Discovering the arrow of time in machine learning

J. Kasmire^{12,*} and Anran Zhao¹²

¹UK Data Service, Essex, UK

²Cathie Marsh Institute, University of Manchester, Manchester, UK

Abstract. Machine learning (ML) is increasingly useful as data increases in both volume and accessibility. Broadly, ML uses computational methods and algorithms to learn to perform tasks, such as categorisation, decision making or anomaly detection through experience and without explicit instruction. ML is most effective in situations where non-computational means or conventional algorithms are impractical or impossible, such as when the data are vast, complex, highly variable and/or full of errors [1, 2]. Thus, ML is useful for analysing natural language, images, or other types of complex and messy data that are now available in ever-growing and impractically large volumes.

Some ML methods are suitable for analysing explicitly ordered or time-dependent data, although these tend to be less tolerant of errors or data asymmetry. Nevertheless, most data has at least an implicit order or temporal context. Selecting an appropriate ML algorithm depends on the properties of the data to analyse and the aims of the project as algorithms vary in the supervision needed, tolerable error levels, and ability to account for order or temporal context, among many other things. Using non-temporal ML algorithms may obscure but not remove the order or temporal features of the data, potentially allowing the hidden ‘arrow of time’ to affect performance.

This research takes the first step in exploring the interaction of ML algorithms and implicit temporal representations in training data. Thus, this research addresses the suitability of ML for analysing the kind of data that is accumulating daily from every social media platform, Internet of Things device, businesses report, transport tracker or other modern data source. Two supervised ML algorithms are selected and described before experiments are run to train those ML algorithms to perform automatic classification tasks under a variety of conditions that balance volume and complexity of data. In this way, the experiments explore whether more data is always better for ML models or whether implicit temporal features of data can influence performance.

The research shows that ML algorithms can be sensitive to subtle or implicit temporal context with consequences for the accuracy in classification tasks. This means that researchers should carefully consider the implications of time within their data when selecting appropriate algorithms, even when the algorithms of choice are not expected to explicitly address order or temporal context.

*e-mail: julia.kasmire@manchester.ac.uk

1 Introduction

Machine learning (ML) is a broad church. Supervised ML is very popular for exploring the relationship between clearly defined input and output variables. ML methods include: linear regression which handles continuous output by finding a line or plane that best fits the data, logistic regression which makes classification predictions for binary output, and support vector machines which classify by finding a hyperplane or boundary that maximises the margin between two or more classes. As an example, a supervised ML algorithm might be trained to classify incoming email based on the sender, subject line or message content (input) that are accurately labelled as SPAM or NOT-SPAM (output). In contrast, unsupervised ML models attempt to draw inferences and find patterns from unlabelled data sets, such as finding neighbourhood clusters within demographic information. Common clustering methods include k-means clustering, which groups input into a pre-determined k number of groups, and hierarchical clustering, which arranges observations into tree-like graphs called dendrograms. For instance, an unsupervised ML algorithm might be asked to classify research articles into disciplines according to their similarity. There are also semi-supervised approaches to machine learning, so the distinction is not always clear cut.

ML algorithms and tasks are diverse, but there are some general features that make ML useful. For example, supervised ML algorithms allow data to be classified when the volume or complexity exceeds the capacity of manual classification or non-learning algorithmic exploration while unsupervised ML algorithms are extremely useful for real-time data exploration and pattern detection in data that changes too quickly or too unpredictably for conventional exploration. ML (supervised and unsupervised) are often understood to perform better with more data. Nevertheless, there are trade-offs as unrepresentative data can lead to 'over-fitting' and poor generalisability, which is when a ML algorithm performs very well on the training data but poorly on other data. ML is also typically understood to be tolerant of errors in the data, with some approaches showing acceptable performance despite error rates up to 39 percent [1, 2]. ML methods usually require substantial calculation, but this can be managed through careful selection of algorithms and setups [3]. For these reasons, ML is understood to be a good tool for the large, complex and messy data that is ballooning too rapidly for manual or conventional means of exploration, classification or analysis.

This paper begins with some background on time, machine learning and the intersection of the two. Following this is a clear description of the research question to be addressed and the methods and data used to address it. Next is the research results and a discussion of how the results relate to the research question. Finally, a brief discussion of what the results might mean for the wider ML research context and what potential next steps for this research topic.

2 Theory and Background

2.1 The "arrow of time"

The "arrow of time" concept comes from 1927 theoretical physics lecture and describes the apparent one-way flow of time observed in irreversible processes at macroscopic levels. Although concluding that the arrow of time is a consequence of entropy, it was also described as a matter of perception, reasoning and statistical approaches [4]. Put another way, the one-way flow of time was understood to have a thermodynamic origin and psychological consequences. Thus, the arrow of time applies to irreversible processes that are both physical (e.g. an egg cannot be unscrambled) and economic or social (e.g. decisions cannot be undone and done differently instead), often known as path dependency or lock-in [5].

More recent work on the arrow of time maintains that people perceive and experience time as flowing in one direction, but highlights how the concept has continued to develop.

Figure 1. Path dependency illustrated in a simple graph; travel between points only moves in one direction (toward the future) after which previously accessible points (alternate futures) are longer available.

Entropy is now linked to information and complexity [6] although the concepts and their applications are not interchangeable [7]. Currently, time's arrow is thought to be a key driver of thermodynamic processes, biochemical processes like evolution via natural selection, and informational processes in both the natural world and cultural systems like language [6, 8].

Setting aside the complicated concepts of information (especially as they related to entropy and thermodynamics), the arrow of time can be found in language; utterances¹ are spoken or written at a specific point in time. Further, language sequentially orders its components (symbols, semantic features, words, etc.). These temporal aspects of language are not always strict, allowing some meaning to be communicated 'out of order'. Still, exactly what is communicated may not be "the same" as what would have been communicated if the utterances had been experienced at a different time or in another order. For example, classic literature is read and reread continuously with frequent new interpretations, all of which are probably unlike the contemporary interpretations from when the literature was originally available. Another example would be how readers may or may not struggle to understand the plot of a book series were they to read the sequel before the original. Likewise, some sentences can be understood more clearly than others when the words are jumbled; 'eats mouse cheese' is relatively clear while 'eats man shark' remains ambiguous.

The difference in interpretations or in interpretability described above derives from shared knowledge or common ground, which accumulates gradually as a consequence of experience [9]. Once learned, common ground allows communicators to recruit previously gained knowledge into helping interpret current utterances. Thus, books in a series can be read out of order if they deal with well-known topics, plots, tropes, events or characters, even if they are not always interpreted in the same way. Likewise, individuals could use common ground to assume that the mouse is eating the cheese in the first jumbled sentence (rather than the cheese eating the mouse) but might not be able to make assumptions about the second without more context (as humans and sharks are both capable of eating each other).

But common ground is also built between individuals through repeated interactions [9] and within a single interaction as participants take turns to speak and check their understanding [10]. Meaning is even built within a single utterance as linguistic input is incrementally processed word-by-word [10], which can create problems when an incremental understanding is violated by a 'garden path' sentence [11]. As a consequence, language use and comprehension are affected by prior exposure to language, most obviously in childhood language development where studies show neurological differences as a result of language exchanges and linguistic turn-taking [12].

¹Utterances are usually understood to be exclusively spoken, but this research counts tweets as utterances.

2.2 Time in ML

ML includes supervised and unsupervised methods and can perform tasks such as prediction, classification, clustering, search and retrieval or pattern discovery [13]. Familiar examples of applying ML to temporal or otherwise ordered data include predictive text functions that suggest complete words based on the sequence of letters entered so far and storm warning systems based on weather data. Comparative studies have found that different ML algorithms perform differently on time-series classification [14] and explicitly temporal tasks [15]. Neural network ML models generally outperformed classic linear analysis techniques and an analysis of 13 more modern ML time series methods showed multilayer perceptron and then Gaussian processes were the best across a range of categories and data features [16]. Clearly, ML can be very sensitive to time, with some ML algorithms working with time better than others. Consequently, there are well-established data sets, benchmarks and approaches for temporal data ML research.

Specifically time-sensitive ML approaches have additional problems or weaknesses that other ML methods can avoid. The problems of over-fitting and generalisability are more complicated because the training data must be representative in relation to time as well as all the other ways that it must be representative. The added dimension makes time-sensitive ML approaches less tolerant of errors or missing values than other ML methods and with greater risks of over-fitting to one time frame. This means that time-sensitive ML algorithms typically have additional data requirements, including that it must be organised into uniformly sized time intervals. Although not insurmountable, the stricter data requirements mean that time-sensitive ML is more difficult to apply, and so less popular, unless the data to be analysed or research question are specifically temporal.

Most of the data of interest for ML exploration and analysis is vast, complex and messy. Importantly, that data usually also has a temporal context. First, much of that data is time-stamped or otherwise ordered, although this is not always seen as important. Second, data is unevenly distributed over time as users, platforms and records are added or shared, seriously complicating temporal representativeness. Additionally, features within that data are also unevenly distributed over time. To illustrate, 10-15% of medical diagnoses at any given point in time are estimated to be wrong [17], although which are erroneous is not known at the time. Some incorrect diagnoses are identified through later testing, condition progression, and medical innovations or discoveries, but the re-diagnosis is also subject to the 10-15% misdiagnosis rate. Further complicating the issue, a correct diagnosis according to the diagnostic criteria of the time could be a misdiagnosis according to later criteria. This becomes clear when we recognise that autism is not childhood schizophrenia [18] or that status lymphaticus is not a real condition [19]. Third, and perhaps most importantly, new data is not created independently of existing data. Instead, data is often generated as the result of learning from previously generated data. For example, users might change what content they share on a social media platform after observing what content is popular or well-received. When considered together, all of this suggests that the data that invites ML analysis is likely to be time-asymmetrical in ways to which many ML algorithms are not well suited.

Researchers have begun to ask how time and ML algorithms interact, both to correct for problems caused by temporal features as well as to capitalise on these features. For example, sentiment analysis of Twitter data is well-established as a ML problem of interest, but accounting for temporal aspects of the data allowed researchers to analyse sentiment in 'real time' tweets [20, 21] while combining sentiment analysis of tweets with geo-tagging enabled researchers to build a spatial and temporal map of New York [22]. Further, using social interactions with temporal features allowed researchers to significantly improve stress detection in young people [23], potentially allowed problems to be detected early enough

for effective intervention. Researchers have clearly noticed the ways that ML algorithms are affected by time and are beginning to investigate these interactions and even to put them to practical use. Thus, it is important to understand which algorithms are sensitive to implicit temporal features and in what ways these temporal sensitivities affect the ML performance.

3 Research Question

This research addresses several related questions. First, is ML performance influenced, biased or complicated by data with subtle or implicit representations of the flow of time? Put another way, will ML algorithms that assume time-independence or time-symmetrical data perform differently when that assumption is violated? This could have consequences for how an appropriate ML algorithm is chosen according to the data available or the task to be performed.

A second problem would revolve around identifying, quantifying or managing the ML effects of data sets that only implicitly capture the one-directional flow of time. Of course, researchers could discard or reshape the data to meet the more stringent requirements of explicitly time-sensitive ML methods, although this would lose many of the benefits of ML approaches. For example, there may be ways to automatically weight or present the data during training, which may help balance the risks between over-fitting and implicit temporal distortion. Alternatively, iterative learning could be introduced so that algorithms learn from already-trained models as well as from the data itself. By approximating the way that the people generating the data learn from past data, such iterative learning may account for the time-asymmetry of the data.

The research seeks to answer whether or not the ‘arrow of time’ can be observed in ML algorithm performance under various training regimes on data with only implicit inclusion of time. In effect, the research question asks whether ‘more data is always better for ML performance?’ or whether ‘data that spans significant time frames will capture an implicit arrow of time that can distort ML performance?’. To answer that, this research collects tweets from 9 different time frames and trains ML models to predict which time frame those tweets come from. the ML models are trained on subsets of 3 time frames, 6 time frames and all 9 time frames. If the accepted ML concept that ‘more training data is always better’ is true, then the performance should be equal or better when the models are trained on 6 or 9 time frames in comparison to when they are trained on only 3 time frames. However, if the data does capture some implicit temporal features that are not well accounted for in these ML models and that disrupt its learning, then the performance of the models trained on more time frames should be worse than those trained on fewer time frames.

Importantly, this research is about a classification task and *not* about a time-series prediction task. Thus, the ML models are not explicitly considering temporal features or order, and time is only present within the model in an implicit way, embedded within the data. In essence, this means that the research question can be rephrased as ‘are non-temporal ML models influenced by purely implicit temporal features?’.

4 Research Method

First, two supervised ML algorithms are explored, compared and contrasted. Unsupervised ML algorithms are also interesting, but as a first step in the exploration of implicit temporal representations on ML performance, this research focuses on supervised ML only. Second, the data collection, preparation and application of the two selected methods are described in

detail and a GitHub repository with the data sets and code is linked [24]. Finally, the experimental conditions and results are discussed, with clear descriptions of the specific training and testing regimes for each model and the accuracy and precision for each.

Two ML models will be explored in this research: naive Bayes classifiers (NBC) and recurrent neural networks (RNN). Both are supervised ML models, meaning that they learn to perform their tasks through explicit training. This is done by providing the ML models with one or more training data sets that consists of at least one 'input' and at least one 'output' so that they can learn to correctly predict the output from the input. Whether or not the models have learned to do this correctly is achieved through one or more test data sets that have the same structure as the training data sets but that do not contain the same entries. Generally, supervised ML models achieve this by taking one large data set and dividing it so that the majority is allocated to the training data set and the rest to the test data set with no overlap. 'Correct' learning is generally understood as the most accurate and precise performance over the test data as possible, although some research projects may instead seek to specifically focus on minimising either false-positives or false-negatives.

While both models take the same basic approach of learning from explicit training sets before being tested, exactly how they learn and operate is unique to each model.

4.1 Naive Bayes Classification and Recurrent Neural Networks

NBC models learn to associate the inputs with the outputs by creating and subsequently adjusting simple weighted associations. For example, a NBC might be given a training data set of emails (input) which are tagged as 'SPAM' or 'NOT-SPAM' (output). Through training, the NBC would first associate the words, phrases or structures of the emails with their 'SPAM' or 'NOT-SPAM' tags and would then refine the weights of these associations. By the end of training, the words, phrases and structures that are only or mostly found in 'SPAM' emails (or vice versa) would be used to predict whether a new email (input in the test data set) is likely to be SPAM or not. Many NBC models seek to be as accurate and precise as possible, although some models may instead seek to focus on ensuring that all SPAM is correctly identified as SPAM even if this means that some NOT-SPAM are also labelled as SPAM.

RNN and other neural network models also learn to predict an output by learning from the training data by feeding activation through a network or graph of connected nodes and layers. The input is provided to the first layer which then activates the subsequent layer according to the connections and weights, proceeding in this way with the activity of the final layer being interpreted as the predicted probability of output classes. The model then compares the predicted output to the real output in the training data and uses the difference to adjust the network connections and weights so as to reduce the error. Simple or feed-forward neural networks only accept a single input at a time so the information only flows in one direction but RNN models allow layers to feed their activation back to the previous layer allowing the input at any given step to contain some of the input from the previous step. For example, a feed-forward neural net might be fed a word one letter at a time with the aim of predicting the next letter; if fed 'n', 'e', 'u', 'r', and 'a', it would predict 'n' (statistically the most common letter to follow 'a'). But a RNN with the same task and inputs would predict 'l' because it would look at *all* of the letters in the input sequence rather than only the most recent.

For the research described in this article, the data set has two parts: the cleaned full text of tweets and when the tweet was first made (in year-month format). Both models will be passed the tweets as input and are asked to predict when the tweet was made as output. As a first exploration of the research question, both models will be judged on simple accuracy with no focus on the issue of false positives or false negatives. However, the research question requires more than simply training and testing the models on all of the data. Instead, the ML

models are trained and tested on subsets of the test data that encompass 3 time periods, 6 time periods or 9 time periods. This allows the researchers to see if the models perform differently when the training sets are larger but also contain more implicit temporal features. In this way, the models can help illuminate whether more data is better (regardless of the temporal nature of that data) or whether language change or specific events captured within the time frames distort the way ML models perform.

4.2 Data Collection and preparation

4.2.1 Collection

The data set is created via the Tweepy[25] package and Twitter application programming interface (API). The Tweepy *search_full_archive* and *cursor* methods take a start date, stop date, search keywords and an environment parameters as input and returns up to a specified maximum of truncated tweets with tweet ID, sender ID, tweet sent time and other details as output. The Tweepy *status_lookup* method then iterates over the IDs of the truncated tweets, 100 at a time, updating to a full-text version with the option to add other relevant information (e.g. whether the tweet was in reply to another, whether it was favourited, from what source the tweet was sent, etc.) as needed.

For this research, the Tweepy methods described above were used to create nine batches of 1200 tweets. Each batch was collected with the keywords ‘vaccine’ and ‘UK’ and each covered a six-month interval between April 2017 and April 2021 for a total of 10800 tweets. For example, one batch retrieved 1200 tweets published between 1 April and 20 April 2021 containing the keywords ‘vaccine’ and ‘UK’.

It is important to note that this method requires a Twitter academic research Twitter developer account, which equips the user with the necessary API tokens to access Twitter materials and contents that are older than seven days. Applying for an academic developer account is free but requires explaining the purpose and planned use of the Twitter materials. Further, the *search_full_archive* developer environment within Tweepy must be set and labelled in the Twitter developer account portal.

4.2.2 Preparation

The Tweepy method above returns a specific ‘Twitter search result’ type object within Python and includes many details that are not considered important for this research question or are not in the desired format or order. Thus, the output from each search is transformed into a table that retains only the ‘full_text’ and ‘created_at’ columns, with all others being dropped. The nine individual tables are then concatenated to create a single table with all of the data, after which the ‘created_at’ values are re-coded from date and time (down to the second) to simple year-month values. The combined data is then saved and exported as a .csv in preparation for Natural Language Processing (NLP) steps.

The NLP steps begin with the *clean* method from ‘Preprocessor’[26] which removes URLs, hashtags, mentions, and reserved words (RT for retweet and FAV for favourite), emojis and smileys. Then, the *sub* method from ‘Re’[27] removes punctuation and any words that appear in the English stopwords list from the ‘NLTK.corpus’ package[28] (e.g. determiners, prepositions, pronouns, etc.). The tweets are then passed through the *word_tokenize*, *pos_tag*, *get_wordnet_pos* and *lemmatize* methods from the ‘WordNet’ functions from the ‘NLTK’ package. This changes words to simple lemma versions according to how they were used in the original text so that ‘further’ and ‘farthest’ are both changed to ‘far’, ‘boxes’ is

changed to ‘box’, and ‘recording’ is changed to ‘record’ if it was originally used as a verb but to ‘recording’ if it was originally used as a noun.

Finally, the data set is copied and edited into three new data sets. Data set 1 contains only entries from the three most recent time periods (2020.4 to 2021.4) and data set 2 contains only entries from the six most recent time periods (2018.10 to 2021.4). Data set 3 contains all entries covering the full nine time periods (2017.4 to 2021.4). These three data sets are used to explore how the ML model performance is influenced when the training data is larger but also more temporally complicated. Each is converted to a list within python and separated into test and train data sets with an 80-20 split.

4.3 Modelling

The NBC model is straightforward; lemmatised tweets are converted to textblobs via the *TextBlob* method from the ‘textblob’[29] package. Following this, three empty NBC models are trained by applying the *NaiveBayesClassifier* method (via ‘textblob.classifiers’ tool from ‘textblob’ package), one each on the training portions of Data sets 1, 2, and 3. The trained NBC models are then tested on the corresponding test portions of the same data sets, so that the NBC trained on the train data from Data set 1 is tested on the test data from Data set 1 and so on.

Figure 2. The RNN model has five layers, each of which serves a specific function.

The RNN is more complicated but mostly uses the ‘Keras’[30] and ‘Tensorflow’[31] python libraries. First, the lemmatised data sets are encoded as tensorflow-type data set through *tf.data.Dataset.from_tensor_slices* function and the ‘created_at’ dates are converted to binary features representing each time period through the *get_dummies* method in Pandas library. Following this preparation step, the training and testing data sets are attached to the

RNN with the *batch* and *prefetch* functions and is passed through all 5 layers of the RNN 20 times (epochs = 20). The RNN, and its 5 layers, is depicted in Figure 2. The first layer uses *experimental.preprocessing.TextVectorization* and *adapt* methods from ‘Tensorflow’ to convert the input to a word index sequence that includes the frequency of each word within the specific data set. The second layer transforms the word index sequences to trainable vector sequences which, with sufficient training, turns different word sequences with similar meanings into similar vector sequences. The third layer is a *tf.keras.layers.Bidirectional* wrapper, which propagates the input forward and backwards through the RNN with long short-term memory (LSTM) architecture. The fourth and fifth layers convert the trained vector sequences into prediction through *tf.keras.layers.Dense* functions. These prediction vectors have as many elements as there are classes (e.g. Data set 1 with 3 time periods has a three-element layer) with each element in this layer belonging to one and only one of the possible output values. The element with the highest score is interpreted as the predicted probability of output classes. The fourth layer applies the common *ReLu* activation function with 64 nodes, while the fifth and final layer adopts the *softmax* activation function to perform the final multi-class prediction. All five layers are compiled using a *categorical_crossentropy* loss function and an *adam* optimiser function.

5 Results and Discussion

Figure 3. Word clouds for each time period, labelled by the closing date of that time period.

It may be useful to begin with a sense of the language in each time period. To this end, Figure 3 shows the most frequent words of each time period as word clouds where the frequency of the word within the data is represented by the size of the word within the word cloud. Clearly, each of the time periods is unique as no two produce the same word clouds. Moreover, they show that some of the changes from one time period to the next may be linked to time in ways that are more or less obvious. For example, the most recent time periods feature words like ‘covid’ or ‘covid 19’ that do not feature in older time periods; this

is not surprising given that the covid-19 pandemic has had a very large impact on Twitter discussions around vaccines. At the same time, 'hvp', 'girl' and 'mmr' feature strongly in the 2018.04 time period but are not so prominent in other time periods. A bit of research shows that 2018 saw the HPV vaccine offered to preteen and teenage boys as well as girls, which surely generated some discussion. 2018 also featured unusually high rates of infection for measles, mumps and rubella. Neither of these vaccine-related events in 2018 was so memorable that the authors didn't have to go looking for an explanation, but were apparently sufficient to have an impact on Twitter discussions.

Having explored the language in the data a bit, it is important to note that both the NBC and RNN models demonstrated good performance on predicting the time class of the tweets; the least accurate performances of any ML model in any experimental condition was just below 80%. This means that both of these ML model approaches are capable of capturing the characteristics of the tweets and using those characteristics to predict which time period category the tweets came from. It is important to reiterate that this task treats the time periods as simple categories into which the data must be classified and that such ML classifiers have no explicit approach to time or order. However, there were performance differences which suggest that both ML models are sensitive to the 'arrow of time' that is inevitably present in purely implicit temporal features of the data.

Accuracy for both models (and test loss for the RNN models) are reported in Table 1. Accuracy is explained already, but loss is the difference between the correct output vector and the predicted output vector. As an example, a tweet from the most recent time period in Data set 1 would have a correct output vector of [1 0 0] but the model might be predict an output vector of [.8 .2 .0], showing that it got the right answer but was not fully certain.

The most accurate NBC model performance of over 90% comes from Data set 1 which was the subset with the least data as it had only the three most recent time periods. The NBC accuracy falls as more data is added, with Data set 2 showing 84% accuracy and Data set 3 falling just below 80% accuracy. This contradicts the commonly accepted idea that ML is always improved by more data and instead supports the idea that even very simple and decidedly non-temporal models such as an NBC can be disrupted by implicit temporal features.

The RNN models show a similar pattern as accuracy drops between the data sets as the time periods included in each data set increases, although the drop in accuracy is less pronounced. Data set 1, with only three time periods, was almost 90% accurate while Data set 2 was over 84% accurate and Data set 3 was only just over 81% accurate. The RNN models seem to exhibit higher sensitivity to time and thus better capture the patterns in the more complex cases with more data over longer time frames. As RNN models are trained by passing the data through multiple times, we can see how the model learns through training in Figure 4. The gradual improvement in accuracy and gradual decrease in loss on the training data (blue lines) are close but not identical to the accuracy and loss for test data (orange lines) showing that the models are well fitted to the data without suffering from over- or under-fitting. It also reveals that handling data of higher temporal complexity necessitates more iterations, as for Data set 1, 10 epochs are enough to find the optimised weights, but Data set 3 requires at least 20 epochs.

6 Conclusions and Further Work

The main conclusion of this research is that both ML models examined here contradict generally accepted ideas about ML; that more data is better and that classification tasks are not affected by time. Both models show that performance as measured by accuracy was not improved by making more data available. Instead, the results suggest that when the data contains

Data Set	NBC	RNN	
	Accuracy score	Test Loss	Accuracy score
1	0.90	0.39	0.89
2	0.84	0.17	0.85
3	0.79	0.13	0.81

Table 1. Test loss and accuracy scores for both ML models according to data set

Figure 4. Loss and accuracy functions from the RNN

implicit temporal features, more data disrupts the capacity to classify. Effectively, implicit temporal data seems to disrupt non-temporal ML models.

This research should be understood as a first step or ‘proof of concept’ rather than an exhaustive exploration of how ML and time interact. Still, the results strongly suggest that a ML classification task with no explicit role for time or order still appears to be disrupted by implicit temporal features. Classifying tweets into time periods is not a particularly important task, especially when Twitter data can easily be exported with a time stamp. Nevertheless, there are many important classification tasks that already routinely employ ML models including predicting diagnoses from medical images or identifying risk categories for individuals from their social media activity. These, and indeed most, classification tasks do not typically have any explicit representation of time. Despite this, the training data is always created at a specific point in time (and within the technological and social context of that time). At the same time, the categories into which the data is classified are also created, modified, or

no longer used within clear temporal contexts. Thus, even without any explicit representation of time, such classification tasks are inevitably influenced by implicit temporal features.

As merely a proof of concept exploration, this research suggests only that those using ML methods that do not have explicit representations of time consider whether and how the arrow of time may be disrupting or influencing their data in subtle and implicit ways. In effect, time should be a factor that researchers consider when choosing which ML models to use, just as they currently account for volume, data complexity, generalisability, calculation costs and error tolerance, among other options.

The authors would like to continue exploring this research question through other ML models, different or larger data sets (potentially with more specific geographical filtering), classification tasks that do not rely on predicting when a tweet was created, and other training regimes. One particularly interesting training regime to explore would be ‘curriculum learning’ in which training data is fed into the model in a deliberate order (e.g. easy or straightforward training data first and then hard or more ambiguous training data later). [32]. In this way, additional interactions between ML and implicit representations of the ‘arrow of time’ may be discovered. As it stands, the authors can only say that implicit time seems to matter for some classification tasks so researchers should be careful when selecting a ML model in light of the data and project aims. Researchers already are careful in this way when they balance the characteristics of their data (complexity, volume, messiness, etc.) against the research project aims. Thus, the authors only suggest that they consider temporal features and the potential disruptions to ML performance from a mismatch between non-temporal methods and implicitly temporal data when they select a ML algorithm. With further research, the authors would like to produce some heuristics to help in this selection process.

References

- [1] C.A. Brunk, M.J. Pazzani, in *Machine Learning Proceedings 1991* (Elsevier, 1991), pp. 389–393
- [2] T.M. Kaiser, P.B. Burger, *Molecules* **24**, 2115 (2019)
- [3] Y. Baştanlar, M. Özuysal, in *miRNomics: MicroRNA Biology and Computational Analysis* (Springer, 2014), pp. 105–128
- [4] A. Eddington, *The nature of the physical world: THE GIFFORD LECTURES 1927*, Vol. 23 (BoD–Books on Demand, 2019)
- [5] S.E. Page et al., *Quarterly Journal of Political Science* **1**, 87 (2006)
- [6] G.E. Mikhailovsky, A.P. Levich, *Entropy* **17**, 4863 (2015)
- [7] A. Ben-Naim, *Entropy* **19**, 48 (2017)
- [8] G. Febres, K. Jaffe, *Entropy* **17**, 1606 (2015)
- [9] E. CLARK, *Common Ground*, Vol. 87 (Wiley Online Library, 2015)
- [10] A. Eshghi, C. Howes, E. Gregoromichelaki, J. Hough, M. Purver, *Feedback in conversation as incremental semantic update*, in *Proceedings of the 11th International Conference on Computational Semantics* (2015), pp. 261–271
- [11] F. Ferreira, J.M. Henderson, *Journal of Memory and Language* **30**, 725 (1991)
- [12] R.R. Romeo, J.A. Leonard, S.T. Robinson, M.R. West, A.P. Mackey, M.L. Rowe, J.D. Gabrieli, *Psychological science* **29**, 700 (2018)
- [13] S. Laxman, P.S. Sastry, *Sadhana* **31**, 173 (2006)
- [14] A. Bagnall, A. Bostrom, J. Large, J. Lines, arXiv preprint arXiv:1602.01711 (2016)
- [15] S. Wang, J. Cao, P. Yu, *IEEE Transactions on Knowledge and Data Engineering* (2020)
- [16] N.K. Ahmed, A.F. Atiya, N.E. Gayar, H. El-Shishiny, *Econometric Reviews* **29**, 594 (2010)

- [17] M.L. Graber, BMJ quality & safety **22**, ii21 (2013)
- [18] L. Kanner et al., Nervous child **2**, 217 (1943)
- [19] H. Dodwell, British medical journal **1**, 149 (1954)
- [20] A. Goel, J. Gautam, S. Kumar, *Real time sentiment analysis of tweets using Naive Bayes*, in *2016 2nd International Conference on Next Generation Computing Technologies (NGCT)* (2016), pp. 257–261
- [21] V. Prakruthi, D. Sindhu, D.S. Anupama Kumar, *Real Time Sentiment Analysis Of Twitter Posts*, in *2018 3rd International Conference on Computational Systems and Information Technology for Sustainable Solutions (CSITSS)* (2018), pp. 29–34
- [22] K.Z. Bertrand, M. Bialik, K. Virdee, A. Gros, Y. Bar-Yam, arXiv preprint arXiv:1308.5010 (2013)
- [23] L. Zhao, J. Jia, L. Feng, *Teenagers' stress detection based on time-sensitive micro-blog comment/response actions*, in *IFIP International Conference on Artificial Intelligence in Theory and Practice* (Springer, 2015), pp. 26–36
- [24] A. Zhao, J. Kasmire, *Ictessh-arrow-of-time* (2021), <https://github.com/UKDataServiceOpen/ICTeSSH-Arrow-of-Time>
- [25] J. Roesslein, URL: <https://github.com/tweepy/tweepy> (2020)
- [26] *Python package index - pypi*, <https://pypi.org/>
- [27] G. Van Rossum, *The Python Library Reference*, release 3.8.2 (Python Software Foundation, 2020)
- [28] S. Bird, E. Klein, E. Loper, *Natural language processing with Python: analyzing text with the natural language toolkit* (" O'Reilly Media, Inc.", 2009)
- [29] S. Loria, Release 0.15 **2** (2018)
- [30] F. Chollet et al., *Keras* (2015), <https://github.com/fchollet/keras>
- [31] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G.S. Corrado, A. Davis, J. Dean, M. Devin et al., *TensorFlow: Large-scale machine learning on heterogeneous systems* (2015), software available from tensorflow.org, <https://www.tensorflow.org/>
- [32] Y. Bengio, J. Louradour, R. Collobert, J. Weston, *Curriculum Learning*, in *Proceedings of the 26th Annual International Conference on Machine Learning* (Association for Computing Machinery, New York, NY, USA, 2009), ICML '09, p. 41–48, ISBN 9781605585161, <https://doi.org/10.1145/1553374.1553380>