

Enhancing Interpretable Clauses Semantically using Pretrained Word Representation

Rohan Kumar Yadav, Lei Jiao, Ole-Christoffer Granmo, and Morten Goodwin

Centre for Artificial Intelligence Research, University of Agder, 4879, Grimstad, Norway
{rohan.k.yadav, lei.jiao, ole.granmo, morten.goodwin}@uia.no

Abstract

Tsetlin Machine (TM) is an interpretable pattern recognition algorithm based on propositional logic, which has demonstrated competitive performance in many Natural Language Processing (NLP) tasks, including sentiment analysis, text classification, and Word Sense Disambiguation. To obtain human-level interpretability, legacy TM employs Boolean input features such as bag-of-words (BOW). However, the BOW representation makes it difficult to use any pre-trained information, for instance, word2vec and GloVe word representations. This restriction has constrained the performance of TM compared to deep neural networks (DNNs) in NLP. To reduce the performance gap, in this paper, we propose a novel way of using pre-trained word representations for TM. The approach significantly enhances the performance and interpretability of TM. We achieve this by extracting semantically related words from pre-trained word representations as input features to the TM. Our experiments show that the accuracy of the proposed approach is significantly higher than the previous BOW-based TM, reaching the level of DNN-based models.

1 Introduction

Tsetlin Machine (TM) is an explainable pattern recognition approach that solves complex classification problems using propositional formulas (Granmo, 2018). Text- (Berge et al., 2019), numerical data- (Abeyrathna et al., 2019), and image classification (Granmo et al., 2019) are recent areas of application. In Natural Language Processing (NLP), TM has provided encouraging trade-offs between accuracy and interpretability for various tasks. These include Sentiment Analysis (SA) (Yadav et al., 2021; Saha et al., 2020), Word Sense Disambiguation (WSD) (Yadav et al., 2021), and novelty detection (Bhattarai et al., 2021). Because TM NLP models employ bag-of-words (BOW) that treat each word as independent

features, it is easy for humans to interpret them. The models can be interpreted simply by inspecting the words that take part in the conjunctive clauses. However, using a simple BOW makes it challenging to attain the same accuracy level as deep neural network (DNN) based models.

A key advantage of DNN models is distributed representation of words in a vector space. By using a single-layer neural network, Mikolov et al. introduced such a representation, allowing for relating words based on the inner product between word vectors (Mikolov et al., 2013). One of the popular methods is skip-gram, an approach that learns word representations by predicting the context surrounding a word within a given window length. However, skip-gram has the disadvantage of not considering the co-occurrence statistics of the corpus. Later, Pennington et al. developed *GloVe* – a model that combines the advantages of local window-based methods and global matrix factorization (Pennington et al., 2014). The foundation for the above vector representation of words is the distributional hypothesis that states that “the word that occurs in the same contexts tend to have similar meanings” (Harris, 1954). This means that in addition to forming a rich high-dimensional representation of words, words that are closer to each other in vector space tend to represent similar meaning. As such, vector representations have been used to enhance for instance information retrieval (Manning et al., 2008), name entity recognition (Turian et al., 2010), and parsing (Socher et al., 2013).

The state of the art in DNN-based NLP has been advanced by incorporating various pre-trained word representations such as GloVe (Pennington et al., 2014), word2vec (Mikolov et al., 2013), and fasttext (Bojanowski et al., 2017). Indeed, building semantic representations of the words has been demonstrated to be a vital factor for improved performance. Most DNN-based models utilize the pre-trained word representations to initialize their

word embeddings. This provides them with additional semantic information that goes beyond a traditional BOW.

However, in the case of TM, such word representations cannot be directly employed because they consist of floating-point numbers. First, these numbers must be converted into Boolean form for TM to use, which may result in information loss. Secondly, replacing the straightforward BOW of a TM with a large number of floating-point numbers in fine-grained Boolean form would impede interpretability. In this paper, we propose a novel pre-processing technique that evades the above challenges entirely by extracting additional features for the BOW. The additional features are found using the pre-trained distributed word representations to identify words that enrich the BOW, based on cosine similarity. In this way, TM can use the information from word representations for increasing performance, and at the same time retaining the interpretability of the model.

The rest of the paper is organised as follows. We summarize related work in Section 2. The proposed semantic feature extraction for TM is then explained in Section 3. In Section 4, we present the TM architecture employing the proposed feature extension. We provide extensive experiment results in Section 5, demonstrating the benefits of our approach, before concluding the paper in Section 6.

2 Related Work

Conventional text classification usually focuses on feature engineering and classification algorithms. One of the most popular feature engineering approaches is the derivation of BOW features. Several complex variants of BOW have been designed such as n -grams (Wang and Manning, 2012) and entities in ontologies (Chenthamarakshan et al., 2011). Apart from BOW approaches, Tang et al. demonstrated a new mechanism for feature engineering using a time series model for short text samples (Tang et al., 2020). There are also several techniques to convert text into a graph and sub-graph (Rousseau et al., 2015; Luo et al., 2017). In general, none of the above methods adopt any pre-trained information, hence have inferior performance.

Deep learning-based text classification either depends on initializing models from pre-trained word representations, or on jointly learning both the word- and document level representations. Various studies report that incorporating such word

representations, embedding the words, significantly enhances the accuracy of text classification (Joulin et al., 2017; Shen et al., 2018a). Another approach related to pre-trained word embedding is to aggregate unsupervised word embeddings into a document embedding, which is then fed to a classifier (Le and Mikolov, 2014; Tang et al., 2015).

Despite being empowered with world knowledge through pre-trained information, DNNs such as BERT (Devlin et al., 2019) and XLNet (Yang et al., 2019) can be very hard to interpret. One interpretation approach is to use attention-based models, relying on the weights they assign to the inputs. However, more careful studies reveal that attention weights in general do not provide a useful explanation (Bai et al., 2020; Serrano and Smith, 2019). Researchers are thus increasingly shifting focus to other kinds of machine learning, with the TM being a recent approach considered to provide human-level interpretability (Berge et al., 2019; Granmo, 2018; Yadav et al., 2021). It offers a very simple model consisting of multiple Tsetlin Automata (TAs) that select which features take part in the classification. However, despite promising performance, there is still a performance gap to the DNN models that utilize pre-trained word embedding. Yet, several TM studies demonstrate high degree of interpretability through simple rules, with a marginal loss in accuracy (Yadav et al., 2021; Yadav. et al., 2021; Saha et al., 2020).

A significant reason for the performance gap between TM-based and state-of-the-art DNN-based NLP models is that TM operates on Boolean inputs, lacking a method for incorporating pre-trained word embeddings. Without pre-trained information, TMs must rely on labelled data available for supervised learning. On the other hand, incorporating high-dimensional Booleanized word embedding vectors directly into the TM would significantly reduce interpretability. In this paper, we address this intertwined challenge. We propose a novel technique that boosts the TM BOW approach, enhancing the BOW with additional word features. The enhancement consists of using cosine similarity between GloVe word representations to obtain semantically related words. We thus distill information from the pre-trained word representations for utilization by the TM. To this end, we propose two methods of feature extension: (1) using the k nearest words in embedding space and (2) using words within a given similarity threshold, measured as

cosine angle (θ). By adopting the two methods, we aim to reduce the current performance gap between interpretable TM and black-box DNN, by achieving either higher or similar accuracy.

3 Boosting TM BOW with Semantically Related Words

Here, we introduce our novel method for boosting the BOW of TM with semantically related words. The method is based on comparing pre-trained word representations using cosine similarity, leveraging distributed word representation. There are various distributional representations of words available. These are obtained from different corpora, using various techniques, such as word2vec, GloVe, and fastText. We here use GloVe because of its general applicability.

3.1 Input Feature Extraction from Distributed Word Representation

Distributed word representation does not necessarily derive word similarity based on synonyms but based on the words that appear in the same context. As such, the representation is essential for NLP because it captures the semantically interconnecting words. Our approach utilizes this property to expand the range of features that we can use in an interpretable manner in TM.

Consider a full vocabulary W of m words, $W = [w_1, w_2, w_3 \dots, w_m]$. Further consider a particular sentence that is represented as a Boolean BOW $X = [x_1, x_2, x_3, \dots, x_m]$. In a Boolean BOW, each element $x_r, r = 1, 2, 3, \dots, m$, refers to a specific word w_r in the vocabulary W . The element x_r takes the value 1 if the corresponding word w_r is present in the sentence and the value 0 if the word is absent. Assume that n words are present in the sentence, i.e., n of the elements in X are 1-valued. Our strategy is to extract additional features from these by expanding them using cosine similarity. To this end, we use a GloVe embedding of each present word $w_r, r \in \{z | x_z = 1, z = 1, 2, 3 \dots, m\}$. The embedding for word w_r is represented by vector $w_r^e \in \mathbb{R}^d$, where d is the dimensionality of the embedding (typically varying from 25 to 300).

We next introduce two selection techniques to expand upon each word:

- Select the top k most similar words,
- Select words up to a fixed similarity angle $\cos(\theta) = \phi$.

For example, let us consider two contexts: “very good movie” and “excellent film, enjoyable”. Figs. 1 and 2 list similar words showing the difference between top k words and words within angle $\cos(\theta)$, i.e., ϕ . In what follows, we will explain how these words are found.

3.2 Similar Words based on Top k Nearest Words

We first boost the Boolean BOW of the considered sentence by expanding X with $(k - 1) \times n$ semantically related words. That is, we add $k - 1$ new words for each of the n present words. We do this by identifying neighbouring words in the GloVe embedding space, using cosine similarity between the embedding vectors.

Consider the GloVe embedding vectors $W_G^e = [w_1^e, w_2^e, \dots, w_m^e]$ of the full vocabulary W . For each word w_r from the sentence considered, the cosine similarity to each word $w_t, t = 1, 2, \dots, m$, of the full vocabulary is given by Eq. (1),

$$\phi_r^t = \cos(w_r^e, w_t^e) = \frac{w_r^e \cdot w_t^e}{\|w_r^e\| \cdot \|w_t^e\|}. \quad (1)$$

Clearly, ϕ_r^t is the cosine similarity between w_r^e and w_t^e . By calculating the cosine similarity of w_r to the words in the vocabulary, we obtain m values: $\phi_r^t, t = 1, 2, \dots, m$. We arrange these values in a vector Φ_r :

$$\Phi_r = [\phi_r^1, \phi_r^2, \dots, \phi_r^m]. \quad (2)$$

The k elements from Φ_r of largest value are then identified and their indices are stored in a new set A_r .

Finally, a boosted BOW, referred to as X_{mod} , can be formed by assigning element x_t value 1 whenever one of the A_r contains t , and 0 otherwise:

$$X_{mod} = [x_1, x_2, x_3, \dots, x_m], \quad (3)$$

$$x_t = \begin{cases} 1 & \exists r, t \in A_r \\ 0 & \nexists r, t \in A_r. \end{cases}$$

In addition, the vocabulary size for a particular task/dataset can be changed accordingly, which is usually less than m . Note that implementation-wise, the GloVe library provides the top k similar words of w_r without considering the word w_r itself, having similarity score 1. Hence, using the GloVe library, w_r must also be added to the boosted BOW.

3.3 Similar Words within Cosine Angle Threshold

Another approach to enrich the Boolean BOW of a sentence is thresholding the cosine angle. This is

excellent film, enjoyable					
excellent	similarity	film	similarity	enjoyable	similarity
good	0.703	movie	0.858	entertaining	0.688
superb	0.680	films	0.839	pleasurable	0.660
terrific	0.656	movies	0.716	exhilarating	0.647
quality	0.603	directed	0.683	fun	0.625
wonderful	0.584	documentary	0.656	exciting	0.619
best	0.550	starring	0.651	amusing	0.605
exceptional	0.550	cinema	0.637	satisfying	0.585
perfect	0.535	screenplay	0.632	engrossing	0.578
impressive	0.534	drama	0.622	enlightening	0.573
decent	0.532	comedy	0.615	informative	0.571

Figure 1: Similar words for an example “excellent film, enjoyable” using 300d GloVe word representation.

very good movie					
very	similarity	good	similarity	movie	similarity
extremely	0.872	better	0.765	film	0.858
quite	0.858	really	0.736	movies	0.849
so	0.785	always	0.717	films	0.790
pretty	0.738	you	0.707	hollywood	0.679
too	0.731	well	0.704	starring	0.675
really	0.729	excellent	0.703	comedy	0.658
well	0.720	very	0.696	sequel	0.646
always	0.712	things	0.693	remake	0.624
especially	0.709	think	0.689	drama	0.608
but	0.707	way	0.683	actor	0.599

Figure 2: Similar words for an example “very good movie” using 300d GloVe word representation.

different from the first technique because the number of additional words extracted will vary rather than being fixed. Whereas the first approach always produces $k - 1$ new features for each given word, the cosine angle thresholding brings in all those words that are sufficiently similar. The cosine similarity threshold is given by $\phi = \cos(\theta)$, where θ is the threshold for vector angle, while ϕ is the corresponding similarity score.

As per Eq. (2), we obtain Φ_r , which consists of the similarity scores of the given word w_r in comparison to the m words in the vocabulary. Then, for each given word w_r , the indices of those scores ϕ_r^t that are greater than or equal to ϕ ($\phi_r^t \geq \phi$) are stored in the set A_r . Similar to the first technique, the words in W with the indices in A_r are utilized to create X_{mod} as given by Eq. (3).

4 Tsetlin Machine-based Classification

4.1 Tsetlin Machine Architecture

A TM is composed by TAs that operate with literals – Boolean inputs and their negations – to form conjunctions of literals (conjunctive clauses). A dedicated team of TAs builds each clause, with each input being associated with a pair of TAs. One TA controls the original Boolean input whereas the other TA controls its negation. The TA pair selects

a combination of “Include” or “Exclude” actions, which decide the form of the literal to include or exclude in the clause.

Each TA decides upon an action according to its current state. There are N states per TA action, $2N$ states in total. When a TA finds itself in states 1 to N , it performs the “Exclude” action. When in states $N + 1$ to $2N$, it performs the “Include” action. How the TA updates its state is shown in Fig. 3. If it receives Reward, the TA moves to a deeper state thereby increasing its confidence in the current action. However, if it receives Penalty, it moves towards the centre, weakening the action. It may eventually jump over the middle decision boundary, to the other action. It is through this game of TAs that the TM shapes the clauses into frequent and discriminative patterns.




Figure 3: A TA with two actions: “Include” and “Exclude”.

With respect to NLP, TM heavily relies on the Boolean BOW introduced earlier in the paper. We now make use of our proposed modified BOW $X_{mod} = [x_1, x_2, x_3, \dots, x_m]$. Let l be the number of clauses that represent each class of the TM, covering q classes altogether. Then, the overall pattern recognition problem is solved using $l \times q$ clauses. Each clause C_i^j , $1 \leq j \leq q$, $1 \leq i \leq l$ of the TM is

given by $C_i^j = \left(\bigwedge_{k \in I_i^j} x_k \right) \wedge \left(\bigwedge_{k \in \bar{I}_i^j} \neg x_k \right)$, where

I_i^j and \bar{I}_i^j are non-overlapping subsets of the input variable indices, $I_i^j, \bar{I}_i^j \subseteq \{1, \dots, m\}$, $I_i^j \cap \bar{I}_i^j = \emptyset$. The subsets decide which of the input variables take part in the clause, and whether they are negated or not. The indices of input variables in I_i^j represent the literals that are included as is, while the indices of input variables in \bar{I}_i^j correspond to the negated ones. Among the q clauses of each class, clauses with odd indexes are assigned positive polarity (+) whereas those with even indexes are assigned negative polarity (-). The clauses with positive polarity vote for the target class and those with negative polarity vote against it. A summation operator aggregates the votes by subtracting the total number of negative votes from positive votes, as shown

in Eq. (4).

$$f^j(X_{mod}) = \frac{\sum_{i=1,3,\dots}^{l-1} C_i^j(X_{mod})}{\sum_{i=2,4,\dots}^l C_i^j(X_{mod})}. \quad (4)$$

For q number of classes, the final output y is given by the argmax operator to classify the input based on the highest sum of votes, $\hat{y} = \operatorname{argmax}_j (f^j(X_{mod}))$.

4.2 Distributed Word Representation in TM

Consider two contexts for sentiment classification: “Very good movie” and “Excellent film, enjoyable”. Both contexts have different vocabularies but some of them are semantically related to each other. For example, “good” and “excellent” have similar semantics as well as “film” and “movie”. Such semantics are not captured in the BOW-based input. However, as shown in Fig. 4, adding words to the BOWs that are semantically related, as proposed in the previous section, makes distributed word representation available to the TM.

The resulting BOW-boosted TM architecture is shown in Fig. 5. Here each input feature is first expanded using the GloVe representation, adding semantically related words. Each feature is then transferred to its corresponding TAs, both in original and negated form. Each TA, in turn, decides whether to include or exclude its literal in the clause by taking part in a decentralized game. The actions of each TA is decided by its current state and updated by the the feedback it receives based on its action. As shown in the figure, the TA actions produce a collection of conjunctive clauses, joining the words into more complex linguistic patterns.

There are two types of feedback that guides the TA learning. They are Type I feedback and Type II feedback, detailed in (Granmo, 2018). Type I feedback is triggered when the ground truth label is 1, i.e., $y = 1$. The purpose of Type I feedback is to include more literals from the BOW to refine the clauses, or to trim them by removing literals. The balance between refinement and trimming is controlled by a parameter called specificity, s . Type I feedback guides the clauses to provide true positive output, while simultaneously controlling overfitting by producing frequent patterns. Conversely, Type II feedback is triggered in case of false positive output. Its main aim is to introduce zero-valued literals into clauses when they give false positive output. The purpose is to change them so that they correctly output zero later in the learning process.

Based on these feedback types, each TA in a clause receives Reward, Penalty or Inaction. The overall learning process is explained in detail by Yadav et al. in (Yadav et al., 2021).




Figure 4: (a) BOW input representation without distributed word representation. (b) BOW input using similar words based on distributed word representation.

5 Experiments and Results

In this section, we evaluate our TM-based solution with the input features enhanced by distributed word representation. Here we use Glove pretrained word vector that is trained using CommonCrawl with the configuration of 42B tokens, 1.9M vocab, uncased, and 300d vectors.

5.1 Datasets

We have selected various types of datasets to investigate how broadly our method is applicable: R8 and R52 of Reuters, Movie Review (MR), and TREC-6. • **Reuters** 21578 dataset include two subsets: R52 and R8 (all-terms version). R8 is divided into 8 sections while there are 52 categories in R52. • **MR** is a movie analysis dataset for binary sentiment classification with just one sentence per review (Pang and Lee, 2005). In this study, we used a training/test split from (Tang et al., 2015)¹. • **TREC-6** is a question classification dataset (Li and Roth, 2002). The task entails categorizing a query into six distinct categories (abbreviation, description, entity, human, location, numeric value).

5.2 TM Parameters

A TM has three parameters that must be initialized before training a model: number of clauses l , voting target T , and specificity s . We configure these parameters as follows. For R8, we use 2,500 clauses, a threshold of 80, and specificity 9. The vocabulary size is 5,000. For R52, we employ 1,500 clauses, the voting target is 80, and specificity is 9. Here, we use a vocabulary of size 6,000. For MR, the number of clauses is 3,000, the voting target is

¹<https://github.com/mnqu/PTE/tree/master/data/mr>.

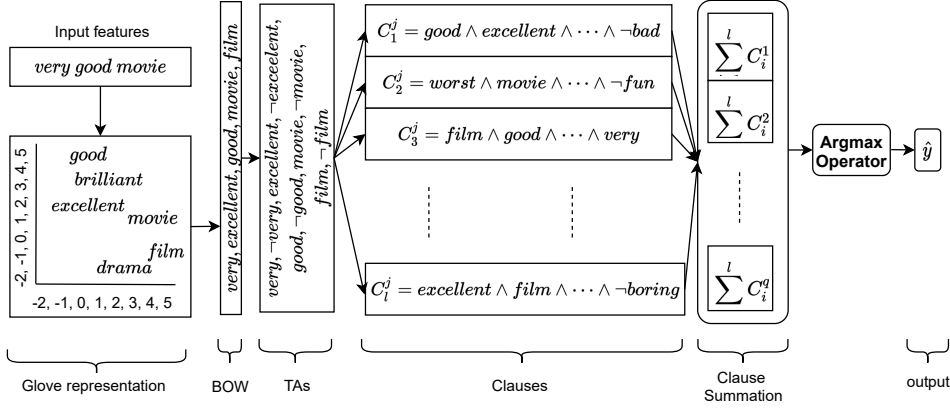


Figure 5: Architecture of TM using modified BOW based on word similarity.

80, and specificity is 9, with a vocabulary of size 5,000. Finally, for TREC, we use 2,000 clause, a voting target of 80, and specificity 9, with vocabulary size 6,000. These parameters are kept static as we explore various k and θ values for selecting similar words to facilitate comparison. The code and datasets are available online ².

5.3 Performance When Using Top k Nearest Neighbors

Here, we demonstrate the performance on each of the datasets, exploring the effect of different k -values, i.e., 3, 5 and 10. The performance of the proposed technique for selected datasets with various values of k is compared in Table 1. It can be seen that by using feature extension, performance is significantly enhanced. Both $k = 3$ and $k = 5$ outperform the simple BOW ($k = 0$). However, for this particular dataset, $k = 10$ performs poorly because extending each word to its 10 nearest neighbors includes many unnecessary contexts that have no significant impact on the classification. In terms of accuracy, $k = 5$ performs best for the R8 dataset. For the R52 dataset, the feature extension with $k = 5$ and $k = 10$ performs poorly compared to using $k = 0$ and $k = 3$. Here, $k = 3$ is the best-performing parameter. The improvement obtained by moving from a simple BOW to a BOW enhanced with semantically similar features is obvious in the case of the R52 dataset. Similarly, in the case of the TREC dataset, the performance of simple BOW ($k = 0$) is markedly outperformed by the feature extension techniques for all the tested k -values, with $k = 5$ and $k = 10$ being good candidates. The advantage of $k = 10$ over $k = 5$ is that $k = 10$ reaches its peak accuracy in an earlier

²<https://github.com/rohanky/Glove-TM>

Parameters	R8	R52	MR	TREC
k=0	96.16	84.62	75.14	88.05
k=3	97.08	88.59	75.21	88.72
k=5	96.80	70.60	76.06	89.16
k=10	87.44	66.94	77.51	89.82

Table 1: Comparison of feature extended TM with several parameters for k .

epoch. Lastly, the performance of the MR is again clear that the feature extension technique outperforms the simple BOW ($k = 0$) with a high margin.

5.4 Performance When Using Neighbors Within a Similarity Threshold

This section demonstrates the performance of our BOW enhancement approach when using various similarity thresholds ϕ for feature extension. Here, ϕ refers to the cosine similarity between a word in the BOW and a target word from the overall vocabulary. Again, similarity is measured in the GloVe embedding space as the cosine of the angle θ between the embedding vectors compared, $\cos(\theta)$. For ϕ , we here explore the values 0.5, 0.6, 0.7, 0.8, and 0.9, whose corresponding angles are 60° , 53.13° , 45.57° , 36.86° , and 25.84° , respectively. The performance of the various ϕ -values for the selected dataset is shown in Table 2. For R8 dataset, feature extension using $\phi = 0.7$, $\phi = 0.8$, and $\phi = 0.9$ outperforms the simple BOW ($\phi = 0$) where $\phi = 0.7$ being the best. In case of the R52 dataset, all of the investigated ϕ -values outperform the simple BOW ($\phi = 0$) where $\phi = 0.5$ and $\phi = 0.8$ performs the best. Similar trend is observed in case of TREC and MR dataset where feature extension outperforms the simple BOW.

Parameters	R8	R52	MR	TREC
$\phi = 0$	96.16	84.62	75.14	88.05
$\phi = 0.5$	88.08	89.14	73.24	90.04
$\phi = 0.6$	90.86	88.05	74.34	87.83
$\phi = 0.7$	96.53	88.51	76.55	89.38
$\phi = 0.8$	96.25	88.94	75.12	88.27
$\phi = 0.9$	96.39	87.50	74.59	87.39

Table 2: Comparison of feature extended TM with several parameters for ϕ .

In most of the cases, however, a too strict similarity threshold ϕ tends to reduce performance because fewer features are added to the BOW. Even though using a looser similarity score thresholds also introduces unnecessary features, these do not seem to impact the formation of accurate clauses. Overall, our experiments show that using ϕ -values from 0.5 to 0.7 peaks performance.

5.5 Comparison with Baselines

We here compare our proposed model with selected text classification- and embedding methods. We have selected representative techniques from various main approaches, both those that leverage similar kinds of pre-trained word embedding and those that only use BOW. The selected baselines are: **•TF-IDF+LR**: This is a bag-of-words model employing Term Frequency-Inverse Document Frequency (TF-IDF) weighting. Logistic Regression is used as a softmax classifier. **•CNN**: The CNN-baselines cover both initialization with random word embedding (CNN-rand) as well as initialization with pretrained word embedding (CNN-non-static) (Kim, 2014). **•LSTM**: The LSTM model that we employ here is from (Liu et al., 2016), representing the entire text using the last hidden state. We tested this model with and without pre-trained word embeddings. **•Bi-LSTM**: Bi-directional LSTMs are widely used for text classification. We compare our model with Bi-LSTM fed with pre-trained word embeddings. **•PV-DBOW**: PV-DBOW is a paragraph vector model where the word order is ignored. Logistic Regression is used as a softmax classifier (Le and Mikolov, 2014). **•PV-DM**: PV-DM is also a paragraph vector model, however with word ordering taken into account. Logistic Regression is used as a softmax classifier (Le and Mikolov, 2014). **•fastText**: This baseline is a simple text classification technique that uses the average of the word embeddings provided

by fastText as document embedding. The embedding is then fed to a linear classifier (Joulin et al., 2017). We evaluate both the use of uni-grams and bigrams. **•SWEM**: SWEM applies simple pooling techniques over the word embeddings to obtain a document embedding (Shen et al., 2018b). **•Graph-CNN-C**: A graph CNN model uses convolutions over a word embedding similarity graph (Defferrard et al., 2016), employing a Chebyshev filter. **• S^2GC** : This technique uses a modified Markov Diffusion Kernel to derive a variant of Graph Convolutional Network (GCN) (Zhu and Kouniusz, 2021). **•LguidedLearn**: It is a label-guided learning framework for text classification. This technique is applied to BERT as well (Liu et al., 2020), which we use for comparison purposes here. **•Feature Projection (FP)**: It is a novel approach to improve representation learning through feature projection. Existing features are projected into an orthogonal space (Qin et al., 2020).

From Table 3, we observe that the TM approaches that employ either of our feature extension techniques outperform several word embedding-based Logistic Regression approaches, such as PV-DBOW, PV-DM, and fastText. Similarly, the legacy TM outperforms sophisticated models like CNN and LSTM based on randomly initialized word embedding. Still, the legacy TM falls of other models when they are initialized by pre-trained word embeddings. By boosting the Boolean BOW with semantically similar features using our proposed technique, however, TM outperforms LSTM (pretrain) on the R8 dataset and performs similarly on R52 and MR. In addition to this, our proposed approach achieves quite similar performance compared to BERT, even though BERT has been pre-trained on a huge text corpus. However, it falls slightly short of sophisticated fine-tuned models like Lguided-BERT-1 and Lguided-BERT-3. Overall, our results show that our proposed feature extension technique for TMs significantly enhances accuracy, reaching state of the art accuracy. Importantly, this accuracy enhancement does not come at the cost of reduced interpretability, unlike DNNs, which we discuss below. The state of the art for the TREC dataset is different from the other three datasets, hence we report results separately in Table 4. These results clearly show that although the basic TM model does not outperform the recent DNN- and transformer-based models, the feature-boosted TM outperforms all of

Model	R8	R52	MR
TF-IDF+LR	93.74	86.95	74.59
CNN-rand	94.02	85.37	74.98
CNN-non-static	95.71	87.59	77.75
LSTM	93.68	85.54	75.06
LSTM (pretrain)	96.09	90.48	77.33
Bi-LSTM	96.31	90.54	77.68
PV-DBOW	85.87	78.29	61.09
PV-DM	52.07	44.92	59.47
fastText	96.13	92.81	75.14
fastText (bigrams)	94.74	90.99	76.24
SWEM	95.32	92.94	76.65
LEAM	93.31	91.84	76.95
Graph-CNN-C	96.99	92.74	77.22
S^2GC	97.40	94.50	76.70
BERT	96.02	89.66	79.24
Lguided-BERT-1	97.49	94.26	81.03
Lguided-BERT-3	98.28	94.32	81.06
TM	96.16± 1.52	84.62± 1.8	75.14± 1.2
TM with k	97.50± 1.12	88.59± 1.2	77.51± 0.6
TM with ϕ	96.39± 1.0	89.14± 1.5	76.55± 0.9

Table 3: Comparison of feature extended TM with the state of the art for R8, R52 and MR. Reported accuracy of TM is the mean of last 50 epochs of 5 independent experiments with their standard deviation.

those models except understandably BAE:BERT (Garg and Ramakrishnan, 2020).

Model	TREC
LSTM	87.19
FP+LSTM	88.83
Transformer	87.33
FP+Transformer	89.51
BAE: BERT	97.6
TM (Dragoş et al., 2021)	87.20
TM	88.05± 1.52
TM with k	89.82± 1.18
TM with ϕ	90.04± 0.94

Table 4: Comparison of feature extended TM with the state of the art for TREC. Reported accuracy of TM is the mean of last 50 epochs of 5 independent experiments with their standard deviation.

5.6 Interpretation

The proposed feature extension-based TM does not only impact accuracy. Perhaps surprisingly, our proposed technique also simplify the clauses that the TM produces, making them more meaningful in a semantic sense. To demonstrate this property, let us consider two samples from the MR dataset: S_1 ="the cast is uniformly excellent and relaxed" and S_2 ="the entire cast is extraordinarily good". Let the vocabulary, in this case, be [cast, excellent,

relaxed, extraordinarily, good, bad, boring, worst] as shown in Fig. 6.

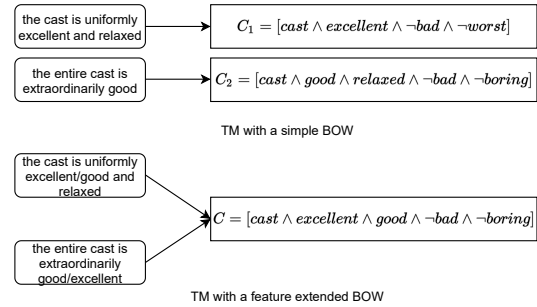


Figure 6: Clause learning semantic for multiple examples compared to simple BOW based TM.

As we can see, that the TM initialized with normal BOW uses two separate clauses to represent two examples. However, augmenting feature on TM uses only one clause that learns the semantic for multiple examples. This indeed makes interpretation of TM more powerful and meaningful as compared to simple BOW based TM.

6 Conclusion

In this paper, we aimed to enhance the performance of Tsetlin Machines (TMs) by introducing a novel way to exploit distributed feature representation for TMs. Given that a TM relies on Bag-of-words (BOW), it is not possible to introduce pre-trained word representation into a TM directly, without sacrificing the interpretability of the model. To address this intertwined challenge, we extended each word feature by using cosine similarity on the distributed word representation. We proposed two techniques for feature extension: (1) using the k nearest words in embedding space and (2) including words within a given cosine angle (θ). Through this enhancement, the TM BOW can be boosted with pre-trained world knowledge in a simple yet effective way. Our experiment results showed that the enhanced TM not only achieve competitive accuracy compared to state of the art, but also outperform some of the sophisticated deep neural network (DNN) models. In addition, our BOW boosting also improved the interpretability of the model by increasing the scope of each clause, semantically relating more samples. We thus believe that our proposed approach significantly enhance the TM in the accuracy/interpretability continuum, establishing a new standard in the field of explainable NLP.

References

- Kuruge Darshana Abeyrathna, Ole-Christoffer Granmo, Xuan Zhang, and Morten Goodwin. 2019. A scheme for continuous input to the Tsetlin machine with applications to forecasting disease outbreaks. In *Advances and Trends in Artificial Intelligence. From Theory to Practice*, pages 564–578. Springer International Publishing.
- Bing Bai, J. Liang, Guanhua Zhang, Hao Li, Kun Bai, and F. Wang. 2020. Why is attention not so interpretable. *arXiv: Machine Learning*.
- Geir Thore Berge, Ole-Christoffer Granmo, Tor Oddbjørn Tveit, Morten Goodwin, Lei Jiao, and Bernt Viggo Matheussen. 2019. Using the tsetlin machine to learn human-interpretable rules for high-accuracy text categorization with medical applications. *IEEE Access*, 7:115134–115146.
- Bimal Bhattarai., Ole-Christoffer Granmo., and Lei Jiao. 2021. Measuring the novelty of natural language text using the conjunctive clauses of a tsetlin machine text classifier. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*, pages 410–417. INSTICC, SciTePress.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Vijil Chenthamarakshan, Prem Melville, Vikas Sindhwani, and Richard D. Lawrence. 2011. Concept labeling: Building text classifiers with minimal supervision. In *IJCAI*, pages 1225–1230.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *ACL: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. ACL.
- Dragoş, Constantin Nicolae, and dragosnicolae. 2021. Question classification using interpretable tsetlin machine. In *International Workshop of Machine Reasoning*. ACM International Conference on Web Search and Data Mining.
- Siddhant Garg and Goutham Ramakrishnan. 2020. [Bae: Bert-based adversarial examples for text classification](#).
- Ole-Christoffer Granmo. 2018. The tsetlin machine - a game theoretic bandit driven approach to optimal pattern recognition with propositional logic. *ArXiv*, abs/1804.01508.
- Ole-Christoffer Granmo, Sondre Glimsdal, Lei Jiao, Morten Goodwin, Christian W. Omlin, and Geir Thore Berge. 2019. The convolutional tsetlin machine. *arXiv*, 1905.09688.
- Zellig S. Harris. 1954. Distributional structure. *WORD*, 10(2-3):146–162.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *EACL: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. ACL.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. ACL.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1188–1196, Beijing, China. PMLR.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *COLING*.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. Recurrent neural network for text classification with multi-task learning. In *IJCAI*, page 2873–2879.
- X. Liu, Song Wang, X. Zhang, Xinxin You, J. Wu, and D. Dou. 2020. Label-guided learning for text classification. *ArXiv*, abs/2002.10772.
- Yuan Luo, Özlem Uzuner, and Peter Szolovits. 2017. Bridging semantics and syntax with graph algorithms - state-of-the-art of extracting biomedical relations. *Briefings in bioinformatics*, 18 1:160–178.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS, Nevada, USA*, volume 26, pages 3111–3119. Curran Associates, Inc.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *ACL*, page 115–124, Michigan, USA. ACL.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP, Doha, Qatar*, page 1532–1543.
- Qi Qin, Wenpeng Hu, and Bing Liu. 2020. Feature projection for improved text classification. In *ACL*, pages 8161–8171, Online. ACL.

- François Rousseau, Emmanouil Kiagias, and Michalis Vazirgiannis. 2015. Text categorization as a graph classification problem. In *ACL (Volume 1: Long Papers)*, pages 1702–1712, Beijing, China. ACL.
- Rupsa Saha, Ole-Christoffer Granmo, and Morten Goodwin. 2020. Mining interpretable rules for sentiment and semantic relation analysis using tsetlin machines. In *Artificial Intelligence XXXVII*, pages 67–78, Cham. Springer International Publishing.
- Sofia Serrano and Noah A. Smith. 2019. Is attention interpretable? In *ACL*, pages 2931–2951, Florence, Italy. ACL.
- Dinghan Shen, Guoyin Wang, Wenlin Wang, Martin Renqiang Min, Qinliang Su, Yizhe Zhang, Chunyuan Li, Ricardo Henao, and Lawrence Carin. 2018a. Baseline needs more love: On simple word-embedding-based models and associated pooling mechanisms. In *ACL Volume 1: Long Papers*, pages 440–450, Melbourne, Australia. ACL.
- Dinghan Shen, Guoyin Wang, Wenlin Wang, Martin Renqiang Min, Qinliang Su, Yizhe Zhang, Chunyuan Li, Ricardo Henao, and Lawrence Carin. 2018b. Baseline needs more love: On simple word-embedding-based models and associated pooling mechanisms. In *ACL (Volume 1: Long Papers)*, pages 440–450, Melbourne, Australia. ACL.
- Richard Socher, John Bauer, Christopher D. Manning, and Andrew Y. Ng. 2013. Parsing with compositional vector grammars. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 455–465, Sofia, Bulgaria. Association for Computational Linguistics.
- Jian Tang, Meng Qu, and Qiaozhu Mei. 2015. Pte: Predictive text embedding through large-scale heterogeneous text networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, page 1165–1174, Sydney, NSW, Australia. Association for Computing Machinery.
- Yichen Tang, Kelly Blincoe, and A. Kempa-Liehr. 2020. Enriching feature engineering for short text samples by language time series analysis. *EPJ Data Science*, 9:1–59.
- Joseph Turian, Lev-Arie Ratinov, and Yoshua Bengio. 2010. Word representations: A simple and general method for semi-supervised learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 384–394, Uppsala, Sweden. Association for Computational Linguistics.
- Sida Wang and Christopher Manning. 2012. Baselines and bigrams: Simple, good sentiment and topic classification. In *ACL (Volume 2: Short Papers)*, pages 90–94, Jeju Island, Korea.
- Rohan Kumar Yadav, Lei Jiao, Ole-Christoffer Granmo, and Morten Goodwin. 2021. Human-level interpretable learning for aspect-based sentiment analysis. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*. AAAI.
- Rohan Kumar Yadav., Lei Jiao., Ole-Christoffer Granmo., and Morten Goodwin. 2021. Interpretability in word sense disambiguation using tsetlin machine. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*, pages 402–409. INSTICC, SciTePress.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Hao Zhu and Piotr Koniusz. 2021. Simple spectral graph convolution. In *International Conference on Learning Representations*.