived from unlabeled data (Søgaard, 2011) or from the Web (Umansky-Pesin et al., 2010; Bansal and Klein, 2011). We further plan to explore ways to decompose the complex Bulgarian morpho-syntactic tags, e.g., as proposed in (Simov and Osenova, 2001) and (Smith et al., 2005). Modeling long-distance syntactic dependencies (Dredze and Wallenberg, 2008) is another promising direction; we believe this can be implemented efficiently using posterior regularization (Graca et al., 2009) or expectation constraints (Bellare et al., 2009).

References

Mohit Bansal and Dan Klein. 2011. Web-scale features for full-scale parsing. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, ACL-HLT '10, pages 693–702, Portland, Oregon, USA.

Kedar Bellare, Gregory Druck, and Andrew McCallum. 2009. Alternating projections for learning with expectation constraints. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, UAI '09, pages 43–50, Montreal, Quebec, Canada.

Thorsten Brants. 2000. TnT – a statistical part-of-speech tagger. In *Proceedings of the Sixth Applied Natural Language Processing*, ANLP '00, pages 224–231, Seattle, Washington, USA.

Eric Brill. 1995. Transformation-based error-driven learning and natural language processing: a case study in part-of-speech tagging. *Comput. Linguist.*, 21:543–565.

Michael Collins and Brian Roark. 2004. Incremental parsing with the perceptron algorithm. In *Proceedings of the 42nd Meeting of the Association for Computational Linguistics, Main Volume*, ACL '04, pages 111–118, Barcelona, Spain.

Michael Collins. 2002. Discriminative training methods for hidden Markov models: theory and experiments with perceptron algorithms. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '02, pages 1–8, Philadelphia, PA, USA.

Walter Daelemans, Jakub Zavrel, Peter Berck, and Steven Gillis. 1996. MBT: A memory-based part of speech tagger generator. In Eva Ejerhed and Ido Dagan, editors, *Fourth Workshop on Very Large Corpora*, pages 14–27, Copenhagen, Denmark.

Veselka Dojchinova and Stoyan Mihov. 2004. High performance part-of-speech tagging of Bulgarian. In Christoph Bussler and Dieter Fensel, editors, *AIMSA*, volume 3192 of *Lecture Notes in Computer Science*, pages 246–255. Springer.

Mark Dredze and Joel Wallenberg. 2008. Icelandic data driven part of speech tagging. In *Proceedings of the 44th Annual Meeting of the Association of Computational Linguistics: Short Papers*, ACL '08, pages 33–36, Columbus, Ohio, USA.

Georgi Georgiev, Preslav Nakov, Petya Osenova, and Kiril Simov. 2009. Cross-lingual adaptation as a baseline: adapting maximum entropy models to Bulgarian. In *Proceedings of the RANLP'09 Workshop on Adaptation of Language Resources and Technology to New Domains*, AdaptLRTtoND '09, pages 35–38, Borovets, Bulgaria.

Jesús Giménez and Lluís Màrquez. 2004. SVMTool: A general POS tagger generator based on support vector machines. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, LREC '04, Lisbon, Portugal.

Joao Graca, Kuzman Ganchev, Ben Taskar, and Fernando Pereira. 2009. Posterior vs parameter sparsity in latent variable models. In Yoshua Bengio, Dale Schuurmans, John D. Lafferty, Christopher K. I. Williams, and Aron Culotta, editors, *Advances in Neural Information Processing Systems 22*, NIPS '09, pages 664–672. Curran Associates, Inc., Vancouver, British Columbia, Canada.

Nizar Habash and Owen Rambow. 2005. Arabic tokenization, part-of-speech tagging and morphological disambiguation in one fell swoop. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, ACL '05, pages 573–580, Ann Arbor, Michigan.

Jan Hajič, Pavel Krbec, Pavel Květoň, Karel Oliva, and Vladimír Petkevič. 2001. Serial combination of rules and statistics: A case study in Czech tagging. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, ACL '01, pages 268–275, Toulouse, France.

Jan Hajič. 1998. Building a Syntactically Annotated Corpus: The Prague Dependency Treebank. In Eva Hajičová, editor, *Issues of Valency and Meaning. Studies in Honor of Jarmila Panevová*, pages 12–19. Prague Karolinum, Charles University Press.

Erhard W. Hinrichs and Julia S. Trushkina. 2004. Forging agreement: Morphological disambiguation of noun phrases. *Research on Language & Computation*, 2:621–648.

Stig Johansson, Eric Atwell, Roger Garside, and Geoffrey Leech, 1986. *The Tagged LOB Corpus: Users' manual*. ICAME, The Norwegian Computing Centre for the Humanities, Bergen University, Norway.

Hristo Krushkov. 1997. *Modelling and building machine dictionaries and morphological processors (in Bulgarian)*. Ph.D. thesis, University of Plovdiv, Faculty of Mathematics and Informatics, Plovdiv, Bulgaria.

Henry Kučera and Winthrop Nelson Francis. 1967. *Computational analysis of present-day American English*. Brown University Press, Providence, RI.

John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the 18th International Conference on Machine Learning*, ICML '01, pages 282–289, San Francisco, CA, USA.

Mohamed Maamouri, Ann Bies, Hubert Jin, and Tim Buckwalter. 2003. Arabic Treebank: Part 1 v 2.0. LDC2003T06.

Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: the Penn Treebank. *Comput. Linguist.*, 19:313–330.

Joakim Nivre, Johan Hall, Jens Nilsson, Atanas Chanev, Gülsen Eryigit, Sandra Kübler, Svetoslav Marinov, and Erwin Marsi. 2007. MaltParser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering*, 13(2):95–135.

Jörgen Pind, Fridrik Magnússon, and Stefán Briem. 1991. The Icelandic frequency dictionary. Technical report, The Institute of Lexicography, University of Iceland, Reykjavik, Iceland.

Robin L. Plackett. 1983. Karl Pearson and the Chi-Squared Test. *International Statistical Review / Revue Internationale de Statistique*, 51(1):59–72.

Joël Plisson, Nada Lavrač, and Dunja Mladenić. 2004. A rule based approach to word lemmatization. In *Proceedings of the 7th International Multiconference: Information Society*, IS '2004, pages 83–86, Ljubljana, Slovenia.

Dimitar Popov, Kiril Simov, and Svetlomira Vidinska. 1998. *Dictionary of Writing, Pronunciation and Punctuation of Bulgarian Language (in Bulgarian)*. Atlantis KL, Sofia, Bulgaria.

Dimityr Popov, Kiril Simov, Svetlomira Vidinska, and Petya Osenova. 2003. *Spelling Dictionary of Bulgarian*. Nauka i izkustvo, Sofia, Bulgaria.

Adwait Ratnaparkhi. 1996. A maximum entropy model for part-of-speech tagging. In Eva Ejerhed and Ido Dagan, editors, *Fourth Workshop on Very Large Corpora*, pages 133–142, Copenhagen, Denmark.

Aleksandar Savkov, Laska Laskova, Petya Osenova, Kiril Simov, and Stanislava Kancheva. 2011. A web-based morphological tagger for Bulgarian. In Daniela Majchráková and Radovan Garabík, editors, *Slovko 2011. Sixth International Conference. Natural Language Processing, Multilinguality*, pages 126–137, Modra/Bratislava, Slovakia.

Helmut Schmid. 1994. Probabilistic part-of-speech tagging using decision trees. In *International Conference on New Methods in Language Processing*, pages 44–49, Manchester, UK.

Ingo Schröder. 2002. A case study in part-of-speech-tagging using the ICOPOST toolkit. Technical Report FBI-HH-M-314/02, Department of Computer Science, University of Hamburg.

Libin Shen, Giorgio Satta, and Aravind Joshi. 2007. Guided learning for bidirectional sequence classification. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, ACL '07, pages 760–767, Prague, Czech Republic.

Kiril Simov and Petya Osenova. 2001. A hybrid system for morphosyntactic disambiguation in Bulgarian. In *Proceedings of the EuroConference on Recent Advances in Natural Language Processing*, RANLP '01, pages 5–7, Tzigov chark, Bulgaria.

Kiril Simov and Petya Osenova. 2004. BTB-TR04: BulTreeBank morphosyntactic annotation of Bulgarian texts. Technical Report BTB-TR04, Bulgarian Academy of Sciences.

Kiril Ivanov Simov, Alexander Simov, Milen Kouylekov, Krasimira Ivanova, Ilko Grigorov, and Hristo Ganev. 2003. Development of corpora within the CLaRK system: The BulTreeBank project experience. In *Proceedings of the 10th conference of the European chapter of the Association for Computational Linguistics*, EACL '03, pages 243–246, Budapest, Hungary.

Kiril Simov, Petya Osenova, and Milena Slavcheva. 2004. BTB-TR03: BulTreeBank morphosyntactic tagset. Technical Report BTB-TR03, Bulgarian Academy of Sciences.

Noah A. Smith, David A. Smith, and Roy W. Tromble. 2005. Context-based morphological disambiguation with random fields. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 475–482, Vancouver, British Columbia, Canada.

Anders Søgaard. 2011. Semi-supervised condensed nearest neighbor for part-of-speech tagging. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, ACL-HLT '10, pages 48–52, Portland, Oregon, USA.

Hristo Tanev and Ruslan Mitkov. 2002. Shallow language processing architecture for Bulgarian. In *Proceedings of the 19th International Conference on Computational Linguistics*, COLING '02, pages 1–7, Taipei, Taiwan.

Kristina Toutanova, Dan Klein, Christopher D. Manning, and Yoram Singer. 2003. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, NAACL '03, pages 173–180, Edmonton, Canada.

Yoshimasa Tsuruoka and Jun'ichi Tsujii. 2005. Bidirectional inference with the easiest-first strategy for tagging sequence data. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, HLT-EMNLP '05, pages 467–474, Vancouver, British Columbia, Canada.

Yoshimasa Tsuruoka, Yusuke Miyao, and Jun'ichi Kazama. 2011. Learning with lookahead: Can history-based models rival globally optimized models? In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, ACL-HLT '10, pages 238–246, Portland, Oregon, USA.

Shulamit Umansky-Pesin, Roi Reichart, and Ari Rappoport. 2010. A multi-domain web-based algorithm for POS tagging of unknown words. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, COLING '10, pages 1274–1282, Beijing, China.