Meaning Representation of English Prepositional Phrase Roles: SNACS Supersenses vs. Tectogrammatical Functors

Wesley Scivetti **Nathan Schneider**
Georgetown University
nathan.schneider@georgetown.edu

Abstract

This work compares two ways of annotating semantic relations expressed in prepositional phrases: semantic classes in the Semantic Network of Adposition and Case Supersenses (SNACS), and tectogrammatical functors from the Prague English Dependency Treebank (PEDT). We compare the label definitions in the respective annotation guidelines to determine expected mappings, then check how well these work empirically using Wall Street Journal text. In the definitions we find substantial overlap in the distributions of the two schemata with respect to participants and circumstantials, but substantial divergence for configurational relationships between nominals. This is borne out by the empirical analysis. Examining the data more closely for participants and circumstantials reveals that there are some unexpected, yet systematic divergences between definitionally aligned groups.

1 Introduction

Broad coverage descriptive frameworks for annotating lexical semantics have proven useful for researchers in the field of computational semantics. Most of these frameworks have a primary focus on verbs and their participants (Baker et al., 1998; Bonial et al., 2014; Kipper et al., 2008; Palmer et al., 2017), though some frameworks extend annotation schema to cover the arguments of nominal phrases (Hajič et al., 2012; Meyers et al., 2004). Relatively few frameworks have focused on comprehensive accounts of prepositions, which can modify both verbal and nominal heads (Schneider et al., 2018; Litkowski and Hargraves, 2005), and can contribute crucial semantic information to sentences despite often being thought of as purely functional elements.

The most recent and comprehensive attempt to cover the semantics of prepositions is the Semantic Network of Adposition and Case Supersenses, or SNACS (Schneider et al., 2015, 2016, 2018), which is a hierarchy of semantic classifications of prepositional modifiers. SNACS contains 52 total preposition semantic classes, or SUPERSENSES, which are arranged into a hierarchy with different levels of granularity at each point in the hierarchy. In English, the SNACS framework has been applied to the reviews section of the English Web Treebank (EWT) corpus (Bies et al., 2012), resulting in the STREUSLE corpus with gold SNACS annotations (Schneider et al., 2018).

For researchers interested in the lexical semantics of prepositions, the STREUSLE corpus is a valuable resource, but is smaller in size compared to corpora that have been annotated for other lexical semantic projects. While some of these other resources do mark some semantic information conveyed by prepositional phrases, it is an open question to what extent these more general semantic frameworks overlap with the preposition-centric hierarchy of SNACS. If there is significant overlap between corresponding classes across different annotation schema, it may be possible to convert the classifications of prepositional phrases in these more general schemata into corresponding SNACS supersenses. This would make it possible to quickly augment the available data annotated within the SNACS hierarchy, and would provide useful comparisons between the coverage of different annotation schemata.

In particular, this research highlights the Prague English Dependency Treebank (PEDT, Hajič et al. 2012) as one resource with potential overlap with SNACS.[1,2] The PEDT contains multiple layers of

[1] PEDT is the English side of the Prague Czech-English Dependency Treebank. One reason to examine this framework and corpus is that if the correspondence proves reliable for English, it might be leveraged to obtain heuristic SNACS annotations of Czech data as well, since the tectogrammatical annotation scheme is also applied in the Czech translation of the Wall Street Journal corpus.

[2] In a comparison of an earlier version of SNACS to PropBank semantic roles, Schneider et al. (2016) found good correspondences between supersenses and PropBank modifiers,

Proceedings of the 4th International Workshop on Designing Meaning Representations, pages 68–73
June, 2023.

Functor	Supersense	Functor	Supersense	Functor	Supersense	Functor	Supersense	Functor	Supersense
TSIN	StartTime	LOC	Locus	MEANS	Instrument,Means	ACT	Agent,Force	EXT	Cost
TTILL	EndTime	DIR1	Source	MANN	Manner	PAT	Theme,Topic	APP	Gestalt
TFHL	Duration	DIR2	Direction,Path	CAUS	Explanation	ORIG	Originator	COMPL	Identity
THL	Duration	DIR3	Goal	AIM	Purpose	ADDR	Recipient	MAT	QuantityItem
THO	Frequency	EXT	Extent			BEN	Beneficiary	RSTR	Characteristic
TPAR	Time					ACMP	Ancillary	CPR	ComparisonRef
TWHEN	Time								

Table 1: Heuristic mapping from PEDT functor to SNACS supersense based on the guidelines. Functors and supersenses without a clear correspondence are omitted. EXT is listed twice because it maps to both Spatial and Configurational supersenses.

syntactic/semantic annotation for the entire WSJ section of the Penn Treebank (Marcus et al., 1993). We focus on the tectogrammatical layer, or t-layer, which describes the deep syntax/semantics of the sentence, and labels nominals with a set of FUNCTORS. Many of these functors seem to show remarkable overlap with SNACS supersenses, though there are some significant divergences. This work investigates the overlap between the SNACS hierarchy and functor labels for prepositional phrases from PEDT, by first qualitatively outlining the similarities between the definitions of semantic classes in the two frameworks, then offering an empirical analysis of their overlapping distributions on a set of WSJ sentences.

2 Definitional Comparison

The SNACS hierarchy v2.6 (Schneider et al., 2022) contains 52 total supersenses organized into 3 main branches: the **CIRCUMSTANCE** branch, the **PARTICIPANT** branch, and the **CONFIGURATION** branch. Recent versions of the SNACS hierarchy assign supersenses to a preposition for both its *scene* role and its *function* role. The *scene* role represents the contextual semantic role of a preposition in combination with the predicate, while the *function* role is more faithful to the lexical semantics of the preposition (Schneider et al., 2018; Hwang et al., 2017). In many instances, the two roles are the same, but in cases where the *scene* and *function* roles differ, the two are represented using the SCENE⇝FUNCTION notation. We hypothesize that the *scene* role SNACS supersenses will more closely align with PEDT functors, and thus focus on *scene* role supersenses unless otherwise specified. Many SNACS supersenses correspond more or less directly to PEDT functors based upon the definitions set forth in their respective guidelines. Table 1 lists PEDT functors with clear corresponding SNACS supersenses. We exclude supersenses without clear corresponding functors, as well as functors which are not directly relevant to the SNACS hierarchy.

We see in Table 1 that most CIRCUMSTANCE supersenses, which add spatial, temporal, or other description to events, usually have corresponding PEDT functors. In Example (1) we see an example of the overlap between the THL ("how long?") functor, and the DURATION supersense. The directional functors DIR1 and DIR3 best correspond to SOURCE and GOAL respectively, and not (despite the terminology) DIRECTION. This is because the start point of movement (which answers the question "where from?") is labeled SOURCE, and the end point of movement (which answers the question "where to?") is labeled as GOAL. Examples of DIR1 and DIR3 are shown in Examples (2) and (3).

(1) Big mainframe computers for business had been around **for**_THL_DURATION years.

(2) All came **from**_DIR1_SOURCE Cray Research.

(3) Despite recent declines in yields, investors continue to pour cash **into**_DIR3_GOAL money funds.

On the other hand, SNACS DIRECTION is used to express the orientation of motion where the end result is not specified. We can observe the distinction in Examples (4) and (5), which are taken from the most recent version of the SNACS annotation guidelines (Schneider et al., 2022). If DIR1 and DIR3 do not generally correspond to DIRECTION, then DIRECTION is exceptional in that it does not have a directly corresponding PEDT functor. DIRECTION, which is a subtype of PATH, is probably most closely related to the more general DIR2.

(4) I headed **to**_GOAL work.

but less uniform correspondences for numbered arguments.

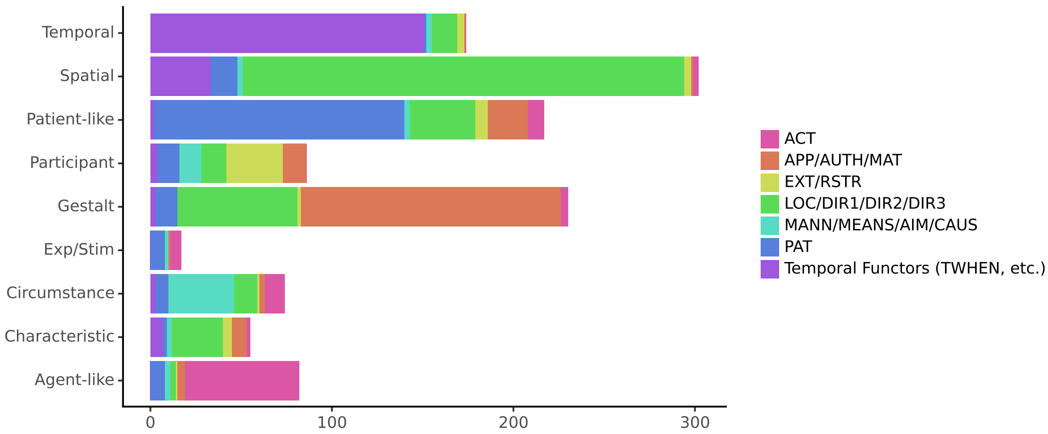

Figure 1: The Overlap of General Supersense Groupings with the PEDT Functors. Supersenses and Functors are combined into groups to show broad overlaps. Here "Circumstance" covers non-spatiotemporal circumstances (Manner, Means, Explanation, and Purpose). "Participant" covers Originator, Recipient, Beneficiary, Instrument, and Cost (excluding the agent-like, patient-like, and experiencer/stimulus participants).

(5) I headed **towards**_DIRECTION work, but never made it there.

PARTICIPANT supersenses, which introduce more canonical participants to events, also often correspond well to PEDT functors. The INSTRUMENT supersense lacks a directly corresponding functor, but is grouped with the MEANS supersense under the scope of the MEANS functor. The ACMP functor at least sometimes corresponds to ANCILLARY, as shown in Example (6). The ACT and PAT functors are potentially problematic, since they mark primarily syntactic roles of arguments, not semantic roles. This means that finer-grained supersenses, such as EXPERIENCER and STIMULUS, are not captured by PEDT functors. Furthermore, COST is perhaps the most problematic of the PARTICIPANT supersenses, with EXT being a marginal match at best.

(6) The U.S., **with**_ACMP_ANCILLARY its regional friends, must play a crucial role in designing its architecture.

CONFIGURATION supersenses, which describe state or property relationships between two nominals, are the least similar to PEDT functors, though there are some clear correspondences shown in Table 1, including the relationship between MAT and QUANTITYITEM as shown in Example (7). SNACS also includes more specific GESTALT subtypes, such as ORG and POSSESSOR, which are finer-grained than what is captured by the APP functor. Some more general configurations such as SPECIES, ENSEMBLE, and SOCIALREL lack corresponding PEDT functors.

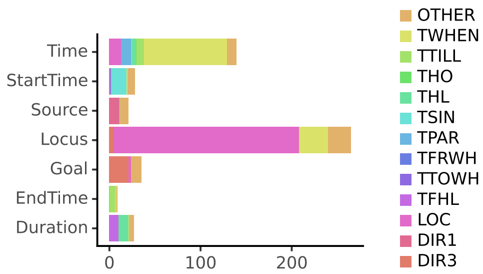

Figure 2: Spatiotemporal Supersense Overlap with PEDT functors

There are also some PEDT functors which are beyond the scope of the SNACS hierarchy: for instance, functors which mark paratactic relations (e.g. CONTRA), express primarily discourse functions (e.g. ATT), or mark types of syntactic information which is not conveyed in prepositional phrases (e.g. APPS). These functors are omitted from further analysis.

(7) About 20,000 sets **of**_MAT_QUANTITYITEM Learning Materials teachers' binders have also been sold in the past four years.

3 Empirical Comparison

3.1 Methodology

Now that we have outlined the overlap in descriptions between the SNACS hierarchy supersenses and various PEDT functors, we wish to quantify how these categories overlap in practice. In or-

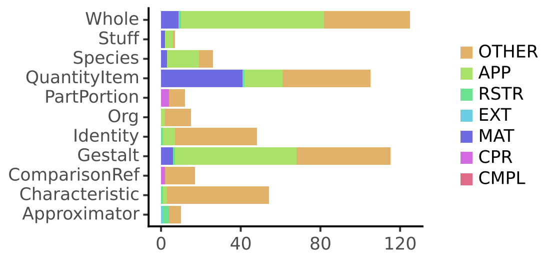

Figure 3: Configuration Supersense Overlap with PEDT functors

Class of Functors	# of Tokens	Percent Overlap
circumstantials	818	50.3
spatials	446	52.7
temporals	212	66.5
other	160	21.9
participants	500	47.0
ACT	113	55.8
PAT	238	58.0
other	149	22.8
configurations	386	36.0

Table 2: For groups of functors, percentage of tokens for which tagger-predicted supersense agrees with the heuristic mapping in Table 1. EXT is only considered in the configurations category.

der to compare the distribution of SNACS supersenses and PEDT functors, we first isolated all sentences containing relevant nominals from a small subsection of English PEDT using the tectogrammatical layer annotations. Specifically, we target nominals introduced by prepositional phrases, which have a formeme in the t-layer of the form "noun+preposition+X". Nominals with formemes of this type are found in a PP in the surface syntax. In total, we extract 838 sentences with 1837 total PPs with functor labels. These sentences were fed into a state-of-the-art SNACS supersense classifier (Arora, 2023), and a predicted supersense label was gathered for each of the target PPs. Since these tags were automatically generated, there is some expected noise in the resulting predictions, particularly for uncommon supersenses. We sampled 100 preposition tokens for manual tagging: focusing on the 71 that were not **of** tokens (as **of** is usually configurational), we found that the predicted classifier agreed with expert judgments roughly 60% of the time. We compare the automatically generated supersense labels with a rule-based heuristic based on our expectations outlined in §2. Generally, our heuristic aligns PEDT functors with the supersense that is most similar in definition. This heuristic was shown to be roughly 52% accurate on the manually tagged sample. After showing the overall distribution of supersenses across different functors, we then isolate examples of divergences between the automatic classifier and rule-based heuristic, finding that divergences come from both tagging errors and meaningful differences in the two frameworks.

3.2 Results

We compare the distribution of SNACS supersenses with PEDT functors in Figures 1 to 3. For all comparisons, supersenses that were predicted less than 5 times were excluded from analysis. Figure 1 shows the general overlap of different coarse groups SNACS supersenses with groupings of PEDT functors. We see here that supersenses grouped around broad semantic domains typically correspond to groups of PEDT functors with similar domains. The most clear correspondences are with the spatiotemporal, "Agent-like" and "Patient-like" supersenses, indicating that despite the syntactic definition of ACT and PAT in PEDT, they still pattern similarly to the semantic based categories in SNACS.

Figure 2 shows the overlap of spatiotemporal supersenses and functors with a higher degree of granularity than in Figure 1. We see that LOCUS and TIME are two of the most frequently predicted supersenses, and generally line up well with the LOC and TWHEN functors. This is in contrast with the overlap for **CONFIGURATION** supersenses, which is shown in Figure 3. We can see here that most of the supersenses seem to be spread over several competing PEDT functors. As expected, APP and MAT have substantial representation in these supersenses, but there is also considerable overlap with other unexpected PEDT functors.

We report the overlap of the predicted classifier supersenses with those predicted by a rule-based heuristic for different functor groupings in Table 2. We see that our expectations for functors and the predictions of the classifier diverge substantially, especially for configurations, though there is substantial divergence even in the spatiotemporal and participant classes.

Since the automatic SNACS classifier has substantial limitations in tagging WSJ data, it is worth considering whether the divergence reported in Table 2 is primarily due to tagging errors, or is due to real differences in annotation distributions for supersenses and functors. In Examples (8–12), we show the classifier-predicted supersense alongside the gold functor. For Examples (8, 9), the predicted supersenses do not align with our expectations due to classification errors. We see in (8) that the classifier mistakenly predicts LOCUS instead of TIME. In this case, the heuristic which matches TWHEN to TIME would get this correct. In (9), the SNACS classifier predicts COST incorrectly, probably because it introduces a monetary amount as its dependent. Throughout the WSJ data, monetary values are often incorrectly classified as COST.

(8) The strong growth followed year-to-year increases of 21% **in**_TWHEN_LOCUS August and 12% in September.

(9) Imports were **at**_PAT_COST $50.38 billion, up 19%.

While classifier errors account for a substantial amount of misalignment between functors and supersenses, there are also systematic divergences. One reason for divergences is that some PEDT functors align more with SNACS *function* roles, rather than *scene* roles (as was expected). This is shown in Examples (10, 11), where both the predicted *scene* and *function* roles are shown. In Example (10), we see that the *scene* role of LOCUS does not align with DIR3, but the *function*, which is GOAL, does align with our expectations. This sentence is an example of *fictive motion*, where a preposition typically indicating motion is used in a static scene (Talmy, 1996; Hwang et al., 2017). In Example (11), we see that the *function* role of ANCILLARY is what we would expect to align with the ACMP functor, though the *scene* role AGENT does not. Problematic cases involving the ANCILLARY supersense have been a focus of prior SNACS research (Hwang et al., 2020), so it is perhaps unsurprising that some divergences arise in this case. Despite such examples, in most cases where *scene* and *function* differ, we observe that the *scene* role is closer to the PEDT functor. More investigation is needed to determine when PEDT functors map to *function* roles instead of *scene* roles in SNACS.

(10) The new plant, located in Chinchon about 60 miles **from**_DIR1_LOCUS⇝GOAL Seoul, will help meet increasing and diversifying demand for control products in South Korea, the company said.

(11) Moscow has settled pre-1917 debts **with**_ACMP_AGENT⇝ANCILLARY other countries in recent years at less than face value.

Beyond the discrepancies between PEDT functors and SNACS supersenses which arise from the *scene* and *function* distinction in SNACS, there are other unexpected divergences between PEDT functors and SNACS supersenses, two of which are shown in Examples (12, 13). In (12), the classifier's prediction of STARTTIME is obviously incorrect, but the expectation that CPR aligns with COMPARISONREF is also incorrect here. Instead, SOURCE is probably most appropriate, but is not predicted by the classifier or from our heuristic. This is one case where the usage of PEDT functors and SNACS supersenses do not overlap. Furthermore, in (13), the PEDT functor DIR3 would typically align with GOAL, but in this sentence the classifier prediction of PURPOSE is actually closer to the correct supersense. In general, it seems that DIR3 is not as clearly aligned with GOAL as anticipated, but also has some overlap with TOPIC, THEME, and PURPOSE. Despite the similar definitions of DIR3 and GOAL, in practice they are used in some non-overlapping situations.

(12) A seat on the Chicago Board of Trade was sold for $350,000, down $16,000 **from**_CPR_STARTTIME the previous sale last Friday.

(13) Then, in the guests' honor, the speedway hauled out four drivers, crews and even the official Indianapolis 500 announcer **for**_DIR3_PURPOSE a 10-lap exhibition race.

4 Conclusion

In this work, we compare SNACS supersenses with PEDT tectogrammatical functors in terms of how they account for English prepositions. We show that the substantial definitional overlap between SNACS supersenses and PEDT functors is reflected in the overlapping distributions of the various semantic classes, particularly for spatial, temporal, and participant related supersenses, with less overlap on the CONFIGURATION branch. However, we also find substantial divergences between the two schemata, due in part to limitations of the automatic SNACS classifier we employed. We observe that a simple heuristic mapping from PEDT functors to SNACS supersenses aligns somewhat with classifier predictions, but also has substantial limitations due to the differences between the two frameworks.

5 Acknowledgements

We thank Jan Hajič for helpful discussions regarding functors in PEDT, and anonymous reviewers for their constructive comments and suggestions. This research was funded in part by NSF award IIS-2144881.

References

Aryaman Arora. 2023. snacs: Models for parsing SNACS datasets. https://github.com/aryamanarora/snacs.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90, Montreal, Quebec, Canada. Association for Computational Linguistics.

Ann Bies, Justin Mott, Colin Warner, and Seth Kulick. 2012. English Web Treebank. LDC2012T13.

Claire Bonial, Julia Bonn, Kathryn Conger, Jena D. Hwang, and Martha Palmer. 2014. PropBank: Semantics of New Predicate Types. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3013–3019, Reykjavik, Iceland. European Language Resources Association (ELRA).

Jan Hajič, Eva Hajičová, Jarmila Panevová, Petr Sgall, Ondřej Bojar, Silvie Cinková, Eva Fučíková, Marie Mikulová, Petr Pajas, Jan Popelka, Jiří Semecký, Jana Šindlerová, Jan Štěpánek, Josef Toman, Zdeňka Urešová, and Zdeněk Žabokrtský. 2012. Announcing Prague Czech-English Dependency Treebank 2.0. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3153–3160, Istanbul, Turkey. European Language Resources Association (ELRA).

Jena D. Hwang, Archna Bhatia, Na-Rae Han, Tim O'Gorman, Vivek Srikumar, and Nathan Schneider. 2017. Double Trouble: The Problem of Construal in Semantic Annotation of Adpositions. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, page 178–188, Vancouver, Canada. Association for Computational Linguistics.

Jena D. Hwang, Nathan Schneider, and Vivek Srikumar. 2020. Sprucing up Supersenses: Untangling the Semantic Clusters of Accompaniment and Purpose. In *Proceedings of the 14th Linguistic Annotation Workshop*, page 127–137, Barcelona, Spain. Association for Computational Linguistics.

Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A Large-scale Classification of English Verbs. *Language Resources and Evaluation*, 42(1):21–40.

Ken Litkowski and Orin Hargraves. 2005. The Preposition Project. In *Proc. of the Second ACL-SIGSEM Workshop on the Linguistic Dimensions of Prepositions and their Use in Computational Linguistics Formalisms and Applications*, page 171–179, Colchester, Essex, UK.

Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: the Penn Treebank. *Computational Linguistics*, 19(2):313–330.

Adam Meyers, Ruth Reeves, Catherine Macleod, Rachel Szekely, Veronika Zielinska, Brian Young, and Ralph Grishman. 2004. The Nombank Project: An Interim Report. In *Proceedings of the Workshop Frontiers in Corpus Annotation at HLT-NAACL 2004*, page 24–31, Boston, Massachusetts, USA. Association for Computational Linguistics.

Martha Palmer, Claire Bonial, and Jena D. Hwang. 2017. VerbNet: Capturing English verb behavior, meaning, and usage. In Susan E. F. Chipman, editor, *The Oxford Handbook of Cognitive Science*, pages 315–336. Oxford University Press.

Nathan Schneider, Jena D. Hwang, Vivek Srikumar, Archna Bhatia, Na-Rae Han, Tim O'Gorman, Sarah R. Moeller, Omri Abend, Adi Shalev, Austin Blodgett, and Jakob Prange. 2022. Adposition and Case Supersenses v2.6: Guidelines for English. arXiv:1704.02134 [cs.CL].

Nathan Schneider, Jena D. Hwang, Vivek Srikumar, Meredith Green, Abhijit Suresh, Kathryn Conger, Tim O'Gorman, and Martha Palmer. 2016. A Corpus of Preposition Supersenses. In *Proceedings of the 10th Linguistic Annotation Workshop held in conjunction with ACL 2016 (LAW-X 2016)*, pages 99–109, Berlin, Germany. Association for Computational Linguistics.

Nathan Schneider, Jena D. Hwang, Vivek Srikumar, Jakob Prange, Austin Blodgett, Sarah R. Moeller, Aviram Stern, Adi Bitan, and Omri Abend. 2018. Comprehensive supersense disambiguation of English prepositions and possessives. In *Proc. of ACL*, pages 185–196, Melbourne, Australia.

Nathan Schneider, Vivek Srikumar, Jena D. Hwang, and Martha Palmer. 2015. A hierarchy with, of, and for preposition supersenses. In *Proc. of The 9th Linguistic Annotation Workshop*, pages 112–123, Denver, Colorado, USA.

Leonard Talmy. 1996. Fictive motion in language and "ception". In Paul Bloom, Mary A. Peterson, Nadel Lynn, and Merrill F. Garrett, editors, *Language and Space*, pages 211–276. MIT Press, Cambridge, MA.