Implementing ASLNet V1.0: Progress and Plans

Colin P. Lualdi[1,2,†] **Elaine Wright**[1,3,4] **Jack Hudson**[1] **Naomi K. Caselli**[5] **Christiane Fellbaum**[4,6]

[1]SignSchool Inc., Madison, Wisconsin, USA
[2]Department of Physics, University of Illinois, Urbana-Champaign, Illinois, USA
[3]Department of Electrical Engineering, Princeton University, Princeton, New Jersey, USA
[4]Program in Linguistics, Princeton University, Princeton, New Jersey, USA
[5]Programs in Deaf Studies, Boston University, Boston, Massachusetts, USA
[6]Department of Computer Science, Princeton University, Princeton, New Jersey, USA
[†]colin@signschool.com

Abstract

We report on the development of ASLNet, a wordnet for American Sign Language (ASL). ASLNet V1.0 is currently under construction by mapping easy-to-translate ASL lexical nouns to Princeton WordNet synsets. We describe our data model and mapping approach, which can be extended to any sign language. Analysis of the 390 synsets processed to date indicates the success of our procedure yet also highlights the need to supplement our mapping with the "merge" method. We outline our plans for upcoming work to remedy this, which include use of ASL free-association data.

1 Background and Motivation

First proposed in 2019 by Lualdi et al., ASLNet is an effort to extend the wordnet model pioneered by the Princeton WordNet (Miller, 1995; Fellbaum, 2010) to the visual-kinesthetic realm of sign languages. This endeavor is in part inspired by the creation of wordnets in dozens of other spoken languages, including those outside the Indo-European language family (Bond and Foster, 2013; Vossen, 2004), as well as images via ImageNet (Deng et al., 2009). It is only natural to develop wordnets for sign languages like American Sign Language (ASL) as they are unique languages in their own right.

There are many benefits to creating a wordnet representation of the ASL lexicon. The semantic relations encoded by a wordnet enable semantically-driven language acquisition (Miller and Fellbaum, 1992), resulting in a powerful first-language (L1) and second-language (L2) pedagogical resource that will also contribute to ASL linguistics. Furthermore, with the Princeton WordNet (PWN) serving as a hub linking multiple wordnets, connecting an ASL wordnet to PWN will bridge ASL to other languages (and ImageNet images), allowing for novel linguistic investigations. Lastly, with wordnets being invaluable to natural language processing (NLP), specifically word sense disambiguation (Navigli, 2009), an ASL wordnet will support burgeoning ASL machine translation efforts (Bragg et al., 2019).

The original ASLNet proposal (Lualdi et al., 2019) examined theoretical questions and strategies for extending the wordnet model to a sign language. The findings were synthesized into a proposed roadmap for creating ASLNet via a hybrid "map" and "merge" approach. ASLNet development would start with mapping straightforward ASL lexical nouns to PWN synsets, followed by creating ASLNet synsets with ASLNet-specific relations to be merged with PWN where appropriate.

In this paper, we report on the latest progress on implementing ASLNet V1.0 according to the prescription set forth by Lualdi et al. (2019). We describe the mapping procedure and evaluate its effectiveness. We also discuss how our work to date is informing the direction of subsequent ASLNet development.

2 ASLNet V1.0 Overview

As recommended by Lualdi et al. (2019), our objective for ASLNet V1.0 is to map ASL signs to their corresponding PWN synsets. For simplicity, we consider only lexical nouns, which often refer to concrete entities and hence are easier to represent; nouns also tend to map better crosslingually than verbs. Furthermore, most of the words in a lexicon (e.g., dictionaries, wordnets, etc.) are generally nouns, resulting in more data to work with. Therefore, ASLNet V1.0 is a table of noun PWN synsets and their mapped ASL sign(s). All semantic structure is directly derived from PWN, considerably simplifying the development work, albeit at the cost of (temporarily) ignoring aspects of ASL not present in English and therefore not encoded by PWN, such as classifier constructions, which lack a clear parallel in the English language.

While the "map" technique is not new to the wordnet community, this work presents a novel challenge in that, by working with a sign language, we are required to employ video exemplars. This stems from the lack of a conventional system for transcribing signs; there is no standardized writing system or even an International Phonetic Alphabet (IPA) for sign languages. Consequently, the signing community has no consensus on how to distinguish phonologically similar signs from one another, which complicates the isolation of particular signed forms for encoding in a sign language wordnet. By leaning on existing PWN synsets and structure during this initial stage of development, we therefore have more bandwidth for implementing an experimental model for organizing the sign data.

Difficulties with encoding signs are significant contributing factors to the resource-scarce nature of sign languages like ASL. They complicate the logistical challenges of gathering and processing video exemplars of signs, especially in the absence of practical computer vision, motion capture, and sign language machine translation technologies. Consequently, sign language lexical databases and corpora tend to be comparatively smaller than those of languages accompanied by robust orthographies. Accordingly, the challenges faced by the ASLNet team are not so different from those of teams working with other under-resourced languages. In fact, many of the techniques utilized in the development of ASLNet V1.0, such as the initial focus on mapping lexical nouns, are similar to those employed by the African Wordnet Project (AWN) in creating wordnets for five resource-scarce African languages (Bosch and Griesel, 2017).

2.1 Sign Data

In the original ASLNet proposal, it was suggested that the ASL sign data be drawn from SignStudy (www.signschool.org), a non-profit ASL research resource. SignStudy is supported by SignSchool Technologies LLC (www.signschool.com), a Deaf-led and owned ASL education company. While SignStudy's database size is respectable with 4,500+ sign videos, ASLNet V1.0 will function best as a wordnet when it possesses multiple clusters with a high density of sign-synset mappings. Furthermore, understanding how PWN structure lends itself to this small subset of ASL data will guide the ASLNet "merge" phase by highlighting any PWN deficiencies (in the context of ASL) that need addressing. So, to improve our ability to create such well-filled regions of PWN semantic hierarchy, we increase the number of documented signs available for mapping by also incorporating signs from two other ASL databases, ASL-LEX 2.0 (Sehyr et al., 2020; Caselli et al., 2017) with ∼2,700 signs and ASL Signbank (Hochgesang et al., 2020) with ∼3,500 signs.

As SignStudy, ASL-LEX, and ASL Signbank contain sign metadata[1], incorporating their signs not only improves ASLNet filling but also makes available linguistic data that will likely prove valuable for implementing ASLNet-specific relations and features during the upcoming "merge" stage of ASLNet development.

2.2 Data Model

To organize the ASLNet V1.0 data, we developed a tripartite model (Fig. 1). The first (lowest) level consists of "Signs", individual sign entries (including metadata) from the three sign databases[2]. Since these databases may overlap in coverage, we introduce a second (middle) level that combines identical-in-form Signs into "Combined Sign" objects. This merges complementary metadata for duplicate signs, resulting in a richly-annotated combined lexical database. Together, Signs and Combined Signs comprise the "Form Level", as they are strictly concerned with sign production; their organization is independent of semantics.

Note that determining whether two Signs should be grouped together in a Combined Sign (i.e., considered identical in form) or kept separate is not always a clear-cut process. One could adopt the strategy of considering signs identical if every phonological component is shared. However, the aforementioned lack of widely-used conventions for cod-

[1]SignStudy: Each sign is annotated with its constituent handshapes (∼70 unique handshapes identified) as well as semantic category (∼40) and subcategory (∼200). ASL-LEX: Each sign is annotated for six phonological properties (sign type, selected fingers, flexion, major and minor location, and movement), four lexical properties (initialization, lexical class, compounding, and fingerspelling), and subjective frequency and iconicity ratings. ASL Signbank: Each sign is identified by a unique "ID gloss" and partially annotated with various phonological, morphological, semantic, and miscellaneous metadata.

[2]In this paper, "Sign" with a capital "S" refers to the sign data object while "sign" with a lowercase "s" refers to the actual sign itself.

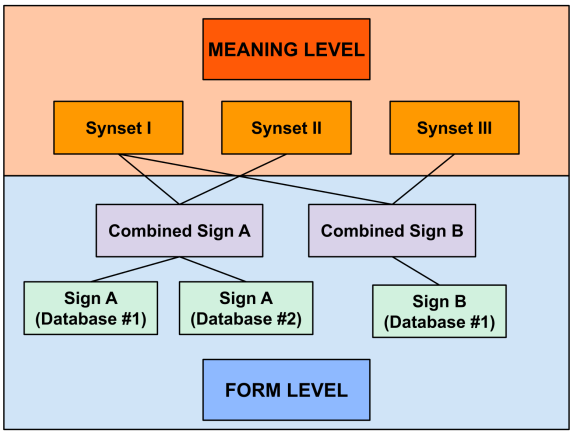

Figure 1: The ASLNet V1.0 data model. Sign A (Database #1) and Sign A (Database #2) are duplicates merged into Combined Sign A. Sign B is a distinct sign with its own Combined Sign object (B). Both Combined Sign A and B are polysemous and map to multiple synsets (I & II and I & III). They are also synonymous for a certain sense (Synset I) and thus both map to it.

ing ASL phonology makes it difficult to distinguish similar-in-form signs from one another. For the time being, we consider signs identical in form if they are treated as indistinguishable in use by native speakers. Any evidence of perceivable difference (e.g., one sign being a known regional variation of the other) is grounds for distinguishability. Since we expect this criteria to evolve as sign encoding conventions develop, our data model is designed to be flexible by allowing for easy rearrangement of Signs under Combined Signs without the need to drastically modify the entire system.

The third (top) level is the "Meaning Level", where we introduce semantics by linking Combined Signs to their corresponding PWN synsets. Due to polysemy, a Combined Sign may link to multiple PWN synsets. Similarly, synonymy results in individual synsets being associated to several Combined Sign objects.

3 Sign-Synset Mapping

To link the "Form" and "Meaning" levels of our data model, we developed a procedure to map signs to synsets with the objective of creating high-density synset clusters in ASLNet V1.0. While we are working with ASL signs, our procedure may be extended to any sign language with available lexical databases and corresponding wordnets.

3.1 Choosing Synset Clusters

With 10^4 signs and 10^5 synsets[3] available, it is challenging to identify the initial synset clusters to build. To condense our options, we imposed two criteria: relevance and efficiency.

We ensure relevance by considering only synsets belonging to common semantic domains (e.g., people, food, etc.) appearing frequently in everyday ASL discourse; their early incorporation will make ASLNet useful sooner.

To help us identify these synsets, we can utilize both the English equivalents of the signs in our combined ASL lexical database and the "Core" WordNet (CPWN), a collection of 5,000 more-frequently used word senses derived from British National Corpus (BNC) frequency data (Boyd-Graber et al., 2006). At high frequencies, BNC only differs by about 10% from the Corpus of Contemporary American English (Davies, 2011–), so the use of CPWN synsets to approximate frequently used word senses in American English is reasonable. While it would be ideal to use ASL frequency ratings, said data is limited due to the sign-coding challenges mentioned previously. However, since a large number of ASL speakers are bilingual ASL-English Americans, it is fair to assume frequency data for ASL and American English are comparable to a degree (Wright, 2020). Indeed, it was found that ASL-LEX subjective frequency data is moderately correlated with English frequency counts (Caselli et al., 2017). Furthermore, the small sizes of the sign databases we utilize imply that the included signs are relatively frequently used.

Therefore, instead of searching the entirety of PWN 3.0 for possible clusters of interest, we constrain ourselves to a smaller subset formed by the union of (A) all CPWN synsets and (B) PWN 3.0 synsets with at least one lemma matching sign data English equivalents[4]. The resulting synsets are favorable to our sign data while also identifying possible gaps worth filling during ASLNet development.

To achieve efficient use of the sign data supplied by SignStudy, ASL-LEX, and ASL Signbank, we chose domains from this subset that were very likely to achieve high sign-synset mapping densities. E.g., if our combined sign database is rich in

[3]PWN 3.0 synsets.

[4]Note that guessing signs' corresponding PWN synsets on the basis of the signs' manually annotated English equivalents is a crude heuristic; the listed translations may not be comprehensive or capture all of a sign's meanings.

"vegetables" signs but lean in "fruits", it is in our interest to perform the mapping work in the former.

We devised a computerized screening process incorporating the two criteria above to identify candidate clusters. First, we generated the union subset. Then, we checked if any synset in this set was a direct PWN hypernym of another set element, and if so, we added the hypernym to the candidate list. Synsets with a common 1st- or 2nd-level hypernym were also identified, with the shared hypernym added to the candidate list. After filtering for duplicates, we generated a list of existing 1st- and 2nd-generation hyponyms for each candidate list synset[5].

Each of these lists were scored by the proportion of constituent synsets with at least one lemma matching signs' English equivalents. The clusters (labeled by their "parent" synset from the candidate list) with the highest scores (closer to 1) were therefore recognized as optimal starting points.

Results from the screening process are summarized in Fig. 2 and Table 1. Smaller clusters, for the most part, score better than larger clusters. Overall, 3606 candidates with nonzero scores were identified, with an average score of 0.33 ($\sigma = 0.24$).

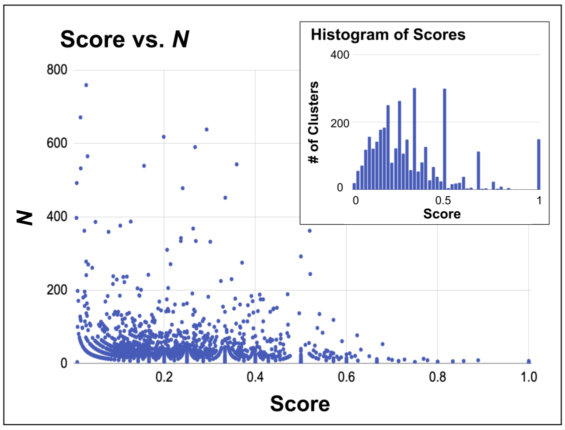

Figure 2: Candidate cluster score versus size (N). High-scoring clusters tend to be smaller. Plot excludes two $N > 1,000$ outliers [(0.002, 2532), (0.19, 1615)]. Inset: Histogram of candidate cluster scores.

3.2 Mapping Protocol and Tool

With the target synset clusters identified, the next step is to map signs to the synsets belonging to these clusters. A team of 3 ASL-English bilingual lexicographers[6] was assembled for the mapping.

[5]CPWN is a disjointed list with no internal navigation functionality. Since we intend to map to well-filled clusters, we elect to generate the hyponym lists from PWN.

Candidate Synset	Score	*N*
contact.n.04	0.89	9
kinsman.n.01	0.67	6
hair.n.01	0.40	43
jewelry.n.01	0.31	36
vegetable.n.01	0.15	75
pasta.n.02	0.08	26

Table 1: Select results of the screening process to determine possible starting clusters for ASLNet V1.0. Score (N) denotes the proportion (number) of 1st- or 2nd-generation hyponyms in the cluster that can likely map to one or more available signs.

As a preliminary step, one of the lexicographers manually generated correspondence tables linking SignStudy, ASL-LEX, and ASL Signbank, grouping duplicate signs into Combined Sign objects. Since combining Signs into Combined Signs can be a difficult task for reasons mentioned previously, Sign objects (rather than Combined Signs) are currently being mapped directly to the synsets; at a later time the Signs will be condensed into Combined Signs on a synset-by-synset basis and the results checked against the manually-prepared Combined Signs.

For the mapping, we developed an online tool ("Synset Mapper") to guide the lexicographers through the mapping protocol as follows:

Step 1: Lexicographer searches for a PWN synset. Search returns a list containing the query synset and all of its existing 1st-generation hyponyms[7].

Step 2: Clicking on any of these synsets opens up its review page displaying the synset's name, definition, status, review state, notes, mapped Signs, computer-suggested Signs, and a Sign search. Each Sign is presented as a user-playable video accompanied by its associated English equivalent(s) and source (i.e., SignStudy, ASL-LEX, or ASL Signbank).

Step 3: Lexicographer reviews the synset definition and selects the appropriate Sign(s) from either the suggestions or a manual gloss search. See Fig. 3 (located at the end of the paper, after the references section) for a screenshot of this step.

[6]The team consisted of two deaf individuals with graduate-level education and one hearing ASL interpreter with a MA in Linguistics.

[7]From the full PWN 3.0.

Step 4: Lexicographer updates the synset status and review state according to the mapping outcome. Clicking on "save" closes the synset review page and brings back the hyponym list from Step 1.

Step 5: Lexicographer repeats Steps 2 - 4 for all existing synsets in the hyponym list.

Step 6: Lexicographer repeats Steps 1-5 with each existing 1st-generation hyponym as the query synset. The cluster is complete once all of its existing 1st and 2nd-generation hyponyms have been reviewed.

As the effectiveness of the Synset Mapper tool and its accompanying mapping protocol is still being evaluated via our preliminary mapping work, the tool is not yet publicly available. However, the video-centered design of Synset Mapper will likely make it very applicable to wordnet-development efforts for languages where a video-based lexical database is an efficient means of documenting individual units of meaning, such as for other signed or spoken languages without robust orthographies. For this reason, we hope to make our tool accessible for this purpose in the near future once it is fully developed.

3.2.1 Synset Status and Review State

The synset status indicates the status of a synset's mapping, with four options to choose from:

- **Unreviewed:** The default status.
- **Incomplete:** Mapping incomplete due to a gap in the sign data.
- **Approved:** Mapping complete; all appropriate Signs have been linked.
- **Deferred:** Mapping is non-trivial, reserve for future analysis.

The review state indicates if the mapping has been finalized by the lexicographers. To ensure consistency and limit individual subjectivity, we adopted a measure-twice-and-cut-once protocol, where each lexicographer's mappings ("Tentative" state) are verified by a second lexicographer ("Final" state). The default state is "not started".

3.2.2 Computer-Suggested Signs

To expedite the mapping task, Synset Mapper can function in a computer-assisted mode by providing "recommended signs" and "corresponding signs".

A Sign appears in "recommended signs" if any of its English equivalents satisfies a "string contains" regular expression match with at least one of the lemmas of the synset under review, or of its hypernym(s) and hyponym(s) (when they exist).

The "corresponding signs" list utilizes the manually-prepared Combined Sign correspondence tables. When a Sign is mapped to a synset, Synset Mapper checks if this Sign belongs to a Combined Sign containing other Signs. Any associated Signs are then added in real time to the "corresponding signs" list for the lexicographer to map (if suitable). This serves as an effective means of verifying our preliminary Combined Sign groupings.

4 Mapping Progress

The lexicographers are currently performing preliminary mapping work to evaluate the strategy described in Sections 2 and 3, paving the way for large-scale development. To date, we have processed 390 synsets (including both "Tentative" and "Final" states) in 14 randomly-selected clusters with scores in vicinity of the average from the cluster screening process (Table 2). As the optimal score threshold for mapping in practice is unknown due to a lack of data, the "in the vicinity of the average" criteria was arbitrarily selected. Variety in the scores of the selected clusters will allow evaluation of correlations between their score and actual completeness upon the conclusion of mapping as a test of our cluster screening procedure.

A total of 271 signs[8] have been mapped to synsets. On average, each cluster contains $\sim$184 synsets with $\sim$50% of its synsets processed (e.g., reviewed by the lexicographers) and $\sim$13% mapped to at least one Sign. These statistics, along with these reported in the remainder of this paper, consider all processed synsets (i.e., both the "Tentative" and "Final" states) unless otherwise noted.

As indicated by Table 3, the fact that the processed synsets are not overly dominated by those with "Incomplete" status is a testament to the success of our cluster screening and the coverage of the combined ASL lexical database. This is further supported by observing that 30% of the "Incomplete Synsets" have at least one mapped sign (i.e., they still need additional ASL forms not present in the sign data to achieve "Approved" status).

Of the synsets with mapped signs, 13 had 1 sign, 18 had 2 signs, and 62 had 3+ signs. The apparent propensity of these synsets to have a large number of mapped signs is likely due to the fact we

[8] 104 from SignStudy, 88 from ASL-LEX, and 79 from ASL Signbank.

Cluster	Size	Score	Fraction Processed	Fraction Mapped	# of Signs
baseball_equipment.n.01	18	0.47	0.67	0.05	2
clock_time.n.01	20	0.42	1.00	0.50	36
hair.n.01	44	0.40	1.00	0.34	10
sports_equipment.n.01	69	0.34	0.46	0.04	4
jewelry.n.01	37	0.31	0.35	0.16	20
head_of_state.n.01	15	0.29	0.40	0.20	7
starches.n.01	43	0.24	0.95	0.14	17
furniture.n.01	83	0.23	0.96	0.12	45
building.n.01	176	0.21	0.01	0.01	4
person.n.01	1616	0.19	0.02	0.01	48
woman.n.01	126	0.18	0.02	0.02	8
vegetable.n.01	79	0.15	0.36	0.07	17
edible_fruit.n.01	136	0.12	0.35	0.07	45
beverage.n.01	112	0.11	0.38	0.11	39

Table 2: The 14 randomly-selected clusters (with scores in vicinity of the average) for preliminary ASLNet V1.0 mapping work. Mapping is underway; "Fraction Processed", "Fraction Mapped", and "# of Signs" indicates the fraction of cluster synsets having been reviewed by lexicographers, the fraction of cluster synsets having at least one mapped Sign, and the number of Signs mapped to the cluster's synsets, respectively. In the ideal case (i.e., where our cluster screening process is indeed reliable), as "Fraction Processed" approaches 1.0 for a given cluster, its "Fraction Mapped" value will approach the cluster's "Score". Based on our progress so far, this seems to be the case. However, our mapping work is still too preliminary to draw a definitive conclusion on the predictive ability of our cluster screening process.

Synset Status	Tentative	Final	Overall
Incomplete	44	69	113
Approved	8	50	58
Deferred	188	31	219
Total	**240**	**150**	**390**

Table 3: Status and review states of processed synsets.

have yet to collapse Signs over Combined Signs, especially since the three sign databases used are known to have some overlap. However, this may also be explained by a high incidence of synonymous signs, which might be an interesting metric to compare against other languages. The actual cause will be revealed when the Signs for each synset are reviewed and condensed into Combined Signs as appropriate.

Comparing the number of synsets in each of the "Tentative" and "Final" review state suggests the presence a processing bottleneck introduced by the measure-twice-cut-once protocol. While this is a worthwhile trade-off for early mapping efforts due to the lexicographers' inexperience, it is not for large-scale work. Mapping quality will instead be maintained via a training regimen for future lexicographers along with the development of a mapping guide with instructions for common cases such as whether to incorporate signs of foreign origin.

Some of the "Deferred" synsets correspond to ASL lexical gaps. Yet it is difficult to disambiguate between gaps and certain signs (e.g., classifier constructions) that differ from basic lexicalized forms. Others are technical concepts (present due to the taxonomic depth of PWN) unfamiliar to our lexicographers. The latter will be addressed by querying relevant experts who are also native ASL signers. Altogether, the non-triviality of the "Deferred" synsets relegate their analysis to future work.

The question of ASL lexical gaps also spotlights a serious limitation of the "map" approach. Despite having $N_{\text{Signs}} << N_{\text{Synsets}}$, we elected to map Signs to synsets rather than vice versa as it is easier for the lexicographers to retrieve Signs corresponding to the definition of a given synset as opposed to searching for a synset matching a given Sign. While suitable for basic mapping work, this precludes identifying PWN gaps for concepts lexicalized in ASL. To find such signs, this deficiency must be addressed in upcoming work.

5 Next Steps

With the mapping infrastructure implemented and its evaluation underway, it is beneficial to identify next steps as we scale up mapping operations.

5.1 Supplementing Mapping with Merging

The challenges pertaining to the "Deferred" synsets and PWN lexical gaps described in Section 4 reveal the limitations of the "map" technique for crosslingual wordnet development. This conclusion is expected, and is similar to that of the AWN team, who realized that mapping PWN to African languages resulted in a translation of predominately European concepts rather than a true African resource (Bosch and Griesel, 2017). One part of the solution is to ramp up the "merge" phase of ASLNet development where a new wordnet is built solely for ASL (and eventually merged with PWN). This affords us the flexibility to include ASL-specific synsets as well as implement the ASLNet-specific structure proposed in (Lualdi et al., 2019). A new wordnet structure and understanding of the nature of ASL-only synsets may guide us in resolving many of the currently "Deferred" synsets.

For the ASL-specific synsets, we propose to start with two basic discovery techniques. First, we will begin by having our lexicographers select specific semantic domains for which they will then supply any ASL signs that come to mind. While some of these will overlap with existing PWN synsets, we anticipate that others will correspond to lexical gaps in English. Second, once the mapping work reaches a stage where a large percentage of the available sign data has been mapped to PWN synsets, the remaining unmapped signs will be reviewed, as chances are high that they represent lexical gaps in English. The signs identified by these techniques will then be incorporated into ASLNet either as a Collaborative Interlingual Index (Bond et al., 2016) synset if a suitable match exists, or as a new synset.

5.2 Free Association

A more involved technique to probe senses and relations unique to ASL is to perform free-association tests on native ASL speakers. The premise is that associated words may be semantically related and therefore inform "merge" ASLNet development.

Free-association has been well studied for the English language (Nelson et al., 2004) and extended to PWN via studies of evocation between synsets (Boyd-Graber et al., 2006). The ASL-LEX team is currently working to collect semantic free associations from native ASL users for all of the signs in ASL-LEX, which will be used to generate a semantic network of the ASL lexicon. Because ASLNet and its sign data will be cross-referenced with ASL-LEX, we will be able to compare the semantic structure of the lexicon as measured in these two different ways (e.g., like Steyvers and Tenenbaum (2005) did for English). Additionally, as has been done for other languages (e.g., (Sinopalnikova, 2004; Ma, 2013)), we will leverage the ASL-LEX semantic associations in building ASLNet (e.g., using the free associates as suggested items in a later version of the Synset Mapper tool, among other possibilities). Accordingly, the "map" ASLNet work will prioritize the linking of ASL-LEX signs in anticipation of ASL-LEX semantic association data.

This work has NLP benefits as well. Spoken-language wordnets are generally thought to model human mental lexicon organization to some extent, hence their utility for word sense disambiguation (Fellbaum, 2010; Navigli, 2009). It is an open question if this premise extends to ASL. By comparing the ASL free-association data against both the "map" and "merge" components of ASLNet, one can verify the suitability of the wordnet model for organizing the ASL lexicon. This has important implications for ASLNet design and its applicability to ASL NLP efforts. Along these lines, one of the major barriers to NLP efforts for sign languages is a lack of the datasets necessary to train models (Bragg et al., 2019). By offering a semantically-structured lexicon, ASLNet could serve as one of the resources for developing such models.

6 Conclusion

Overall, progress is being made with developing ASLNet V1.0, with a focus on mapping easy-to-translate lexical nouns. Our tripartite data model, cluster screening technique, Synset Mapper tool, and mapping protocol all have enabled successful linking of ASL signs to PWN synsets, and in fact can be easily extended to other sign languages. In particular, these tools so far have been helpful in solving the unique challenges of building a sign language wordnet, overcoming the fact that there is no conventional notation system for identifying and disambiguating signs. However, preliminary work has highlighted the need for the "merge" technique to incorporate aspects of ASL overlooked by our current mapping efforts such as ASL-only synsets. Moving forward, the "map" technique used so far will be supplemented by "merge" development work that include the utilization of ASL free-association data.

Acknowledgments

The authors would like to thank lexicographers Miriam Goldberg and Heidi Johnson for their work. Comments and feedback from the 2021 Global Wordnet Conference attendees as well as three anonymous reviewers are also greatly appreciated.

References

Francis Bond and Ryan Foster. 2013. Linking and Extending an Open Multilingual Wordnet. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1352–1362, Sofia, Bulgaria. Association for Computational Linguistics.

Francis Bond, Piek Vossen, John P. McCrae, and Christiane Fellbaum. 2016. CILI: The Collaborative Interlingual Index. In *Proceedings of the Eighth Global WordNet Conference*, pages 50–57.

Sonja E. Bosch and Marissa Griesel. 2017. Strategies for building wordnets for under-resourced languages: The case of African languages. *Literator (Potchefstroom. Online)*, 38:1 – 12.

Jordan Boyd-Graber, Christiane Fellbaum, Daniel Osherson, and Robert Schapire. 2006. Adding Dense, Weighted Connections to Wordnet. In *Proceedings of the Third Global WordNet Conference*, pages 29–35.

Danielle Bragg, Oscar Koller, Mary Bellard, Larwan Berke, Patrick Boudreault, Annelies Braffort, Naomi Caselli, Matt Huenerfauth, Hernisa Kacorri, Tessa Verhoef, Christian Vogler, and Meredith Ringel Morris. 2019. Sign Language Recognition, Generation, and Translation: An Interdisciplinary Perspective. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '19, pages 16–31, New York, NY, USA. Association for Computing Machinery.

Naomi K. Caselli, Zed Sevcikova Sehyr, Ariel M. Cohen-Goldberg, and Karen Emmorey. 2017. ASL-LEX: A lexical database of American Sign Language. *Behavior Research Methods*, 49(2):784–801.

Mark Davies. 2011–. Most frequent 100,000 word forms in English (based on data from the COCA corpus). Available online at https://www.wordfrequency.info/.

J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. ISSN: 1063-6919.

Christiane Fellbaum. 2010. WordNet. In Roberto Poli, Michael Healy, and Achilles Kameas, editors, *Theory and Applications of Ontology: Computer Applications*, pages 231–243. Springer Netherlands, Dordrecht.

Julie A. Hochgesang, Onno Crasborn, and Diane Lillo-Martin. 2020. ASL Signbank. New Haven, CT. Haskins Lab, Yale University.

Colin Lualdi, Jack Hudson, Christiane Fellbaum, and Noah Buchholz. 2019. Building ASLNet, a Wordnet for American Sign Language. In *Proceedings of the Tenth Global WordNet Conference*, pages 315–322.

Xiaojuan Ma. 2013. Evocation: analyzing and propagating a semantic link based on free word association. *Language Resources and Evaluation*, 47(3):819–837.

George A. Miller. 1995. WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41.

George A. Miller and Christiane Fellbaum. 1992. WordNet and the Organization of Lexical Memory. In *Intelligent Tutoring Systems for Foreign Language Learning*, NATO ASI Series, pages 89–102, Berlin, Heidelberg. Springer.

Roberto Navigli. 2009. Word sense disambiguation: A survey. *ACM Computing Surveys*, 41(2):10:1–10:69.

Douglas L. Nelson, Cathy L. McEvoy, and Thomas A. Schreiber. 2004. The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, 36(3):402–407.

Zed Sevcikova Sehyr, Naomi K. Caselli, Ariel M. Cohen-Goldberg, and Karen Emmorey. 2020. Under review.

Anna Sinopalnikova. 2004. Word Association Thesaurus As a Resource for Building WordNet. In *Proceedings of the Second Global WordNet Conference*, pages 199–205.

Mark Steyvers and Joshua B. Tenenbaum. 2005. The Large-Scale Structure of Semantic Networks: Statistical Analyses and a Model of Semantic Growth. *Cognitive Science*, 29(1).

Piek Vossen. 2004. EuroWordNet: A multilingual database of autonomous and language-specific wordnets connected via an inter-lingual-index. *International Journal of Lexicography*, 17(2):161–173. Publisher: Oxford Academic.

Elaine Wright. 2020. From Intractable to Tractable: Simplifications and Strategies for Implementing ASLNet. Princeton, NJ. Princeton University Undergraduate Independent Work.

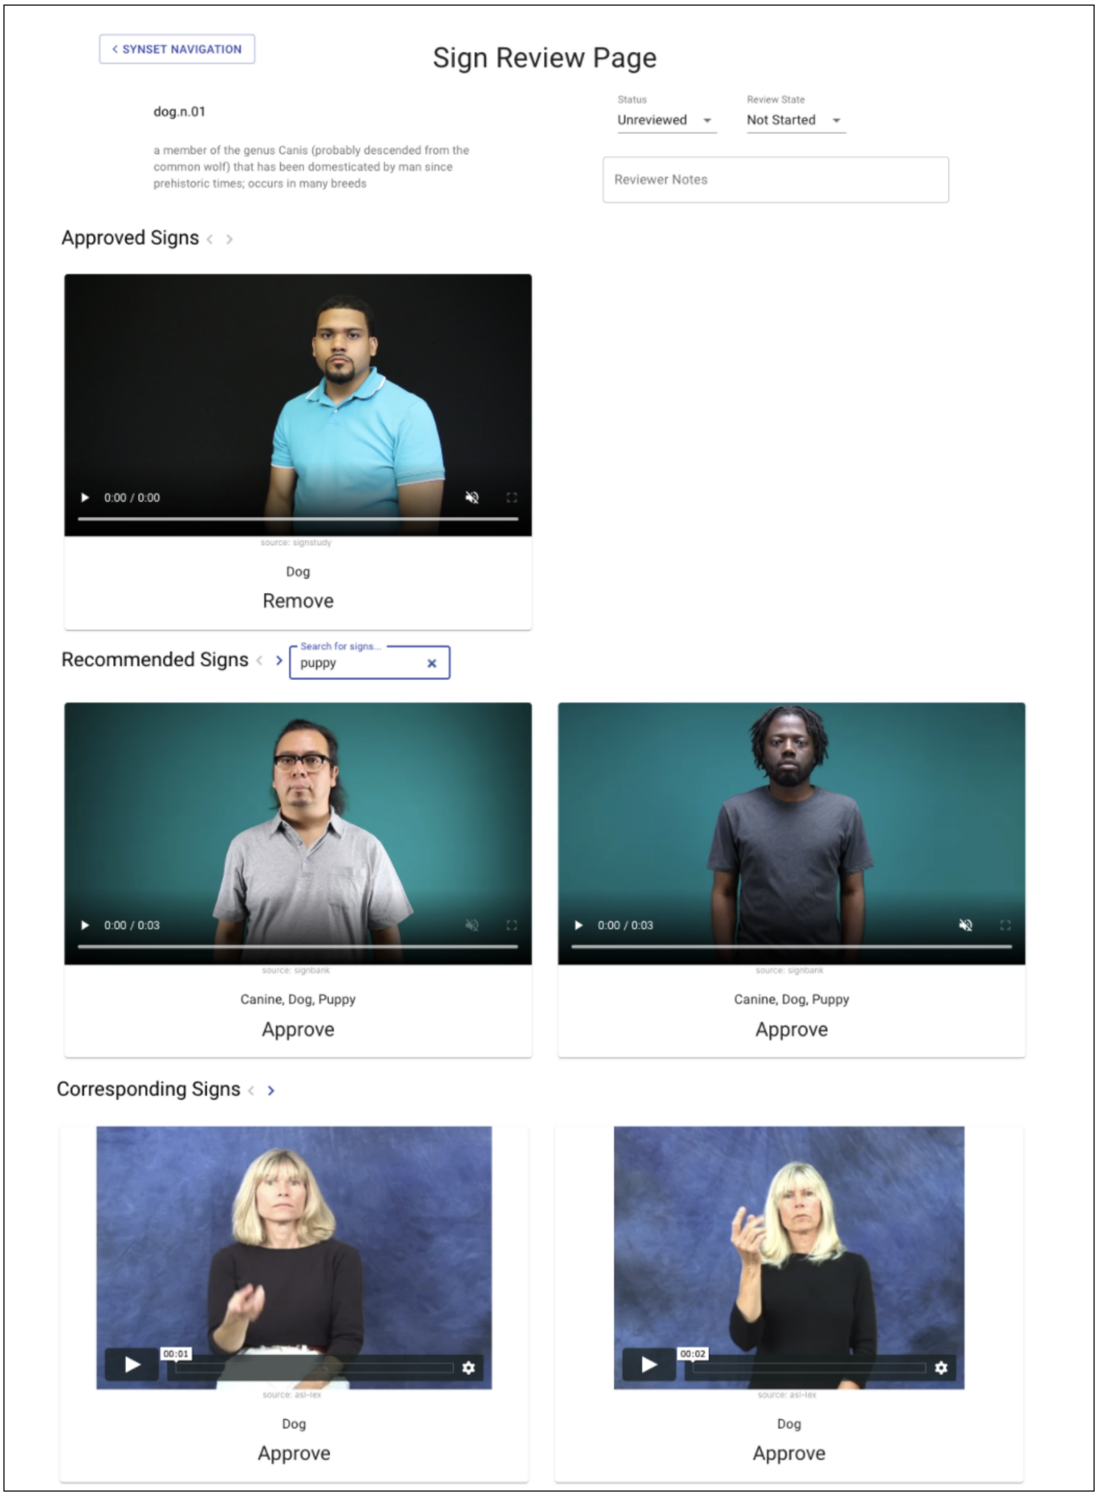

Figure 3: Synset Mapper tool Step 3: Synset review page displaying the synset's name, definition, status, review state, notes, mapped Signs, "recommended signs", "corresponding signs", and a Sign search.