YARN: Spinning-in-Progress

Pavel Braslavski
Ural Federal University
Yekaterinburg, Russia
pbras@yandex.ru

Dmitry Ustalov
Ural Federal University
Yekaterinburg, Russia
dmitry.ustalov@urfu.ru

Mikhail Mukhin
Ural Federal University
Yekaterinburg, Russia
mfly@sky.ru

Yuri Kiselev
Yandex
Yekaterinburg, Russia
yurikiselev@yandex-team.ru

Abstract

YARN (Yet Another RussNet), a project started in 2013, aims at creating a large open WordNet-like thesaurus for Russian by means of crowdsourcing. The first stage of the project was to create noun synsets. Currently, the resource comprises 100K+ word entries and 46K+ synsets. More than 200 people have taken part in assembling synsets throughout the project. The paper describes the linguistic, technical, and organizational principles of the project, as well as the evaluation results, lessons learned, and the future plans.

1 Introduction

The Global WordNet Association website lists 76 wordnets for 70 different languages[1], including multilingual resources. Although the table mentions as many as three wordnets for Russian, unfortunately no open Russian thesaurus of an acceptable quality and size is still available.

The Yet Another RussNet (YARN) project[2] started in 2013. It aims at creating a comprehensive and open thesaurus for Russian. From the linguistics point of view, the proposed thesaurus has rather a traditional structure: it consists of synsets—groups of near-synonyms corresponding to a concept, while synsets are linked to each other, primarily via hierarchical hyponymic/hypernymic relations.

YARN intends to cover Russian nouns, verbs and adjectives. Following the divide and conquer approach, we treat synset assembly and relationship establishing separately.

The main difference between YARN and the previous projects is that YARN is based on crowdsourcing. We hope that the crowdsourcing approach will make it possible to create a resource of a satisfactory quality and size in the foreseeable future and with limited financial resources. Our optimism is based both on the international practice and the recent examples of successful Russian NLP projects fueled by volunteers. Another important distinction is that the editors do not build the thesaurus from scratch; instead, they use "raw data" as the input. These "raw data" stem from pre-processed dictionaries, Wiktionary, Wikipedia, and text corpora. More than 200 people have taken part in the synset assembly in the course of the project. Currently, the resource comprises 100K+ word entries and 46K+ synsets that are available under CC BY-SA license.

The paper describes the main linguistic and organizational principles of YARN, the tools developed, and the results of the current content evaluation. We also point to some pitfalls of the chosen crowdsourcing methodology and discuss how we could address them in the future.

2 Related Work

In this section, we briefly survey projects aimed at creation of WordNet-like semantic resources for Russian, describe peculiarities of other thesauri for Slavic languages, and systematize different crowdsourcing approaches to building lexicographic resources.

[1] http://globalwordnet.org/wordnets-in-the-world/

[2] http://russianword.net/en/, not to be confused with a Hadoop subsystem.

2.1 Russian Thesauri

The **RussNet** project[3] was launched in 1999 at Saint-Petersburg university (Azarova et al., 2002). According to the RussNet developers, the resource currently contains about 40K word entries, 30K synsets, and 45K semantic relations. However, this data is not encoded in a uniform format and cannot be published or used in a NLP application in its current form.

RuThes is probably the most successful WordNet-like resource for Russian (Loukachevitch, 2011). It has been developing since 2002, and now contains 158K lexical units constituting 55K concepts. RuThes is a proprietary resource; however a subset of it was published recently[4]. The main hurdle for a wider use of the resource is a restrictive license and the fact that the data in XML format can be obtained by request only.

Another resource—**RussianWordNet**—was a result of a fully automatic translation of the Princeton WordNet (PWN) into Russian undertaken in 2003 and is freely available[5] under the PWN license. The approach based on bilingual dictionaries, parallel corpora, and dictionaries of synonyms resulted in the translation of about 45% of the PWN entries. The thesaurus contains 18K nouns, 6K adverbs, 5.5K verbs, and 1.8K adverbs; no systematic quality assessments of the obtained data were performed (Gelfenbeyn et al., 2003). Another attempt to translate the PWN into Russian, in this case—in a semi-automatic fashion—is the **Russian Wordnet** project (Balkova et al., 2004) started in 2003, but its deliverables are not available to the general public.

Russian Wiktionary[6] can be seen as an ersatz of a proper thesaurus, since along with definitions it contains—though marginally—semantic relations. Wikokit project[7] allows handling Wiktionary data as a relational database (Krizhanovsky and Smirnov, 2013). Russian Wiktionary contains about 190K word entries and 70K synonym relations as of September, 2015.

The **Universal Networking Language**[8] project is dedicated to the development of a computer language that replicates the functions of natural languages. The Russian version of its semantic network—the Universal Dictionary of Concepts—contains approximately 62K universal words (UWs) and 90K links between them and is available[9] under CC BY-SA license.

One of the recent trends is the creation of semantic resources in a fully automatic manner, where collaboratively created resources like Wikipedia and Wiktionary are used as the input. A striking example of this approach is **BabelNet**, a very large automatically generated multilingual thesaurus (Navigli and Ponzetto, 2012); the Russian part of BabelNet consists of 2.37M lemmas, 1.35M synsets, and 3.7M word senses[10]. The data is accessible through an API under CC BY-NC-SA 3.0 license. No evaluation of the Russian data has been performed yet.

As can be seen from the survey, no open human-crafted wordnet for Russian is available so far. Automatically created resources are freely available and potentially have very good coverage, but their quality is disputed.

2.2 Thesauri of Other Slavic Languages

Slavic languages are highly inflectional and have a rich derivation system. The survey of wordnets for Czech (Pala and Smrž, 2004), Polish (Maziarz et al., 2014) and Ukrainian (Anisimov et al., 2013) shows that in each case a special attention is paid to dealing with the morphological characteristics. For instance, plWordNet features a versatile system of relations with dozens of subtypes of relations between synsets and lexical units, many of which reflect derivational relations.

2.3 Crowdsourcing Language Resources

Crowdsourcing, a human-computer technique for collaborative problem solving by online communities, has gained high popularity since its inception in the mid 2000's (Kittur et al., 2013). Creation and expansion of linguistic resources using crowdsourcing became a trend in recent years as shown by Gurevych and Kim (2013).

Despite the ongoing unabated discussions about the types, merits and limitations of crowdsourcing (Wang et al., 2013), we consider the following genres of crowdsourcing: *wisdom of the crowds* (WOTC), *mechanized labor* (MLAB) and *games with a purpose* (GWAPS).

[3] http://project.phil.spbu.ru/RussNet/
[4] http://labinform.ru/pub/ruthes/
[5] http://wordnet.ru/
[6] http://ru.wiktionary.org/
[7] https://github.com/componavt/wikokit
[8] http://www.undl.org/
[9] https://github.com/dikonov/Universal-Dictionary-of-Concepts
[10] http://babelnet.org/stats

In the WOTC genre, the resource is constructed *explicitly* by a crowd of volunteers that collaborates in an online editing environment. Their participation is mostly altruistic and a participant's benefit is either self-exaltation or self-promotion of any kind. Successful examples of this genre are Wikipedia and Wiktionary. The primary issues of such resources are vandalism and "edit wars", which are usually resolved by edit patrolling and edit protection.

In the MLAB genre, the resource is created *implicitly* by the workers who submit answers to simple tasks provided by the requester. This genre is proven to be effective in many practical applications. For instance, Lin and Davis (2010) extracted ontological structure from social tagging systems and engaged workers in evaluation. Rumshisky (2011) used crowdsourcing to create an empirically-derived sense inventory and proposed an approach for automated assessment of the obtained data. Biemann (2013) described how workers can contribute to thesaurus creation by solving simple lexical substitution tasks. Most of these studies have been conducted on the commodity platforms like Amazon Mechanical Turk[11] (MTurk) and CrowdFlower[12]. Unfortunately, MTurk can hardly be used for tasks implying the knowledge of Russian because: (1) there are virtually no workers from Russia presented on the platform (Pavlick et al., 2014), and (2) a requester must have a U.S. billing address to submit tasks[13]. Having no access to the global online labor marketplaces is a serious obstacle to paying the workers due to the requirements of the local legislation of Russia. However, projects like OpenCorpora are trying to work around this problem by developing custom crowdsourcing platforms and effectively appealing to *altruism* instead of money reward (Bocharov et al., 2013). Since such altruistic mechanized labor does not imply money reward, it is not prone to spam, where an unfair worker may permanently submit random answers instead of sensible ones.

In the GWAPS genre, the crowdsourcing process is embedded into a multi-player game, in which the players have to accomplish various goals by creating new data items to win the game. Although such games are attractive and entertaining, game development is an expensive and complex kind of activity that may be feasible only for large-scale annotation projects. The examples here are Phrase Detectives[14] and JeuxDeMots[15].

3 YARN Essentials

YARN is conceptually similar to Princeton WordNet (Fellbaum, 1998) and its followers: it consists of synsets—groups of quasi-synonyms corresponding to a concept. Concepts are linked to each other, primarily via hierarchical hyponymic/hypernymic relationships.

3.1 YARN Structure

Each single-word entry in YARN is characterized by the grammatical features (the types of POS and inflection) according to Zaliznyak's dictionary (1977). Synsets may include single-word entries {суффикс (*suffix*)}, multi-word expressions {подводная лодка (*submarine*)}, and abbreviations {ПО (программное обеспечение, *software*)}. Synsets may contain a definition (*gloss* in terms of PWN). Additionally, definitions can be attached to individual words in a synset—these definitions are inherited from the dictionary data and specify a word meaning, but cannot serve as a good definition for the whole synset. "Empty synsets" (i.e. containing no words) that correspond to a non-lexicalized concept are legitimate and help to create a more harmonious hierarchy of synsets.

Each word in a synset can be accompanied by one or more usage examples. Words within synsets can attach labels from the five categories: *emotional*, *stylistic*, *chronological*, *domain/territorial*, and *semantic* (28 labels in total). This list is a result of the systematization of large and diverse Wiktionary label set. One of the synset words can be marked as the head word. Its sense is stylistically neutral, and it encompasses the meanings of the whole synset, e.g. {армия (*army*), войска (*troops*), вооружённые силы (*armed forces*)}. Each synset may belong to a domain, e.g. {кино (*movie*), кинофильм (*movie picture*), фильм (*film*)} → "Arts", {думать (*to think*), размышлять (*to ponder*)} → "Intellect".

The vertical, hypo-/hypernymic relations between synsets are decisive for the hierarchical

[11] https://www.mturk.com/mturk/welcome

[12] http://crowdflower.com/

[13] https://requester.mturk.com/help/faq#can_international_requesters_use_mturk

[14] https://anawiki.essex.ac.uk/phrasedetectives/

[15] http://www.jeuxdemots.org/

macrostructure of the thesaurus. The root of the YARN hierarchy is {предмет (*entity*), объект (*object*), вещь (*thing*)}; the second level is represented by {физическое явление (*physical phenomenon*)}, {отвлечённое понятие, абстрактное понятие, абстракция (*an abstraction*)}, {совокупность, набор (*set*), группа (*group*)}, {воображаемое, представляемое (*imaginary*)}. We elaborated 4–5 top levels for each part of speech.

The vertical links in YARN are also formed by the meronymy relations (the part-whole relations): ноздря (*nostrill*)—нос (*nose*)—лицо (*face*)—голова (*head*). The antonymy relationship connects specific words in the context of corresponding synsets. For example, the verb прибыть (*to arrive*) is the antonym of the verb отбыть (*to depart*), but not of направиться (*to head somewhere*) and the other words in the synset.

In the future, YARN will reflect the cross-POS relations between derivates: {двигаться (*to move*), движение (*movement*)}, {лес (*forest*), лесной ($forest_{adj}$)}. It will be significant for the word pairs with a minimum difference in senses.

3.2 Raw Data

As the "raw data" for the thesaurus construction we employed existing resources such as Wiktionary (which constituted the core of the input data), Wikipedia (redirects), the aforementioned result of the automatic translation of the PWN, the Universal Dictionary of Concepts, and the data from two dictionaries in the public domain. We also implicitly use the data from the Russian National Corpus (RNC) so that the corpus statistics influence the queue of words presented to the editors. Wikipedia and RNC were also used to compile the list of multi-word expressions to be included in the resource.

3.3 User Interface

Our initial approach to synset building is based on the WOTC inspired by the highly successful examples of Wikipedia and Wiktionary: our editors assemble synsets using word lists and definitions from dictionaries as the "raw data". Technically, virtually everybody can edit the YARN data—one needs only to login using a social network account. However, the task design implies minimal lexicographical skills and is more complicated than an average task offered for instance to MTurk workers. Our target editors are college or university students, preferably from the linguistics departments, who are native Russian speakers. It is desirable that students receive instructions from a university teacher and may seek their advice in complex cases. YARN differentiates the two levels of contributors—line editors and moderators. Moderators are authorized to approve thesaurus elements thus excluding them being modified by line editors.

The current synset editing interface can be accessed online[16]; its main window is presented in Figure 1. The "raw data" are placed on the left-hand side of the interface: definitions of the initial word and examples, and possible synonyms for each of the meanings, with definitions and examples for each of the synonyms. The right-hand part represents the resulting synsets including words, definitions, and examples. In principle, an editor can assemble a "minimal" synset from the dictionary "raw data" simply with several mouse clicks, without any typing.

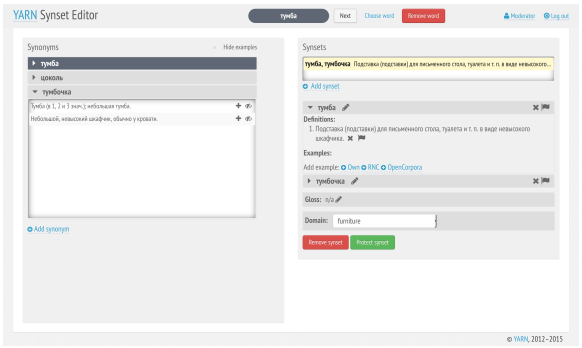

Figure 1: YARN synset assembly interface (the interface captions are translated into English for the convenience of the readers; originally all interface elements are in Russian).

Synset assembly begins with a word, or "synset starter". The editor selects an item from the list of words ranked by decreasing frequency; the already processed words are shaded. The editor can go through the words one after another or choose an arbitrary word using the search box. The top-left pane displays definitions of the initial word and usage examples if any. The possible synonyms of the initial word are listed on the bottom-left pane; they in turn contain their definitions and examples. The top-right pane displays a list of synsets containing the initial word. The editor can copy definitions and usage examples of the initial word

[16] http://russianword.net/editor

```
<synsetEntry id="s9439" author="122" version="29" timestamp="2014-11-17T07:49:46Z">
  <word ref="w9244">
    <definition source="ru.wiktionary" url="http://ru.wiktionary.org/wiki/суп">
      Жидкое кушанье, обычно представляющее собой отвар с приправами и употребляемое как первое блюдо.
    </definition>
    <example source="'Путешествие в седьмую сторону света', 2000, НКРЯ.">
      Был обед - овощной суп и курица на второе.
    </example>
  </word>
  <word ref="w40078"/>
  <word ref="w2893"/>
</synsetEntry>
```

Figure 2: XML representation of the synset {суп, бульон, похлёбка (*soup*)}.

from the top-left pane of the interface to the current synset by clicking the mouse. From the synonyms pane one can transfer words along with their definitions and examples. The editor can add a new word to the list of synonyms; it will appear with dictionary definitions and examples if presented in the parsed data. If the editor is not satisfied with the collected definitions, they can create a new one—either from scratch or based on one of the existing descriptions. Using search in the Russian National Corpus[17] and OpenCorpora[18], the editor can add usage examples. Additionally, a word or a definition within a synset can be flagged as "main", and be provided with labels. All synset edits are tracked and stored in the database along with the timestamps and the editor ID.

As a pilot study showed, editors spent about two minutes on average to compile a non-trivial synset, i.e. containing more than a single word. The top contributors demonstrated a learning effect: the average time per synset tended to decrease as the editor proceeded through the tasks, see Braslavski et al. (2014) for details.

Our next goal is to lower the threshold of participation in the data annotation and thus—to increase the number of participants. To do this, we are developing a mobile application in the MLab genre that is aimed at gathering "raw synsets": users are presented with a series of sentences with highlighted words and lists of possible contextual substitutes. This approach is similar to the experiment described in (Biemann, 2013).

3.4 Implementation Details

The YARN data are stored in a centralized database that can be accessed through a web interface. In addition, distributed teams can work directly with the database through an API. The database is periodically exported to XML format. Although the original dictionaries and thesauri were coming in different formats, we decided to develop a custom XML schema for data export[19]. We believe that XML format provides sufficient flexibility and preserves the connection to the internal data representation. The developed format is modular, as different types of objects (lexical units, synsets, and relationships) are described separately. The proposed format is somewhat similar to the Lexical Markup Framework (LMF)[20] approach, although the YARN format does not refer to the latter directly. All editing actions (in fact, aggregated "action chunks") are stored in the database. The YARN format stores the revision history analogously to the OpenStreetMap XML format[21]. A synset structure is illustrated in Figure 2.

The YARN software is implemented using Ruby on Rails framework. All data are stored in a PostgreSQL database. The user interface is implemented as a browser JavaScript application, which interacts with the back-end via JSON API. User authentication is performed through an OAuth endpoint provided by Facebook, VK and GitHub. The entire source code of the project is available in a GitHub repository[22].

3.5 Current State and Problems

The current version of the the YARN (September 2015) contains 44K synsets that consist of 48K words and 5.4K multi-word expressions; 838 words carry labels; 2.6K words are provided with at least one usage example (there are 4.2K examples in total). The resource contains 2.5K synset-level and 8.3K word-level definitions. The synset size distribution is presented in Figure 3.

[17] http://ruscorpora.ru/en/
[18] http://opencorpora.org/
[19] https://github.com/russianwordnet/yarn-formats/
[20] http://www.lexicalmarkupframework.org/
[21] http://www.openstreetmap.org/
[22] https://github.com/russianwordnet

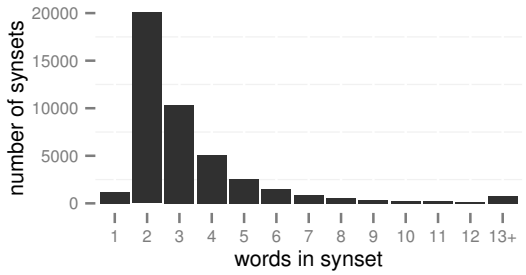

Figure 3: Synset distribution by size.

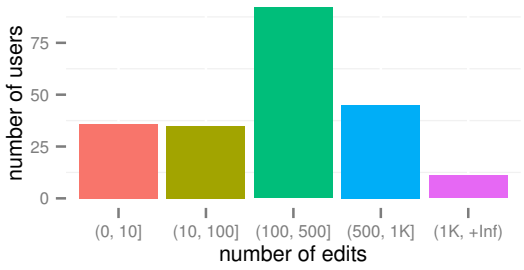

Figure 4: Distribution of users by edit count.

More than 200 people have taken part in editing YARN in the course of the project; the distribution of users by activity is shown in Figure 4. Whereas we consider the early experiment under a controlled crowd to be successful, we found the three significant problems replicating over time: organization issues, synset duplication and hyponymy/synonymy confusion.

Organization Issues. The number of synsets was growing rapidly and moderators were not able to assess all the incoming edits. In order to work around this problem, we are experimenting with MLAB workflows.

Synset Duplication. Participants do not consult the other people's work, which results in creation of duplicate synsets like {авто (*auto*), автомобиль (*automobile*), машина (*car*)} and {машина (*car*), тачка (*ride*)}.

Hyponymy Confusion. In some cases the participants mix hyponymy and synonymy, which results in strange synsets like {мультфильм (*cartoon*), мультик (*cartoon*), аниме (*anime*)}.

4 Evaluation

We compared YARN with other Russian thesauri (Kiselev et al., 2015), which have been described in Section 2.1 (Table 1). Besides YARN, the only resource available for use is RuThes-lite, the commercial use of which requires licensing. It should be noted that although the lexicon of YARN represents 100K+ words, only half of them are included in synsets. Thus, we provide the latter number.

The number of concepts indicates that crowdsourcing is a promising approach for thesauri creation for the Russian language. Interestingly, YARN contains more concepts than RussNet, a project started in 1999. However, when comparing YARN and RuThes-lite, one may notice, that they have an approximately equal number of concepts, yet the number of words in the latter is twice bigger than in YARN. This implies the hypothesis that expert-built thesauri include richer lexis that could be covered by non-expert users. Hence, the YARN synset quality requires more thorough evaluation.

4.1 Synset Quality

Since YARN is created using crowdsourcing, it seems reasonable to apply this technique for evaluation purposes, too. In our experiments we used an open source engine for MLAB workflows (Ustalov, 2015). In order to estimate the quality of the current YARN synsets, we retrieved the 200 most frequently edited synsets. We asked four experts to assess the quality of each synset by rating them on the following scale: *Excellent*—the synset completely represents a concept, *Satisfactory*—the synset is related to the concept, but some words are missing or odd words are present, and *Bad*—the synset is either ambiguous or it does not represent any sensible concept.

We aggregated the 800 obtained answers using the majority voting strategy, where the ties are resolved by choosing the worst of two answers, e.g. given the same number of votes for both *Good* and *Bad*, the latter will be selected. This resulted in 103 synsets of *Excellent*, 70 of *Satisfactory* and 27 of *Bad* quality. The results are shown in Table 2. Values in column **MV** are the numbers of synsets per each of the three grades, values in the last three columns are the numbers of synsets grouped by answer diversity—*all the answers are the same* in **1**, *two different answers present* in **2**, and *the expert opinions divided* in **3**.

We also computed the alpha annotator reliability coefficient for ordinal values to estimate the

Table 1: Russian thesauri comparison.

	# of concepts	# of relations	# of words	Availability	Commercial Usage
RussNet	5.5K	8K	15K	No	No
Russian Wordnet	157K	—	124K	No	No
RuThes	55K	210K	158K	No	No
RuThes-lite	26K	108K	115K	Yes	No
YARN	44K	0	48.6K	Yes	Yes

Table 2: YARN synset quality.

	MV	1	2	3
Excellent	103	37	62	21
Satisfactory	70	3	43	11
Bad	27	0	12	11
Total	200	40	117	43

inter-rater reliability (Krippendorff, 2013). The Krippendorff's alpha is $\alpha = 0.202$ due to the skewness of the answer distribution: more than half of the answers (434) are *Excellent*, the numbers of *Satisfactory* and *Bad* answers are 253 and 113 correspondingly. Given these results, we treat the top 200 YARN synsets as sufficiently good. These evaluation results define the upper bound for the average quality of the resource in its current state. Ustalov (2014) showed that revision count is a good proxy for quality in the Russian Wiktionary that is created in a similar fashion.

4.2 Duplicate Synsets

Sometimes users create new synsets without investigating the current synsets presented in YARN. The main problem with this is the presence of multiple entries for the same concept in the resource. Detecting such concepts requires special effort because they are not described with identical synsets but with similar ones.

Hence, we had to develop a method for automatically retrieving duplicate synsets. It was based on the heuristics suggesting that any two synonyms uniquely define a concept. This is not always true, but it lets us discover duplicate synsets with a very good recall. To estimate it, we compared the senses of random 200+ synsets having two or more common words. It turned out that more than in 85% of the cases these pairs described the same sense.

However, we found out that non-linguists do not recognize subtle nuances of meaning that are noticeable to experts, so the non-linguists cannot significantly improve the quality of duplicate extraction. Thus, this method—considering any synsets having more than two common words as duplicates—allows to detect and merge identical concepts with a quality that is comparable to what can be achieved by volunteers.

5 Conclusion

The deliverables of YARN are available under the CC BY-SA 3.0 license on the project website[23] in XML, CSV, and RDF formats. So far, we have the following plans for the future work.

- Creating verb and adjective synsets.
- Establishing hierarchical links between synsets through validation of the relationships imported from Wiktionary and other resources.
- Development of automatic methods for generating hypotheses based on Wikipedia and large text corpora.
- Development of automatic methods for preparing "raw data", as well as for post-processing of annotation results produced by the crowd.
- Widening the audience of the project's participants through mobile applications and simpler tasks.
- Development of crowd management methods, such as automatic methods for evaluation of workers, task difficulty, and annotation results, the system of incentives, etc.

Acknowledgments

This work is supported by the Russian Foundation for the Humanities project no. 13-04-12020 "New Open Electronic Thesaurus for Russian". We are grateful to Yulia Badryzlova for proofreading the text. We would also like to thank the three anonymous reviewers, who offered very helpful suggestions.

[23] http://russianword.net/data

References

Anatoly Anisimov, Oleksandr Marchenko, Andrey Nikonenko, et al. 2013. Ukrainian WordNet: Creation and Filling. In *Flexible Query Answering Systems*, volume 8132 of *Lecture Notes in Computer Science*, pages 649–660. Springer Berlin Heidelberg.

Irina Azarova, Olga Mitrofanova, Anna Sinopalnikova, Maria Yavorskaya, and Ilya Oparin. 2002. RussNet: Building a Lexical Database for the Russian Language. In *Proc. of Workshop on WordNet Structures and Standardisation, and How These Affect WordNet Applications and Evaluation*, pages 60–64, Gran Canaria, Spain.

Valentina Balkova, Andrey Sukhonogov, and Sergey Yablonsky. 2004. Russian WordNet. In *Proceedings of the Second International WordNet Conference—GWC 2004*, pages 31–38, Brno, Czech Republic. Masaryk University Brno, Czech Republic.

Chris Biemann. 2013. Creating a system for lexical substitutions from scratch using crowdsourcing. *Language Resources and Evaluation*, 47(1):97–122.

Victor Bocharov, Svetlana Alexeeva, Dmitry Granovsky, et al. 2013. Crowdsourcing morphological annotation. In *Computational Linguistics and Intellectual Technologies: papers from the Annual conference "Dialogue"*, volume 12(19), pages 109–124, Moscow, Russia. RGGU.

Pavel Braslavski, Dmitry Ustalov, and Mikhail Mukhin. 2014. A Spinning Wheel for YARN: User Interface for a Crowdsourced Thesaurus. In *Proceedings of the Demonstrations at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 101–104, Gothenburg, Sweden. Association for Computational Linguistics.

Christiane Fellbaum. 1998. *WordNet: An Electronic Database*. MIT Press.

Ilya Gelfenbeyn, Artem Goncharuk, Vlad Lekhelt, et al. 2003. Automatic translation of WordNet semantic network to Russian language. In *Proceedings of Dialog-2003*.

Iryna Gurevych and Jungi Kim, editors. 2013. *The People's Web Meets NLP*. Springer Berlin Heidelberg.

Yuri Kiselev, Sergey V. Porshnev, and Mikhail Mukhin. 2015. Current Status of Russian Electronic Thesauri: Quality, Completeness and Availability. *Programmnaya Ingeneria*, (6):34–40.

Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, et al. 2013. The Future of Crowd Work. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, CSCW '13, pages 1301–1318, New York, NY, USA. ACM.

Klaus Krippendorff. 2013. *Content Analysis: An Introduction to Its Methodology*. SAGE, Thousand Oaks, CA, USA, 3rd edition.

Andrew A. Krizhanovsky and Alexander V. Smirnov. 2013. An approach to automated construction of a general-purpose lexical ontology based on wiktionary. *Journal of Computer and Systems Sciences International*, 52(2):215–225.

Huairen Lin and Joseph Davis. 2010. Computational and Crowdsourcing Methods for Extracting Ontological Structure from Folksonomy. In *The Semantic Web: Research and Applications*, volume 6089 of *Lecture Notes in Computer Science*, pages 472–477. Springer Berlin Heidelberg.

Natalia Loukachevitch. 2011. *Thesauri in information retrieval tasks*. Moscow University Press, Moscow, Russia.

Marek Maziarz, Maciej Piasecki, Ewa Rudnicka, and Stan Szpakowicz. 2014. plWordNet as the Cornerstone of a Toolkit of Lexico-semantic Resources. In *Proceedings of the Seventh Global Wordnet Conference*, pages 304–312, Tartu, Estonia.

Roberto Navigli and Simone Paolo Ponzetto. 2012. BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*, 193:217–250.

Karel Pala and Pavel Smrž. 2004. Building Czech Wordnet. *Romanian Journal of Information Science and Technology*, 7(1–2):79–88.

Ellie Pavlick, Matt Post, Ann Irvine, Dmitry Kachaev, and Chris Callison-Burch. 2014. The Language Demographics of Amazon Mechanical Turk. *Transactions of the Association for Computational Linguistics*, 2:79–92.

Anna Rumshisky. 2011. Crowdsourcing Word Sense Definition. In *Proceedings of the 5th Linguistic Annotation Workshop*, LAW V '11, pages 74–81, Stroudsburg, PA, USA. Association for Computational Linguistics.

Dmitry Ustalov. 2014. Words Worth Attention: Predicting Words of the Week on the Russian Wiktionary. In *Knowledge Engineering and the Semantic Web*, volume 468 of *Communications in Computer and Information Science*, pages 196–207. Springer International Publishing.

Dmitry Ustalov. 2015. A Crowdsourcing Engine for Mechanized Labor. *Proceedings of the Institute for System Programming*, 27(3):351–364.

Aobo Wang, Cong Duy Vu Hoang, and Min-Yen Kan. 2013. Perspectives on crowdsourcing annotations for natural language processing. *Language Resources and Evaluation*, 47(1):9–31.

Andrey Zaliznyak. 1977. *Grammatical dictionary of Russian*. Russky yazyk, Moscow, USSR.