Constructing a Lexicon of Relational Nouns

Edward Newell, Jackie C.K. Cheung
edward.newell@mail.mcgill.ca
jcheung@cs.mcgill.ca
McGill University

Abstract

Relational nouns refer to an entity by virtue of how it relates to another entity. Their identification in text is a prerequisite for the correct semantic interpretation of a sentence, and could be used to improve information extraction. Although various systems for extracting relations expressed using nouns have been developed, there are no dedicated lexical resources for relational nouns. We contribute a lexicon of 6,224 labeled nouns which includes 1,446 relational nouns. We describe the bootstrapped annotation of relational nouns, and develop a classifier that achieves 70.4% F1 when tested on held out nouns that are among the most common 2,500 word types in Gigaword. We make the lexicon and classifier available to the scientific community.

Keywords: relational noun, relation extraction, possessive, lexicon

1. Introduction

The extraction of relations is a fundamental aspect of natural language understanding, playing a central role in knowledge base construction, question answering, and recognizing textual entailment. Semantic relations can be expressed via a syntactically diverse set of constructions. Consider the following approximate paraphrases:

(1a) Jack **befriended** Jill 5 years ago.
(1b) Jack and Jill's **friendship** is now 5 years old.
(1c) Jack has been Jill's **friend** for 5 years.

Whereas **(1a)** establishes the relation by means of a verb, **(1b)** and **(1c)** do so by means of nouns. In example **(1c)**, *friend* is part of a semantically motivated class of nouns called *relational nouns*, which we focus on in this work.

Relational nouns refer to an entity by virtue of how it relates to something else (Barker, 2011). In the above example, *friend* establishes a relation between its referent, *Jack*, and the external entity, *Jill*, which could be depicted using a two-place predicate such as `friend(Jack, Jill)`. This differs from how a sortal (i.e. non-relational) noun like *person* is interpreted. The sentence *Jill is a person* could be depicted as a unary predicate, `person(Jill)`. In other words, there is an absolute set of people, but there is no absolute set of friends, without first specifying the person with whom they are friends.

Despite this semantic difference, relational nouns behave syntactically like other nouns. Semantic parsers such as Boxer that are trained on CCGBank do not currently distinguish between relational and non-relational nouns, leading to errors in sentences that contain them (Bos, 2008). Relation extraction systems such as RENOUN (Yahya et al., 2014) and RELNOUN (Pal and Mausam, 2016) rely on automatically extracted patterns, and also do not make a distinction between relational and non-relational nouns. Various information extraction systems could be improved by the identification of relational nouns. In this work, we create a high-quality lexicon of relational nouns using bootstrapped manual annotation.

Relational nouns pose a conceptually difficult annotation task, in part because the meaning of most nouns involves some relation at least indirectly. Part of our contribution is a series of decisions about what should be included or excluded from the class; decisions which are made with applications to relation extraction in mind. We describe these decisions and the approach to annotation.

Based on our annotation results, approximately one out of every 15 nouns, by type, is relational. This represents a substantial semantic class, but means that a significant amount of annotation effort will be expended on annotating the negative class. To more efficiently deploy annotation effort, we bootstrap annotation using a relational noun classifier.

We obtain a dataset of 6,224 labelled nouns, containing 1,446 relational nouns and 4,778 sortal nouns. We also develop a classifier that achieves 70.4% F1 when classifying held-out nouns from the top 2,500 most common nouns. We release the dataset[1] and classifier[2] to the scientific community so they can be used in relation extraction systems and other NLP applications.

2. Related work

There exists a line of work examining the properties of relational nouns using a formal theoretical framework (De Bruin and Scha, 1988; Partree, 2008; Laczkó et al., 2009; Barker, 2011). There, the focus is on developing a consistent theoretical treatment rather than lexical resources. Many of our decisions and heuristics used in annotation and classification are based on this literature.

Other work focuses on the development of resources and systems for nominal semantics, which is complementary to our work. This includes NomBank (Meyers et al., 2004), and other work on nominal semantic role labelling (Padó et al., 2008; Gerber and Chai, 2010). These resources and systems are concerned with the argument structure of nouns (relational nouns are often treated as taking arguments). But these resources do not focus on relational nouns in particular. NomBank, by its extension to NomLex (Macleod et al., 1998), includes a short list of 331 relational nouns. By focusing specifically on relational nouns,

[1] http://cgi.cs.mcgill.ca/~enewel3/publications/relational-nouns-lrec-2018

[2] https://github.com/enewe101/relational-nouns-lrec-2018

we provide a greater than fourfold increase in the number of labelled relational nouns (1,446). Within the context of Open Information Extraction, earlier work such as ReVerb focused on extracting relations from verbs (Fader et al., 2011). Mausam et al. (2012) examined the role of nouns and adjectives as bearers of predicates as well, showing that doing so increases coverage. Yahya et al. (2014) developed the RENOUN system, which focuses on extracting information about rarer attributes expressed using nouns.

The work most related to ours is the RELNOUN Open IE system, most recently augmented by Pal and Mausam (2016). This work extracts relations expressed using nouns, including relational nouns, using a combination of deterministic patterns and lexical resources, yielding 209 correct extractions from 2,000 newswire sentences.

Our current work complements automated extraction systems like RENOUN and RELNOUN by assembling a lexicon of relational nouns. Rather than only relying on automatically extracted patterns, we use manual annotation and bootstrapping to create a high-quality lexicon. This lexicon of relational nouns can be used as a semantic resource in relation extraction and other NLP tasks.

3. Operationalizing Relational Nouns

Our operationalization of relational nouns follows the theoretical treatment in the context of possessive constructs fairly closely. However, we deviate in certain cases (as will be noted) to prioritize relation extraction applications.

Applying the definition of relational nouns during annotation is quite difficult, in part because relational nouns carry no strong syntactic characteristics that distinguish them. It helps to decompose the definition of relational nouns into two criteria: (1) relational nouns must provide intrinsic lexical evidence that a relation is being expressed, and (2) relational nouns must refer to one of the members in the relationship expressed. This is equivalent to the prior definition in which a relational noun identifies an entity by virtue of how it relates to something else. But the decomposition makes it more obvious that there are two ways that a noun can fail to be relational.

Criterion 1 eliminates sortal nouns such as *car*. Consider, by way of example, the sentences *that is Jill's car* and *Jack is Jill's brother*. There we see that the referents of *car* and *brother* are both participating in relations. In contrast, if Jill utters the sentence *Jack is a brother* one would infer `brother(Jack, Jill)`, whereas uttering *that is a car* indicates no corresponding relation. Lacking specific context, we see that *car* does not provide intrinsic lexical evidence for a relation.

Criterion 2 eliminates nouns that indicate relationships yet which are not relational nouns, such as event and result nominals. Consider *disagreement* and *reconciliation*: although these both indicate a relation, notice that instead of referring to one of the participants of the relation, these nouns refer to the relation itself. So too do nouns like *friendship*. Although these nouns do indicate relations, they are not proper relational nouns and are correctly eliminated by criterion 2. (Note that agent and patient deverbal nouns, such as *employee* and *employer* adhere to criterion 2, and so are considered relational.)

In preparation for designing the annotation task, we collected many relational nouns given as examples in the literature, and organized them under the broad types of relations they expressed, shown in **Table 1**. In our experience, annotators perceive these classes as quite different, so it is useful to decompose the notion of relational nouns in terms of the subclasses. We found that sequentially introducing the subclasses simplified annotator training. We now review the subclasses:

Kinship. Kinship nouns like *brother*, describe family relations, and are the most common example in the literature.

Social non-kin. This includes informal roles, like *friend*, and formal or organizational roles like *mayor*, *CEO*, or *goalie*. Nouns that depict roles without providing lexical evidence for a relation, such as *butcher*, are excluded.

Operational. This includes non-social relations: purpose, cause/effect, function, representation, etc. Theories of possessive constructs dictate that relational nouns can occupy postnominal possessive constructions (Barker, 2011), but we include nouns like *cure* even though it is more natural to say *the cure **for** the disease*, than it is to say *the cure **of** the disease*.

Relative parts. This includes nouns designating a physical region based on a spatial or temporal relationship, such as *corner* or *intro*. These are typically reified by the relation itself: the corner of a desk exists by virtue of being the corner and cannot exist apart from the desk.

In creating these subclasses, we have specifically excluded two others containing nouns normally considered relational (Partee and Borschev, 2003; Cresswell, 1996; Laczkó et al., 2009; Barker, 2011; Lichtenberk et al., 2011; Partree, 2008; De Bruin and Scha, 1988):

Body parts. All body part nouns are traditionally considered to be relational (Laczkó et al., 2009). However, the relation between a body part and the whole body does not seem to be an essential part of the meaning of body part nouns. Supose one is practicing drawing ears, and is asked *what are you drawing?* In the response *that is an ear*, the word *ear* does not seem to indicate any sort of relation. Even in the gruesome case of a real disembodied ear, no relation is implied when one says *that is an ear*. The same cannot be said for other nouns we consider relational. With that said, in many cases a body part noun *can* have a relational meaning, when it specifically represents the relation that the body part has to the rest of the body, or a generalized analogical meaning, e.g. *head* as in *head of the department* or *leg* as in *leg of the journey*. We do consider these nouns relational (belonging either the operational or relative part subclasses), but do not admit all body part nouns wholesale.

Properties. Nouns such as *height*, or *colour* are also often treated as relational nouns, but we distinguish between nouns that identify the properties of individual entities from those establishing relationships *between* entities. We consider properties to warrant a separate annotation effort (in part undertaken in (De Bruin and Scha, 1988)).

Excluding body part nouns and properties, the operational

Subclass	Literature examples	Additional examples
Kinship	**Brother**^{6}, **sister**^{5}, **mother**^{4}, **child**^{4†}, **grandmother**^{3}, **husband**^{2}, **wife**, **spouse**, **father**, **daughter**, **aunt**, **uncle**, **cousin**, **family**, **relative**	**Stepdaughter**, **brother-in-law**, **kin**, **kindred**
Social non-kin	**Enemy**^{4}, **friend**^{3}, **pet**^{2}, **stranger**^{2}, **neighbour**^{2}, **mayor**, **governor**, **commander**, **co-author**, **employee**, **tutor**	**Captain**†, **guitarist**†, **investor**†, **linebacker**, **lawyer**†, **spokesperson**, **entourage**, **confidante**
Operational	**Name**^{2}, **birthday**^{2}, **picture**^{2}, **portrait**^{2†}, **tracks**†, **sake**, rumor, **description**, **reputation**	**Passenger**, **solution**, **cure**, **hole**
Relative part	**Edge**^{2}, **mantel**†, **side**, **corner**, **middle**	**Top**, **tip**†, **base**†, **leg**†, **eye**†, **body**†, **face**†,
Body part	**Leg**^{3†}, hand^{2}, **head**†, **eye**†, **body**†, **face**†, bone, blood, voice, ulcer, nose, sweat, hairdo, tears, finger	Epithelium, duodenum
Properties	Height^{2}, speed^{2}, distance, rating, length, readiness, color, weight, shape, temperature, gesture, habit, fear, posture, shadow	Happiness, badness, capacity, range, clarity, piety, extravagance

Table 1: relational noun examples drawn from the literature (Lichtenberk et al., 2011; Partree, 2008; Laczkó et al., 2009; Cresswell, 1996; De Bruin and Scha, 1988; Barker, 2011; Asudeh, 2005; Partee and Borschev, 2003), grouped by subclass, along with additional examples we provide. superscripts indicate the number of distinct authors that used the noun as an example. the nouns which we consider relational under our operationalization are shown in bold, with a dagger (†) indicating that the noun also has prominent non-relational meaning(s).

subclass catches any other non-social, non-spatio-temporal relational noun, making this classification exhaustive.

3.1. Annotation

We conducted annotation using 3 experts (natural language processing and linguistics researchers) and 13 non-expert annotators. We used 392 expert-annotated seed examples drawn from the literature and randomly sampled from Gigaword (Graff and Cieri., 2003) as a source of training examples and quality-control test questions. In an initial training phase, non-expert annotators learned about each relational noun subclass, and annotated examples from each with immediate feedback. Then, in a quiz phase, annotators completed 25 test (seed) examples before proceeding to the task. During the task, 2 of every 25 examples was randomly chosen to be a test example to track annotation quality. We only include data from annotators that sustained more than 70% correct test questions throughout the quiz and annotation.

Many nouns have both relational and non-relational meanings. To better serve as an input to NLP systems, we adopt a recall-oriented stance, and consider a noun to be relational if it has any relational meanings. However, during annotation, annotators were instructed to label a noun as *usually relational* if they judge relational sense(s) to be strictly more likely in general usage, and *occaisionally relational* if the noun has relational sense(s) that they judge not to be strictly more likely. Nouns lacking any a relational sense were to be labelled *almost never relational*.

Each noun was presented without context, and was annotated either by 1 expert or at least 3 non-experts (5 if initial non-expert annotations were not unanimous), except for a subset of 250 nouns which was annotated by all annotators (expert and non-expert) to test agreement. We reconcile multiple annotations by taking the *mode label* that received the most votes. Ties are resolved by applying the label *occasionally relational*. For nouns with expert and non-expert annotations, we only consider the expert annotations.

Consistent with our recall-oriented stance, we consider both the *occasionally relational* and *usually relational* as relational, but the distinction is preserved in the dataset so that analysts can use it as an indication of confidence that a given word in context is relational. Interested readers can view a complete reproduction of the annotation guidelines[3].

As mentioned, 250 nouns were annotated by all annotators to measure agreement. Agreement (Krippendorff's α, squared error metric) was 0.53 among expert annotators, 0.47 among non-expert participants, and 0.43 among all annotators. These modest agreement levels reflect the fact that the task is conceptually extremely challenging. They were the highest agreement levels we achieved after several iterations of task design and pilot studies. Reviewing disagreements, in many cases, annotators failed to recall a given noun's alternative relational meaning(s). In part this contributed to our decision to default to *occasionally relational* to resolve ties.

3.2. Bootstrapping

In choosing nouns to annotate, we sought to balance the desire to annotate common words, the desire to cover rare phenomena, and the desire limit the amount annotation effort expended on labelling negative examples.

To focus on examples likely to be positive, we bootstraped the sampling of nouns using a relational noun classifier trained on partial data, corresponding to the *feature-rich* model to be described in §4.2. We began with the expert-annotated seed set of 392 examples. During each round of bootstrapped annotation, we trained a model based on the existing annotations, and use the model to select new nouns for the next round of annotation, while oversampling those predicted by the classifier to be relational. We still included some nouns predicted by the classifier to be negative, to

[3]http://cgi.cs.mcgill.ca/~enewel3/publications/relational-nouns-lrec-2018

avoid excessive drift or bias. Sampling nouns for annotation was done in three rounds, with approximately 2,000 nouns sampled in each round. Nouns were sampled by taking, among nouns not yet annotated: 800 nouns scoring highest in the classifier's decision function, 200 nouns scoring lowest, 800 most common nouns in Gigaword, and 800 nouns uniformly randomly selected from Gigaword (but occurring at least 5 times). These sets overlap to some extent, resulting in each round contributing somewhat fewer than 2,000 nouns (1944 on average). The classifier's positive samples were enriched in relational nouns by factor of 8, being 49.8% relational.

4. Relational noun detection

As mentioned, the annotations provide three-class labels: *usually relational*, *occasionally relational*, and *almost never relatoinal*. But, motivated to create a more recall-oriented model, we collapse *usually relational* and *occasionally relational* into a single positive class.

Automatically detecting relational nouns is quite difficult, given (a) relational nouns behave grammatically the same way as sortal nouns, (b) the annotation of relational nouns is quite difficult even for people, and (c) even after merging the *usually relational* and *occasionally relational* classes, there is considerable class imbalance.

4.1. Baseline

Our baseline classifier draws on the observation of Barker (2011) that relational nouns characteristically arise under possessive constructions, specifically the prenominal (*Jill's brother*) and the postnominal (*the brother of Jill*) constructions. We collected counts for the number of times that nouns arose in each grammatical context in Gigaword (Graff and Cieri., 2003), as a fraction of total occurrences, and used these as features in shallow learners described below.

4.2. Feature-rich model

To create a more performant classifier, we added the following features:

Dependency tree features. We note the frequency with which a noun arises in particular dependency tree contexts, according to Stanford CoreNLP's dependency parse (Manning et al., 2014) of Gigaword. The dependency tree context was encoded by starting from the noun, and following a non-intersecting path along dependency tree relations for up to three hops, while noting the sequence (and direction) of the relations, and the part of speech (POS) tags of the nodes.

Sequence features. We record counts for the POS, lemma, and surface form, and bigrams thereof, of the surrounding 10 tokens along with their relative position.

Morphological features. We record the suffix of the noun, based on a list of common suffixes (see supplementary material[4]).

Derivational features. We record whether a noun is derivationally related to words of other POSs, based on WordNet's `derivationally_related_forms` (Fellbaum, 1998).

Semantic features. 300-dimensional Google word embeddings (Mikolov et al., 2013).

Hand-crafted features. 31 theoretically and empirically motivated features based on surrounding lemmas, POS-tags, named entities, and dependency-tree relations in Gigaword (Graff and Cieri., 2003) (see supplementary material[4]).

4.3. Feature transformations

We tried alternative encodings of dependency tree and sequence features: raw counts, frequencies (i.e. counts normalized by the number of occurrences of the noun), log-frequencies, and binarized features based on a threshold of the pth percentile for feature's frequency, $p \in \{25, 50, 75\}$. We pruned these features to the top k with highest mutual information, for $k \times 10^{-3} \in \{10, 20, 40, 80\}$. **Table 2** shows the performance of classifiers using each representation, showing that the frequency representation performed best.

4.4. Learning algorithms

Using the scikit-learn package (Pedregosa et al., 2011), we optimized over several learners for both the baseline and feature rich models: support vector machine (SVM), logistic regression (LR), naive Bayes (NB), and random forest (RF). For learners with linear decision surfaces (SVM, LR, and NB), we shifted the surface to optimize F1 on the training set due to class imbalance. As shown in **Table 3** SVM performed best.

4.5. Feature Ablation

The various features mentioned above were ablated to determine the best-performing combination of features. As shown in **Table 4**, The best performance is achieved by including all except the morphological features.

4.6. Classifier performance

Because we are interested in using the classifier both for bootstrapping and for classification of unannotated nouns, we report two evaluation metrics: the average precision, and the F1 score. Average precision considers the ranking of nouns by order of decreasing likelihood of being relational according to the classifier's decision function, rather than the actual classification, which indicates the usefulness of the classifier in bootstrapping. F1 provides a measure of the binary prediction recall and precision, which is more useful in NLP applications when classifying un-annotated nouns.

Optimizing F1 over the space of learners, hyperparameters, features, and feature representations, the best performance on the dev set was achieved by excluding morphological features, using frequencies, $k = 80,000$, while using an SVM with radial basis function, and $C = 100$ and $\gamma = 0.001$[5].

The best-performing classifier was one built using al features except morphological (suffix) features, and using

[4]http://cgi.cs.mcgill.ca/~enewel3/publications/relational-nouns-lrec-2018

[5]see Pedregosa et al. (2011) for the meaning of γ and C.

Feature representation	AP (%)	F1 (%)
raw counts	56.3	58.6
frequency	61.1	63.2
log-frequency	53.2	56.0
threshold-25%	56.2	59.3
threshold-50%	58.4	60.3
threshold-75%	54.8	58.6

Table 2: Performance of relational noun classifiers using various feature representations. Support vector machine is the learner for each model shown. AP = Average Precision.

classifier	AP (%)	F1 (%)
Support Vector Machine	79.8	76.8
Random Forest	75.1	69.9
Logistic Regression	79.6	74.8
Naive Bayes	53.8	67.6

Table 3: Performance of relational noun classifiers using a variety of learners. All models shown use all features except the morphological (suffix-based) features and use the frequency encoding. AP = Average Precision.

SVM as the learner. This classifier achieved average precision $AP = 76.1\%$ and $F1 = 70.4\%$ on the 2500 most common nouns in Gigaword (Graff and Cieri., 2003) and $AP = 49.5\%$ and $F1 = 46.7\%$ on randomly sampled nouns occurring at least 5 times in Gigaword.

While these performances are modest, for the purposes of bootstrapping, the set of nouns predicted to be relational turned out to be enriched by eight fold over the baseline rate of occurrence, from approximately 1 in 15 words to 1 in 2. This was crucial to increase the number of relational nouns discovered during the study.

As previously mentioned, the subclasses of relational nouns we identified were perceived very differently by annotators, and seemed to differ greatly in difficulty. Based on post-hoc manual classification of 300 nouns from the held-out test set, according to relational noun subclass, we can see that classifier performance also differs greatly by subclass (**Table 5**). Whereas classifier recall is high for kinship and social non-kin nouns, it is very low for operational and relative part nouns.

Ablated feature	AP (%)	F1 (%)
baseline	80.2	75.2
dependency	80.3	76.1
hand-picked	79.9	75.2
lemma	80.2	75.2
surface	80.3	74.9
POS	80.1	75.2
derivational	80.3	75.5
google-vectors	80.2	75.2
suffix	79.8	76.8

Table 4: Performance of relational noun classifiers using various feature sets. Each classifier is built using all features except the ablated feature; lower scores indicate greater importance of the feature. All models used support vector machine with the frequency representation of count-based features. AP = Average Precision.

Subclass	Recall (%)	Fraction of unique token types (%)
Kinship	80.0	5.1
Social non-kin	85.5	46.9
Operational	37.0	33.0
Relative part	34.3	15.0

Table 5: Classifier recall for various subclasses of relational noun, and the size of the subclasses in terms of the fraction of unique tokens.

The kinship nouns constitute the smallest subclass by number of unique token types, so it is at first surprising that recall was very high. However, this stands to reason when one considers that they tend to have fewer alternate non-relational meanings, and are the more common class by number of occurrences (rather than by token type). These facts mean that their corpus-derived features are based on a larger number of occurrences in Gigaword and may have less diversity in their contexts (and hence features).

On the other hand, the operational and relative part nouns had very low recall, despite acounting for an intermediate number of examples (by number of unique token types). This may be due to the fact that these words tend to labelled *occasionally relational* more frequently than the other relational noun subclasses. This results both from polysemy and annotator disagreement, both of which lead to greater difficulty for the classifier—in the first case due to more diverse contexts in Gigaword (and presumably poor clustering in feature space), and in the second case due to noisier training data.

Overall, recall was highest for the social non-kin nouns, which may be due to a combination of having fewer non-relational meanings and being a much larger class compared to the other three (by number of unique token types).

5. Conclusions

In this paper we presented the first effort dedicated to building a high-quality lexicon of relational nouns, and contribute the largest lexicon of relational nouns to date.

The lexicon is based on an operationalization of relational nouns designed with applications to relation extraction in mind. We provide the dataset of 6,224 nouns, including 1,446 relational nouns, along with a relational noun classifier to the research community for inclusion in relation extraction systems.

Delineating the class of relational nouns, whether automatically or manually, is a very challenging task. We hope this first effort will provide a starting point for future work to develop additional lexical resources for relational nouns and relation extraction.

6. Acknowledgements

We wish to acknowledge many helpful conversations with Professor Brendan Gillion of the Linguistics Department at McGill University. Funding for this research was provided by the Natural Sciences and Engineering Research Council of Canada.

7. Bibliographical References

Asudeh, A. (2005). Relational nouns, pronouns, and resumption. *Linguistics and Philosophy*, 28(4):375–446.

Barker, C. (2011). Possessives and relational nouns. *Semantics: An international handbook of natural language meaning*, 2:1109–1130.

Bos, J. (2008). Wide-coverage semantic analysis with Boxer. In *Proceedings of the 2008 Conference on Semantics in Text Processing*, pages 277–286. Association for Computational Linguistics.

Cresswell, M. J., (1996). *Relational Nouns*, pages 21–36. Springer Netherlands, Dordrecht.

De Bruin, J. and Scha, R. (1988). The interpretation of relational nouns. In *Proceedings of the 26th annual meeting on Association for Computational Linguistics*, pages 25–32. Association for Computational Linguistics.

Fader, A., Soderland, S., and Etzioni, O. (2011). Identifying relations for open information extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545. Association for Computational Linguistics.

Fellbaum, C. (1998). *WordNet*. Wiley Online Library.

Gerber, M. and Chai, J. Y. (2010). Beyond NomBank: A study of implicit arguments for nominal predicates. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1583–1592. Association for Computational Linguistics.

Graff, D. and Cieri., C. (2003). English Gigaword LDC2003T05.

Laczkó, T., Butt, M., and King, T. H. (2009). Relational nouns and argument structure: Evidence from hungarian.

Lichtenberk, F., Vaid, J., and Chen, H.-C. (2011). On the interpretation of alienable vs. inalienable possession: A psycholinguistic investigation. *Cognitive Linguistics*, 22(4):659–689.

Macleod, C., Grishman, R., Meyers, A., Barrett, L., and Reeves, R. (1998). Nomlex: A lexicon of nominalizations. In *Proceedings of EURALEX*, volume 98, pages 187–193.

Manning, C. D., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S. J., and McClosky, D. (2014). The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.

Mausam, Schmitz, M., Bart, R., Soderland, S., and Etzioni, O. (2012). Open language learning for information extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 523–534. Association for Computational Linguistics.

Meyers, A., Reeves, R., Macleod, C., Szekely, R., Zielinska, V., Young, B., and Grishman, R. (2004). The nombank project: An interim report. In *HLT-NAACL 2004 workshop: Frontiers in corpus annotation*, volume 24, page 31.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Padó, S., Pennacchiotti, M., and Sporleder, C. (2008). Semantic role assignment for event nominalisations by leveraging verbal data. In *Proceedings of the 22nd International Conference on Computational Linguistics-Volume 1*, pages 665–672. Association for Computational Linguistics.

Pal, H. and Mausam. (2016). Demonyms and compound relational nouns in nominal open ie. *Proceedings of AKBC*, pages 35–39.

Partee, B. H. and Borschev, V. (2003). Genitives, relational nouns, and argument-modifier ambiguity. *Modifying adjuncts*, 4:67–112.

Partree, B., (2008). *Genitives: A Case Study*, pages 182–189. Blackwell Publishing Ltd.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Yahya, M., Whang, S., Gupta, R., and Halevy, A. Y. (2014). Renoun: Fact extraction for nominal attributes. In *EMNLP*, pages 325–335.