Expanding Russian PropBank: Challenges and Insights for Developing New SRL Resources

Skatje Myers[*]**, Roman Khamov**[*]**, Adam Pollins**[*]**, Rebekah Tozier**[*]**,**
Olga Babko-Malaya[†]**, Martha Palmer**[*]

[*]University of Colorado
{firstname.lastname}@colorado.edu

[†]BAE Systems
olga.babko-malaya@baesystems.com

Abstract

Semantic role labeling (SRL) resources, such as Proposition Bank (PropBank), provide useful input to downstream applications. In this paper we present some challenges and insights we learned while expanding the previously developed Russian PropBank. This new effort involved annotation and adjudication of *all* predicates within a subset of the prior work in order to provide a test corpus for future applications. We discuss a number of new issues that arose while developing our PropBank for Russian as well as our solutions. Framing issues include: distinguishing between morphological processes that warrant new frames, differentiating between modal verbs and predicate verbs, and maintaining accurate representations of a given language's semantics. Annotation issues include disagreements derived from variability in Universal Dependency parses and semantic ambiguity within the text. Finally, we demonstrate how Russian sentence structures reveal inherent limitations to PropBank's ability to capture semantic data. These discussions should prove useful to anyone developing a PropBank or similar SRL resources for a new language.

Keywords: Semantic role labeling, Semantically annotated resources, Russian semantics

1. Introduction

The ability to identify the semantic elements of a sentence (*who* did *what* to *whom*, *where* and *when*) is crucial for machine understanding of natural language and downstream tasks such as information extraction (MacAvaney et al., 2017), question-answering systems (Yih et al., 2016), text summarization (Mohamed and Oussalah, 2019), and machine translation (Rapp, 2022). The process of automatically identifying and classifying the predicates in a sentence and the arguments that relate to them is called semantic role labeling (SRL).

Using the PropBank schema (Palmer et al., 2005) (Pradhan et al., 2022), a Russian-language lexicon and corpus was manually annotated, called Russian PropBank (which we will refer to as RuPB1) (Moeller et al., 2020). In this paper, we present our work expanding RuPB1 (we refer to the expanded version as RuPB2), the challenges encountered, and our proposed solutions. We present this discussion to benefit future work for new PropBanks and semantic representations in other languages, many of which may encounter similar challenges during annotation and in representing the semantics of target languages.

In particular, we have been creating new frames [1] and expanding double-annotated and adjudi-cated coverage of the verbs, as well as expanding the scope of annotation to include participles and both relativizers and their head words. Our efforts have resulted in a smaller but more thorough dataset. This paper first provides a general overview of our project's source material and goals as well as related projects that facilitated the process in Section 2. Next, we distinguish the respective scopes of RuPB1 and RuPB2 in Section 3. Section 4 covers changes made to RuPB1's frames and the issues faced when adding frames to RuPB2. We provide an overview of our infrastructure and annotation process in Section 5. In Section 6, we discuss sources of disagreement between annotators and the guidelines we devised to resolve them. Finally, we review how Russian's dropped copulas provide a challenge for accurate, detailed semantic representation in Section 7.

2. Background

Proposition Bank (PropBank) takes a verb-oriented but very generalizable approach to representing semantics. The list of permissible semantic roles is defined by the sense of each verb using numbered labels, ARG0 through ARG6. Typically an ARG0 is similar to a Proto-agent (per Dowty (1991)), and is the Agent or Experiencer, while ARG1 is usually the Patient or Theme of the predicate, similarly to a Proto-patient. By

---

[1] https://github.com/cu-clear/RussianPropbank/

generalising the arguments in this way, automatic semantic role labelers can produce useful information even if they misidentify the frame. Additionally, there are adjunct-like arguments, called argument modifiers (ARGM), to incorporate other semantically relevant information such as location (ARGM-LOC) and direction (ARGM-DIR).

The standard approach to developing a PropBank for a new language is to begin by defining a valency lexicon, known as a set of PropBank Frame Files, that defines the predicate-argument structure for all predicates to be annotated (Xue and Palmer, 2009; Zaghouani et al., 2010; Palmer et al., 2006; Bhatt et al., 2009). Once a sufficient number of frames has been defined, the annotation process begins, with the annotators referring to the frames for guidance for each individual predicate. In order to maintain complete annotation coverage for each sentence, additional frames are typically added during the annotation process. Double-blind annotation is recommended, followed by an expert adjudication pass. It is expected that the annotation process will reveal various ways in the which the original frame definitions need to be revised, sometimes resulting in follow-on revisions to previous annotations.

The Low Resource Languages for Emergent Incidents (LORELEI) project [2] sought to explore techniques for rapidly developing natural language processing technologies for low-resource languages. The dataset released as part of this project consists of parallel corpora for 23 low resource languages across many genres, such as newswire, phrasebooks, and weblogs. A subset of the English data was manually annotated with PropBank SRL.

The RuPB1 corpus (Moeller et al., 2020) project constructed 364 frames and annotated PropBank-style semantic roles on a portion of the Russian *newswire* and *phrasebook* sentences that paralleled the English dataset. This consists of 91 *newswire* sentences (2,228 tokens) and 496 *phrasebook* sentences (2,471 tokens). The previous work focused on annotating high-frequency verbs, which resulted in most sentences in the corpus having partial annotation. Our work has focused on filling in missing predicate annotations to produce fully labeled sentences in order to facilitate use of this corpus for training and testing SRL models and for evaluating how well annotation projection methods, such as those used by the Universal PropBanks project (UPB) (Jindal et al., 2022), map to SRL designed for the target language. The latter requires fully-annotated sentences to determine which predicates have been missed, added,

	RuPB1	RuPB2
# frames	364	497
# sentences	587	257
# predicates	431	331

Table 1: Comparison of annotation coverage between the partial annotation of RuPB1 and the smaller but completely annotated sentences of RuPB2.

or misplaced by the projection. See Table 1 for more details.

Russian PropBank is not the only resource for Russian SRL. Russian FrameBank (Lyashevskaya and Kashkin, 2015) is a project to develop FrameNet-style (Baker et al., 1998) frames designed for Russian and annotate examples of those frames from the Russian National Corpus [3]. Their annotation scheme uses 96 distinct semantic roles, such as Result or Beneficiary, organised in an hierarchical graph. Frames for approximately 4,000 target verbs, adjectives, and nouns were constructed, and over 50,000 examples of these frames were annotated. There is a fundamental difference in the approach of both resources: Russian FrameBank is rooted more in lexical semantics, while PropBanks are more focused on the syntax-semantics interface (Levin, 1993). As a result, RuPB offers a coarser-grained, more general SRL schema. Instead of having 96 specific semantic roles, PropBank uses the numbered arguments described above. For instance, an ARG0 can be either an Agent or an Experiencer depending on the predicate. Additionally, while FrameNet accounts for peripheral arguments and modifiers, Russian FrameBank does not; its annotations focus only on the core arguments of a given example predicate. In contrast, for each predicate, RuPB labels both the core arguments and modifiers (PropBank's equivalent of peripheral arguments). Additionally, RuPB2's goal is to annotate every predicate in a given sentence, instead of only annotating a specific, example predicate. Unfortunately, this means there is no automatic way for RuPB to take advantage of the 50,000 annotated example sentences in Russian FrameBank without extensive manual review, since the latter's annotations only provide partial coverage of the predicates in a sentence. For the same reason, Russian FrameBank does not provide an appropriate evaluation corpus for UPB.

3. Scope

As discussed above, while RuPB1 prioritised developing frames in the order of verb frequency, our

---

aim was to ensure that sentences have *complete* annotations, so that this resource can also be used as a test dataset. As a result, RuPB2 produced 200 sentences of *phrasebook* and 57 sentences of *newswire* with all predicates annotated.

Besides the additional verb annotation, we also extended the scope of RuPB from only verbs to include participles for 36 verbs and 9 relative-head pairs, such as обнаруженных 'discovered'. These are annotated with the same frames that they would be as verbs (обнаружить 'to discover').

RuPB2 also expanded annotations to include R-ARGs, in alignment with EnPB guidelines. Previously, the relativizer was the only argument annotated (such as который 'who' in this example):

(1) мальчик который любит кошек
    mal'čik kotoryj ljubit košek
    boy   who   loves   cats
    -        ARG0   pred   ARG1

RuPB2 now captures both the relativizer and the head noun:

(2) мальчик который любит кошек
    mal'čik kotoryj ljubit košek
    boy   who   loves   cats
    ARG0   R-ARG0   pred   ARG1

Our scope is still narrower than than of the current EnPB, which extensively annotated nominalizations and predicative adjectives. Some of these additional parts of speech may be added to RuPB in the future, such as nominalizations and eventive nouns, depending on applications. Some of the eventive nouns were added to EnPB during projects that focused on disasters, such as tornado.01, which captures arguments for things such as death toll and Fujita scale.

Another type of predication that EnPB includes, but RuPB does not, are adjectival predicates, such as blue.01: "He was blue from the cold."

4. Framing

As discussed above, the development of a high-quality, comprehensive valency lexicon is the cornerstone of the PropBanking process. Thanks to RuPB1, we began with a pre-existing set of Russian Frame Files. Our goal with RuPB2 was twofold: 1) to add enough frames to get full sentence coverage; 2) and to expand the scope of the predicates being annotated.

In addition to the expansion, 134 new frames were added, and many previous frames were re-examined. During the initial stages of RuPB2, we ran into the issue of using different terms when discussing framing decisions, and settled on the following clarifications for the terms: alias, roleset, and predicates.

An alias is a grammatical or syntactic form of a verb. Both *drank* and *drunk* are aliases of the verb *drink*. A roleset is a particular sense of a verb as well as a list of its core arguments according to their semantic roles. Rolesets also include all aliases of the verb in question.

Roleset id: drink.01	
'ingest liquid'	
ARG0	drinker, agent
ARG1	liquid
ARG2	source of liquid

Table 2: Roleset for drink.01

Predicates are collections of rolesets. Many verbs are polysemous, and each sense or meaning of the verb (predicate) is captured by different rolesets. The predicate *drink* can have two rolesets, *drink*.01 and *drink*.02, as in 'I drank water from a well' vs 'I drink to your health'. See Tables 2 and 3.

Roleset id: drink.02	
'salute'	
ARG0	drinker, agent
ARG1	thing saluted

Table 3: Roleset for drink.02

Determining whether a given token warrants its own predicate or roleset, or is simply an alias of an existing roleset, can be challenging, especially in morphologically rich languages. For additional examples and details of our framing process, please refer to the RuPB2 Framing Guidelines on the website.

As discussed by Moeller et al. (2020), Russian verbs can undergo many morphological processes that sometimes change the verb's aspect but can change semantic meaning as well.

For example, the reflexive affix -ся can simply change a verb's grammar (new alias) but can also add a new sense (new roleset). The verb молить, 'to beg' (Table 4), becomes 'to pray', молиться (Table 5), when the reflexive affix is added. By comparison, хотеть and хотеться, 'to want', have no semantic difference.

Often these kinds of differences lead to discussions of semantic domains and analysing frequencies of arguments in the literature. The RuPB2

Roleset id: молить.01	
molit' 'to beg'	
ARG0	asker, agent
ARG1	person being begged
ARG2	thing asked for

Table 4: Roleset for молить.01.

Roleset id: молиться.01	
molit'sja 'to pray'	
ARG0	pray-er, agent
ARG1	prayer
ARG2	deity

Table 5: Roleset for молиться.01.

Figure 1: Anna can wait here or there

Figure 2: Anna can read books

guidelines err on the side of making different rolesets as opposed to different aliases when the framer is unsure. This can be referred to as *splitting* as opposed to *lumping*.[4] There are two main factors that led to this decision. Firstly, having a clear golden rule speeds up the process of creating new frames. Secondly, it means an ill-judged decision is easily reversible. Should a framer make two separate rolesets instead of separate aliases, it is always easier to go back through annotations and deterministically merge two different tags into a single tag. This is simple to implement and much easier than the reverse: deciding that two aliases should be separate rolesets and manually re-annotating every occurrence according to the new senses.

In contrast with EnPB, other PropBank projects have set a precedent of splitting over lumping, as seen in the Turkish PropBank (Ak et al., 2018), where very rich morphological processes result in lots of very similar rolesets. Verbs with negative or modal affixes are given their own frames despite having identical rolesets.

Our suggestion is that any potential PropBank should have explicit guidelines on splitting vs. lumping. On the one hand, a liberal approach to splitting may result in the amount of frames ballooning drastically. Yet a conservative approach may result in much time and effort spent on reversing previous decisions. Both linguistic and computational factors must be considered.

In the case of RuPB2, aside from some minor edits to existing frames (such as typos and confusing example sentences), there were a few decisions that resulted in different annotations compared to RuPB1. Some involved removing frames entirely.

The first case was the frame мочь.03 'can, may', which has no core arguments. This differs from мочь.01 'can, have ability', which has an ARG1 (*agent with ability*) and an ARG2 (*ability itself*).

мочь.03 can be seen in the following sentence in Figure 1:

In RuPB1, может would have been marked as мочь.03. This sense is contrasted with a sentence such as Figure 2:

In Figure 2, annotators should mark может as мочь.01, with ARG1 being 'Anna', and ARG2 be-

---
[4]Splitting and lumping have long been used by lexicographers to illustrate a bias in favor of either more coarse-grained senses or more fine-grained senses.
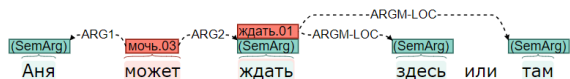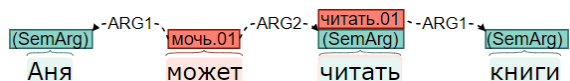
ing 'read'. In RuPB2, the может should be marked as ARGM-MOD for 'wait' not as its own predicate, since it is a modal indicating possibility (see Figure 3). In RuPB2, мочь.03 is removed entirely.
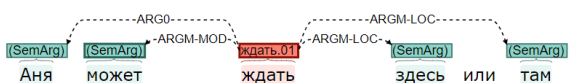

Figure 3: Anna can wait here or there

Likewise, we removed the roleset давай(те).09, which can be translated into English as 'let's', as in "let's look at a few examples". Instead of having a dedicated roleset, the verb will simply be marked ARGM-MOD, since it is essentially a hortative, modal verb.

Although we initially added быть.08, which was modeled on EnPB be.03 (the auxiliary verb 'will/was/were'), we eventually opted to remove this frame.

(3) Мы будем есть
     My budem est'
     We will eat.

EnPB set out to annotate semantic components including temporal relations as an ARGM (Kingsbury and Palmer, 2002). One could argue that быть.08 should be included in RuPB2 to adhere more closely to its English counterpart. Ultimately, быть.08 seemed to perform more of a functional, placeholder role; annotators would label this sense to avoid confusion with other быть senses. Because быть.08 lacks significant lexical information, we opted for its discontinuation.

5. Annotation Process

All RuPB2 annotation and adjudication was completed through the text-annotation platform INCEpTION (Klie et al., 2018). INCEpTION's interface streamlines corpus creation, annotation, and adjudication (the INCEpTION term for this phase is 'curation'). Our project required an environment
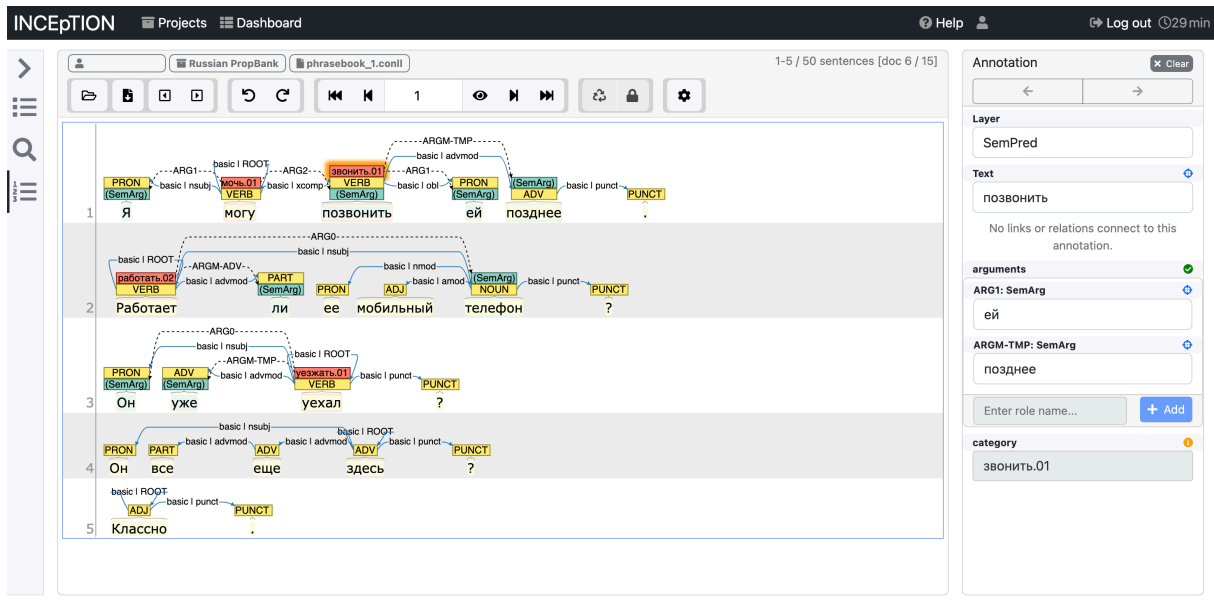
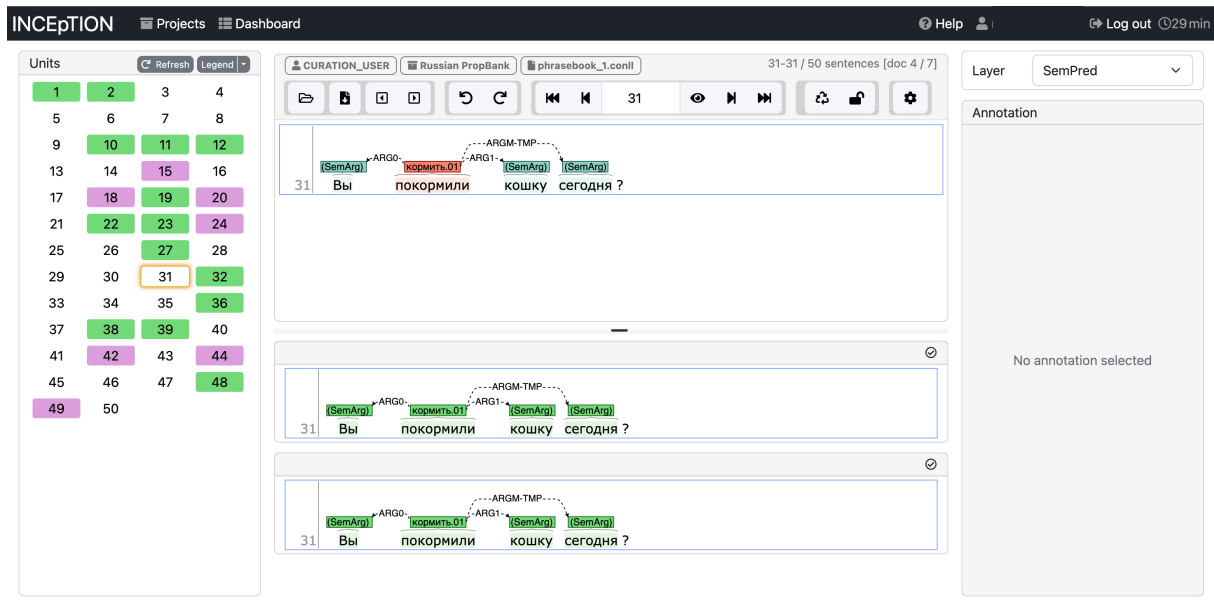Figure 4: INCEpTION Annotation mode.

Figure 5: INCEpTION Curation mode.

that would allow multiple users to annotate a semantic layer of predicates and arguments. This annotation process was additionally assisted by being able to simultaneously view dependency parse and part of speech layers. We automatically parsed the sentences using UDPipe (Straka and Straková, 2017) and provided these parses as our initial data in INCEpTION.

Figure 4 provides an example of the RuPB2 sentences using the annotation feature of the INCEpTION platform. Looking more closely at sentences 4 and 5 in Figure 4, these are examples of INCEpTION feature layers before RuPB2 annotation was

complete. No semantic roles could be annotated for sentences 4 and 5 due to the lack of predicates and arguments. The figure displays the layers that assisted the annotators: the sentences are written in Cyrillic text and further organized by each word's part of speech (yellow boxes). In addition, each sentence is syntactically parsed (e.g., subjects, objects, and sentence roots). In contrast with 4 and 5, sentences 1 through 3 have semantic predicates (red boxes) and arguments (various ARG arrows and green SemArg boxes) annotated. Observe that each predicate takes a specific verb frame, such as мочь.01, and core (numbered) ar-

guments are distinguished from ARGMs (e.g., adverbial, ARG-ADV or temporal, ARG-TMP) When clicking on a predicate, a pane on the right of the platform also shows the details of the semantic predicate layer.

Upon completion of a document, annotators submitted their finished work to an adjudicator for adjudication (the 'curation' pass). The adjudicator compared the finished annotations between users to assess inter-annotator agreement, with discrepancies highlighted. Additionally, the adjudication process resulted in the adjudicator creating a final, gold standard, fully annotated sentence. This process is illustrated in Figure 5; the top sentence reflects the adjudicator's gold standard annotated sentence, whereas the lower two sentences are the annotators'. For simplicity, only semantic predicate and argument layers are shown in Figure 5, but the layers included in Figure 4 are also available during the adjudication process.

Since the annotation guidelines and framing decisions evolved concurrently with the annotation process itself, all members of the RuPB2 group participated in the adjudication step. This thorough process allowed all aspects, framing, annotation challenges, and future work to be discussed. Annotation challenges and future work are presented more thoroughly in the following sections.

6. Annotation Challenges

RuPB annotators rely on an underlying Universal Dependencies (UD) syntactic parse to resolve ambiguity (de Marneffe et al., 2021). This parse itself sometimes introduces new ambiguity. Unlike EnPB annotators, who tag arguments as spans of words, RuPB annotators must identify and tag the word that corresponds to the argument's head. The automatic UD parser's choice of head is not always intuitive or consistent, and we observed it caused annotator disagreement most frequently in part-whole constructions and phrases that comprise more than one temporal modifier. Phrases containing locative modifiers were another source of disagreement. The counts of these phenomena that occurred in the RuPB2 sentences are totaled in Table 6, largely occurring in the more complex *newswire* sentences.

Pseudopartitives	7
Temporal Doublets	4
Locative Modifiers	4

Table 6: Cases of Challenging Annotation

6.1. Part-Whole Constructions

The head of a quantified nominal phrase is usually the inner nominal, which refers to the whole entity quantified (e.g., две тысячи **людей** / *two thousand **people***). By contrast, the head of a partitive construction is the outer nominal or part (***tons** of rice*). When the parser labels a quantifier as a noun instead of a numeral, the quantifier becomes the head of that phrase (**тысячи** людей / ***thousands** of people*) because the construction appears syntactically partitive. The part of speech of the quantifier thus changes the head of the phrase in the parse, though it does not affect the phrase's lexical meaning.

These constructions are *pseudopartitives*, and should not be analyzed as having the same syntax as partitives (Falco and Zamparelli, 2019). Their prevalence varies from one language to another, but they appear more frequently in Russian UD parses than in English. Compare English *a million **residents*** (numeral) and ***millions** of residents* (pseudopartitive) with Russian **миллион** жителей lit. 'million of residents' and **миллионы** жителей 'millions of residents' (both pseudopartitive). Annotators must take care to choose the head of each argument when working with a dependency parse that does not handle constructions such as pseudopartitives.

6.2. Temporal Doublets

Temporal modifiers occasionally appear in a series. However, the parser does not always treat the modifiers either as a single oblique nominal or as two separate ones, as seen in Figure 6. At first, annotators tagged according to the parse, but then noticed these inconsistencies and needed a different solution. In EnPB practice, arguments that comprise conjuncts are treated as a single argument and never tagged twice. With this in mind, we chose to treat these constructions as asyndetic coordination, in which the first element is the head per UD guidelines.[5]

6.3. Locative Modifiers

Annotators encountered difficulty as they decided whether to tag locative modifiers as arguments of the verb or to consider them as modifying a noun and thus not tagged. Straightforward cases of both the former type ('People died **in the village**') and the latter ('The head of the program **in Bangladesh** expressed his fears') appeared, in which annotators agreed.

Yet in ambiguous cases, annotators diverged. For example, in тысячи людей **в Индии и Бан-**
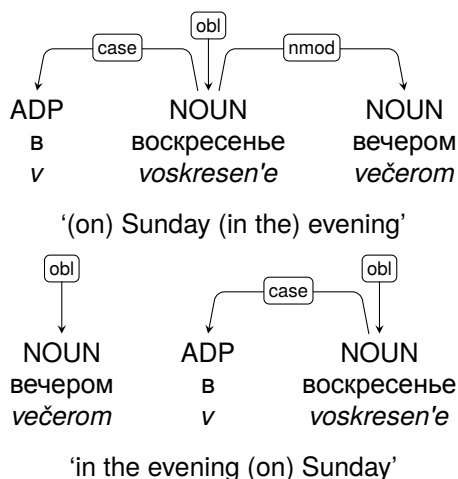

Figure 6: Two temporal modifier arguments, identical save word order, with arbitrarily different parses

гладеше до сих пор обращаются … 'Thousands of people **in India and Bangladesh** are still seeking …', they did not agree as to whether this locative was an argument of 'seeking' or modified 'people'. We resolved this by tagging the locative as a modifier of the verb in each ambiguous case, as we preferred to be thorough and prevent omissions.

7. Implicit Predicates

Russian usually drops the present tense linking verb, есть.[6] A lexical unit is omitted in instances similar to those in English where *be* would appear as *am*, *is* or *are*. For instance, one does not say "I am a student" in Russian but literally я студент-ка "I student." Ten percent of the sentences in our dataset were affected by dropped copulas (i.e., sentences had one less predicate or were completely unable to be annotated), predominately in the *phrasebook* portion.

This issue of dropping copulas is not limited to Russian; the World Atlas of Language Structures Online (Stassen, 2013) reports that 45% of the languages accounted for in their database (175 of 386 languages) allow zero copula constructions with nominal predicates.

8. Conclusion

We have presented numerous issues that were encountered during our endeavor to expand and complete PropBank annotation for RuPB2, a Russian SRL dataset for training and testing purposes. In

the future, RuPB2 can be further expanded to include nominalizations and light verb constructions to provide better coverage. We described our approach to constructing frames for Russian, which can provide a precedent for other morphologically rich languages and others with similar characteristics. More particularly, we analyzed the complexity in differentiating between predicates and modal verbs. We discussed our solutions to frequent cases of annotator disagreement, as well as the importance of the parse in settling ambiguities. In the final section, we discussed the challenges of implicit predicates that can be found in zero-copula sentences, which we expand on below. Throughout the development of RuPB2, there has been an aim to stay true to this schema and maintain parity with EnPB, all the while reflecting the semantics of Russian with the highest accuracy possible. These discussions should prove useful to anyone building a new PropBank for another language.

Although there are many benefits of the PropBank schema, it is important to also consider limitations when constructing a semantic corpus for a new language. PropBank can capture shallow semantic information about who did what to whom, but a deeper complete sentence representation that includes discourse relations and modality can be more effective. Uniform Meaning Representations (UMRs), (Van Gysel et al., 2021), provide a cross-lingual approach to such a representation. UMRs are based on the popular Abstract Meaning Representations project (Banarescu et al., 2013) which directly incorporates English PropBank for predicate argument structures. The ability of AMR-UMR to represent implicit predications yields a strategy for capturing semantics that is not covered by PropBank alone. Our Russian PropBank provides an essential foundational element for this type of richer, more nuanced semantics. The discussion and suggestions for how to develop guidelines and frame files for a Slavic language that are contained in this paper should provide a road-map for anyone else undertaking such an endeavor.

9. Acknowledgements

---

[6]The phenomenon where a subject and predicate are not overtly connected through a linking verb is known as a zero (or null) copula.

authors and do not necessarily reflect the views of NSF, DARPA or the U.S. government.

10.   Bibliographical References

References

K. Ak, C. Toprak, V. Esgel, and O. T. Yildiz. 2018. Construction of a turkish proposition bank. *Turkish Journal of Electrical Engineering and Computer Sciences*, 26(1):570–581.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *COLING 1998 Volume 1: The 17th International Conference on Computational Linguistics*.

L. Banarescu, Claire Bonial, Shu Cai, M. Georgescu, K. Griffitt, U. Hermjakob, K. Knight, P. Koehn, Martha Palmer, and N. Schneider. 2013. Abstract meaning representation for sembanking. *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186.

Rajesh Bhatt, Bhuvana Narasimhan, Martha Palmer, Owen Rambow, Dipti Sharma, and Fei Xia. 2009. A multi-representational and multi-layered treebank for Hindi/Urdu. In *Proceedings of the Third Linguistic Annotation Workshop (LAW III)*, pages 186–189, Suntec, Singapore. Association for Computational Linguistics.

Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal Dependencies. *Computational Linguistics*, 47(2):255–308.

David Dowty. 1991. Thematic proto-roles and argument selection. *language*, 67(3):547–619.

Michelangelo Falco and Roberto Zamparelli. 2019. Partitives and partitivity. *Glossa: a journal of general linguistics*, 4(1).

Ishan Jindal, Alexandre Rademaker, Michał Ulewicz, Ha Linh, Huyen Nguyen, Khoi-Nguyen Tran, Huaiyu Zhu, and Yunyao Li. 2022. Universal Proposition Bank 2.0. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1700–1711, Marseille, France. European Language Resources Association.

Paul Kingsbury and Martha Palmer. 2002. From TreeBank to PropBank. In *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC'02)*, Las Palmas, Canary Islands - Spain. European Language Resources Association (ELRA).

Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. The inception platform: Machine-assisted and knowledge-oriented interactive annotation. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9. Association for Computational Linguistics. Event Title: The 27th International Conference on Computational Linguistics (COLING 2018).

Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press.

Olga Lyashevskaya and Egor Kashkin. 2015. FrameBank: A database of russian lexical constructions. In *Communications in Computer and Information Science*, pages 350–360. Springer International Publishing.

Sean MacAvaney, Arman Cohan, and Nazli Goharian. 2017. GUIR at SemEval-2017 task 12: A framework for cross-domain clinical temporal information extraction. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1024–1029, Vancouver, Canada. Association for Computational Linguistics.

Sarah Moeller, Irina Wagner, Martha Palmer, Kathryn Conger, and Skatje Myers. 2020. The Russian PropBank. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5995–6002, Marseille, France. European Language Resources Association.

Muhidin Mohamed and Mourad Oussalah. 2019. Srl-esa-textsum: A text summarization approach based on semantic role labeling and explicit semantic analysis. *Information Processing & Management*, 56(4):1356–1372.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics*, 31(1):71–106.

Martha Palmer, Shijong Ryu, Jinyoung Choi, Sinwon Yoon, and Yeongmi Jeon. 2006. Korean PropBank. *LDC Catalog No.: LDC2006T03 ISBN*, pages 1–58563.

Sameer Pradhan, Julia Bonn, Skatje Myers, Kathryn Conger, Tim O'gorman, James Gung, Kristin Wright-bettner, and Martha Palmer. 2022. PropBank comes of Age—Larger, smarter, and more diverse. In *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics*, pages 278–288, Seattle, Washington. Association for Computational Linguistics.

Reinhard Rapp. 2022. Using semantic role labeling to improve neural machine translation. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3079–3083, Marseille, France. European Language Resources Association.

Leon Stassen. 2013. Zero copula for predicate nominals (v2020.3). In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Zenodo.

Milan Straka and Jana Straková. 2017. Tokenizing, pos tagging, lemmatizing and parsing ud 2.0 with udpipe. In *Proceedings of the CoNLL 2017 shared task: Multilingual parsing from raw text to universal dependencies*, pages 88–99.

Jens EL Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Tim O'Gorman, Andrew Cowell, William Croft, Chu-Ren Huang, et al. 2021. Designing a uniform meaning representation for natural language processing. *KI-Künstliche Intelligenz*, 35(3-4):343–360.

Nianwen Xue and Martha Palmer. 2009. Adding semantic roles to the Chinese TreeBank. *Natural Language Engineering*, 15(1):143–172.

Wen-tau Yih, Matthew Richardson, Chris Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–206, Berlin, Germany. Association for Computational Linguistics.

Wajdi Zaghouani, Mona Diab, Aous Mansouri, Sameer Pradhan, and Martha Palmer. 2010. The revised arabic propbank. In *Proceedings of the Fourth Linguistic Annotation Workshop*, pages 222–226. Association for Computational Linguistics.