

論文の内容の要旨

論文題目 : Multilingual computational lexicography: frame semantics meets distributional semantics
(多言語計算辞書学: フレーム意味論と分布意味論の接点)

氏名 : ロージャズ アンナ Anna Rogers

This thesis discusses several issues in multilingual frame-based computational lexicography from the combined perspectives of frame semantics and distributional semantics. The case study under consideration is three posture verbs (*sit*, *lie*, *stand*) in English, and their counterparts in Russian and Japanese. Frame-based representation of these verbs presents several challenges for the current FrameNets (Baker, Fillmore, & Lowe, 1998; Ohara et al., 2004), and I explore potential theoretical and practical solutions.

While many languages have lexemes corresponding to English verbs *sit*, *lie* and *stand*, they might differ with respect to the aspect and inchoativity/causativity. For example, consider (1):

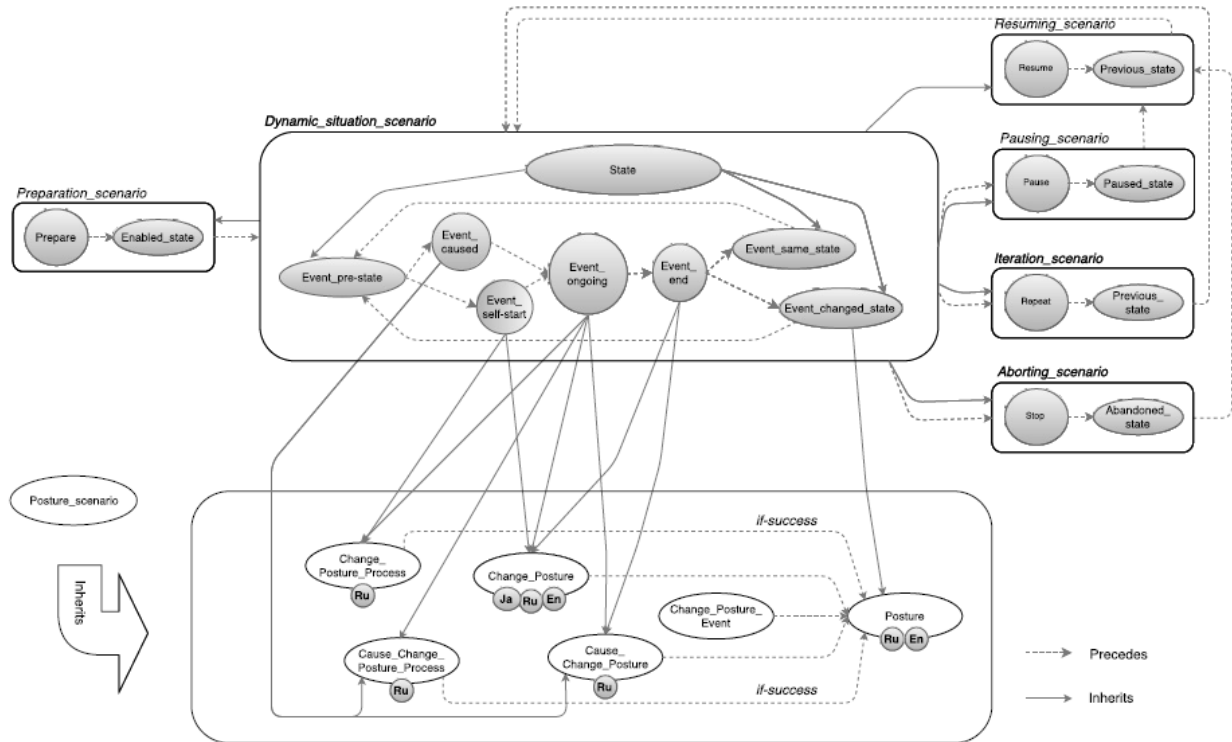
(1) I sat on the chair.

In English (1) is ambiguous; it could mean either “I sat myself down on the chair” or “I was sitting on a chair (and didn't move)” (Newman, 2002, p. 4). In other languages the posture verbs may have other aspectual properties, lexicalized or expressed with constructions.

In **chapter 3** I argue that such high-level features characterizing verb classes can and should be consistently represented in frame-based resources, and that doing so requires a more typological perspective. I develop a model (Fig.1) that grounds these features in the event structure representation. This model can be implemented in the current Berkeley FrameNet architecture, it is better equipped to deal with multilingual data, and would encourage consistency with respect to aspect and inchoativity/causativity. It is also more inference-friendly and thus would be more useful from the point of view of applications of FrameNets in Natural Language Processing. Furthermore, I outline a mechanism for establishing translational equivalence between a frame evoked by a lexical unit in one language, and a complex expression – in another.

Since implementing this model would require too large a number of frame-to-frame relations, I explore the possibilities of identifying verb classes with supervised machine learning over vector space models. Such a mechanism would ease identifying words that are candidates for inheritance from an abstract frame, and it would also be helpful for conducting consistency checks in a FrameNet database. Furthermore, it lends strong support to the distributional hypothesis. While most previous work on identifying verb classes from distributional data relied on syntactically parsed corpora, **chapter 4** of this thesis shows that relatively abstract high-level verb features can also be derived from simple bag-of-words models.

Fig. 1. Compositional account of POSTURE_SCENARIO based on event structure.



The proposed approach successfully recovers Russian imperfective/perfective verbs, including some that are not registered in current dictionaries (Drozd, Gladkova, & Matsuoka, 2015a, 2015b). I also demonstrate the possibility of identifying pairs of Japanese intransitive/transitive verbs with word analogies. Work on English and Russian data (Gladkova, Drozd, & Matsuoka, 2016) showed that success of solving word analogies varies considerably for different linguistic relations. The method relying on supervised learning (Drozd, Gladkova, & Matsuoka, 2016) is overall more successful, and in this thesis I show that it can be easily augmented with pattern-based rules.

In **chapter 5** I consider the problem of representing selectional preferences in a frame-based lexical resource. The challenge is that languages differ with respect to whether verbs that indicate human posture can also be used for indicating location of objects (Ameka & Levinson, 2007; Viberg, 2013). When they do, they retain their original spatial characteristics to various extent, and they also vary with respect to which objects they are compatible with. For example, in English the verb *sit* is often used to indicate location of an object, irrespective of its spatial orientation (2), but *lie* or *stand* are more tied to the spatial context:

(2) The cup sat on the desk.

In Russian, on the other hand, it is possible to combine a posture verb with “cup” – but it would “stand” rather than “sit”. As for Japanese, it normally does not allow for any such combinations. In both cases, the restrictions are motivated by certain properties of the real-world entities that are denoted by the arguments of posture verbs, and, as such, should be accounted for by frame semantics.

I argue against the current Berkeley FrameNet model in which semantic types of frame elements are represented as independent ontological categories such as “sentients” or “physical objects”. This view is supported by the data of a survey of 72 native speakers of English, Russian and Japanese (24 speakers per language). The task for the participants was to indicate whether either or any combination of posture verbs were acceptable to describe various combinations of head nouns and locations (e.g. “cat sitting/lying/standing on a table”), or some other expression would be preferable. In total there were 372 combinations of posture verbs, their nouns and locations.

The survey was designed to encourage rather than suppress variation, and no claims are made about its results being representative of the language norm. Still, Russian and Japanese exhibited fairly strong constraints on selectional preferences of posture verbs that few participants were willing to violate. English exhibited greater variability than either Russian or Japanese. Only two verb-noun combinations were universally accepted by speakers of all three languages, and only 39 (out of 372) were accepted by 60% majority. However, the distribution of the answer patterns is overall similar for the three languages, with

“man” and “doll” having the highest acceptance, and saliently vertical objects such as boxes and houses tending to be accepted more.

I attribute the observed variation between speakers of the same language to dialectal, idiolectal differences, as well as differences in individual experience and imagination. I argue that the best way to account for variation in posture verb use both within and between languages is with a model based on extensions from the anthropocentric prototype. I outline 8 semantic constraints on extensions of this prototype (Table 1).

In **chapter 6** I consider an alternative approach to modeling selectional preferences, namely supervised classification over word vectors (Bergsma, Lin, & Goebel, 2008). For each English posture verb I train a logistic regression classifier that would predict to what extent a given word is acceptable as its subject.

In a series of experiments I show that a more linguistically informed training dataset (clean-up of syntactic parser output, balanced subsampling of positive examples) can yield up to 50% increase in accuracy of such classification over the commonly used approach with random and unbalanced samples. I also explore the possibility of more semantically relevant negative examples. Usually negative examples are sampled from corpus vocabulary at random (e.g. using *blue* as a negative example of things that are not able to *stand*), so more relevant negative examples should provide the classifier with better semantic knowledge. However, in my experiments inclusion of relevant negative examples (e.g. *snake* as an example of things that are not able to *stand*) did not yield consistent advantage. This result points at a larger problem of identifying narrow semantic classes for selectional preferences (“standability”) that in distributional space are hard to distinguish from bigger classes such as “names” or “animals”.

In **chapter 7** I outline several theoretical challenges to frame semantics and distributional semantics, and discuss the possibilities of closer interaction between linguistics and natural language processing. I also consider the methodological challenges of constructing a multilingual frame-based resource. The primary tenet of frame semantics is that meanings are relativized to scenes, and scenes come from experience. This could be taken as an argument for universality of frames, since humans have a lot of shared experience.

However, in fact the experiential basis of linguistic frames renders them susceptible to considerable variability from society to society, and even from person to person. Even something as basic as sitting is different for Japanese culture, where sitting on the floor is typical, and American culture, where it is less so. The actual sitting posture also differs in these two situations, ranging from extremely stiff and formal *seiza* pose to comfortable sinking in a huge sofa. Does that mean that English and Japanese have different *sitting* frames? If so, how do we establish the correspondence? I argue that the solution lies in shifting from the corpus-based averages and the idea of frame identity to the idea of functional equivalence of frames, and their overlaps between the lexicons of different speakers.

I also argue that a multilingual frame-based resource is more feasible if we do not attempt to project frames from one language onto others (as is done in many current FrameNets that are based on the English frame database). A more typological and systematic perspective is required from the outset, especially for defining top-level frames and scenarios. Such a resource would require analyzing more data than can be done manually, and thus closer collaboration with natural language processing is essential.

I conclude with discussing the practical contributions, theoretical implications, and limitations of this work. I also outline a list of questions it raised with regards to theoretical apparatus of frame semantics, the nature of distributed meaning representations, representations of event structure and selectional preferences, compositionality, and multilingual frame-based resources.

Table 1. Constraints on selectional preferences of posture verbs in English, Russian and Japanese

Schematic principles	English	Russian	Japanese
(1) Anthropomorphism	<i>sit, lie, stand</i>	<i>sit, lie, stand</i>	<i>sit, lie, stand</i>
(2) Physical ability to assume postures visually and/or functionally similar to human postures	<i>sit, stand</i>	<i>lie, sit, stand</i>	<i>sit, lie, stand</i>
(3) Salient vertical orientation	<i>stand</i>	<i>stand</i>	<i>Stand</i>
(4) Salient horizontal orientation	<i>lie</i>	<i>lie</i>	---
(5) A base or something resembling legs	<i>stand</i>	<i>stand</i>	---
(6) Enclosure in a container	---	<i>lie</i>	---
(7) Birds in a perched position	<i>stand</i>	<i>sit</i>	---
(8) “Fallback” option for unclear cases	<i>sit</i>	<i>lie</i>	

References

- Ameka, F. K., & Levinson, S. C. (2007). Introduction: The typology and semantics of locative predicates: posturals, positionals, and other beasts. *Linguistics*, 45(5), 847–871. <https://doi.org/10.1515/LING.2007.025>
- Baker, C. F., Fillmore, C. J., & Lowe, J. B. (1998). The Berkeley Framenet project. In *Proceedings of the 17th international conference on Computational Linguistics* (Vol. 1, pp. 86–90). Association for Computational Linguistics.
- Bergsma, S., Lin, D., & Goebel, R. (2008). Discriminative learning of selectional preference from unlabeled text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 59–68). Association for Computational Linguistics. Retrieved from <http://dl.acm.org/citation.cfm?id=1613725>
- Droz, A., Gladkova, A., & Matsuoka, S. (2015a). Discovering aspectual classes of Russian verbs in untagged large corpora. In *Proceedings of 2015 IEEE International Conference on Data Science and Data Intensive Systems (DSDIS)* (pp. 61–68). <https://doi.org/10.1109/DSDIS.2015.30>
- Droz, A., Gladkova, A., & Matsuoka, S. (2015b). Python, performance, and Natural Language Processing. In *Proceedings of the 5th Workshop on Python for High-Performance and Scientific Computing* (p. 1:1–1:10). New York, NY, USA: ACM. <https://doi.org/10.1145/2835857.2835858>
- Droz, A., Gladkova, A., & Matsuoka, S. (2016). Word embeddings, analogies, and machine learning: beyond king - man + woman = queen. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (pp. 3519–3530). Osaka, Japan, December 11-17: ACL. Retrieved from <https://www.aclweb.org/anthology/C/C16/C16-1332.pdf>
- Gladkova, A., Droz, A., & Matsuoka, S. (2016). Analogy-based detection of morphological and semantic relations with word embeddings: what works and what doesn't. In *Proceedings of the NAACL-HLT SRW* (pp. 47–54). San Diego, California, June 12-17, 2016: ACL. Retrieved from <https://www.aclweb.org/anthology/N/N16/N16-2002.pdf>
- Newman, J. (Ed.). (2002). *The linguistics of sitting, standing and lying* (Vol. 51). Philadelphia, PA: J. Benjamins Pub.
- Ohara, K. H., Fujii, S., Ohori, T., Suzuki, R., Saito, H., & Ishizaki, S. (2004). The Japanese FrameNet project: An introduction. In *Proceedings of the Workshop on Building Lexical Resources from Semantically Annotated Corpora at LREC 2004*.
- Viberg, Å. (2013). Posture verbs: A multilingual contrastive study. *Languages in Contrast*, 13(2), 139–169. <https://doi.org/10.1075/lic.13.2.02vib>