

Preparing SNACS for Subjects and Objects

Adi Shalev

Hebrew University of Jerusalem
adi.bitan@mail.huji.ac.il

Jena D. Hwang

IHMC
jhwang@ihmc.us

Nathan Schneider

Georgetown University
nathan.schneider@georgetown.edu

Vivek Srikumar

University of Utah
svivek@cs.utah.edu

Omri Abend Ari Rappoport

Hebrew University of Jerusalem
{oabend, arir}@cs.huji.ac.il

Abstract

Research on adpositions and possessives in multiple languages has led to a small inventory of general-purpose meaning classes that disambiguate tokens. Importantly, that work has argued for a principled separation of the *semantic role* in a scene from the *function* coded by morphosyntax. Here, we ask whether this approach can be generalized beyond adpositions and possessives to cover all scene participants—including subjects and objects—directly, without reference to a frame lexicon. We present new guidelines for English and the results of an interannotator agreement study.

1 Introduction

Studies of verbal argument structure have established some clear semantic correlations of syntactic relations like subject and object, and there are various approaches to expressing these generalizations using categorical semantic roles (Fillmore, 1968, 1982; Levin, 1993) or bundles of proto-properties (Dowty, 1991; Reisinger et al., 2015) that generalize across verbs. A parallel line of work (§2) has looked at the meanings coded by grammatical phrase-markers such as prepositions and possessives and how to disambiguate them. These inquiries necessarily overlap because many prepositions mark verb arguments or modifiers. Consequently, insights from the study of prepositions/case may improve the meaning representation of core syntactic arguments, or vice versa.

In this paper, we investigate whether SNACS (Schneider et al., 2018b), an approach to semantic disambiguation of adpositions and possessives, can be adapted to cover syntactically core grammatical relations (subjects and objects). We believe this may have several practical advantages for NLP.

First, many of the semantic labels in SNACS derive from VerbNet (Kipper et al., 2008) role labels.

However, VerbNet and other frame-semantic approaches like FrameNet (Fillmore and Baker, 2009) and PropBank (Palmer et al., 2005) assume a *lexicon* as a prerequisite for semantic role annotation. This can be an obstacle to comprehensive corpus annotation when out-of-vocabulary predicates are encountered. But is a lexicon really necessary for role annotation? A general-purpose set of role labels with detailed criteria for each can potentially bypass coverage limitations of lexicon-based approaches, while still supporting some degree of generalization across grammatical paraphrases.

Second, the nonreliance on a lexicon potentially *simplifies the annotation process* in some respects. For example, no explicit predicate disambiguation step is necessary, and the annotator does not need to consult frame-specific role definitions.¹

Third, the semantic criteria for SNACS labels are designed to be *language-neutral*, and investigations thus far suggest that they can be generalized to languages besides English (Hwang et al., 2017; Zhu et al., 2019). While this paper focuses on English, we see the future opportunity for cross-lingual extension without the construction of new lexicons as a major advantage.

Finally, SNACS is unique in allowing *two* semantic labels per target, one reflecting a level of meaning closer to the grammatical coding, and the other at a deeper level associated with the predicate scene type (§3). We show below that the SNACS analysis, while designed for PPs, can be extended to subjects and objects, to the extent that the coarse-grained inventory distinguishes roles in the scene.

We summarize SNACS in §2, and in §3 propose a strategy for adapting SNACS for English subjects

¹On the other hand, consulting a frame-specific set of core roles may simplify the role labeling task for an annotator, producing higher-quality annotations. In the future it may be worth exploring a hybrid solution that maps lexicon-defined roles to supersenses and asks the annotator to apply supersenses directly only for out-of-vocabulary predicates.



Figure 1: Modified inventory of supersenses based on the SNACS hierarchy (additions and removals in bold).

and objects. This involves minor changes to the label inventory and new annotation guidelines for a variety of challenging phenomena. We conduct a pilot interannotator agreement study on Wikipedia articles (§4) and release the annotations.²

2 Background

The SNACS³ hierarchy is a taxonomy of coarse-grained supersenses developed to mark semantic relations as expressed by adpositions (prepositions + postpositions) and possessives (Schneider et al., 2018b). The complete SNACS hierarchy is shown in figure 1 with our modifications highlighted.

SNACS includes the usual thematic relations (e.g., **AGENT**, **THEME**, **RECIPIENT**) and adjunct relations (e.g., **TIME**, **LOCUS** including locations, **PURPOSE**) used by most resources designed for SRL annotation. SNACS diverges from the general predicate-argument labeling standards in its inclusion of non-standard roles such as **ORIGINATOR** in creation (creator), transfer (giver) and communication (speaker) events, and labels regarding properties involved in a static relationship to one another (e.g., **POSSESSION** in “owner **of** the car”).

Unlike labels used by efforts such as PropBank and FrameNet, SNACS labels are highly coarse-grained and generalize across various scenes and situations. This approach also differs from frame-alternation-based lexicons like VerbNet, which defines classes of verbs whose members exhibit similar syntactic alternations involving the same subset

of roles. Instead, SNACS places the burden of semantics directly on a fixed set of **supersenses**, forgoing the use of frame (or class) definitions. The supersenses can be thought of as disambiguating coarse-grained adposition senses. The supersense labels effectively encapsulate—at a highly abstract/schematic level—various basic scenarios that are important to language and grammar, such as transitive action, motion, unidirectional transfer/communication, and psychological events, as well as stative relations like possession, quantity, comparison, and identity. SNACS does not formalize a semantic core/non-core or argument/adjunct distinction, though roles in the **PARTICIPANT** hierarchy are typically core and roles in the **CIRCUMSTANCE** hierarchy are typically non-core in predicate-argument annotation schemes like PropBank and FrameNet.

SNACS further adopts a device called *construal* (Hwang et al., 2017), explained below.

3 Applying SNACS

We adopt the SNACS labels originally developed for disambiguating adpositions and possessives as exemplified in (1) and extend their use to annotate the subject and object of a verb as seen in (2).

- (1) a. The bagel was eaten **by**_{AGENT} Jane.
b. Jane dined **on**_{THEME} a bagel.
- (2) [Jane]_{AGENT} ate [a bagel]_{THEME}.

Following the construal approach, which is illustrated in table 1 for adpositions, we separate two semantic dimensions of an annotation target: **Scene Role**: What semantic role is most closely associated with the type of scene (typically indicated by the verb/predicate)? **Function**: What semantic role is most salient in the morphosyntactic coding of the phrase (with a grammatical relation like subject or object, or overt marking with closed-class morphology like adpositions and case)? Consider the following examples. Construal is notated by **SCENE ROLE**→**FUNCTION**.

- (3) [Jane]_{RECIPIENT→AGENT} bought [the book]_{POSSESSION→THEME}.
- (4) [Bingley]_{SOCIALREL→THEME} married [Jane]_{SOCIALREL→THEME}.

The scene role indicates the participation role of the target in the scene described by the verb. Jane is the **RECIPIENT** in a transfer scene in (3), and she is in a certain social relationship with Bingley (i.e., **SOCIALREL**) given the marriage scene in (4). The function label, on the other hand, captures the

²https://github.com/adishalev/SNACS_DMR_IAA

³Semantic Network of Adposition and Case Supersenses

Phrase	Scene Role	Coding	Function	Congruent?
The ball was hit by the batter	AGENT	by	AGENT	✓
Put the book on the shelf	GOAL	on	LOCUS	✗
Put the book onto the shelf	GOAL	onto	GOAL	✓
I talked to her	RECIPIENT	to	GOAL	✗
I heard it in my bedroom	LOCUS	in	LOCUS	✓
I heard it from my bedroom	LOCUS	from	SOURCE	✗
John's death	THEME	's	GESTALT	✗
the windshield of the car	WHOLE	of	WHOLE	✓

Table 1: SNACS for adpositions/possessives (Schneider et al., 2018b,a). The scene role and function annotations are labels from figure 1 and are often but not always congruent for a particular token. The function annotation reflects the semantics of the morphosyntactic coding (such as the choice of adposition). Note that, especially for adnominal PPs and genitives, the governor sometimes does not lexically denote an event or state; rather, a semantic relation such as possession or part-whole is indicated by the morphosyntax.

orthogonal dimension of agency which is more closely tied to syntactic realization: Jane is the **AGENT** of the buying action, while the book is the **THEME** in (3); Jane and Bingley are the **THEMES** of the marriage in (4). Further examples are shown in figure 2. In many cases, the scene role will be identical to the function. These are called **congruent construals**. But in other cases, they can differ, as illustrated in table 1.

In the rest of the section, we discuss a few difficult cases while assessing SNACS labels for the annotation of subject and objects, and decisions made regarding these challenges including slight deviations from the latest SNACS standards.

Scene role prioritization. In some cases, multiple supersenses are equally applicable at the scene level. In such cases, we give highest priority to more complex and less frequent scene types such as transfer (**ORIGINATOR**, **RECIPIENT**) or employee-organization (or social) relations (**ORGMEMBER**, **ORG**, **SOCIALREL**). The causal roles (**AGENT**, **INSTRUMENT**, **THEME**), if appearing in the scene position, are prioritized next. The highly frequent locative scenes (**LOCUS**, **SOURCE**, **GOAL**) are given the lowest priority. In example (10), the subject “I” could be considered either a metaphorical source location of the recommendation (i.e., **SOURCE**) or can be considered the speaker in a communication event (i.e., **ORIGINATOR**). The latter scene is prioritized, and the scene roles **ORIGINATOR**, **TOPIC** (i.e., the message), and **RECIPIENT** reflect the prioritized choice.

- (10) [I]_{ORIGINATOR→AGENT} recommended [the book]_{TOPIC→TOPIC} [to him]_{RECIPIENT→GOAL}.

Transfer of possession often implies change of location, and being a part of something often implies being located in it. If both are salient, for the scene role annotation, we prioritize the more

complex scene over the locative semantics:

- (11) [Jane]_{RECIPIENT→AGENT} took the book from me.
(12) I relinquished the book [to Jane]_{RECIPIENT→GOAL}.
(13) At the play, he spotted Mary [in the cast]_{ORG→LOCUS}.

In (11, 12), Jane is arguably a **GOAL** of motion by virtue of being a **RECIPIENT** of something physical. We do not use **GOAL** as the scene role, however, if **RECIPIENT** applies. In (13), Mary can be understood as part of the cast (which is an organization) or as located within the cast. We prioritize the former. Other pairs that tend to overlap include: **RECIPIENT/BENEFICIARY**, closely correlated when someone is given or told something for their benefit or harm—we prioritize **RECIPIENT** for the scene role; and **STIMULUS/TOPIC**, closely correlated when a thought or message triggers an emotional reaction—we prioritize **STIMULUS** for the scene role.

If two equally prioritized scenes are in conflict with one another, we rely on the semantics of the predicate to disambiguate the scene. Note that in (14), CJ is likely an employee of the White House. However, CJ is not considered the **ORGMEMBER** as the verb “brief” does not intrinsically conventionalize the employee-organization relationship in its semantics. With a predicate that conventionally encodes employment or some other stable relationship—*employ, hire, work for/at*, etc.—**ORGMEMBER** would be annotated.

- (14) [CJ]_{ORIGINATOR→AGENT} briefs the press [for the White House]_{ORG→BENEFICIARY}.

Role duplication. The latest version of SNACS we adopt for our study does not allow participant

- (5) a. [Rachel]_{AGENT→AGENT} opened [the door]_{THEME→THEME} [with the remote control]_{INSTRUMENT→INSTRUMENT}.
 b. [The remote control]_{INSTRUMENT→INSTRUMENT} opened [the door]_{THEME→THEME}.
 c. [I]_{AGENT→AGENT} used [the remote control]_{INSTRUMENT→INSTRUMENT} [to open the door]_{PURPOSE→PURPOSE}.
 d. [The door]_{THEME→THEME} opened.
 (6) [Rachel]_{THEME→THEME} sneezed.
 (7) [Rachel]_{EXPERIENCER→AGENT} watched [the children playing]_{STIMULUS→THEME}.
 (8) [Rachel]_{EXPERIENCER→THEME} heard [the noise]_{STIMULUS→THEME}.
 (9) [Rachel]_{AGENT→AGENT} spent [\$5]_{COST→COST} [on coffee]_{POSSESSION→THEME}.

Figure 2: Annotated examples from our guidelines.

labels such as **AGENT** or **THEME** to appear multiple times in a given scene, opting for the use of a “Co-” label for the second participant sharing the same role (e.g., **Co-AGENT**). In applying SNACS guidelines for subjects and objects, this became untenable, as “Co-” prefixation could apply to a good majority of the **PARTICIPANT** labels, threatening a quick proliferation of the supersenses. E.g., (4) would require **CO-SOCIALREL**, (16) would require **CO-EXPERIENCER**, and so forth. In an effort to keep the supersense inventory limited, we diverge from the latest SNACS standards to allow role duplication in a scene. This is allowed even when targets assigned the same role are not fully symmetric or are qualitatively distinct as in (17).

- (15) [A reception]_{THEME→THEME} will precede [the dinner]_{THEME→THEME}.
 (16) [He]_{EXPERIENCER→THEME} heard the news [with a stranger]_{EXPERIENCER→ACCOMPANIER}.
 (17) Replace [the old one]_{THEME→THEME} [with the new one]_{THEME→ACCOMPANIER}.

Thematic hierarchy. As discussed above, the function label generally reflects **AGENT-THEME** relations of a proposition. More specifically, we annotate all subjects and *direct* objects with a function in the following thematic hierarchy: {**AGENT**, **CAUSER**} > {**INSTRUMENT**, **MEANS**} > {**THEME**, **TOPIC**, **COST**}. In a transitive clause, the supersense of the subject cannot be ranked lower than the direct object (e.g., a subject construed as a **THEME** cannot have a direct object construed as an **AGENT**). Indirect objects in the English double object construction⁴ are treated as **RECIPIENT** construals.

- (18) I sent [John]_{RECIPIENT→RECIPIENT} a cake.
 (19) I sent a cake [to John]_{RECIPIENT→GOAL}.
 (20) I baked [John]_{RECIPIENT→RECIPIENT} a cake.
 (21) I paid [John]_{RECIPIENT→RECIPIENT} [\$10]_{COST→COST}.

⁴If there is a single overt object, we treat it as an indirect object if what would normally be a direct object is implicit: e.g., *John* is treated as an indirect object in *I told John*. (where the content of what is told is implicit).

Copular sentences. These are treated differently from non-copular sentences. The English copula relates a subject to an object in what is semantically an identificational (22a) or predicational (22b) relationship. To these cases we assign **IDENTITY-IDENTITY** or **GESTALT-CHARACTERISTIC** at the scene level, respectively. Roughly speaking, **IDENTITY** indicates the identified or identifying category or referent, and **CHARACTERISTIC** indicates a property being ascribed to the **GESTALT**:

- (22) a. [John]_{IDENTITY→IDENTITY} is [a man]_{IDENTITY→IDENTITY}.
 b. [John]_{GESTALT→THEME} is [tall]_{CHARACTERISTIC→CHARACTERISTIC}.

Open issues. The unresolved problem of causatives and caused-motion constructions is discussed in appendix A.

4 Interannotator Agreement Study

Data. We piloted our guidelines using a sample of 100 scenes from the English UCCA-annotated Wiki corpus⁵ as detailed by Abend and Rappoport (2013). UCCA is a scheme for annotating coarse-grained predicate-argument structure such that syntactically varied paraphrases and translations should receive similar analyses. It captures both static and dynamic scenes and their participants, but does not mark semantic roles.

Annotators. Four annotators (A, B, C, D), all authors of this paper, took part in this study. All are computational linguistics researchers.

Datasets. Prior to development of guidelines for subjects and objects, one of the annotators (Annotator A) sampled 106 Wiki documents (44k tokens) and tagged all 10k instances of UCCA Participants⁶ with a supersense based on the existing guidelines

⁵<http://cs.huji.ac.il/~oabend/uca.html>; the Wikipedia corpus contains 369 documents (biographies of entertainment industry figures) with 159k tokens and 36k Participant units.

⁶The UCCA category Participant is broader than the **PARTICIPANT** supersense, also including locations, for example.

Subjects/Objects (N=57)					PPs (N=42)				
κ	A	B	C	D	κ	A	B	C	D
A		.75	.38	.72	A		.68	.68	.68
B	.64		.42	.83	B	.54		.79	.84
C	.50	.63		.54	C	.57	.64		.92
D	.68	.83	.65		D	.60	.75	.75	
Scene Role					Scene Role				
Function					Function				

Table 2: Cohen’s Kappa scores for interannotator agreement for all pairings of four annotators.

for adpositions. This preliminary dataset was used to stimulate discussion for developing guidelines for subjects and objects. Once the new guidelines were written, four annotators first annotated a practice dataset of 48 UCCA Participant tokens, then adjudicated disagreements through discussion and clarified certain policies before annotating a final sample of 100 tokens for measuring agreement.⁷ Participant units were sampled based on the preliminary annotation, stratified across subtypes of **PARTICIPANT** and **CONFIGURATION** to ensure diversity. In the final sample, the syntactic distribution is as follows: 31 subjects (including 4 passive subjects and 6 copular subjects); 26 objects (including 1 indirect object and 2 copular complements); 42 PPs; and 1 possessive.

Coverage of the hierarchy. Under the **PARTICIPANT** tree in the hierarchy, there are 12 supersenses, of which 11 were used as scene roles and 9 as functions. (By design, **PARTICIPANT** itself is never used and exists only to organize the hierarchy.) The **CONFIGURATION** tree includes 19 supersenses, of which 14 were used as scene roles and 10 as functions. In the **CIRCUMSTANCE** tree—which primarily applies to syntactic adverbials—**GOAL**, **LOCUS**, **SOURCE**, **MANNER**, **MEANS**, and **CIRCUMSTANCE** were all used as functions, and all but **SOURCE** also appeared as a scene role.

Quantitative IAA results. We first compare agreement on two subsamples: the subject/object Participants, and the prepositional phrase Participants. Pairwise Cohen’s κ scores appear in table 2.

Subjects/objects: For the scene role, all annotators agree on 46% of items (26/57), and at least 3 annotators on 84%. For the function, 51% have total agreement, and 86% have a majority. Average pairwise κ is 0.66 for scene and 0.61 for function.

PPs: At the scene level, 48% (20/42) have total agreement, and 71% have a majority. For the function, 64% have total agreement, and 88% have a majority. Average pairwise κ is 0.64 for scene and 0.77 for function.

⁷4 tokens where annotators noticed a problem with the UCCA annotation were discarded and replaced.

Thus subjects/objects (SOs) receive higher scene role agreement than PPs—somewhat surprising given that the labels were originally designed for prepositions! This may be an artifact of the particular sample, or may indicate that the scene role is more intuitive for SOs than for PPs. PPs have higher agreement than SOs with respect to function; this may be due to some difficulty deciding between **AGENT** and **THEME** for the function of SOs, plus the availability of extensive guidelines/examples for prepositional SNACS annotation.⁸

Disagreements involving agentivity. We found it can be difficult to choose between **AGENT** and **THEME** for the function of a subject with borderline agentivity, e.g., in scenes of *befriending* someone or *forming* a musical group with others. Likewise, the line between **AGENT** and **THEME** for the function can be unclear in cognition/perception scenes like *[She] enjoyed the fame* and *[She] saw the social scene as tedious and superficial*. We decided the annotator should consider whether the scene involves judgment or is more of a passive experience; **EXPERIENCER**~**THEME** would thus apply to the first example and **EXPERIENCER**~**AGENT** to second.⁹ Finally, the line between **CAUSER** and **INSTRUMENT** can be unclear in sentences like *I was hit [by a car]* and *I was quoted [by a magazine]*.

UCCA issues. We found a handful of UCCA annotation errors—primarily where two verbs were analyzed as separate scenes but the first ought to be considered a light verb. A more interesting case was the relation between the two bolded expressions in *William S. Paley set terms that included... ownership of the negative at the end of the contract*. The UCCA annotation treats *William S. Paley* as a Participant of *ownership* (i.e., the owner). Though **POSSESSOR** is a natural scene role for the owner of something, we concluded that this was an indirect inference not suitable for annotating with a function.

5 Conclusion

We explored whether a system for semantic relation annotation can be extended beyond prepositions and possessives to cover English subjects and objects. While initial annotation results are promising, further work is needed to substantiate the approach on a larger scale, and ideally in multiple languages.

⁸Raw agreements are higher than kappa, but the same trends hold.

⁹There is precedent for this distinction in FrameNet’s *Perception_active* vs. *Perception_experience* frames.

Acknowledgments

We would like to thank Jakob Prange and anonymous reviewers for their feedback. This research was supported in part by NSF award IIS-1812778 and grant 2016375 from the United States–Israel Binational Science Foundation (BSF), Jerusalem, Israel.

References

- Omri Abend and Ari Rappoport. 2013. [Universal Conceptual Cognitive Annotation \(UCCA\)](#). In *Proc. of ACL*, pages 228–238, Sofia, Bulgaria.
- David Dowty. 1991. [Thematic proto-roles and argument selection](#). *Language*, 67(3):547–619.
- Charles J. Fillmore. 1968. [The case for case](#). In Emon Bach and Robert Thomas Harms, editors, *Universals in Linguistic Theory*, pages 1–88. Holt, Rinehart, and Winston, New York.
- Charles J. Fillmore. 1982. Frame Semantics. In *Linguistics in the Morning Calm*, pages 111–137. Han-shin Publishing Co., Seoul, South Korea.
- Charles J. Fillmore and Collin Baker. 2009. [A frames approach to semantic analysis](#). In Bernd Heine and Heiko Narrog, editors, *The Oxford Handbook of Linguistic Analysis*, pages 791–816. Oxford University Press, Oxford, UK.
- Adele E. Goldberg. 2006. *Constructions at work: the nature of generalization in language*. Oxford University Press, Oxford.
- Jena D. Hwang, Archana 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 *Proc. of *SEM*, pages 178–188, Vancouver, Canada.
- 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.
- Beth Levin. 1993. *English verb classes and alternations: a preliminary investigation*. University of Chicago Press, Chicago.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. [The Proposition Bank: An annotated corpus of semantic roles](#). *Computational Linguistics*, 31(1):71–106.
- Drew Reisinger, Rachel Rudinger, Francis Ferraro, Craig Harman, Kyle Rawlins, and Benjamin Van Durme. 2015. [Semantic proto-roles](#). *Transactions of the Association for Computational Linguistics*, 3:475–488.
- Nathan Schneider, Jena D. Hwang, Archana Bhatia, Na-Rae Han, Vivek Srikumar, Tim O’Gorman, Sarah R. Moeller, Omri Abend, Austin Blodgett, and Jakob Prange. 2018a. [Adposition and Case Supersenses v2: Guidelines for English](#). *arXiv:1704.02134 [cs.CL]*.
- Nathan Schneider, Jena D. Hwang, Vivek Srikumar, Jakob Prange, Austin Blodgett, Sarah R. Moeller, Aviram Stern, Adi Bitan, and Omri Abend. 2018b. [Comprehensive supersense disambiguation of English prepositions and possessives](#). In *Proc. of ACL*, pages 185–196, Melbourne, Australia.
- Yilun Zhu, Yang Liu, Siyao Peng, Austin Blodgett, Yushi Zhao, and Nathan Schneider. 2019. [Adpositional supersenses for Mandarin Chinese](#). In *Proc. of SCiL*, volume 2, pages 334–337, New York, NY, USA.

A Open Issues

Sometimes a sentence will construe a scene as involving more arguments than a predicate normally licenses, as in the following causative or caused-motion examples (Goldberg, 2006):

- (23) [Rachel]_{THEME→?} sneezed (implicit: blew) [the napkin]_{THEME→THEME} [off the table]_{PATH→SOURCE}.
- (24) [Rachel]_{AGENT→?} (implicit: caused) jumped [the horse]_{AGENT→THEME} [over the fence]_{PATH→PATH}.

So far, we have posited the scene to be the situation or event described by the predicate. The problem is that in addition to the scene evoked by the verb (sneezing in (23) and jumping in (24)), there is an added caused-motion scene whose semantics derives from the construction. Should there be an indication that the sneezer is also the causer of motion, and that the ultimate causer of jumping is separate from the impelled jumper? One possible solution would be to add implicit predicates so the verb- and construction-triggered scenes would be annotated separately. A different solution may be to relax the definition of what constitutes a scene to allow for non-predicate-driven scenes as well.