

Living Machines: A study of *atypical* animacy

This paper proposes a new approach to animacy detection, the task of determining whether an entity is represented as animate in a text. In particular, this work is focused on *atypical* animacy and examines the scenario in which typically inanimate objects, specifically machines, are given animate attributes. To address it, we have created the first dataset for atypical animacy detection, based on nineteenth-century sentences in English, with machines represented as either animate or inanimate. Our method builds upon recent innovations in language modeling, specifically BERT contextualized word embeddings, to better capture fine-grained contextual properties of words. We present a fully unsupervised pipeline, which can be easily adapted to different contexts, and report its performance on an established animacy dataset and our newly introduced resource. We show that our method provides a substantially more accurate characterization of atypical animacy, especially when applied to highly complex forms of language use.

Conceptualization

Mariona Coll Ardanuy^{1,5}
Daniel CS Wilson^{1,5}

Methodology

Mariona Coll Ardanuy
Federico Nanni¹
Kasra Hosseini¹

Implementation

Federico Nanni
Kasra Hosseini
Mariona Coll Ardanuy
Kaspar Beelen^{1,5}

Reproducibility

Kasra Hosseini
Federico Nanni

Interpretation

Kaspar Beelen
Mariona Coll Ardanuy
Katherine McDonough^{1,5}
Daniel CS Wilson
Ruth Ahnert⁵
Jon Lawrence⁴
Giorgia Tolfo²

Historical Analysis

Daniel CS Wilson
Katherine McDonough
Kaspar Beelen
Jon Lawrence

Data Curation

Kaspar Beelen
Mariona Coll Ardanuy
Federico Nanni
Giorgia Tolfo

Annotation

Giorgia Tolfo
Ruth Ahnert
Kaspar Beelen
Mariona Coll Ardanuy
Jon Lawrence
Katherine McDonough
Federico Nanni
Daniel CS Wilson

Writing and Editing

Mariona Coll Ardanuy
Federico Nanni
Ruth Ahnert
Kaspar Beelen
Kasra Hosseini
Jon Lawrence
Katherine McDonough
Barbara McGillivray^{1,3}
Daniel CS Wilson

Supervision

Barbara McGillivray
Ruth Ahnert

Project Management

Barbara McGillivray
Ruth Ahnert
Mariona Coll Ardanuy

¹The Alan Turing Institute, ²The British Library, ³University of Cambridge,
⁴University of Exeter, ⁵Queen Mary University of London

Contacts of the corresponding authors: {mcollardanuy,dwilson}@turing.ac.uk

Work for this paper was produced as part of "Living with Machines". This project, funded by the UK Research and Innovation (UKRI) Strategic Priority Fund, is a multidisciplinary collaboration delivered by the Arts and Humanities Research Council (AHRC), with The Alan Turing Institute, the British Library and the Universities of Cambridge, East Anglia, Exeter, and Queen Mary University of London. This work was also supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

1 Introduction

Animacy is the property of being alive. Although the *perception* of a given entity as animate (or not) tends to align with its biological animacy, discrepancies are not uncommon. These may arise from differences in how we unconsciously perceive entities, or from the deliberate use of animate expressions to describe inanimate entities (or vice versa). Machines sit at the fuzzy boundary of animacy and inanimacy (Turing, 1950; Yamamoto, 1999). In this paper, we examine how machines have been imagined over the nineteenth century from lifeless mechanical objects to human-like agents that feel, think, and even love. We focus on nineteenth-century Britain, a society being transformed by industrialization, as a good candidate for studying this transition.

This paper applies state-of-the-art contextualized word representations, trained using the BERT architecture (Devlin et al., 2018), to animacy detection. In contrast to previous research, this paper provides an in-depth exploration of the ambiguities and figurative aspects that characterize animacy in natural language, and analyzes how context shapes animacy. Context is constitutive of meaning (Wittgenstein, 1921, 3.3), an observation acknowledged by generations of scholars, but which is still difficult to apply to its full extent in computational models of language. We show how the increased sensitivity of BERT-based models to contextual cues can be exploited to analyze how the same entity (e.g., a machine) can be at once represented as animate or inanimate depending on the purpose of the writer.

This paper makes several contributions: we present an unsupervised method to detect animacy that is highly sensitive to the context and therefore suited to capture not only typical animacy, but especially atypical animacy. Additionally, we provide the first benchmark for atypical animacy detection based on a dataset of nineteenth-century sentences in English with machines represented as animate and inanimate. We conduct an extensive quantitative evaluation of our approach in comparison with supervised and unsupervised baselines on both an established animacy dataset and on our newly introduced resource, and demonstrate the generalizability of our approach. Finally, we discuss the distinction between animacy and humanness, and provide preliminary quantifiable insights into the linguistic representation of the historical process of dehumanization by mechanization.

Atypical observations are rare by definition. Because of this, addressing them is often an ungratifying undertaking, as they can only marginally improve the accuracy of general natural language processing systems on existing benchmarks, if at all. And yet, precisely because of this, atypical observations tend to acquire a certain salience from a qualitative and interpretative point of view. For the humanities scholar and the linguist, such deviations prove particularly interesting objects of study because they flout expectations.

2 Related work

Animacy and its relation to cognition has been extensively studied in a range of linguistic fields, from neurolinguistics and language acquisition research (Gao et al., 2012; Opfer, 2002) to morphology and syntax (Rosenbach, 2008; McLaughlin, 2014; Vihman and Nelson, 2019). There is evidence that animacy is not a fixed property of lexical items but is subject to their context of use (Nieuwland and van Berkum, 2005). This points to a more nuanced and graduated view of animacy than a binary distinction between “animate” and “inanimate” (Peltola, 2018; Bayanati and Toivonen, 2019), which results in a hierarchy of entities that reflects notions of agency, closeness to the speaker, individuation, and empathy (Comrie, 1989; Croft, 2002; Yamamoto, 1999). Yamamoto (1999) identifies modern machines as one of the most prominent examples at the frontier area between animacy and inanimacy.

The distinction between animate and inanimate is a fundamental aspect of cognition and language, and has been shown to be a useful feature in natural language processing (NLP), in tasks such as coreference and anaphora resolution (Lee et al., 2013; Orasan and Evans, 2007; Poesio et al., 2008; Raghunathan et al., 2010), word sense disambiguation (Chen et al., 2006; Øvrelid, 2008), and semantic role labeling (Connor et al., 2010). Earlier approaches to animacy detection relied on semantic lexicons (such as WordNet, Fellbaum (1998)) combined with syntactic analysis (Evans and Orasan, 2000), or developed machine-learning classifiers that use syntactic and morphological features (Øvrelid, 2008). More recently, Karsdorp et al. (2015) focused on Dutch folk tales and trained a classifier to identify animate

entities based on a combination of linguistic features and word embeddings trained using a skip-gram model. They showed that close-to-optimal scores could be achieved using word embeddings alone. Jahan et al. (2018) developed a hybrid classification system which relies on rich linguistic text processing, by combining static word embeddings with a number of hand-built rules to compute the animacy of referring expressions and co-reference chains. Previous research (Karsdorp et al., 2015; Jahan et al., 2018) has acknowledged the importance of context in atypical animacy, but it has not explicitly tackled it, or attempted to quantify how well existing methods have handled such complexities.

Whereas static word representations such as Word2Vec (Mikolov et al., 2013) have been shown to perform well in typical animacy detection tasks, we argue that they are not capable of detecting atypical cases of animacy, as by definition animacy in the latter case must arise from the context, and not the target entity itself. The emergence of contextualized word representations has yielded significant advances in many NLP tasks (Peters et al., 2017; Radford et al., 2018; Devlin et al., 2018). Unlike their static counterparts, they are optimized to capture the representations of target words in their contexts, and are therefore more sensitive to context-dependent aspects of meaning. BERT (Bidirectional Encoder Representations from Transformers, Devlin et al. (2018)) incorporates the latest improvements in language modeling and, through its deep bidirectionality and its self-attention mechanism, has become one of the most successful attempts to train context-sensitive language models. BERT is pre-trained on two tasks: Masked Language Model (MLM), which tries to predict a masked token based on both its left and right context, and Next Sentence Prediction (NSP), which tries to predict the following sentence through a binary classification task. This dual learning objective ensures that the contextual representations of words are learned, also across sentences. Its simple and efficient fine-tuning mechanism allows BERT to easily adapt to different tasks and domains.

3 Method

In this section, we describe our approach to determine the animacy of a target expression in its context. The intuition behind our method is the following: an entity becomes animate in a given context if it occurs in a position in which one would typically expect a living entity. More specifically, given a sentence in which a target expression has been adequately masked, we rely on contextualized masked language models to provide ranked predictions for the masked element, as shown in example 1:

- (1) **Original sentence:** And why should one say that the machine does not live?
Masked sentence: And why should one say that the [MASK] does not live?
Predictions with scores: man (5.0788), person (4.4484), other (4.1866), child (4.1600), king (4.1510), patient (4.1249), one (4.1141), stranger (4.1067), ...

We then determine the animacy of the masked expression by averaging the animacy of the top κ tokens that have been predicted to fill the mask in the sentence. While this may seem a circular argument at first glance, the fundamentally probabilistic nature of language models means that we are in fact replacing the masked element with tokens that have a high probability of occurring in this context. Our method rests on the assumption that, given a context requiring an animate entity, a contextualized language model should predict tokens corresponding to conventionally animate entities.

We use a BERT language model to predict a number of possible fillers given a sentence with a masked token, with their corresponding probability scores. We then use WordNet,¹ a lexical database that encodes relations between word senses (Fellbaum, 1998), to determine whether the predicted tokens correspond to typically animate or inanimate entities. Tokens can be ambiguous: the same token can be used for several different word senses, some of which may correspond to living entities and some others not. For example, the word ‘dresser’ has several meanings, including a profession – typically animate – and piece of furniture – typically inanimate. We disambiguate each predicted token to its most relevant sense in WordNet by measuring the similarity between the original sentence and the gloss of each WordNet sense. Inspired by previous research on distributional semantic models for word sense disambiguation

¹<https://wordnet.princeton.edu/>.

(Basile et al., 2014), we have implemented a BERT-adapted version of the Lesk algorithm, which leverages recent advancements in transformer-based sentence representations (Reimers and Gurevych, 2019).

WordNet organizes nouns according to hierarchies, which eventually converge at the root node *entity*. Senses of nouns that correspond to living entities fall under the *living_thing* node, which is the common parent of the person, animal, plant, and microorganism classes, among others. Therefore, we determine whether each predicted token corresponds to an animate or inanimate entity based on whether its disambiguated sense is a descendant of the *living_thing* node. Finally, we produce a single animacy score between 0 and 1 for the masked element, by averaging the animacy values (i.e. 0 if inanimate, 1 if animate) of the predicted tokens, weighted by their probability score. We find the optimal animacy threshold τ and cutoff κ (i.e. number of predicted tokens) through experimentation.²

4 Data

We use two datasets to evaluate the performance of our algorithm: the first is derived from the data released by Jahan et al. (2018), while the second has been created by us for specifically testing detection of unconventional animacy, in nineteenth-century English texts in particular. Both datasets are described in sections 4.1 and 4.2 respectively, and summarized and discussed in section 4.3.

4.1 The *Stories* animacy dataset

In their paper, Jahan et al. (2018) used a collection of stories (i.e. Russian folktales, Islamist extremist stories, and Islamic Hadiths) translated into English and already provided with several layers of linguistic annotations (Finlayson et al., 2014; Finlayson, 2017). The authors enriched the texts with animacy annotations at the level of coreference chains and of their referring expressions. The authors reported a near-perfect inter-annotation agreement (Cohen’s $\kappa = 0.99$). Given that our method works at the sentence level, we reformatted their data to make it compatible with our approach. This process resulted in a new dataset (henceforth *Stories* dataset) consisting of 5,835 sentences, each of which contains a target expression annotated with animacy (see some examples in Table 1).³

Target expression	Sentence	Animacy
this very day	That furrow can be seen to this very day ; it is fourteen feet high.	0
a long time or a short time	The bear kept her with him, and after some time, a long time or a short time , she had a son by him.	0
the window	Very early next morning Vasilisa awoke, after Baba Yaga had arisen, and looked out of the window .	0
herself	But the pike knew quite well what he was thinking about, and laid herself right across the sea.	1
Fox	The Fox took it home, stuck what remained of his coins behind the hoop, and brought it back to the tsar.	1
a tsar	In a kingdom, in a far-away land, there lived, there were a tsar and his tsaritsa	1

Table 1: Examples of sentences from the *Stories* dataset, with the target expressions and their animacy values, derived from the work by Jahan et al. (2018).

4.2 The *19thC Machines* animacy dataset

The *Stories* dataset is largely composed of target expressions that correspond to either typically animate or typically inanimate entities. Even though some cases of unconventional animacy can be found (folktales, in particular, are richer in typically inanimate entities that become animate), these account for a very small proportion of the data.⁴ We decided to create our own dataset (henceforth *19thC Ma-*

²We will share the code, data, and experimental results openly on Github, with an appropriate software DOI, upon publication of the paper.

³Unfortunately, we were not able to reproduce the same number of target expressions as are reported in Jahan et al. (2018), but we will provide the code we used to generate our datasets for future studies. See table 3 for a summary of the *Stories* dataset.

⁴Note that this is based on observation, as the number of atypical cases is not provided in Jahan et al. (2018).

chines dataset) to gain a better sense of the suitability of our method to the problem of atypical animacy detection, with particular attention to the case of animacy of machinery in nineteenth-century texts.

We extracted sentences containing nouns that correspond to types of machines from an open dataset of nineteenth-century books (from now on *19thC BL Books*).⁵ Even though the OCR quality is relatively good, some noise can still be found in the dataset. In order to extract sentences which contain machine-related words, we manually selected words that occurred close to the combined vector of ‘machine’ and ‘machines’ in Word2vec models trained on BL books from before and after 1850 (to make sure the selection is not biased towards a particular half of the nineteenth century). We refined this list in multiple iterations, adding new words and recomputing the combined vector. The result was a stable list of generic words referring to machines across the period under investigation.⁶ In most sentences, machines are treated as inanimate objects. We therefore employed a pooled strategy⁷ to identify meaningful sentences for annotation: we specified four animacy bands (0.0-0.25, 0.25-0.50, 0.50-0.75, and 0.75-1.00) and we used the different methods described in section 6 to obtain a fixed number of sentences for each band. This way, we obtained a large pool of sentences capturing a variety of different types of animate and inanimate contexts present in the corpus.

4.2.1 Preliminary annotations

For human annotators, even history and literature experts, language subtleties made this task extremely subjective. In order to gain a better understanding of the problem, we started with two preliminary annotation tasks. A first set of 100 sentences derived from the pooling process was distributed among the annotators.⁸ In the first task, we masked the target word (i.e. the machine) in each sentence and asked the annotator to fill the slot with the most likely entity between ‘human’, ‘horse’, and ‘machine’, representing three levels in the animacy hierarchy: human, animal, and object (Comrie, 1989, 185). We asked the annotators to stick to the most literal meaning and stay away from metaphoric interpretations when possible. Interestingly, even though the original masked expressions contained only instances of the lemma ‘machine’, the annotators selected ‘machine’ as the most likely option in only 62% of the total number of annotations. However, the agreement was low, with a Krippendorff α of 0.32. This indicates that, at least in some contexts, machines seem to be interchangeable with humans and animals, and that annotators may disagree about when one is preferred over the other.

The second task was more straightforwardly related to determining the animacy of the target entity. We asked the annotators to provide a score between -2 and 2, with -2 being definitely inanimate, -1 possibly inanimate, 1 possibly animate, and 2 definitely animate. Neutral judgements were not allowed. The agreement for this second task was low as well (Krippendorff α of 0.43). Neither collating the annotations into positive (scores 1 and 2) and negative groups (scores -2 and -1) nor collating slightly animate and slightly inanimate together improved inter-annotation agreement significantly. We explored the cases in which annotators disagreed, and found that the same sentence would often be annotated as highly animate by one annotator and as highly inanimate by another. This was especially the case of sentences containing similes or metaphors that liken machines to humans, animals, or systems.

Preliminary annotations helped us to understand the data and improve our experimental design. Annotators were asked to leave comments and provide feedback, and agreed that both tasks were more challenging than expected, mostly due to the high incidence of figurative language, as in example 2.

- (2) (a) *He is himself but a mere **machine**, unconscious of the operations of his own mind.*
(b) *Our servants, like mere **machines**, move on their mercenary track without feeling.*

⁵This nineteenth-century book dataset was digitized by the British Library. It contains $\approx 48,200$ volumes with $\approx 4.9\text{B}$ tokens. Data available via <https://data.bl.uk/digbks/db14.html>, DOI <https://doi.org/10.21250/db14> (British Library Labs, 2014).

⁶A publication is forthcoming which explains the lexicon expansion procedure in detail. We will make the Word2vec models available upon publication of the paper. The curated words are: ‘machine’, ‘machines’, ‘machinery’, ‘engines’, ‘engine’, ‘locomotive’, ‘locomotives’, ‘turbine’, ‘turbines’, ‘boiler’, ‘boilers’, ‘dynamo’, ‘dynamos’, ‘motor’, ‘motors’, ‘apparatus’, ‘apparatuses’, ‘accumulator’, ‘accumulators’, ‘compressor’, and ‘compressors’.

⁷Pooling is an established evaluation strategy for information retrieval systems (Spark-Jones, 1975; Buckley et al., 2007).

⁸A combination of computational linguists, historians, literary scholars, and data scientists.

- (c) *My companions treated me as a **machine**, and never in any way repaid my services.*
 (d) *A master who looks upon thy kind, not as mere **machines**, but as valued friends.*

These kinds of sentences present a very particular type of interpretative openness. In each case a human or group of humans (animate beings) are likened to a machine to suggest that they have been reduced somehow in their agency or animacy. Some of the annotators deduced an implied inanimacy of the machine, which would have the rhetorical effect of suggesting that the humans too are rendered inanimate. Conversely, for other annotators, the comparison conjured a kind of automaton, a human-machine-hybrid, and therefore an animate machine.

4.2.2 Final annotations

A subgroup of five annotators collaboratively wrote the guidelines based on their experience annotating the first batch of sentences, taking into account the most common discrepancies. After discussion, it was decided that a machine would be tagged as animate if it is described as having traits or characteristics distinctive of biologically animate beings or human-specific skills, or portrayed as having feelings, emotions, or a soul. Sentences like the ones in example 2 would be considered animate, but an additional annotation layer would be provided to capture the notion of **humanness** (or lack thereof, i.e. dehumanization through mechanization).⁹ A new batch of 400 unseen sentences was sent to the annotators. The Krippendorff α of this annotation task was of 0.74 for animacy and 0.50 for humanness. The gold standard was produced by one of the annotators and author of the guidelines, who assigned the final labels by adjudication, taking into account the agreements and disagreements between annotators and their comments. We provide examples of annotations in table 2.

Target	Sentence	Animacy	Humanness
engine	In December, the first steam fire engine was received, and tried on the shore of Lake Monona, with one thousand feet of hose.	0	0
engine	It was not necessary for Jakie to slow down in order to allow the wild engine to come up with him; she was coming up at every revolution of her wheels.	1	1
locomotive	Nearly a generation had been strangely neglected to grow up un-Americanized, and the private adventurer and the locomotive were the untechnical missionaries to open a way for the common school.	1	1
machine	The worst of it was, the people were surly; not one would get out of our way until the last minute, and many pretended not to see us coming, though the machine , held in by the brake, squeaked a pitiful warning.	1	1
machines	Our servants, like mere machines , move on their mercenary track without feeling.	1	0
machinery	We have everywhere water power to any desirable extent, suitable for propelling all kinds of machinery .	0	0

Table 2: Examples of sentences from the *19thC Machines* dataset with their target expression and corresponding annotations in terms of animacy and humanness.

4.3 Datasets summary and discussion

Table 3 summarizes the main differences between the two datasets. The *Stories* dataset is larger and more varied in terms of unique target expressions, and has a nearly-perfect inter-annotation agreement (Cohen’s κ of 0.99).¹⁰ The *19thC Machines* corpus consists of 393 sentences¹¹ with 13 unique target expressions, which can be either animate or inanimate, depending on the context. As discussed, the disagreement was quite high in comparison, proving that detecting atypical animacy can be a very semantically complex problem (in particular in highly figurative language). There are 183 sentences in which the machine has been tagged as animate, out of which 134 are also instances of humanness.

⁹For us, ‘animacy’ encompasses ‘humanness’: all sentences that are tagged as cases of humanness are also annotated as animate, but not all sentences that are annotated as animate are instances of humanness. The term ‘humanness’ is used differently in other research fields, for instance when evaluating the performance of a chatbot (Svenningsson and Faraon, 2019).

¹⁰Previous work on animacy in Dutch reported a similarly high κ agreement score of 0.95 (Karsdorp et al., 2015).

¹¹13 sentences were discarded because they mentioned a human boiler instead of a boiler engine.

	#Sentences	IAA	UniqueExpressions	Train/Test	#Animate
<i>Stories</i>	5835	Cohen’s κ 0.99	4277	4084/1751	2080
<i>19thC Machines</i>	393	Krippendorff’s α 0.74	13	99/294	183

Table 3: Comparison between the *Stories* and *19thC Machines* animacy datasets.

5 Language Models

In our experiments, we used the ‘BERT base uncased’ model and tokenizer as contemporary models,¹² hereinafter referred to as *BERT-base*. Besides, in order to investigate pattern changes over time, we also fine-tuned *BERT-base* on the *19thC BL Books* dataset, split into four time periods (before 1850, between 1850 and 1875, between 1875 and 1890, and between 1890 and 1900), each containing ≈ 1.3 B words per period, except for the 1890-1900 time period which had ≈ 940 M words.¹³ The fine-tuning was done in four sequential steps. The *BERT-base* model was first fine-tuned on the oldest time period (i.e., books published before 1850). We then used the resulting language model and further fine-tuned it on the next time period. This procedure of fine-tuning a language model on the subsequent time period was repeated for the other two time periods. For each time period, we preprocessed all books¹⁴ and tokenized them using the original BERT-base tokenizer as implemented by HuggingFace¹⁵ (Wolf et al., 2019). We did not train new tokenizers for each time period. This way, the resulting language models can be compared easily with no further processing or adjustments. The tokenized sentences are then fed to the language model fine-tuning tool in which only the masked language model (MLM) objective is optimized.¹⁶

We do not aim at modeling animacy change diachronically in this paper. Instead, we treat the different fine-tuned models as four different snapshots of time that we can then use for comparison. Table 4 shows how our four fine-tuned language models differ in predicting the same masked element in a sentence. While this is a cherry-picked example, it serves as illustration of the importance of having language models that adequately reflect the language that is contemporary to our data.

FT BERT model	Predicted tokens
≤ 1850	<i>man</i> (5.3291), <i>prisoners</i> (4.9758), <i>men</i> (4.885), <i>book</i> (4.6477), <i>people</i> (4.556), <i>one</i> (4.4271), <i>slave</i> (4.4034), <i>air</i> (4.1329), <i>water</i> (4.1148)
1850–1875	<i>men</i> (10.7655), <i>people</i> (9.497), <i>miners</i> (9.249), <i>engine</i> (8.0428), <i>women</i> (8.0126), <i>company</i> (7.7261), <i>machine</i> (7.6021), <i>labourers</i> (7.5987), <i>machines</i> (7.5012)
1875–1890	<i>men</i> (10.2048), <i>miners</i> (7.6654), <i>machines</i> (7.4062), <i>people</i> (7.2991), <i>engine</i> (7.232), <i>labourers</i> (7.0957), <i>engines</i> (6.7786), <i>engineers</i> (6.5642), <i>machine</i> (6.4712)
1890–1900	<i>mercury</i> (8.0446), <i>machinery</i> (7.4067), <i>machine</i> (7.2903), <i>mine</i> (7.274), <i>mill</i> (7.057), <i>men</i> (7.0257), <i>engine</i> (6.9966), <i>lead</i> (6.9177), <i>miners</i> (6.7764)

Table 4: This table illustrates the differences between language models fine-tuned over time: the *Predicted tokens* column contains the tokens predicted as most probable to fill the [MASK] gap in the sentence “They were told that the [MASK] stopped working”, according to each model. The more recent models clearly predict machine-related words more often than the older ones.

¹²<https://github.com/google-research/bert>.

¹³While the data distribution for fine-tuning was decided on largely by the number of tokens, these periods also work well in representing distinctive cultural eras. For example, the pre-1850 dataset sets apart the first industrial revolution from later developments in Britain. Likewise, 1890-1900 is seen as distinct, especially in literary terms, for the emergence of ‘modernist’ sensibilities and the questioning of class and gender hierarchies associated with the term ‘*fin de siècle*’.

¹⁴We normalized white spaces, removed accents and repeated “.” (as they are common in the OCR’d texts), added a white space before and after punctuation signs, and finally split token streams into sentences using syntok library: <https://pypi.org/project/syntok/>.

¹⁵<https://github.com/huggingface/transformers>.

¹⁶We used a batch size of 5 per GPU and fine-tuned for 1 epoch over the books in each time-period. The choice of batch size was dictated by the available GPU memory (we used $4 \times$ NVIDIA Tesla K80 GPUs in parallel). Similar to the original BERT pre-training procedure, we used the Adam optimization method (Kingma and Ba, 2014) with learning rate of $1e-4$, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and L_2 weight decay of 0.01. In our fine-tunings, we used a linear learning-rate warmup over the first 2,000 steps. A dropout probability of 0.1 was used in all layers.

6 Experiments and results

6.1 Baselines

We provide two different types of baselines: a masking approach using static word representations and a classification approach. We also provide the performance of the most frequent class, which is the inanimate class both in the *Stories* and the *19thC Machines* datasets.

Masking Approach. In order to understand the added value of relying on contextualized word representations for predicting the masked entity, we compare our results with a simpler alternative based on the use of traditional static word embeddings (hereinafter *MaskPredict: WordEmb*). It predicts words which are semantically similar to the masked expression (via cosine similarity of their word embeddings), without any additional information on the context in which the word is mentioned. We determine the animacy value of each predicted token and compute the combined animacy score of all predicted tokens using the same WordNet-based approach as in our method. The optimal cutoff (number of predicted tokens) and animacy threshold are found by maximizing F-Score on the training set.

Classification Approach. One alternative approach used in Karsdorp et al. (2015) and Jahan et al. (2018) is to treat animacy as a classification problem by training supervised classifiers on examples annotated with a binary label. We report the performance of three different classifiers: two SVMs using either tf-idf or word embeddings as feature vectors, and a BERT classifier (from now on *SVM TFIDF*, *SVM WordEmb*, and *BERTClassifier*, respectively).¹⁷ All classifiers are trained on the *Stories* training set (over 4000 instances), either at the target expression level (*targetExp*), or at the context level (i.e. trained on the target expressions and n words to the left and to the right, where n is 3, as in Karsdorp et al. (2015)), either including the target expression itself (*targetExp + ctxt*) or replacing it with a mask (*maskedExp + ctxt*). We find the optimal animacy threshold for each classifier and dataset through parameter-tuning on the respective training set, by maximizing F-Score. While such approaches act as “skylines” when compared with our unsupervised masking methods on the *Stories* dataset, examining their performance on the *19thC Machines* dataset highlights their drawbacks when used out of domain.

6.2 Evaluation metrics

Since our datasets are not always balanced and we want to give equal importance to each class, we report on macro precision and recall, and macro average F-score. For reference, we also provide mean average precision (Map), a popular metric in information retrieval which highlights how well the ranking of the animacy score correlates with the labels.

6.3 Experimental results

Table 5 reports the performance of the different baselines and methods on the *Stories* and *19thC Machines* datasets.¹⁸ Classifiers based on the target expression alone are the best performing methods in the *Stories* dataset. Interestingly, their performance becomes worse when more context is added, and even more so when the target expression itself is masked. Unlike the baseline classifiers, our method (*MaskPredict: BERT-base*) does not use the target expression as a feature at all: it relies solely on the context. In fact, adding context (i.e. one sentence to the left and to the right, *MaskPredict: BERT-base + ctxt*) helps improve its performance (from 0.77 to 0.84 in F-Score). This analysis shows that target expression is the most indicative feature of conventional animacy. And yet, the good performance of our context-based method proves that animacy is not only entity-level, but that it is informed by the context as well.

Classifier baselines perform strikingly worse on the *19thC Machines* dataset.¹⁹ Both baselines and our method produce comparatively worse results in the *19thC Machines* dataset than in the *Stories* dataset.

¹⁷The SVMs use both a linear kernel with standard parameters, from the Scikit-learn package. For consistency, we have used the following Scikit-learn wrapper for training the BERT classifier, again with default parameters: <https://github.com/charles9n/bert-sklearn>.

¹⁸Due to space limitations, we only provide the most representative results. We will publish an appendix with the results for all baselines and optimal threshold and cutoff parameters per scenario upon acceptance of the paper.

¹⁹As mentioned, classifiers were trained on the *Stories* dataset. All thresholds (and cutoffs) were found by maximizing F-Score on the *19thC Machines* training set (99 sentences).

	<i>Stories</i>				<i>19thC Machines</i>			
	Precision	Recall	F-Score	Map	Precision	Recall	F-Score	Map
Most frequent class	0.317	0.5	0.388	0.418	0.268	0.5	0.349	0.505
SVM TFIDF: targetExp	0.878	0.867	0.873	0.896	0.268	0.5	0.349	0.505
SVM WordEmb: targetExp	0.9	0.907	0.903	0.932	0.428	0.442	0.435	0.462
BERTClassifier: targetExp	0.931	0.94	0.935	0.944	0.453	0.49	0.471	0.648
SVM TFIDF: targetExp + ctxt	0.776	0.772	0.774	0.781	0.504	0.502	0.503	0.504
SVM WordEmb: targetExp + ctxt	0.781	0.783	0.782	0.79	0.58	0.557	0.569	0.579
BERTClassifier: targetExp + ctxt	0.91	0.918	0.914	0.929	0.559	0.538	0.549	0.559
SVM TFIDF: maskedExp + ctxt	0.671	0.66	0.666	0.642	0.548	0.548	0.548	0.511
SVM WordEmb: maskedExp + ctxt	0.651	0.655	0.653	0.616	0.561	0.526	0.543	0.528
BERTClassifier: maskedExp + ctxt	0.76	0.766	0.763	0.77	0.583	0.583	0.583	0.539
MaskPredict: WordEmb	0.709	0.709	0.709	0.689	0.441	0.487	0.463	0.407
MaskPredict: BERT-base	0.767	0.767	0.767	0.742	0.741	0.74	0.741	0.764
MaskPredict: BERT-base +ctxt	0.834	0.845	0.839	0.845	0.762	0.76	0.761	0.841
MaskPredict: 19thcBERT +ctxt	–	–	–	–	0.783	0.774	0.779	0.853

Table 5: Evaluation results on the *Stories* and *19thC Machines* dataset.

This is probably due to the higher complexity of detecting atypical animacy (as suggested by the comparatively higher disagreement between annotators) and the noisier nature of this dataset, due to OCR errors. As opposed to the other approaches, our models yield consistent performances in both datasets, showing the advantages of its unsupervised context-dependent architecture.²⁰

The *19thC Machines* dataset is composed of sentences from the selected four time periods. As shown in table 3, the appropriate fine-tuned BERT model of the period to which each sentence belongs to (i.e. *MaskPredict: 19thcBERT +ctxt*) provides better results than the contemporary model, especially in terms of mean average precision, i.e. the ranking generated by the animacy score. Even though this difference is not found to be statistically significant in our dataset, a more in-depth analysis reveals interesting trends and patterns in the predictions of the different language models (see section 7.2).

7 Discussion and interpretation

Researchers across many disciplines have long debated the relation between language and the social worlds in which it exists. Studying the linguistic forms used to depict machines *as if* they were alive raises important questions about the relation between humans and machines which go beyond language. Animacy and its related concept of agency (Yamamoto, 2006) are important markers of social and political power: when ascribed to non-human actors they indicate the shifting perception of human agency in distinction to that of machines, as recorded in these common turns of phrase. The forms that ‘machine language’ take are, therefore, unlikely to be timeless; that is to say, their quantity and also their quality appear differently at different periods. The nature of these differences will be of great interest to historians seeking to investigate aspects of life, for example, during Britain’s rapid industrialization in the nineteenth century. For these reasons, understanding the linguistic patterns of ‘living machines’ can help make sense of how humans have been living with machines more generally. In sections 7.1 and 7.2 we present a preliminary investigation of these issues.

7.1 Animacy and humanness

As discussed in section 4.2.2, we consider entities as animate if they are given attributes and physical faculties that are characteristic of living entities. They are attributed the further subfeature of *humanness* if they are portrayed as sentient and capable of specifically human emotions.²¹ The latter is loosely

²⁰Training complex sequential models on small datasets is something we explore in future work. Our experiments suggested that without additional tweaks, they don’t perform well on atypical cases.

²¹Example 2 shows examples of sentences in which the machine is animate but lacking humanness. An example of a sentence displaying humanness is: ‘He bore the movement well [...] and make one wonder why the poor crazy **machine** is thought worthy of being put together again, with infinite pains, and wonderful science.’

tied to the idea of an anthropocentric hierarchy in animacy (Comrie, 1989; Croft, 2002; Yamamoto, 1999), which ranks entities most capable of human perception as the most animate, reflecting notions of agency, closeness to the speaker, and speaker empathy among others. All baselines and methods are worse in predicting humanness than in predicting more general animacy.²² The lower agreement between annotators in detecting humanness (Krippendorff α of 0.50) suggests a higher subjectivity of the task. In addition, our WordNet approach to determine animacy of predicted tokens is insensitive to animacy hierarchy: any living entity is considered equally animate. Interestingly, the performance of our method does not improve if we consider as animate only entities under the *person* node in WordNet, instead of those under *living_thing*.

We analyzed BERT’s predictions in sentences where machines are attributed or negated humanness. Table 6 shows the top predicted tokens by BERT for each case, and exposes some social biases embedded in nineteenth-century language that are captured in the language models. While ‘man’ remains the most predicted token replacing machine (and ‘woman’ is not far behind), the appearance of ‘slave(s)’ and ‘savage’ in contexts of negated humanness reflects the tendency to use these words in discourses that confer diminished human rights and qualities on those people.

Humanness	Most predicted tokens
Attributed	<i>man, child, dog, gentleman, woman, horse, fellow, person, boy, soldier</i>
Negated	<i>man, soldier, slave, woman, men, coward, beings, slaves, child, savage</i>

Table 6: Most predicted tokens in which the machine is attributed humanness or lack thereof, according to the pre-1850 BERT model.

7.2 Exploration of historical models

A language model is a probabilistic representation of a given language. The meaning and usage of words change over time due to linguistic, but also cognitive, social, and contextual factors (Hamilton et al., 2016; Kutuzov et al., 2018; Giulianelli, 2019). Social and technological changes are paralleled by changes in the language used to describe them. New terms arise or are created and new meanings come to infuse old terms (Schatzberg, 2018). The way we think and talk about machines has necessarily changed in line with the widespread adoption of new technology over time. We started by inspecting which living entities are replaced by machines or, put differently, what BERT predictions tell us about the characteristics of machines when they are portrayed as being *alive*.

In the nineteenth century, who (or what) was performing work was changing dramatically. Children were entering and exiting the labor pool at different times, and servants and slaves were similarly key parts of the workforce. Here we use the historical language models to explore the way that such groups were related to machines (and vice versa). We followed a simple procedure: given sentences with a masked machine-related concept and two lists of words related either to children or servants, we compute the mean reciprocal rank between the lists and BERT’s predictions. A high score would suggest that terms related to a target concept (e.g., children) rank highly among BERT predictions. In figure 1, we show changes in the relevance of the concepts *child* and *servant*²³ among the predictions for the masked machine. We plotted results as a function of time for 13,538 sentences classified as animate from the *19thC BL Books* corpus and ran the experiment on both the pre-1850 and post-1890 language models. The timelines in both cases show an increased substitution of machines for children over the course of the century, while predictions of servant-related words decrease. Children-related predictions overtake servant-related predictions at different points in time depending on the language model, potentially signaling a change in perception of both these groups of people and of machines. The pre-1850 model suggests that the relative probabilities of children and servant terms replacing machines are reversed and diverge slightly after 1860, while the post-1890 model shows an even greater divergence.

²²Due to space limitation, we could not include the evaluation on *humanness* in this paper, we will publish it in an appendix in GitHub. Performance is lower overall, with F-Score in our best-performing method decreasing from 0.78 to 0.68.

²³Terms related to *child* are ‘child’, ‘boy’, ‘girl’, ‘children’, ‘youth’, ‘infant’, ‘boys’; terms related to *servant* are ‘slave’, ‘servant’, ‘service’, ‘slaves’, ‘servants’, ‘maid’.

Does something change in the 1850s to cause this change, e.g., in ongoing debates about factory legislation? Although still experimental, these plots show how the method we propose in this paper could assist historians locate and explore longitudinal trends.




Figure 1: Changing relevance (five-yearly rolling average) of the concepts ‘child’ (red) and ‘servant’ (blue) for sentences originally containing machine-related concepts, according to the BERT predictions of the pre-1850 language model (left) and the post-1890 language model (right), as a function of time. Y-axis is average mean-rank over all sentences.

Contextualized word embeddings have been used in the past to identify cultural and social biases that permeate language (Kurita et al., 2019). In future work, we will explore how biases are reflected differently in language models from different periods, potentially revealing more granular changes in the way that writers in specific genres use the trope of the animate machine. This is relevant not only to nineteenth-century discourses of industrialization, but also to contemporary discussion of the impact of technology in our society, highlighting, for example, threats to social hierarchies or transformations of work environments.

8 Conclusion and further work

We have introduced a new method for animacy detection based on contextualized word embeddings, which efficiently handles atypical animacy. Our case study explores how machines were portrayed in nineteenth-century texts and is motivated by the ubiquitous trope of the living machine; both in the historical discourse of industrialization, and also in today’s discussion of AI and robotics, prefigured by Alan Turing’s famous provocation: ‘Can machines think?’ (Turing, 1950). This work opens many avenues for future research. We intend to explore strategies to derive an animacy value from BERT’s predictions by inspecting the embedding space; study the contextual cues which grant animacy (and how these relate to the neighboring concepts of humanness and agency); and explore the extent to which such atypicalities are conveyed through figurative language. Finally, we will apply all of the above in addressing the historical questions raised in this paper.

References

- Pierpaolo Basile, Annalina Caputo, and Giovanni Semeraro. 2014. An enhanced Lesk word sense disambiguation algorithm through a distributional semantic model. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1591–1600.
- Shiva Bayanati and Ida Toivonen. 2019. Humans, animals, things and animacy. *Open Linguistics*, 5(1):156–170.
- Chris Buckley, Darrin Dimmick, Ian Soboroff, and Ellen Voorhees. 2007. Bias and the limits of pooling for large collections. *Information retrieval*, 10(6):491–508.
- Jinying Chen, Andrew Schein, Lyle Ungar, and Martha Palmer. 2006. An empirical study of the behavior of active learning for word sense disambiguation. In *Proceedings of the Main Conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 120–127. Association for Computational Linguistics.

- Bernard Comrie. 1989. *Language Universals and Linguistic Typology*. University of Chicago Press.
- Michael Connor, Yael Gertner, Cynthia Fisher, and Dan Roth. 2010. Starting from scratch in semantic role labeling. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 989–998. Association for Computational Linguistics.
- William Croft. 2002. *Typology and universals*. Cambridge University Press.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805*.
- Richard Evans and Constantin Orasan. 2000. Improving anaphora resolution by identifying animate entities in texts. In *Proceedings of the Discourse Anaphora and Reference Resolution Conference (DAARC2000)*, pages 154–162.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Mark A Finlayson, Jeffry R Halverson, and Steven R Corman. 2014. The N2 corpus: A semantically annotated collection of Islamist extremist stories. In *LREC*, pages 896–902.
- Mark A Finlayson. 2017. ProppLearner: Deeply annotating a corpus of Russian folktales to enable the machine learning of a Russian formalist theory. *Digital Scholarship in the Humanities*, 32(2):284–300.
- Tao Gao, Brian Scholl, and Gregory McCarthy. 2012. Dissociating the detection of intentionality from animacy in the right posterior superior temporal sulcus. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, (32):14276–14280.
- Mario Giulianelli. 2019. Lexical semantic change analysis with contextualised word representations. Master’s thesis, University of Amsterdam - Institute for logic, Language and computation.
- William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic word embeddings reveal statistical laws of semantic change. *arXiv preprint arXiv:1605.09096*.
- Labiba Jahan, Geeticka Chauhan, and Mark Finlayson. 2018. A new approach to animacy detection. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1–12.
- Folger Karsdorp, Marten van der Meulen, Theo Meder, and Antal van den Bosch. 2015. Animacy detection in stories. In *Proceedings of the 6th Workshop on Computational Models of Narrative*, pages 82–97.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. Measuring bias in contextualized word representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy, August. Association for Computational Linguistics.
- Andrey Kutuzov, Lilja Øvrelid, Terrence Szymanski, and Erik Velldal. 2018. Diachronic word embeddings and semantic shifts: a survey. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1384–1397, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Heeyoung Lee, Angel Chang, Yves Peirsman, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. 2013. Deterministic coreference resolution based on entity-centric, precision-ranked rules. *Computational Linguistics*, (39).
- Brittany Dael McLaughlin. 2014. *Animacy in Morphosyntactic Variation*. University of Pennsylvania.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mante S. Nieuwland and Jos J.A. van Berkum. 2005. When peanuts fall in love: N400 evidence for the power of discourse. *Journal of Cognitive Neuroscience*, (18):1098–1111.
- John Opfer. 2002. Identifying living and sentient kinds from dynamic information: The case of goal-directed versus aimless autonomous movement in conceptual change. *Cognition*, (86):97–122.
- Constantin Orasan and Richard J Evans. 2007. NP animacy identification for anaphora resolution. *Journal of Artificial Intelligence Research*, 29:79–103.

- Lilja Øvrelid. 2008. Linguistic features in data-driven dependency parsing. In *Proceedings of the Conference on Computational Natural Language Learning (CoNLL 2008)*, pages 25–32.
- Rea Peltola. 2018. Interspecies identification in nature observations: Modal expressions and open reference constructions with non-human animate reference in finnish. *Open Linguistics*, 4(1):453–477.
- Matthew E Peters, Waleed Ammar, Chandra Bhagavatula, and Russell Power. 2017. Semi-supervised sequence tagging with bidirectional language models. *arXiv preprint arXiv:1705.00108*.
- Massimo Poesio, Ron Artstein, et al. 2008. Anaphoric Annotation in the ARRAU Corpus. In *LREC*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding with unsupervised learning. *Technical report, OpenAI*.
- Karthik Raghunathan, Heeyoung Lee, Sudarshan Rangarajan, Nathanael Chambers, Mihai Surdeanu, Dan Jurafsky, and Christopher Manning. 2010. A multi-pass sieve for coreference resolution. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 492–501. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using siamese BERT-networks. *arXiv preprint arXiv:1908.10084*.
- Anette Rosenbach. 2008. Animacy and grammatical variation. findings from English genitive variation. *Lingua*, (118):151–171.
- Eric Schatzberg. 2018. *Technology: critical history of a concept*. University of Chicago Press.
- Karen Spark-Jones. 1975. Report on the need for and provision of an ‘ideal’ information retrieval test collection. *Computer Laboratory*.
- Nina Svenningsson and Montathar Faraon. 2019. Artificial intelligence in conversational agents: A study of factors related to perceived humanness in chatbots. In *Proceedings of the 2019 2nd Artificial Intelligence and Cloud Computing Conference*, pages 151–161.
- Alan Turing. 1950. Computing machinery and intelligence. *Mind*, 59(236):433–460.
- Virve-Anneli Vihman and Diane Nelson. 2019. Effects of animacy in grammar and cognition: Introduction to special issue. *Open Linguistics*, 5(1):260–267.
- Ludwig Wittgenstein. 1921. *Tractatus Logico Philosophicus*. Simon and Schuster.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv*, abs/1910.03771.
- Mutsumi Yamamoto. 1999. *Animacy and reference: A cognitive approach to corpus linguistics*, volume 46. John Benjamins Publishing.
- Mutsumi Yamamoto. 2006. *Agency and Impersonality. Their Linguistic and Cultural Manifestations*. John Benjamins Publishing Company.