

Possessors Change Over Time: A Case Study with Artworks

Dhivya Chinnappa and Eduardo Blanco

Human Intelligence and Language Technologies Lab

University of North Texas

Denton, TX, 76203

dhivyainfantchinnappa@my.unt.edu, eduardo.blanco@unt.edu

Abstract

This paper presents a corpus and experimental results to extract possession relations over time. We work with Wikipedia articles about artworks, and extract possession relations along with temporal information indicating when these relations are true. The annotation scheme yields many possessors over time for a given artwork, and experimental results show that an LSTM ensemble can automate the task.

1 Introduction

All languages have a way to express possessive relationships (Aikhenvald and Dixon, 2012). Possession is an asymmetric semantic relation between two entities, where one entity (the possessee) *belongs* to the other entity (the possessor) (Stassen, 2009). When it comes to defining possession, *belongs* includes a wide range of relationships, including kinship (e.g., $[my]_{\text{possessor}}$ *oldest* $[son]_{\text{possessee}}$), part-whole (e.g., *the* $[car]_{\text{possessor}}$'s $[dashboard]_{\text{possessee}}$), possession of something intangible (e.g., $[John]_{\text{possessor}}$ *got* *the* $[flu]_{\text{possessee}}$ *last year*), proximity (e.g., *The* $[shelf]_{\text{possessor}}$ *has* *a* $[glass\ sculpture]_{\text{possessee}}$), physical and temporary possession (e.g., $[I]_{\text{possessor}}$ *have* *John's* $[book]_{\text{possessee}}$), and ownership (e.g., $[John]_{\text{possessor}}$ *bought* *a* $[house]_{\text{possessee}}$ *last year*).

Possession relations can be divided into alienable (also referred to as acquired, transferable, non-intimate, etc.) and inalienable (also referred to as inherent, inseparable, intimate, etc.). Possesseees that can be separated from their possessors are alienable, and possesseees that cannot normally be separated from their possessors are inalienable (Heine, 1997). For example, $[John]_x$'s $[condo]_y$ is alienable, and $[John]_x$'s $[arm]_y$ is inalienable (some previous works would call the latter a part-whole relation instead). Tham (2004)

In 1530 the painting was inherited by Margaret's niece Mary of Hungary, who [...]. It is clearly described in an inventory taken after her death in 1558, when it was inherited by Philip II of Spain.
--

<i>Mary of Hungary</i> : 1530–1558 <i>Philip II of Spain</i> : after 1558
--

Figure 1: Excerpt from the Wikipedia article about *Arnolfini Portrait* and its possessors over time.

defines control possession as a relation in which the possessor has temporary control of the possessee, but does not necessarily alienably possess it (e.g., $[John]_x$ *borrowed* *the* $[car]_y$ *for the weekend*). Following the aforementioned works, possession goes beyond ownership of property.

Virtually all possesseees change possessors over time, especially if possession relationships are understood in a broad sense as outlined above. Consider the excerpt of the Wikipedia article about the *Arnolfini Portrait* in Figure 1. From this excerpt, we know that the painting had at least two possessors (*Mary of Hungary* and *Philip II of Spain*), and that they were the possessors from 1530 to 1558 and after 1558 respectively.

In this paper, we track possessors of selected possesseees over time. Unlike most previous works (Section 3), we (a) start with a document relevant to the possessee of interest, (b) select plausible possessors and years without syntactic restrictions and including inter-sentential pairs, and then (c) determine whether the plausible possessors are actual possessors with respect to the years. The main contributions of this paper are: (a) 88 Wikipedia articles about artworks annotated with their possessors over time;¹ (b) a detailed corpus analysis (e.g., unique possessors, years and possessor-year pairs); and (c) experimental results showing that an LSTM ensemble outperforms SVM.

¹Available at dhivyachinnappa.com

2 Previous Work

We briefly summarize work on possession relationships from a theoretical perspective, and then move to work in computational linguistics.

2.1 Possession relations

The very definition of possession is not set in stone. Aikhenvald (2013) distinguishes three core meanings for possessive noun phrases that occur across languages: ownership (of property), whole-part (often referred to as part-whole), and kinship. Following a cross-linguistic perspective, she discusses possessions and time (present and former possession relationships, e.g., *my tooth* vs. *my former axe*), temporary and permanent possession (e.g., *borrow* vs. *acquire*) and others. Heine (1997) classifies possession relationships depending on the possessor and possessee. First, he makes a distinction between human (e.g., $[I]_x$ *have a [house]_y*) and non-human possessors (e.g. $[This\ house]_x$ *has [two\ bedrooms]_y*). Second, he differentiates three kinds of possession depending on the possessee: concrete possession (e.g., $[I]_x$ *have [two\ cats]_y*), social possession (e.g., $[I]_x$ *have [two\ sisters]_y*), and abstract possession (e.g., $[I]_x$ *have [an\ idea]_y*). Miller and Johnson-Laird (1976) differentiate between three kinds of possession: inherent, accidental, and physical; and provide the following example: *He owns an umbrella (inherent), but she's borrowed it (accidental), though she doesn't have it with her (physical)*.

Possession relations have also been defined in terms of parameters. For example, Stassen (2009) considers two parameters (permanent contact and control) and Heine (1997) defines five parameters (human possessor, concrete possessee, spatial proximity, temporal permanence, and control).

While we do not closely follow any of these previous works, we borrow from them the broad definition of possession relations, and the motivation to work with possessions over time.

2.2 Computational Linguistics

Within computational linguistics, possession relations have been mostly studied as one of the many relations encoded in a given syntactic construction. For example, Tratz and Hovy (2013) extract semantic relations within English possessives. They propose a set of 18 relations, e.g. *temporal* (e.g., $[today]_x$'s $[rates]_y$), *extent* (e.g., $[6\ hours]_y$ ' $[drive]_x$). Their *controller / owner*

/ user relation (one relation with three aliases) is the closest relation to the possession relations we target in this paper. Extracting semantic relations between noun compounds (Nakov and Hearst, 2013; Tratz and Hovy, 2010) usually includes extracting possession relations, e.g., $[family]_x$ $[estate]_y$. These previous works extract all semantic relations—including possessions—between arguments that follow a syntactic construction.

In our previous work (Chinnappa and Blanco, 2018), we identify possession relations between a deterministically chosen person (possessor) and a concrete object (possessee) within a sentence. If a possession relation exists, we also identify the possession type (alienable or control). Finally, we temporally anchor the possession relation with respect to the verb of which the possessor is the subject. In this paper, we take a complementary approach. We start with text relevant to the possessee of interest—specifically, its Wikipedia article—and then extract its possessors without any restrictions beyond considering as possessors only named entities. Furthermore, we specify in which years the possessions were true.

To the best of our knowledge, the work by Banea et al. (2016) is the only one on extracting possession relations without imposing syntactic constraints. Banea and Mihalcea (2018) build a corpus working with personal blogs, and present results on automatic extraction of possession using a naive bayes approach. They consider as possessors the author of a blog, and as possesseees concrete nouns in blog posts. Regarding time, they annotate possessions at the time of the utterance (when the blog posts were published). Unlike them, we work with one possessee per Wikipedia article (i.e., the artwork the article is about), and then find possessors in the article. Additionally, we extract when a possessor-possessee relation is true with respect to the years in the article, and present results using SVM and end-to-end neural networks.

3 Annotating Possessions Over Time

In this section, we detail the methodology to create a corpus of possession relations over time. We first discuss the selection of source documents and possesseees of interest. Then, we detail what is considered as a potential possessor, and how these possessors are paired with years. Finally, we describe

		#	%
x, potential possessor	PERSON	1,152	48.0
	ORG	986	30.5
	GPE, LOC	692	21.5
	All	3,230	100.0
y, year	All	940	100.0
(x, y) pairs	x is PERSON	6,304	48.8
	x is ORG	3,840	29.7
	x is GPE, LOC	2,769	21.4
	All	12,913	100.0

Table 1: Counts of potential possessors (x) and years (y), and (x, y) pairs selected for annotation.

the annotation process (is the potential possessor an actual possessor with respect to the years?) and analyze the resulting corpus.

3.1 Selecting Source Documents

Our goal is to target possessors of a given possessee over time. A natural choice is to work with documents about specific objects, as they are likely to describe the history and key events involving the objects. We decided to work with Wikipedia articles about important artworks. The methodology presented here, however, is not limited to artworks, and we believe it is applicable to any article about an object of interest.

We selected 100 artworks using online content, including Google queries for *famous artwork* and *famous paintings*, and online lists.² Then, we downloaded the full content of the corresponding Wikipedia articles. Some of the selected artworks are *The Third of May 1808*, *Philosopher in Meditation*, and *Saturn Devouring His Son*. The final corpus has 88 articles because we discarded articles if we could not select at least three (potential possessor, year) pairs (see below).

3.2 Selecting Potential Possessors and Years

Once possesseees and their Wikipedia articles were selected, we identified potential possessors and years following the five steps below for each section in each Wikipedia article:

1. Run the named entity recognizers in spaCy³ and Stanford CoreNLP (Manning et al., 2014).

²<http://en.most-famous-paintings.com>,
<http://remliel.com/2016/07/08/100-greatest-paintings-of-all-time>

³<https://spacy.io/>

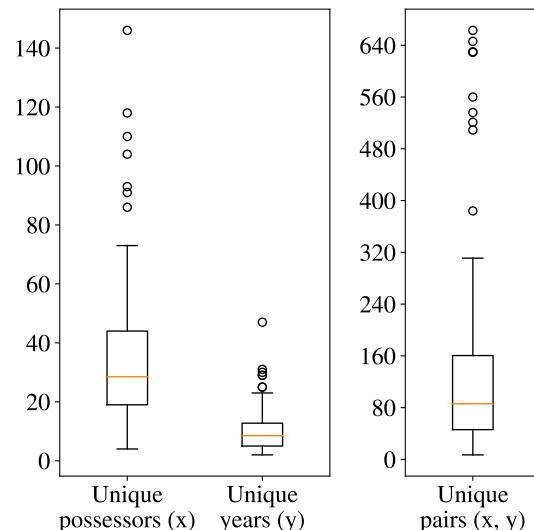


Figure 2: Distribution of unique potential possessors, unique years and unique pairs per article. Each box-plot displays the minimum, first quartile, median, third quartile and maximum.

2. Select as potential possessors all instances of the following named entities: PERSON, ORGANIZATION, LOCATION, and GPE.
3. Select as years all sequences of four digits inside a DATE named entity.
4. Remove all duplicate potential possessors and years from steps 2 and 3.
5. Generate all pairs of potential possessors and years.

Table 1 presents basic counts and percentages of the potential possessors and years after removing duplicates (Step 4), and the pairs generated (Step 5) for all documents. There are 3,230 potential possessors and 940 years, and 12,913 (potential possessor, year) pairs. The most common named entity of potential possessors is PERSON (48%), followed by ORG (30.5%) and GPE / LOC (21.5%). The percentage of (potential possessor, year) pairs depending on the named entity of the potential possessor almost follows an identical distribution (48.8%, 29.7% and 21.4%).

Figure 2 shows the distributions of unique potential possessors, years and (potential possessor, year) pairs generated per article (or equivalently, per possessee). While the distributions are far from uniform, the boxplots show that most articles have a substantial number of potential possessors, years and pairs. The minimum number of potential possessors is 4, of years 2, and of pairs 7. But

	Cohen's κ
Before	0.69
During	0.59
After	0.77
All	0.70

Table 2: Inter-annotator agreement (Cohen's kappa). Values over 0.6 are considered *substantial*, over 0.8 are considered *perfect* (Artstein and Poesio, 2008).

over 75% of articles have at least 19 unique potential possessors, 5 years and 46 pairs; and over 50% of articles have at least 28 unique potential possessors, 8 years and 86 pairs. In other words, our corpus takes into account many potential possessors and years for the vast majority of articles.

3.3 Validating Possessors and Years

After (potential possessor, year) pairs were generated, they were validated manually. To do so, we asked the following questions to annotators:

- Did a possession relation exist between the potential possessor and the possessee at any point of time *before year*?
- Did a possession relation exist between the potential possessor and the possessee at any point of time *during year*?
- Did a possession relation exist between the potential possessor and the possessee at any point of time *after year*?

In all questions, *possessee* refers to the artwork the Wikipedia article is about. Annotators had to choose from two answers: *yes* or *no*, where *no* indicates all cases in which there is not enough information to determine that a possession relation exists at any point of time before / during / after year. In other words, *no* does not mean that the potential possessor did not possess the possessee, and it may mean that there is no information about whether a possession relations exists.

The annotation interface showed the title of the article and the section to which the potential possessor and year belong to (section title + text). Annotators were instructed to first read the section and then answer all questions. Thus, annotators reveal possession information involving possessors and years that are potentially far away (different clauses, sentences, etc.). Recall that all potential possessors and years within a section are paired, thus we allow to cross sentence boundaries.

Annotation Quality. Annotations were done in-house by two graduate students. Both of them

annotated 25% of the articles individually. Table 2 shows inter-annotator agreements (Cohen's kappa) for each question. Overall, inter-annotator agreement is 0.70 (values between 0.60 and 0.80 are considered substantial (Artstein and Poesio, 2008)). Agreements are higher for *Before* and *After* than *During* (0.69 and 0.77 vs. 0.59). The remaining articles were annotated once.

Annotation Examples. Figure 3 shows the annotations for one paragraph of the Wikipedia article about *Girl with a Pearl Earring* (more specifically, from the section titled *Ownership and display*). The figure shows the annotations on top of a screenshot of the article for clarity purposes, but the annotation interface only showed one section at a time along with all the generated pairs (Section 3.3, equivalent to pre-drawing edges).

Five potential possessors and two years were selected, thus ten (potential possessor, year) pairs were generated. The annotations reveal the intuitive possession information contained within the paragraph. First, *Victor de Stuers* was an advisor to *Arnoldus Andries des Tombe*, so there is no evidence that he was a possessor at any point of time (missing label edges). Second, *Vermeer* is the artist who made *Girl with a Pearl Earring*, so there are possession relations before 1881 and 1902. Third, *Arnoldus* purchased the piece in *The Hague* in 1881, and in 1902 it was donated to *Mauritshuis*. So *Arnoldus* was a possessor in 1881 and after 1881 (until 1902), *The Hague* in 1881 (recall that non-humans can be possessors, spatial proximity is also considered possession, Section 2.1), and *Mauritshuis* during and after 1902. We discuss the limitation of the annotation approach in Section 3.5.

3.4 Annotation Analysis.

Counts of *yes* labels for the three questions (before, during and after) are rather low (17%, 9% and 19%, Figure 4). This is not surprising, as any PERSON, ORG, LOC and GPE named entity is considered as potential possessors. We note, however, that we annotated a possession relation (*yes* label) in 35% of (potential possessor, year) pairs generated (either before, during or after year).

Figure 5 depicts the distribution of labels per article for (potential possessor, year) pairs generated from the same and different sentences. It is worth noting a couple of interesting patterns. First, the annotations contain many more possessions because we pair potential possessors and years that

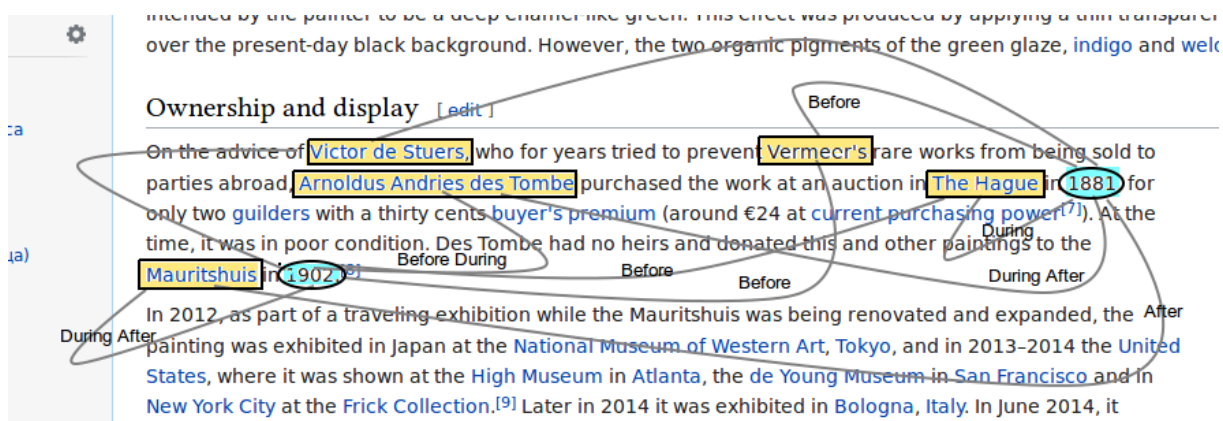


Figure 3: Excerpt of the Wikipedia article of *Girl with a Pearl Earring* and annotations. Edges indicate (potential possessor, year) pairs generated. Edge labels indicate *yes* label for *before*, *during* and *after*; missing edge labels indicates all annotations are *no* (i.e., the potential possessor was invalid).




Figure 4: Label percentages for each question (left). While most labels are *no*, one of the three questions (before, during or after) was answered with *yes* in 35% of (potential possessor, year) pairs (right).

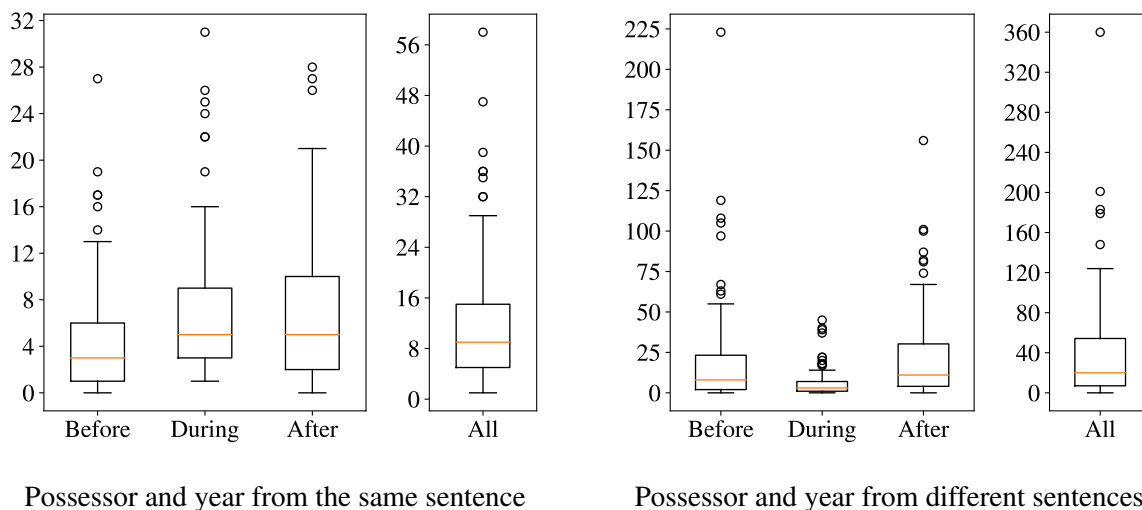


Figure 5: Distribution of *yes* label per article. We provide distributions for each temporal anchor (at some point of time before, during or after *year*) and for all anchors, and distinguish between possessors and years belonging to the same sentence (left) or different sentences (right). Each boxplot shows the minimum, first quartile, median, third quartile, and maximum.

belong to different sentences (note the different scales in the y-axis). Second, for pairs generated from different sentences, *yes* label for *during* questions is much less likely than for the labels *before* and *after*.

Finally, Figure 6 shows the distribution of *yes* label for all generated pairs. At a minimum, each possessee has at least two possession relations. There are a few outliers articles in which annotators identified over 150 possessors in time. The

<p><i>From the possessor and year:</i> the concatenation of tokens; binary flags for each token; the syntactic head (token, lemma and part-of-speech tag); and the named entity type.</p>
<p><i>From the sentences to which the possessor and year belong to:</i> for (a) a window of 4 tokens to the left and right, (b) all the verbs to the left and right, (c) all the verbs that are ancestors or children in the dependency tree, and (d) all the left and right siblings in the dependency tree, the tokens, lemmas and part-of-speech tags.</p>
<p><i>Other and Wikipedia article:</i> whether the possessor and year belong to the same sentence, whether the possessor appears before or after the year, the Wikipedia article title (concatenation of tokens and binary flags for each token), and the section title (concatenation of tokens and binary flags for each token).</p>

Table 3: Feature set used with Support Vector Machines.

first quartile is 14, the median 30, and the third quartile 65 (All, right most boxplot). Thus the procedure presented here reveals a substantial amount of possession relations along with temporal information anchored in the form of years.




Figure 6: Distribution of *yes* label per article for all possessors and years.

3.5 Limitations

While the proposed procedure successfully identifies possession relations over time, we acknowledge limitations in both the possessors and temporal information considered.

First, we only consider named entities as potential possessors, so it is possible we miss some possessors (e.g., pronouns, *the artist*, *his son*). Because of the source documents we work with (Wikipedia articles about artwork) and the fact that we pair all potential possessors and years within a section, this is not a big issue: most Wikipedia sections do not have mentions that cannot be resolved to a named entity within the same section. We note, however, that coreference resolution (Pradhan et al., 2011) would alleviate this problem.

Second, we only consider four digits within a DATE named entity as temporal information. This means that temporal information encoded in relative dates (e.g., four years later) or historical events (e.g. after World War II) is disregarded. Additionally, we cannot distinguish between several possessors within a year, finer-grained times would be required to do so. To address these issues, temporal parsers (Lee et al., 2014; Strötgen and Gertz, 2015) and anchoring events in time (Reimers et al., 2016) are required.

4 Experiments and Results

We experiment with traditional Support Vector Machines and neural networks. We divided the articles (and the corresponding (potential possessor, year) pairs) into train (80%) and test (20%), and report results obtained with the test split. Note that splitting pairs randomly would be unsound, as possession relations for the same possessee would be in the train and test splits. We build three classifiers with both SVMs and neural networks (one per question: before, during and after), and all of them predict two labels: *yes* or *no*.

4.1 Support Vector Machines

We trained the three classifiers using the SVM implementation in scikit-learn (Pedregosa et al., 2011), and tuned hyper-parameters C and γ using 10-fold cross-validation with the train split. We used features extracted from the possessor, the year, and the sentences they belong to. Additionally, we also included the Wikipedia article title and the section title from which the possessor and year were selected. The full feature set is described in Table 3 and we do not elaborate further. Our motivation to try SVMs is to establish a strong supervised baseline, and to compare with neural networks that take as input only plain text.

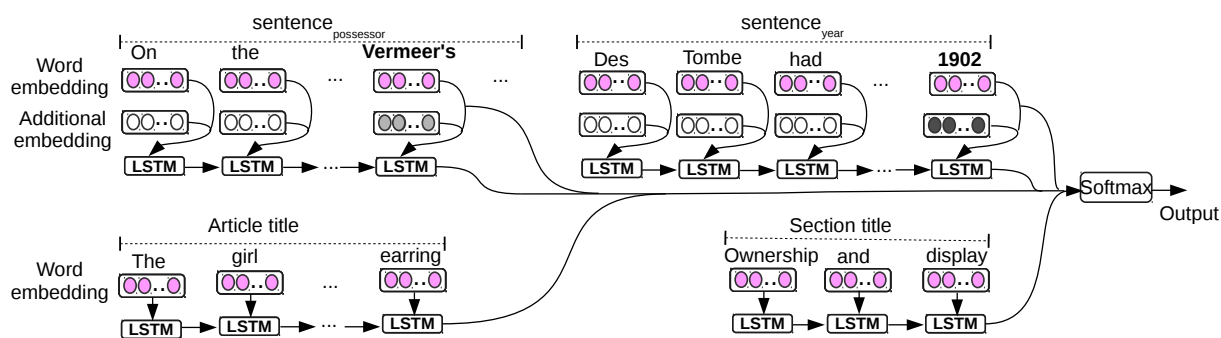


Figure 7: Neural network architecture to predict whether a *potential possessor* is a possessor before, during or after *year*. We exemplify the input to the network with the (*Vermeer*, 1902) pair from Figure 3.

4.2 Neural Networks

We use the implementations provided by the Keras neural network API (Chollet et al., 2015) with TensorFlow backend (Abadi et al., 2015). Additionally, we use GloVe embeddings with 300 dimensions (Pennington et al., 2014)⁴ to transform words into their distributed representations, the Adam optimizer (Kingma and Ba, 2014) and categorical cross entropy as a loss function. We train the network with batch size 16 for up to 200 epochs, but stop earlier if no improvement is observed in the validation set for 5 epochs. We reserve 20% of the train split for validation.

The neural network is composed of four Long Short-Term Memory networks (Hochreiter and Schmidhuber, 1997) with 200 units. The outputs of the LSTMs are concatenated along with the embeddings of the possessor and year, and the final output is calculated with a Softmax layer. Each LSTM has as its input a different chunk of text:

- The first LSTM takes as input the sequence of tokens in the sentence containing the *possessor* (top left in Figure 7). Each token is represented by the corresponding word embedding, and an additional embedding (also with 300 dimensions) for the possessor and all other tokens (there are only two unique additional embeddings, white and light gray in Figure 7). Unlike the word embedding from GloVe, the additional embeddings are initialized randomly and are updated during the training process. Our rationale to add the additional embeddings is to provide the LSTM with information to learn which tokens surrounding the *possessor* are more important.

- The second LSTM takes as input the sentence containing the *year* (top right in Figure 7). The input representation is very similar to the one used in the first LSTM, the only difference is that the additional embeddings (white and dark grey) indicate the *year* and any other token. Again, our rationale for the additional embeddings is to provide the LSTM with information to learn which tokens surrounding the *year* are more important.
- The third LSTM (bottom left in Figure 7) takes as input the Wikipedia article (i.e., the name of the possessee). The input words are represented with their GloVe embeddings.
- The fourth LSTM (bottom right) takes as input the section title from which the *possessor* and *year* were selected. The input words are also represented with their GloVe embeddings and no additional information. Our rationale is that some sections are less likely to contain valid possessors (e.g., *Cultural Impact* (low likelihood) vs. *Ownership and display* (high likelihood)).

4.3 Results

Results obtained with the test set are provided in Table 4. F-measures are always higher for *no* than *yes*, but recall that only *yes* label allows us to extract valid possession relations.

Baselines. The majority baseline always predicts *no* label for all temporal tags (before, during and after, see percentages in Figure 4), thus it fails to extract any possession information.

SVMs. SVMs obtain higher-than-chance results, but F-scores with *yes* label are relatively low (before: 0.33, during: 0.31 and after: 0.44).

Neural Networks. The full neural network always outperforms SVMs, but the difference in F-score

⁴Available at <https://nlp.stanford.edu/projects/glove/>, file `glove/glove.6B.300.txt`

		Before			During			After		
		P	R	F	P	R	F	P	R	F
SVM	yes	0.47	0.26	0.33	0.33	0.29	0.31	0.42	0.46	0.44
	no	0.86	0.94	0.90	0.93	0.94	0.93	0.88	0.87	0.87
NN Full	yes	0.41	0.38	0.40	0.22	0.56	0.32	0.44	0.65	0.53
	no	0.87	0.89	0.88	0.95	0.80	0.87	0.92	0.82	0.87
NN Full - addt_embeds.	yes	0.40	0.40	0.40	0.21	0.46	0.29	0.35	0.75	0.47
	no	0.88	0.87	0.88	0.94	0.82	0.88	0.93	0.69	0.79
NN Full - LSTM _{article_title}	yes	0.33	0.42	0.37	0.21	0.65	0.31	0.33	0.77	0.46
	no	0.87	0.82	0.85	0.95	0.74	0.83	0.93	0.66	0.77
NN Full - LSTM _{section_title}	yes	0.41	0.50	0.45	0.25	0.46	0.33	0.38	0.80	0.48
	no	0.89	0.85	0.87	0.94	0.86	0.90	0.94	0.67	0.78

Table 4: Results obtained in the test set using SVMs (all features), and the neural network (the full architecture and after disabling some components). Recall that *yes* is the only label that enables us to extract possession relations, results with *no* are mostly irrelevant.

with *yes* label is minimal for *during* (before: 0.40 vs 0.33, +21.2%; during: 0.32 vs. 0.31, +3.2%; after: 0.53 vs. 0.44, +20.5%).

We also experimented with modifications of the full neural network to provide insights into which components are more useful. Specifically, we report results not using the additional embeddings for the possessor and year, and disabling the LSTMs for the article title and section title. Note that while *yes* F-scores for *during* barely vary regardless of the modifications to the network, we found interesting patterns for *before* and *after*. All F-scores discussed below are for *yes* label, the only label that is useful to extract possession relations.

- First, the additional embeddings for the possessor and year are beneficial for *after* (0.47 vs 0.53, +12.8%) and *during* (0.29 vs 0.32, +10.3%), and neutral for *before*. This leads to the conclusion that the LSTM learns the contexts surrounding the possessor and year successfully only for *after* and *during*. Note that the additional embeddings provides information regarding the position of the possessor and year within their sentences.
- Second, the LSTM that takes as its input the article title is beneficial for *before* (0.37 vs. 0.40, +8.1%) and *after* (0.46 vs 0.53, +15.2%), and barely for *during* (0.31 vs. 0.32, +3.1%). Thus we can conclude that the article title contains useful information to determine the existence of possession relations, and that pretrained word embeddings capture this information.
- Third, the LSTM that takes as its input the

section title is beneficial for *after* (0.48 vs 0.53, +10.4%), detrimental for *before* (0.45 vs. 0.40, -11.1%) and barely detrimental for *during* (0.32 vs. 0.33, -3.0%). These results lead to the conclusion that the section title only contains useful information to determine possession relations in future years with respect to the years mentioned in the section.

5 Conclusions

Possession is an asymmetric semantic relation between two entities, where one entity (the possessee) *belongs* to the other entity (the possessor). Following theoretical works, we understand *belongs* in a broad sense, including physical, temporal, and control possessions.

In this paper, we track possession relations over time. Specifically, we work with Wikipedia articles about artworks, and extract their possessors as well as temporal information with respect to the years explicitly mentioned (before, during or after). We have presented an approach to extract potential possessors and pair them with years, and an annotation scheme to validate them. Overall inter-annotator agreement (Cohen’s kappa) is 0.70, and the resulting corpus has substantial information regarding possessors over time: in 75% of articles we validate at least 14 (possessor, year) pairs, and in 50%, at least 30 pairs.

Experimental results show that the task can be automated, although we obtain moderate results. We present an LSTM ensemble that outperforms a traditional SVM. Disabling certain components of the full network show that the article title and sec-

tion title benefit different temporal tags, and that the additional embeddings for the possessor and year are beneficial for *during*.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. [TensorFlow: Large-scale machine learning on heterogeneous systems](https://www.tensorflow.org/). Software available from [tensorflow.org](https://www.tensorflow.org/). <https://www.tensorflow.org/>.
- A.Y. Aikhenvald. 2013. Possession and ownership: a cross-linguistic perspective. In A.Y. Aikhenvald and R.M.W. Dixon, editors, *Possession and Ownership: A Cross-Linguistic Typology*, Oxford University Press, Oxford, chapter 1, pages 1–64.
- A.Y. Aikhenvald and R.M.W. Dixon. 2012. *Possession and Ownership: A Cross-Linguistic Typology*. Explorations in Linguistic Typology. OUP Oxford.
- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Comput. Linguist.* 34(4):555–596.
- Carmen Banea, Xi Chen, and Rada Mihalcea. 2016. Building a dataset for possessions identification in text. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, Helene Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*. European Language Resources Association (ELRA), Paris, France.
- Carmen Banea and Rada Mihalcea. 2018. [Possession identification in text](https://doi.org/10.1017/S1351324918000062). *Natural Language Engineering* page 122. <https://doi.org/10.1017/S1351324918000062>.
- Dhivya Chinnappa and Eduardo Blanco. 2018. Mining possessions: Existence, type and temporal anchors. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, New Orleans, Louisiana, USA, pages 496–505.
- François Chollet et al. 2015. Keras. <https://github.com/fchollet/keras>.
- B. Heine. 1997. *Possession: Cognitive Sources, Forces, and Grammaticalization*. Cambridge Studies in Linguistics. Cambridge University Press.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](https://doi.org/10.1162/neco.1997.9.8.1735). *Neural Comput.* 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](http://arxiv.org/abs/1412.6980). *CoRR* abs/1412.6980. <http://arxiv.org/abs/1412.6980>.
- Kenton Lee, Yoav Artzi, Jesse Dodge, and Luke Zettlemoyer. 2014. [Context-dependent semantic parsing for time expressions](http://www.aclweb.org/anthology/P14-1135). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Baltimore, Maryland, pages 1437–1447. <http://www.aclweb.org/anthology/P14-1135>.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*. pages 55–60.
- G.A. Miller and P.N. Johnson-Laird. 1976. *Language and perception*. Belknap Press. Belknap Press of Harvard University Press.
- Preslav I. Nakov and Marti A. Hearst. 2013. Semantic interpretation of noun compounds using verbal and other paraphrases. *ACM Trans. Speech Lang. Process.* 10(3):13:1–13:51.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](http://www.aclweb.org/anthology/D14-1162). In *Empirical Methods in Natural Language Processing (EMNLP)*. pages 1532–1543. <http://www.aclweb.org/anthology/D14-1162>.
- Sameer Pradhan, Lance Ramshaw, Mitchell Marcus, Martha Palmer, Ralph Weischedel, and Nianwen Xue. 2011. Conll-2011 shared task: Modeling unrestricted coreference in ontonotes. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task*. Association for Computational Linguistics, Portland, Oregon, USA, pages 1–27.
- Nils Reimers, Nazanin Dehghani, and Iryna Gurevych. 2016. [Temporal anchoring of events for the time-bank corpus](https://arxiv.org/abs/1610.06092). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages

2195–2204. <http://www.aclweb.org/anthology/P16-1207>.

- L. Stassen. 2009. *Predicative Possession*. Oxford Studies in Typology and Linguistic Theory. OUP Oxford.
- Jannik Strötgen and Michael Gertz. 2015. [A baseline temporal tagger for all languages](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, pages 541–547. <http://aclweb.org/anthology/D15-1063>.
- Shiao Wei Tham. 2004. *Representing Possessive Predication: Semantic Dimensions and Pragmatic Bases*. Ph.D. thesis, Stanford University.
- Stephen Tratz and Eduard Hovy. 2010. A taxonomy, dataset, and classifier for automatic noun compound interpretation. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Stroudsburg, PA, USA, ACL '10, pages 678–687.
- Stephen Tratz and Eduard H. Hovy. 2013. Automatic interpretation of the english possessive. In *ACL (1)*. The Association for Computer Linguistics, pages 372–381.