

Human-Level Interpretable Learning for Aspect-Based Sentiment Analysis

Rohan K Yadav, Lei Jiao, Ole-Christoffer Granmo, Morten Goodwin

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

Abstract

This paper proposes a human-interpretable learning approach for aspect-based sentiment analysis (ABSA), employing the recently introduced Tsetlin Machines (TMs). We attain interpretability by converting the intricate position-dependent textual semantics into binary form, mapping all the features into bag-of-words (BOWs). The binary-form BOWs are encoded so that the information on the aspect and context words are retained for sentiment classification. We further adopt the BOWs as input to the TM, enabling learning of aspect-based sentiment patterns in propositional logic. To evaluate interpretability and accuracy, we conducted experiments on two widely used ABSA datasets from SemEval 2014: Restaurant 14 and Laptop 14. The experiments show how each relevant feature takes part in conjunctive clauses that contain the context information for the corresponding aspect word, demonstrating human-level interpretability. At the same time, the obtained accuracy is on par with existing neural network models, reaching 78.02% on Restaurant 14 and 73.51% on Laptop 14.

Introduction

Sentiment analysis, which identifies people’s opinion on specific topics, is a classic problem in natural language processing (NLP). Under the umbrella of sentiment analysis, aspect-based sentiment analysis (ABSA), which is a fine-grained evaluation framework for sentiment classification (Zhang and Liu 2017), has become a hot research topic (Pontiki et al. 2014). Among various tasks in ABSA, this paper focuses on the sentiment polarity (positive, neutral, negative) of a target word in given comments or reviews. For example, let us consider a review: “*Certainly not the best **sushi** in New York, however, it is always fresh and the **place** is very clean, sterile*”. The target word “*sushi*” is closely associated with its context words “*not best*”, assorting it as a negative polarity. The target word, “*place*”, is associated with its context words “*clean*” and “*sterile*”, classifying it as a positive sentiment. Such a complex form of sentiment classification is highly dependent on where the word appears in the sentence. To address this challenge, several recent approaches to ABSA have been based on attention mechanisms (Ma et al. 2017). Although the accuracy

of attention-based ABSA approaches are progressively improved, the interpretability of these models is still questionable, making them less trust-worthy. Not surprisingly, little research has been done on ABSA learning techniques that are interpretable at a human level (Do et al. 2019).

Recently, interpretable AI has taken a big leap in industrial application (Samek et al. 2019). Indeed, the scientific community has performed extensive research on ways to interpret neural networks. In a modern neural network, one can use the fact that the variants of attention (Bahdanau, Cho, and Bengio 2015) assign soft weights to the input representations, and then extract highly weighted tokens as rationales. However, these attention weights do not provide faithful explanations for classification (Serrano and Smith 2019; Wiegrefe and Pinter 2019; Brunner et al. 2020; Vashishth et al. 2019). On the other hand, certain classic models, like Decision Trees, are particularly easy to understand, yet still compromise on accuracy compared with neural networks. Hence, an effective trade-off between accuracy and interpretability has still not been achieved.

In this article, we propose a Tsetlin Machine (TM) (Granmo 2018) based ABSA that employs a binary representation of the input features. The resulting architecture is *interpretable* and achieves competitive accuracy compared with state-of-the-art techniques. The ABSA task has two important inputs: a context word and an aspect word. Such aspect-based classification usually relies heavily on the position of the aspect word in the context. Such position information can be easily embedded in the neural network models. However, in TM, as all patterns and outputs are expressed in bits, learning and classification depend on bit manipulation, making it a challenging task to embed all the information into binary form. We therefore also aim to propose an extensive pre-processing approach for the ABSA inputs so that the binary form retains as much useful information as possible for the classification.

Our main contributions can be summarized as follows:

- We propose a novel pre-processing scheme to convert the ABSA inputs into binary form with limited information loss.
- We design an interpretable learning architecture using TM. The architecture offers human-level interpretable results with comparable classification accuracy.

- We employ additional knowledge from SentiWordnet (Esuli and Sebastiani 2006) to enhance the accuracy of the architecture. It provides additional knowledge to the model and has significant impact on accuracy as explained later.

The remainder of the paper is organized as follows: We summarize related work in Section 2. The proposed pre-processing and TM architecture along with its learning process are described in Section 3. In Section 4, we report the experiment results and the comparisons with state-of-the-art. The interpretability of trained models is demonstrated in Section 5 before we conclude this work in Section 6.

Related Work

Sentiment analysis operates at three levels: document level, sentence level and aspect level. This work focuses on aspect level. Most of traditional supervised approaches depend heavily on handcrafted features to identify the sentiment of a word based on its context (Jiang et al. 2011; Kiritchenko et al. 2014). However, these models fail to capture the semantic relatedness between the aspect word and its context. This problem gives rise to the attention-based models that are able to capture such a relationship (Shen et al. 2018; Bahdanau, Cho, and Bengio 2015; Tang, Qin, and Liu 2016; Liu and Zhang 2017). Furthermore, it is shown in (Tang et al. 2016) how an attention layer captures the weightage of the context words for predicting the sentiment of an aspect word. However, existing models cannot leverage the syntactic structure of the sentence, thereby making it difficult to distinguish various sentiments for multiple aspects of the sentence. To address this challenge, the RepWalk neural network model was recently proposed (Zheng et al. 2020). It performs a replicated random walk on a syntax graph, effectively focusing on the descriptive contextual words.

Despite the fact that neural network-based models with attention, including BERT and contextualized embedding (Bahdanau, Cho, and Bengio 2015; Devlin et al. 2019; Peters et al. 2018), capture the semantic relatedness among words in the context, they still lack interpretability. This arguably makes them black box models (Rudin 2018). Many applications of attention mechanisms show, however, that a model can be interpreted based on the weight assigned by the attention vector to each input, but they do not provide a faithful explanation of classification (Serrano and Smith 2019; Wiegrefe and Pinter 2019). Many researchers have attempted to replicate human learning behavior in neural networks (Lei et al. 2019), but have failed to answer the question of making the learning interpretable. In order to overcome the issue of interpretability in NLP, we explore the recently introduced Tsetlin Machine (TM), which recognizes patterns in the form of propositional logic (Granmo 2018; Zhang et al. 2020). TMs have demonstrated promising results in various classification tasks involving numerical data, image data, text data, and board games (Granmo et al. 2019; Berge et al. 2019).

In this paper, we aim to reduce the gap between interpretability and accuracy with a significant margin on the ABSA task. To the best of our knowledge, this is the first




Figure 1: Representation of an aspect word and its surrounding words.

study using TM to explore how each word in the context includes or excludes themselves to form conjunctive clauses for sentiment classification. Once the model is trained, clauses in the TM hold the information about which individual features in the context take part in the sentiment classification of the aspect word.

Methodology

Input Binarization

For both datasets, the ABSA tasks have a context word and an aspect word whose polarity is to be classified. Usually, the sentiment of the aspect word is reflected by its surrounding words in a sentence, as shown in Figure 1. In this example, the aspect word “price” has positive sentiment due to the word “reasonable” in the context. Similarly, for “service”, the context word “poor” describes its negative sentiment. This reveals that the sentiment of aspect words heavily relies on its position in a sentence and thus position embedding (Gu et al. 2018) is necessary. Such embedding creates a probability distribution of the sentence based on the aspect word. Recently, position-aware modelling has shown promising results on ABSA tasks (Zhou et al. 2020).

Since TM requires binary inputs, to utilize TM for interpretability, the inputs must be binarized. It is challenging to incorporate the required position-based word relations in binary form, to allow for ABSA. In particular, since a TM does not employ any world knowledge like Word2vec (Mikolov et al. 2013), Elmo (Peters et al. 2018) or BERT (Devlin et al. 2019), so as to retain the interpretability of the model, we reduce the size of vocabulary by replacing the sentiment carrying words with a common token. Understandably, without pre-trained embeddings, a model cannot find the similarity between two semantically related words such as “excellent” and “good”. Hence, we adopt Opinion Lexicon (Hu and Liu 2004), which is a list of English positive and negative sentiment words. In more details, we replace every possible word in the dataset by the common token “positive” or “negative”, as shown in Figure 2. Such external knowledge also helps to reduce the vocabulary size thereby decreasing the sparsity of BOW representations.

Once the vocabulary size is determined, the context word and the aspect word can be converted into binary form, named as $BOW_{context}$ and BOW_{aspect} respectively. Since BOW in binary form does not consider the frequency of the replaced common tag (i.e., “positive” and “negative”), it becomes a rough representation of those tokens. In order to determine the location of these sentiment-carrying tokens,

The **price** is reasonable although the **service** is poor.



The **price** is **positive** although the **service** is **negative**.

Figure 2: Replacement of sentiment-carrying words with a common sentiment token using Opinion Lexicon.




Figure 3: 3-bit input feature representing the location of common sentiment-carrying tokens: negative, no sentiment, and positive.

the sentence is split into two parts, divided by the aspect word. More specifically, we create additional binary vectors LOC^1_{vec} and LOC^2_{vec} , representing the location of the common tokens. The dimension of LOC^1_{vec} and LOC^2_{vec} is three (the 1st bit: negative, the 2nd bit: no sentiment, the 3rd bit: positive) as shown in Figure 3. LOC^1_{vec} represents the presence of the common tokens “positive” or “negative” in the first part. If there are no sentiment tags, this is represented by “no sentiment”. Similarly, LOC^2_{vec} represents the presence of the common tokens in the second part.

After the pre-processing of inputs, we use SentiWordNet to obtain the sentiment score (SC) of the 1st part and the 2nd part of the split sentence. This involvement of such additional knowledge enrich the input information. We adopt the sentiment score in a 3-D binary form for each part of the sentence. The SC vector SC^1_{vec} for the 1st part of the context is given by Eq. (1). Similarly, vector SC^2_{vec} is utilized for the second part of the context.

$$SC^1_{vec} = \begin{cases} [0, 0, 1](positive), & \text{if } SC > 0, \\ [1, 0, 0](negative), & \text{if } SC < 0, \\ [0, 1, 0](no\ sentiment), & \text{if } SC = 0. \end{cases} \quad (1)$$

After processing all these binary representations, we concatenate them all to make a final input vector of size $(2n + 12)$ as shown in Figure 4.

The Tsetlin Machine Based ABSA

TM is a recent classification method that manipulates expressions in propositional logic based on a team of Tsetlin Automata (TA) (Granmo 2018). TA is a fixed structure deterministic automaton that learns the optimal action from a set of actions suggested by the environment. In TM, each input bit corresponds to two TAs, i.e., TA and TA' . TA controls the original bit of the input sample whereas TA' controls its negation. Here we use TA to represent a general Tsetlin automata that can be a TA or a TA' . Each TA corresponds to




Figure 4: Construction of binary input by concatenating all the pre-processed features.




Figure 5: The two-action TA and its transition in TM.

one literal. A literal here indicates an input bit or its negation. For example, if the bit represents the word “food”, TA controls “food” itself and then TA' handles “not food”. Any TA employed by a TM has two actions with $2N$ states in total, as shown in Figure 5. When it operates in states from 1 to N , action “exclude” is selected while action “include” is adopted for states from $N + 1$ to $2N$. For each iteration, a TA performs “include” or “exclude” based on the current state. This in turn triggers a reward or penalty. If a reward is received, the TA moves to the deeper side of the action whereas if it obtains a penalty, it moves towards the center and eventually jumps to the other side of the action. Clearly, a TA, through its actions, decides whether to include or exclude its corresponding literal.

TM has a novel game theoretic strategy that regulates a decentralized team of TAs. This strategy guides the TAs to learn an arbitrarily complex propositional formula by including or excluding certain literals. More specifically, the included literals, by the operation of conjunction, formulate clauses. Each clause, after training, is expected to capture a sub-pattern. The overall pattern is decided by summing up the output of all clauses for any unknown input. The architecture for ABSA using TM is shown in Figs. 6 and 7.

Let us consider the input feature as a vector with a vocabulary size of n words, which is represented in BOW as $X_s = [x_1, x_2, x_3, \dots, x_n, \dots, x_{2n}, x_{2n+1}, x_{2n+2}, \dots, x_{2n+12}]$ with $x_k \in \{0, 1\}$ and $k \in \{1, \dots, 2n + 12\}$. Here, $[x_{2n+1}, x_{2n+2}, x_{2n+3}]$ and $[x_{2n+4}, x_{2n+5}, x_{2n+6}]$ represent LOC^1_{vec} and LOC^2_{vec} respectively. Similarly, $[x_{2n+7}, x_{2n+8}, x_{2n+9}]$ and $[x_{2n+10}, x_{2n+11}, x_{2n+12}]$ represent SC^1_{vec} and SC^2_{vec} respectively. Let q be the number of classes ($q = 3$ in ABSA task: positive, neutral and negative). If a pattern has m sub-patterns, the pattern can be captured using $q \times m$ conjunctive clauses $C_i^j, 1 \leq j \leq q, 1 \leq i \leq m$:

$$C_i^j = \left(\bigwedge_{k \in I_i^j} x_k \right) \wedge \left(\bigwedge_{k \in \bar{I}_i^j} \neg x_k \right), \quad (2)$$

where I_i^j and \bar{I}_i^j are non-overlapping subsets of the input variable indices, $I_i^j, \bar{I}_i^j \subseteq \{1, \dots, 2n + 12\}, I_i^j \cap \bar{I}_i^j = \emptyset$. 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




Figure 6: TA team forms a Clause C_i^j by either including or excluding the input features.




Figure 7: (a). The sum of the votes for the clauses offers a score for a particular class. (b). Argmax operator decides the output class based on the score of the clauses in each class.

included as is, while the indices of input variables in \bar{I}_j^i correspond to the negated ones. Among m clauses in each class, clauses with odd indexes are assigned to positive polarity (+) whereas those with even indexes are assigned to negative polarity (-). The clauses with positive polarity vote for the target class and those with the negative vote against it. A summation operator aggregates them by subtracting the total number of negative votes from positive votes, as shown in Eq. (3).

$$f^j(X_s) = \sum_{i=1,3,\dots}^{m-1} C_i^j(X_s) - \sum_{i=2,4,\dots}^m C_i^j(X_s). \quad (3)$$

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, as shown in Eq. (4).

$$y = \operatorname{argmax}_j (f^j(X_s)). \quad (4)$$

The Learning Process of TM Based ABSA

In this section, we will detail the learning process of TM for the ABSA task. We explain the learning process with a walk-through of a specific sample context: “The food is very good and the place is clean as well”, using the aspect word “place” whose sentiment is to be predicted. The context is first changed to “The food is very **positive** and the place is **positive** as well.” For ease of explanation, we use the text word as a feature instead of the index in its binary form. For additional features, we will use the index of the binary input so as to differentiate the features that take part in classification. The indexes for additional features are $LOC_{vec}^1 =$




Figure 8: TAs with 100 states per action that learn whether to exclude or include a specific word (or its negation), location of common token (or its negation) and the sentiment score information (or its negation) in a clause at time step 1.




Figure 9: TAs with 100 states per action that learn whether to exclude or include a specific word (or its negation), location of common token (or its negation) and the sentiment score information (or its negation) in a clause at time step t .

$[2n + 1, 2n + 2, 2n + 3]$, $LOC_{vec}^2 = [2n + 4, 2n + 5, 2n + 6]$. Since the sentiment scores for both the first part of the context (“The food is very good and the”) and that for the second part (“is clean as well”) are greater than zero, we have $SC_{vec}^1 = [2n + 7, 2n + 8, 2n + 9] = [0, 0, 1]$, and $SC_{vec}^2 = [2n + 10, 2n + 11, 2n + 12] = [0, 0, 1]$, according to Eq. (1).

Figures 8, 9, and 10 show the learning process of the ABSA task with the TM model. The subscripts c and a in the figures represent the word from context and aspect respectively. The TA or TA' that received reward will move away from the center while those that received penalty will move towards the center. In this way, the TA (or TA') can be trained to either “include” or “exclude” a word (or its negation), helping the clauses, which are composed by the literals, learn different subpatterns. Consequently, the TM, composed by clauses, will gradually converge to the intended pattern. The feedback (reward or penalty) given to the TM follows two types: Type I and Type II feedback. Based on these feedback types, rewards or penalties are fed to the TA for the training samples. Type I Feedback is activated when a given input feature is either correctly assigned to the target sentiment (true positive) or mistakenly ignored (false negative). This feedback provides two countering effects: (1) involving more literals from the sample to refine the clauses; (2) trimming of the clauses by a factor specificity s that makes all clauses eventually evaluate to 1. The s -parameter is also responsible for avoiding overfitting. Type II Feedback is activated when an input feature is wrongly assigned to the target sentiment (false positive). It is responsible for introducing literals that make the clause evaluate to false, every time a false positive occurs. Type I Feedback and Type II Feedback are summarized in Tables 1 and 2 respectively.

Let us consider an example: a clause $C_1^1 = [food_c \wedge$

Input features	time	Clause formed and its learning in each step	Clause Output	Pred Class for true class = 1	Feedback type
the_c $food_c$ $very_c$ $positive_c$ $place_c$ $well_c$ $place_a$ $2n + 3$ $2n + 12$	1	$C_1^1 = food_c \wedge \neg positive_c \wedge place_c \wedge (2n + 3) \wedge \neg(2n + 12)$	$C_1^1 = 0$	$y = 0$	Feedback I
		$C_1^1 = 1 \wedge 0 \wedge 1 \wedge 1 \wedge 0$			
	2	$C_1^1 = food_c \wedge place_c \wedge (2n + 3)$	$C_1^1 = 1$	$y = 0$	Feedback I
		$C_1^1 = 1 \wedge 1 \wedge 1$			
	3	$C_1^1 = food_c \wedge positive_c \wedge place_c \wedge (2n + 3) \wedge (2n + 12)$	$C_1^1 = 1$	$y = 1$	Feedback I
		$C_1^1 = 1 \wedge 1 \wedge 1 \wedge 1 \wedge 1$			
⋮	⋮	⋮	⋮	⋮	⋮
t	$C_1^1 = food_c \wedge positive_c \wedge place_c \wedge place_a \wedge (2n + 3) \wedge (2n + 12)$	$C_1^1 = 1$	$y = 1$	Feedback I	
	$C_1^1 = 1 \wedge 1 \wedge 1 \wedge 1 \wedge 1 \wedge 1$				

Figure 10: The illustration of the clause update until reaching to an intended pattern at time step t .

Input	Clause Literal	1		0	
		1	0	1	0
Include Literal	P(Reward)	$\frac{s-1}{s}$	NA	0	0
	P(Inaction)	$\frac{1}{s}$	NA	$\frac{s-1}{s}$	$\frac{s-1}{s}$
	P(Penalty)	0	NA	$\frac{1}{s}$	$\frac{1}{s}$
Exclude Literal	P(Reward)	0	$\frac{1}{s}$	$\frac{1}{s}$	$\frac{1}{s}$
	P(Inaction)	$\frac{1}{s}$	$\frac{s-1}{s}$	$\frac{s-1}{s}$	$\frac{s-1}{s}$
	P(Penalty)	$\frac{s-1}{s}$	0	0	0

Table 1: The Type I Feedback (Granmo 2018).

Input	Clause Literal	1		0	
		1	0	1	0
Include Literal	P(Reward)	0	NA	0	0
	P(Inaction)	1.0	NA	1.0	1.0
	P(Penalty)	0	NA	0	0
Exclude Literal	P(Reward)	0	0	0	0
	P(Inaction)	1.0	0	1.0	1.0
	P(Penalty)	0	1.0	0	0

Table 2: The Type II Feedback (Granmo 2018).

$\neg positive_c \wedge place_c \wedge (2n + 3) \wedge \neg(2n + 12)$] that is formed at time step $t = 1$, as shown in Figure 8. Here, the time step indicates the instant the clause is updated during training iterations. The clause is composed by a combination of literals that are “included” by its associated TAs. At the current step, the excluded literals in this case (i.e., $2n + 12$, $place_a$, the_c , $well_c$) are controlled by TA , and the negated literals (i.e., $place_a$, $very_c$, $well_c$) are governed by TA' . Clearly, this clause evaluates to 0, thereby contributing to predict class 0 despite the true class being 1, as shown in Figure 10. This indeed triggers the Type I feedback. With Type I feedback, the reward or penalty for each literals is decided by Table 1. Since the literal $\neg positive_c$ is included, its feature is 0 (-1) and the clause output is 0. Therefore, it receives penalty for being included with the probability of $\frac{1}{s}$, making it slowly move towards the center and eventually jumping to the side with action “exclude”. Similarly, the literal $\neg(2n + 12)$ also receives the penalty with prob-

ability $\frac{1}{s}$, making it slowly moving towards the center, as well, eventually jumping to exclude action. Once this happens, the clause C_1^1 becomes $[food_c \wedge place_c \wedge (2n + 3)]$ as shown in time step $t = 2$ that outputs 1, making a prediction of class 0 as depicted in Figure 10. Table 1 shows if the clause output is 1, the literals are of value 1, and the actions of the literals are “excluded”, such literals obtain inaction or penalty with probability $\frac{1}{s}$ or $\frac{s-1}{s}$ respectively, making them slowly move towards the center and eventually jump to “include” action. Once it happens, the clause becomes $C_1^1 = [food_c \wedge positive_c \wedge place_c \wedge (2n + 3) \wedge (2n + 12)]$ as shown in time step $t = 3$. Based on reward and penalty, TM reaches to the intended pattern at time step t by the arrangement of literals controlled by their respective TAs, as shown in Figure 9. The final clause is given by $C_1^1 = [food_c \wedge positive_c \wedge place_c \wedge (2n + 3) \wedge place_a \wedge (2n + 12)]$. The clause will still obtain Type I feedback when more training samples are given and they reinforce the true positive oc-

Algorithm 1 Training Process of TM based ABSA

Require: Given Input = [Context sentence, Aspect Word, Sentiment Score]

- 1: Pre-processed Input = $\text{Concat}(BOW_{context}, BOW_{aspect}, LOC_{vec}^1, LOC_{vec}^2, SC_{vec}^1, SC_{vec}^2)$
- 2: Final Input: $X_{s,tr} = [x_1, \dots, x_n, \dots, x_{2n}, \dots, x_{2n+12}]$ and y $\triangleright y$ is the label of the input sample
- 3: Output: trained TM.
- 4: **for** Each training sample **do**
- 5: $\hat{y} = \text{TM}(X_{s,tr}, T, s)$ \triangleright Current sentiment estimate for the input sample
- 6: **if** $y = 1$ **then:**
- 7: **for** each clause C_i^j with odd index **do**
- 8: Use Type I Feedback($X_{s,tr}, \hat{y}, T, C_i^j, s$) to update all Tsetlin automata in C_i^j .
- 9: **end for**
- 10: **else** \triangleright if $y = 0$
- 11: **for** each clause C_i^j with odd index **do**
- 12: Use Type II Feedback($X_{s,tr}, \hat{y}, T, C_i^j, s$) to update all Tsetlin automata C_i^j .
- 13: **end for**
- 14: **end if**
- 15: **end for**
- 16: **return** Trained TM.

Algorithm 2 Testing Process of TM based ABSA

Require: Given Input = $X_{s,te}$

- 1: Output: predicted class
- 2: $f^j(X_{s,te}) = 0$, for all j
- 3: **for** all j **do** \triangleright For all classes
- 4: **for** all i in class j **do** \triangleright For all clauses in this class
- 5: $f^j(X_{s,te}) = f^j(X_{s,te}) + (-1)^{i+1} C_i^j(X_{s,te})$
- 6: **end for**
- 7: **end for**
- 8: **return** $\text{argmax}_j f^j(X_{s,te})$

currences until the sum of the votes by these clauses reaches a threshold parameter T .

The overall training and testing processes of TM-based ABSA are summarized in Algorithm 1 and Algorithm 2 respectively. For conciseness, we present, in Algorithm 1, the training procedure for the clauses with positive polarity, i.e., the clauses with odd index number. Clearly, the feedback types for the negative ones are just opposite. The complete training approach of a TM can be found in (Granmo 2018).

Once the class is predicted, we can explore its clauses for interpretability. The clauses that are triggered (i.e., $C_i^j(X_{s,te}) = 1$) are explored and their literals are converted into the original words for interpretation with the help of the additional information like LOC_{vec}^1 , LOC_{vec}^2 , SC_{vec}^1 , or SC_{vec}^2 .

Dataset	Positive	Negative	Neutral	Total
res14 (train)	2164	807	637	3608
res14 (test)	728	196	196	1120
lap14 (train)	994	870	464	2238
lap14 (test)	341	128	169	638

Table 3: The statistics of SemEval-2014 dataset.

Experiment Results

Datasets

The datasets are obtained from SemEval-2014 Task 4. The task has two domain-specific datasets, namely, Restaurant 14 (res14) and Laptop 14 (lap14). These datasets are provided with training and testing data. The statistics of the two datasets is shown in Table 3. The code and the datasets are available online¹.

Baselines

In our experiment, we evaluate the proposed method and compare it with related approaches for ABSA as baselines.

- **ContextAvg** averages the word embedding to form a context embedding (Tang, Qin, and Liu 2016).
- **LSTM** uses the last hidden vector of the LSTM for classification (Hochreiter and Schmidhuber 1997).
- **TD-LSTM** utilizes two LSTMs to learn the language model from the left and the right contexts of the aspect (Tang, Qin, and Liu 2016).
- **ATAE-BiLSTM** is an attention-based LSTM with Aspect Embedding model (Wang et al. 2016).
- **MemNet** integrates the content and the position of the aspect word into a deep neural network (Tang, Qin, and Liu 2016).
- **RAM** is a multi-layer architecture where each layer consists of attention-based aggregation of word features and a GRU cell (Chen et al. 2017).
- **IAN** is an Interactive Attention Network model that calculates the attention weights of the word in its sentiment and aspect interactively (Ma et al. 2017).
- **PRET+MULT** uses two approaches of transfer knowledge from document level using pretraining and multitask training (He et al. 2018).
- **HCSN** proposes a Human-like Semantic Cognition network for the ABSA task, motivated by the human beings' reading cognitive process (Lei et al. 2019). We show that performance of our proposed scheme is quite similar to this technique with high interpretability.
- **TNet** employs a CNN layer instead of attention layer to extract features from the transformed word representations originated from a bi-directional RNN layer (Li et al. 2018).

¹https://github.com/rohanky/tm_absa

- **AGDT** is an Aspect-Guided Deep Transition model that uses the given aspect to direct the sentence encoding from scratch with specially designed deep transition architecture. This model generates the aspect-based sentence representation and hence predicts sentiment more accurately (Liang et al. 2019).

Results

In our experiment, the main selling-point of the architecture is transparent learning and interpretability rather than accuracy. Better accuracy may be achieved when grid search is adopted. As we have used the integer weighted TM (Abeyrathna, Granmo, and Goodwin 2020), the parameters available are the number of clauses, the threshold T , and the specificity s , which are configured as 700, 90×100 , and 15 respectively for both datasets. For pre-processing of text, we substitute the short form to its full form, such as “isn’t” to “is not”. Additionally, we stem the words to reduce the vocabulary size created due to spelling mistakes and variants of words². The remaining pre-processing procedure has already been explained before. We train the TM model on both the datasets for 100 epochs each.

Since the output sentiment label has imbalanced training samples, we use two evaluation metrics: Accuracy and Macro-F1 (Manning and Schütze 2002). Following most of the related studies within the ABSA task, we report the best reproducible results by running the ABSA TM for 100 epochs, as shown in Table 4. We have reported the highest reproducible accuracy along with its mean and standard deviation obtained during 5 experiments. As we can see, Context2vec and LSTM perform quite poorly as they do not consider the aspect information when deciding the sentiment polarity. However, due to the consideration of left and right context information, TD-LSTM performs slightly better than LSTM. The variants of attention perform consistently better than LSTM and TD-LSTM. This is due to the fact that attention captures important information with regard to the aspect word. Other methods like RAM and MemNet perform slightly better because of the integrated memory in sentiment modeling. Another kind of the neural network-based model is HCSN. HCSN utilizes a human-being-like cognitive network for ABSA, which is motivated by the principles of human beings’ reading cognitive processes. Its pre-reading, active reading, and post-reading technique mimics the human behavior, which is then fed to the GRU network. As interesting as it seems, the involvement of the neural network still brings this below human-level interpretation on what drives the model to make the decision. Our model, which offers a transparent view of the learning process, obtains quite similar or higher accuracy compared to HCSN and PRET+MULT techniques. However, the TNet architecture with a CNN layer, which extracts salient features from transformed word representation, achieves higher accuracy compared to TM. AGDT is a model that uses Aspect guided GRU along with Max pooling to obtain Aspect Concatenated Embedding. It obtains quite similar accuracy compared to TM on Restaurant 14, whereas

²In this work, we adopt the Porter Stemmer.

accuracy is lower on Laptop 14. Note that we do not use any pre-trained word2vec or glove embedding for TM and our model still performs better than LSTM, TD-LSTM as well as attention based BiLSTM for both datasets. The Macro-F1 score shows that TM does not only greedily learn a particular class but also creates a set of features for each and every class. Even though the performance of our proposed model does not outperform the state-of-the-arts models, it reaches to comparable accuracy and Macro-F1 with transparent learning and interpretable prediction.

In addition to the above comparisons, we demonstrate here the necessity of including both *LOCs* and *SCs* vectors. First we only used the *LOCs* in the model and observed that the accuracy of the model reaches 76.51%. Secondly, we replaced *LOCs* with *SCs* and the performance of the model decreased to 75%. This shows that both vectors add useful information when employed together thereby reaching the stated accuracy of 78.02%.

To compare the performance of TM with classical interpretable models such as Logistic Regression (LR), we use our preprocessed *BOW* as input to LR. We observed that the TM performs better than LR in terms of accuracy. LR obtains the accuracy of 75.38% as compared with TM’s 78.02% on the Restaurant 14 dataset. Indeed, those two approaches operate based on different concepts. LR is trained by adjusting weights and bias. TMs, on the other hand, relates words using propositional logic to represent a class. Employing propositional logic for knowledge representation provides rules rather than a mathematical computation. This crucial difference between a rule-based approach and regression methods is explored in (Haghighi et al. 2016). One can analyze why a LR model assigns a particular class to an input by inspecting the weights and bias. However, assigning them meanings requires understanding of the mathematical computation that LR carries out. Since TM creates a list of patterns for a particular class based on the interaction of aspect words and the sentiment words in the context, its conjunctive clauses hold information of words in a rule-based form. It is well-known that evaluating a conjunctive clause is particularly easy for humans, making them natively interpretable and easier to explain than LR.

Interpretability and Analysis

Characteristics of Clauses

In this section, we will explain one phenomenon of a TM after training with the datasets. When analyzing the clauses after training, we noticed that the TM employs more negated literals to form a clause. This is a bit counter intuitive as there should be, intuitively, more literals in their original forms than their negations in a clause. To explain this behavior, let us study two general sentences having positive sentiments for aspect word “laptop”:

- This laptop is in excellent condition.
- Battery life of this laptop is better compared to other brands.

In this example, we assume a vocabulary containing both positive and negative sentiment words of size 6, $V = [\text{excel-}$

Methods	Restaurant 14		Laptop 14	
	Accuracy	Macro-F1	Accuracy	Macro-F1
ContextAvg	71.5	58.0	61.5	53.9
LSTM	74.3	63.0	66.5	60.1
TD-LSTM	75.6	64.5	68.1	63.9
ATAE-BiLSTM	77.6	65.3	68.7	64.2
MemNet	76.9	66.4	68.9	62.8
RAM	78.5	68.5	72.1	68.4
IAN	78.6	NA	72.1	NA
PRET+MULT	79.1	69.7	71.2	67.5
HCSN	77.8	70.2	76.1	72.5
TNet	80.79	70.84	76.01	71.47
AGDT	78.85	NA	71.50	NA
TM based ABSA	78.02 (76.40 ± 1.0)	67.85 (64.01 ± 0.8)	73.51 (71.47 ± 0.9)	70.82 (67.48 ± 1.5)

Table 4: Experiment results of various approaches for SemEval-2014 dataset. The upper results show the best reproducible accuracy and lower ones represent the mean and standard deviation of the last 50 epochs when running the model for five times.

lent, bad, condition, worst, costly, better, laptop]. When literals in their original forms are utilized to compose a clause, the two sentences require two clauses to follow the sentiment, i.e., $C_1 = [laptop \wedge excellent \wedge condition]$ and $C_2 = [battery \wedge better \wedge laptop \wedge others]$. On the contrary, when negated form of literals are employed, only one clause is sufficient to satisfy the sentiment in both sentences: $C_1 = [laptop \wedge \neg bad \wedge \neg worst \wedge \neg costly]$. As the negation is a more efficient way to represent a pattern in NLP, the trained TM employs naturally more negated literals to form a clause.

A Case Study for Interpretability

In this case study, we demonstrate the interpretable result from trained model. We randomly select a sentence from the dataset as an example and demonstrate its literals that are responsible to form the clause. The selected sentence is “Looks nice, but has a horribly cheap feel.” with a aspect word “looks” whose sentiment prediction of TM is positive. The sentence after pre-processing becomes “Looks positive, but has a negative negative feel.” Among various clauses that are triggered by the given input, we randomly select a clause for interpretation. The clause is given by:

- $C_i^j = positive \wedge (2n+6) \wedge \neg(\text{words not in the sentence and aspect})$.

The above clause can be interpreted as: the aspect word “looks” has positive sentiment because it has words “positive” and it lies in the second part of the sentence (indicated by $2n + 6$, i.e. $LOC_{vec}^2 = [0, 0, 1]$) when split from aspect word “looks”. Similarly, if the aspect word in the sentence is “feel” then its sentiment is predicted to be negative and a randomly selected clause is:

- $C_i^j = negative \wedge (2n + 1) \wedge \neg(\text{words not in the sentence and aspect})$.

This clause means that the sentiment is negative because it has words like “negative” and it lies in the first part of the sentence (indicated by $2n + 1$, i.e. $LOC_{vec}^1 = [1, 0, 0]$)




Figure 11: Interpretation of a randomly selected sample from ABSA task.

when split from aspect word “feel”. In both the cases, $\neg(\text{words not in the sentence and aspect})$ represents the words in negated form that are presented in the input features. Finally, reversing back all the information and binarization to the original form of the words, we can obtain interpretation that shows the influence of words in the classification as in Figure 11.

Conclusions

In this paper, we aim to reduce the gap between the interpretability and the accuracy of aspect based sentiment analysis (ABSA) by employing the recently introduced Tsetlin Machine (TM). Our proposed model embeds the aspect-based inputs into binary form for classifying the sentiment of a particular word in a sentence. Such binary representations are then fed to a TM architecture where the learning process is transparent, which gives a clear picture of what actually drives the TM to learn the particular sentiment for a given input. Additionally, we show the involvement of words carrying the sentiment for the aspect words in the case study. In short, the proposed model successfully provides an human-interpretable learning approach on ABSA task with comparable accuracy.

References

- Abeyrathna, K. D.; Granmo, O.-C.; and Goodwin, M. 2020. Extending the Tsetlin Machine With Integer-Weighted Clauses for Increased Interpretability. *arXiv* 2005.05131.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *ICLR*. California, USA.
- Berge, G. T.; Granmo, O.-C.; Tveit, T. O.; Goodwin, M.; Jiao, L.; and Matheussen, B. V. 2019. Using the Tsetlin Machine to Learn Human-Interpretable Rules for High-Accuracy Text Categorization With Medical Applications. *IEEE Access* 7: 115134–115146.
- Brunner, G.; Liu, Y.; Pascual, D.; Richter, O.; Ciaramita, M.; and Wattenhofer, R. 2020. On Identifiability in Transformers. In *ICLR*. Addis Ababa, Ethiopia.
- Chen, P.; Sun, Z.; Bing, L.; and Yang, W. 2017. Recurrent Attention Network on Memory for Aspect Sentiment Analysis. In *EMNLP*, 452–461. Copenhagen, Denmark: ACL.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*, 4171–4186. Minneapolis, Minnesota: ACL.
- Do, H. H.; Prasad, P.; Maag, A.; and Alsadoon, A. 2019. Deep Learning for Aspect-Based Sentiment Analysis: A Comparative Review. *Expert Systems with Applications* 118: 272–299.
- Esuli, A.; and Sebastiani, F. 2006. SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining. In *LREC*. Genoa - Italy.
- Granmo, O.-C. 2018. The Tsetlin Machine - A Game Theoretic Bandit Driven Approach to Optimal Pattern Recognition with Propositional Logic. *ArXiv* abs/1804.01508.
- Granmo, O.-C.; Glimsdal, S.; Jiao, L.; Goodwin, M.; Omlin, C. W.; and Berge, G. T. 2019. The Convolutional Tsetlin Machine. *arXiv* 1905.09688.
- Gu, S.; Zhang, L.; Hou, Y.; and Song, Y. 2018. A Position-aware Bidirectional Attention Network for Aspect-level Sentiment Analysis. In *COLING*, 774–784. Santa Fe, New Mexico, USA: ACL.
- Haghighi, M.; Johnson, S.; Qian, X.; Lynch, K.; Vehik, K.; and S. Huang, a. T. S. G. 2016. A Comparison of Rule-based Analysis with Regression Methods in Understanding the Risk Factors for Study Withdrawal in a Pediatric Study. *Scientific Reports* 6.
- He, R.; Lee, W. S.; Ng, H. T.; and Dahlmeier, D. 2018. Exploiting Document Knowledge for Aspect-level Sentiment Classification. In *ACL*, 579–585. Melbourne, Australia: ACL.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation* 9: 1735–1780.
- Hu, M.; and Liu, B. 2004. Mining and Summarizing Customer Reviews. In *ACM SIGKDD*, 168–177. New York, NY, USA: ACM.
- Jiang, L.; Yu, M.; Zhou, M.; Liu, X.; and Zhao, T. 2011. Target-dependent Twitter Sentiment Classification. In *ACL*, 151–160. Portland, Oregon, USA: ACL.
- Kiritchenko, S.; Zhu, X.; Cherry, C.; and Mohammad, S. 2014. NRC-Canada-2014: Detecting Aspects and Sentiment in Customer Reviews. In *International Workshop on Semantic Evaluation (SemEval 2014)*, 437–442. Dublin, Ireland: ACL.
- Lei, Z.; Yang, Y.; Yang, M.; Zhao, W.; Guo, J.; and Liu, Y. 2019. A Human-Like Semantic Cognition Network for Aspect-Level Sentiment Classification. In *AAAI*. Hawaii, USA.
- Li, X.; Bing, L.; Lam, W.; and Shi, B. 2018. Transformation Networks for Target-Oriented Sentiment Classification. In *ACL*, 946–956. Melbourne, Australia: ACL.
- Liang, Y.; Meng, F.; Zhang, J.; Xu, J.; Chen, Y.; and Zhou, J. 2019. A Novel Aspect-Guided Deep Transition Model for Aspect Based Sentiment Analysis. In *EMNLP-IJCNLP*, 5569–5580. Hong Kong, China: ACL.
- Liu, J.; and Zhang, Y. 2017. Attention Modeling for Targeted Sentiment. In *EACL*, 572–577. Valencia, Spain: ACL.
- Ma, D.; Li, S.; Zhang, X.; and Wang, H. 2017. Interactive Attention Networks for Aspect-Level Sentiment Classification. In *IJCAI*, 4068–4074. Melbourne, Australia.
- Manning, C. D.; and Schütze, H. 2002. Foundations of Statistical Natural Language Processing. In *SGMD*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*. Stateline, United States: Curran Associates, Inc.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep Contextualized Word Representations. In *NAACL*, 2227–2237. New Orleans, Louisiana: ACL.
- Pontiki, M.; Galanis, D.; Pavlopoulos, J.; Papageorgiou, H.; Androutsopoulos, I.; and Manandhar, S. 2014. SemEval-2014 Task 4: Aspect Based Sentiment Analysis. In *International Workshop on Semantic Evaluation (SemEval 2014)*, 27–35. Dublin, Ireland: ACL.
- Rudin, C. 2018. "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead". *Nature Machine Intelligence* 1: 206–215.
- Samek, W.; Montavon, G.; Vedaldi, A.; Hansen, L.; and Müller, K. 2019. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Springer.
- Serrano, S.; and Smith, N. A. 2019. Is Attention Interpretable? In *ACL*, 2931–2951. Florence, Italy: ACL.
- Shen, T.; Zhou, T.; Long, G.; Jiang, J.; Pan, S.; and Zhang, C. 2018. DiSAN: Directional Self-Attention Network for RNN/CNN-free Language Understanding. In *AAAI*. New Orleans, USA.

- Tang, D.; Qin, B.; Feng, X.; and Liu, T. 2016. Effective LSTMs for Target-Dependent Sentiment Classification. In *COLING*, 3298–3307. Osaka, Japan: ACL.
- Tang, D.; Qin, B.; and Liu, T. 2016. Aspect Level Sentiment Classification with Deep Memory Network. In *EMNLP*, 214–224. Austin, Texas: ACL.
- Vashishth, S.; Upadhyay, S.; Tomar, G. S.; and Faruqui, M. 2019. Attention Interpretability Across NLP Tasks. *arXiv* 1909.11218.
- Wang, Y.; Huang, M.; Zhu, X.; and Zhao, L. 2016. Attention-based LSTM for Aspect-level Sentiment Classification. In *EMNLP*, 606–615. Austin, Texas: ACL.
- Wiegrefe, S.; and Pinter, Y. 2019. Attention is not not Explanation. In *EMNLP-IJCNLP*, 11–20. Hong Kong, China: ACL.
- Zhang, L.; and Liu, B. 2017. *Sentiment Analysis and Opinion Mining*, 1152–1161. Boston, MA: Springer US.
- Zhang, X.; Jiao, L.; Granmo, O.-C.; and Goodwin, M. 2020. On the Convergence of Tsetlin Machines for the IDENTITY- and NOT Operators. *arXiv* 2007.14268.
- Zheng, Y.; Zhang, R.; Mensah, S.; and yi Mao, Y. 2020. Replicate, Walk, and Stop on Syntax: An Effective Neural Network Model for Aspect-Level Sentiment Classification. In *AAAI*. New York, USA.
- Zhou, J.; Chen, Q.; Huang, X.; Hu, Q.; and He, L. 2020. Position-aware hierarchical transfer model for aspect-level sentiment classification. *Information Sciences* 513: 1–16.