

Living Machines: A study of *atypical* animacy

Mariona Coll Ardanuy
The Alan Turing Institute
Queen Mary University of London
mcollardanuy@turing.ac.uk

Federico Nanni
The Alan Turing Institute
fnanni@turing.ac.uk

Kaspar Beelen
The Alan Turing Institute
Queen Mary University of London
kbeelen@turing.ac.uk

Kasra Hosseini
The Alan Turing Institute
khosseini@turing.ac.uk

Ruth Ahnert
Queen Mary University of London
r.r.ahnert@qmul.ac.uk

Jon Lawrence
University of Exeter
j.lawrence3@exeter.ac.uk

Katherine McDonough
The Alan Turing Institute
Queen Mary University of London
kmcDonough@turing.ac.uk

Giorgia Tolfo
The British Library
giorgia.tolfo@bl.uk

Daniel CS Wilson
The Alan Turing Institute
Queen Mary University of London
dwilson@turing.ac.uk

Barbara McGillivray
The Alan Turing Institute
University of Cambridge
bmcgillivray@turing.ac.uk

Abstract

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 on 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.

1 Introduction

Animacy is the property of being alive. Although the *perception* of a given entity as animate (or not) tends to align correctly with its biological animacy, discrepancies are not uncommon. These may arise from the ways different entities are unconsciously perceived, or from the deliberately figurative use of animate expressions to describe inanimate entities (or vice versa). Machines sit at the fuzzy boundary between animacy and inanimacy (Turing, 1950; Yamamoto, 1999). In this paper, we provide a basis

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

for examining how machines were imagined during the nineteenth century as everything from lifeless mechanical objects to living beings, or even human-like agents that feel, think, and love. We focus on texts from nineteenth-century Britain, a society being transformed by industrialization, as a good candidate for studying the broader issue.

This paper applies state-of-the-art contextualized word representations, trained using the BERT architecture (Devlin et al., 2019), 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 intent 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. 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 strong 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 which, we believe, has the potential to shed light on the related (but separate) historical process of dehumanization through the language of mechanization, but which lies beyond the scope of the present paper.

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 (human) cognition has been extensively studied in a range of linguistic fields, from neurolinguistics and language acquisition research (Opfer, 2002; Gao et al., 2012) 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 (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; Yamamoto, 1999; Croft, 2002). 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 co-reference and anaphora resolution (Orasan and Evans, 2007; Poesio et al., 2008; Raghunathan et al., 2010; Lee et al., 2013), 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 this 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., 2019). Unlike their static counterparts, they are optimized to capture the representations of target words in their contexts, and are more sensitive to context-dependent aspects of meaning. BERT (Bidirectional Encoder Representations from Transformers, Devlin et al. (2019)) incorporates the latest improvements in language modeling. Through its deep bidirectionality and self-attention mechanism, it 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 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 expression¹ by averaging the animacy of the top κ tokens that have been predicted to fill the mask in the sentence. While this may initially seem to be a circular argument, 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 expression and their corresponding probability scores. We then use WordNet,² 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 a 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 (Basile et al., 2014), we use a BERT-adapted version of the Lesk algorithm, which adopts 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

¹Note that this can be a multi-word expression.

²<https://wordnet.princeton.edu/>.

its disambiguated sense is a descendant of the *living_thing* node.³ 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.⁴

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 English-translated stories that already contained several layers of linguistic annotation (Finlayson et al., 2014; Finlayson, 2017). The authors enriched the texts with animacy annotations at the level of co-reference chains and of their referring expressions. They reported a near-perfect inter-annotation agreement (Cohen’s $\kappa = 0.99$). We restructured their data to make it compatible with our approach, which takes full sentences as inputs. This process resulted in a new dataset (henceforth *Stories* dataset) consisting of 18,708 sentences, each of which contains a target expression annotated with animacy (see some examples in Table 1).⁵

Target expression	Sentence	Animacy
noise of the leaves	“No it was only the noise of the leaves. ”	0
need	He whistled and shouted with a mighty hiss, in a truly powerful voice: “Magic steed, horse of my need! ”	0
long time or a short time	Nikita fought him for a long time or a short time ; in any event, he defeated him.	0
merchant’s daughter	The merchant’s daughter took the letter, burst into tears, and prepared to go, taking also her maidservant; and no one could distinguish this maid from the merchant’s daughter , they were so like each other.	1
twelve-headed dragon	They began to fight, and after a long struggle or a short struggle, Prince Ivan defeated the twelve-headed dragon , slashed open his trunk, and found the coffer in his right side.	1
She	She wept, but still could not awaken Prince Ivan.	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.⁶ We decided to create our own dataset (henceforth *19thC Machines* 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 machines 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*).⁷ 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’

³Pronouns are not represented in WordNet. Because of this, we also automatically assign animacy to a predicted token if it is included in a closed list of pronouns typically referring to persons, such as ‘she’, ‘him’, and ‘yourselves’.

⁴Code and experimental results are available at <https://github.com/Living-with-machines/AtypicalAnimacy>.

⁵Unfortunately, we were not able to reproduce the same number of target expressions as are reported in Jahan et al. (2018), but we provide the code used to generate our datasets from the original corpus. See Table 3 for a summary of the *Stories* dataset.

⁶Note that this is based on observation, as the number of atypical cases is not provided in Jahan et al. (2018).

⁷This 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).

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.⁸ In most sentences, machines are treated as inanimate objects. We therefore employed a pooled strategy⁹ 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 six annotators.¹⁰ 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 annotators to stick to the most literal meaning and avoid metaphorical 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, especially in a historical context.

The second task was more straightforwardly related to determining the animacy of the target entity, given the same 100 sentences. We asked 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 and negative groups 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 true for sentences containing similes or metaphors that liken machines to humans, animals, or systems.

These 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 the sentences 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.*
- (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) is likened to a machine to suggest that their agency or animacy has been reduced. Some annotators detected an implied inanimacy of the machine, which had the rhetorical effect of rendering the humans inanimate. Conversely, for others, the comparison conjured a kind of automaton, a human-machine-hybrid, and therefore a machine made animate through its association with the human.

⁸The 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’. We are preparing a publication on the lexicon expansion procedure that will allow us to assess its impact on downstream tasks.

⁹Pooling is an established evaluation strategy for information retrieval systems (Spark-Jones, 1975; Buckley et al., 2007).

¹⁰A combination of computational linguists, historians, literary scholars, and data scientists.

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 common discrepancies. After discussion, it was decided that a machine would be tagged as animate if it is described as having traits 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**, which would be true if the machine is portrayed as sentient and capable of specifically human emotions, and false if it used to suggest some degree of dehumanization.¹¹ The focus of this paper is on detecting atypical animacy, and we plan to further study the phenomenon of humanness (or lack thereof) in animate objects in future work.

A new batch of 600 unseen sentences was sent to the same annotators, who were asked to document their decisions, especially the ones they were most uncertain of. A subset of 100 sentences were annotated by five domain experts, for which the Krippendorff α was of 0.74 for animacy and 0.50 for humanness. The remaining annotations were produced by subgroups of two to three annotators. 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 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. *Stories* 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).¹² *19thC Machines* consists of 593 sentences with 13 unique target expressions, which can be either animate or inanimate, depending on the context.¹³ As discussed, the disagreement was quite high in comparison, indicating that detecting atypical animacy can be a very semantically complex problem, in particular in highly figurative language.

5 Language Models

In our experiments, we used the ‘BERT base uncased’ model and tokenizer as contemporary model,¹⁴ hereinafter referred to as *BERT-base*. In addition, in order to investigate pattern changes over time, we

¹¹For us, ‘animacy’ encompasses ‘humanness’: all sentences that are tagged as cases of humanness are also annotated as animate, but not the other way around. The term ‘humanness’ is used differently in other research fields, for instance when evaluating the performance of a chatbot (Svenningsson and Faraon, 2019).

¹²Previous work on animacy in Dutch reported a similarly high κ agreement score of 0.95 (Karsdorp et al., 2015).

¹³Some sentences were discarded because they mentioned a human boiler instead of a boiler engine. Only 13 of the curated machine words appear in the annotations, due to the comparatively lower frequency of the other word forms in the corpus. There are 201 sentences in which the machine has been tagged as animate, of which 81 are also instances of humanness. The *19thC Machines* dataset is publicly available at <https://doi.org/10.23636/1215>.

¹⁴<https://github.com/google-research/bert>.

	#Sentences	IAA	UniqueExpressions	Train/Test	#Animate
<i>Stories</i>	18,708	Cohen’s κ 0.99	5,340	13,096/5,612	11,666
<i>19thC Machines</i>	593	Krippendorff’s α 0.74	13	178/415	201

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

also fine-tuned *BERT-base* on the *19thC BL Books* dataset, split into four 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.¹⁵ 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 books¹⁶ and tokenized them using the original BERT-base tokenizer as implemented by HuggingFace¹⁷ (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.¹⁸

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 compare. 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 illustrates the importance of having language models that adequately reflect language 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.

6 Experiments and results

6.1 Baselines and Evaluation Metrics

We provide two different types of competitive baselines: a classification approach and a sequence tagging approach. As a reference, we also provide the results obtained by selecting the most frequent class, which

¹⁵While the data distribution for fine-tuning was decided on largely by the number of tokens, these periods also work well in representing distinct cultural eras. For example, the pre-1850 dataset sets apart the first industrial revolution from later developments in Britain. Likewise, 1890-1900 is set off, especially in literary terms, by the emergence of ‘modernist’ sensibilities and the questioning of class and gender hierarchies associated with the term ‘*fin de siècle*’.

¹⁶We 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 the *syntok* library: <https://pypi.org/project/syntok/>.

¹⁷<https://github.com/huggingface/transformers>.

¹⁸We 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.

is the animate class in the *Stories* dataset and the inanimate class in the *19thC Machines* dataset.

- **Classification approach.** Karsdorp et al. (2015) and Jahan et al. (2018) treat animacy as a classification problem. We report the performance of three different supervised classifiers trained on examples annotated with a binary label: two SVMs, one using tf-idf and the other using word embeddings as features, and a BERT classifier (from now on *SVM TFIDF*, *SVM WordEmb*, and *BERTClassifier*, respectively).¹⁹ All classifiers are trained both at the target expression level (*targetExp*), and 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*). For the *Stories* dataset, we used the training set of 13,096 instances to find the optimal hyperparameters and thresholds, which we then applied to the test set. For the *19thC Machines* dataset, due to the small amount of domain-specific data, we trained the classifiers on the concatenation of the *Stories* and *19thC Machines* training sets (resulting in a total of 13,274 instances) and used the *19thC Machines* training set (178 instances) to find the optimal thresholds.
- **Sequential tagging approach:** Moreover, we experimented with LSTMs (Hochreiter and Schmidhuber, 1997; Akbik et al., 2019) to predict the animacy of target tokens (*SeqModel: LSTM*). In this scenario, we enclosed the target words with ‘[TARGET]’ token — so the classifier knows which words to tag — and labeled them as 0 or 1. We assigned ‘O’ to all tokens outside of the ‘[TARGET]’ tokens. We used embeddings extracted from a BERT model as input for a Sequential LSTM tagger. The model is trained and tested on the *Stories* dataset. To assess performance on the *19thC Machines* dataset, we further fine-tuned the *Stories* model on the sentences from the *19thC Machines* training set.

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.

6.2 Experimental results

Table 5 reports the performance of the different baselines and methods on the *Stories* and *19thC Machines* datasets. The sequence tagging approach (*SeqModel: LSTM*) and classifiers (especially *BERTClassifier*) based on the target expression alone (*targetExp*) are the best performing methods in the *Stories* dataset. Interestingly, the classifiers’ performances become worse when more context is added (*targetExp + ctxt*), and even more so when the target expression is masked (*maskedExp + ctxt*). Our method (*MaskPredict: BERT-base*) does not use the target expression as a feature at all: it relies solely on context. In fact, adding context (one sentence to the left and to the right, *MaskPredict: BERT-base + ctxt*) improves the F-Score from 0.72 to 0.81. This analysis shows that the target expression is a very indicative feature of conventional animacy. And yet, the good performance of our context-based method proves that animacy is not only determined at the level of the target expression, but informed by contextual cues as well.

All baselines perform worse on the *19thC Machines* dataset. As mentioned, in this case we trained the classifiers on the concatenation of the *Stories* and the *19thC Machines* training sets, and optimal classification thresholds were found by maximizing F-Score on the *19thC Machines* training set. The sequential model was trained on the *Stories* training set and fine-tuned on the *19thC Machines* training set. Our approach yields consistent performance on both datasets, showing the advantage of its unsupervised context-dependent architecture, as opposed to the other competitive baselines, which heavily rely on the availability of in-domain training data.

We show in Table 5 the results of using different BERT models in our masking approach: *fit19thCBERT* uses the appropriate fine-tuned BERT model of the period to which each sentence belongs (e.g. a sentence from a book published in 1856 is processed using the BERT model fine-tuned on

¹⁹Both SVMs use a linear kernel with standard parameters, from the *Scikit-learn* package. We have used the following *Scikit-learn* wrapper for training the BERT classifier, again with default parameters: <https://github.com/charles9n/bert-sklearn>.

	<i>Stories</i>				<i>19thC Machines</i>			
	Precision	Recall	F-Score	Map	Precision	Recall	F-Score	Map
Most frequent class	0.31	0.5	0.383	0.623	0.336	0.5	0.402	0.318
SVM TFIDF: targetExp	0.911	0.893	0.902	0.928	0.696	0.713	0.704	0.474
SVM WordEmb: targetExp	0.927	0.919	0.923	0.954	0.694	0.711	0.702	0.499
BERTClassifier: targetExp	0.951	0.948	0.949	0.985	0.698	0.715	0.706	0.51
SVM TFIDF: targetExp + ctxt	0.734	0.739	0.737	0.859	0.688	0.71	0.699	0.651
SVM WordEmb: targetExp + ctxt	0.758	0.742	0.75	0.876	0.728	0.531	0.614	0.481
BERTClassifier: targetExp + ctxt	0.931	0.926	0.929	0.978	0.695	0.721	0.708	0.721
SVM TFIDF: maskedExp + ctxt	0.674	0.677	0.675	0.804	0.592	0.6	0.596	0.498
SVM WordEmb: maskedExp + ctxt	0.674	0.678	0.676	0.809	0.518	0.52	0.519	0.339
BERTClassifier: maskedExp + ctxt	0.855	0.852	0.854	0.951	0.687	0.696	0.692	0.603
SeqModel: LSTM	0.952	0.948	0.95	0.949	0.697	0.719	0.708	0.482
MaskPredict: BERT-base	0.739	0.703	0.72	0.848	0.719	0.742	0.73	0.74
MaskPredict: BERT-base +ctxt	0.839	0.774	0.806	0.892	0.758	0.778	0.768	0.795
MaskPredict: fit19thBERT +ctxt	–	–	–	–	0.758	0.775	0.766	0.777
MaskPredict: early19thBERT +ctxt	–	–	–	–	0.799	0.773	0.786	0.784

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

books published between 1850 and 1875), whereas *early19thCBERT* uses the BERT model fine-tuned on the first half of the nineteenth-century. While the latter provides the best results, the differences are not statistically significant.

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 depictions of machines *as if* they were alive raises important questions about the relationship between humans and machines that 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 discussed in section 4.2.2, we consider entities to be animate if they are given attributes and faculties that are characteristic of living entities; they are attributed *humanness* if they are portrayed as sentient and capable of specifically human emotions.²⁰ The latter is loosely 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, or speaker empathy.

All baselines and methods are clearly worse in predicting humanness than animacy.²¹ In the case of the supervised baselines, this is possibly due in part to the very limited amount of annotated data for this concept (81 instances of animacy with humanness, 120 instances of animacy without humanness). Besides, the lower agreement between annotators in detecting humanness (Krippendorff α of 0.50) suggests that the task is more subjective. Table 6 reports the top predicted tokens by the pre-1850 BERT model in sentences where animated machines are attributed or negated humanness. This experiment shows how language models can hint a social bias embedding in nineteenth-century discourse.²² 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 nineteenth-century tendency to use these words in discourses that confer diminished human qualities on those people.

²⁰Example 2 shows 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[s] one wonder why the poor crazy **machine** is thought worthy of being put together again, with infinite pains, and wonderful science.’

²¹Performance for ‘humanness’ is lower overall, with the F-Score of our best-performing method decreasing from 0.79 to 0.63. Due to space limitations, we could not include the evaluation on *humanness* annotations in this paper, but it can be found on our Github repository: <https://github.com/Living-with-machines/AtypicalAnimacy>.

²²Contextualized word embeddings have been used in the past to identify cultural and social biases that permeate language (Kurita et al., 2019).

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

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

Language is not static, it evolves continuously (Hamilton et al., 2016; Kutuzov et al., 2018). Social and technological changes are paralleled by changes in the language used to describe them. In the nineteenth century, who or what was performing work was changing dramatically, with different groups of people entering and exiting the (free and unfree) labor pool and becoming key parts of the workforce at different times.

Sentence	≤ 1850 BERT	Contemporary BERT
Our sewing machines stood near the wall where grated windows admitted sunshine, and their hymn to Labour was the only sound that broke the brooding silence.	<i>girls, women, men, boys, children, machines, parties, females, mothers, hands</i>	<i>machines, tables, stalls, machine, rooms, sheds, carts, girls, boxes, women</i>

Table 7: Top predictions for the word ‘*machines*’, using the pre-1850 and contemporary BERT models.

Such changes are often implicitly captured by a language model. Table 7 gives a concrete example, related to the activity of ‘sewing’, which was originally performed by humans (often women and children) but increasingly mechanized. This shift is reflected in the difference between the predictions of the pre-1850 and the contemporary language model. In future work, we will explore how language models encode biases and attitudes, and therefore can be used to scrutinize subtle but critical shifts in the understanding of technology and mechanization. This is relevant not only to nineteenth-century discourses of industrialization, but also to contemporary discussion of the impact of technology in our society.

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 time-sensitive 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.

Acknowledgements

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 grant AH/S01179X/1), 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 (EPSRC grant EP/N510129/1).

References

Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. 2019. FLAIR: An easy-to-use framework for state-of-the-art NLP. In *Proceedings of the 2019 Conference of the North Amer-*

- ican Chapter of the Association for Computational Linguistics (*Demonstrations*), pages 54–59, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- 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. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- 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, Jeffrey 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.
- William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, Berlin, Germany, August. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- 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.
- 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.