



SentiQ: A Probabilistic Logic Approach to Enhance Sentiment Analysis Tool Quality

Wissam Mammar Kouadri, Salima Benbernou, Mourad Ouziri, Themis Palpanas,
Iheb Benamor

► To cite this version:

Wissam Mammar Kouadri, Salima Benbernou, Mourad Ouziri, Themis Palpanas, Iheb Benamor. SentiQ: A Probabilistic Logic Approach to Enhance Sentiment Analysis Tool Quality. WISDOM 2020 : The 9th KDD Workshop on Issues of Sentiment Discovery and Opinion Mining, Aug 2020, San Diego, United States. ⟨hal-03299105⟩

HAL Id: hal-03299105

<https://hal.science/hal-03299105v1>

Submitted on 26 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

SentiQ: A Probabilistic Logic Approach to Enhance Sentiment Analysis Tool Quality

Wissam Mammar Kouadri, Wissam Kouadri, Salima Benbernou, Mourad Ouziri, Themis Palpanas, Iheb Benamor

► To cite this version:

Wissam Mammar Kouadri, Wissam Kouadri, Salima Benbernou, Mourad Ouziri, Themis Palpanas, et al.. SentiQ: A Probabilistic Logic Approach to Enhance Sentiment Analysis Tool Quality. WISDOM 2020: The 9th KDD Workshop on Issues of Sentiment Discovery and Opinion Mining, Aug 2020, San Diego, United States. 10.1145/1122445.1122456 . hal-03299105

HAL Id: hal-03299105

<https://hal.archives-ouvertes.fr/hal-03299105>

Submitted on 26 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SentiQ: A Probabilistic Logic Approach to Enhance Sentiment Analysis Tool Quality

Wissam Maamar Kouadri
wissam.maamar_kouadri@u-paris.fr
Université de Paris

Salima Benbernou
salima.benbernou@u-paris.fr
Université de Paris

Mourad Ouziri
mourad.ouziri@u-paris.fr
Université de Paris

Themis Palpanas
themis.palpanas@u-paris.fr
Université de Paris & Institut
Universitaire de France (IUF)

Iheb Ben Amor
Iheb.Benamor@imbaconsulting.fr
IMBA Consulting

ABSTRACT

The opinion expressed in various Web sites and social-media is an essential contributor to the decision making process of several organizations. Existing sentiment analysis tools aim to extract the polarity (i.e., positive, negative, neutral) from these opinionated contents. Despite the advance of the research in the field, sentiment analysis tools give *inconsistent* polarities, which is harmful to business decisions. In this paper, we propose SentiQ, an unsupervised Markov logic Network-based approach that injects the semantic dimension in the tools through rules. It allows to detect and solve inconsistencies and then improves the overall accuracy of the tools. Preliminary experimental results demonstrate the usefulness of SentiQ.

CCS CONCEPTS

• **Machine learning** → **Text labelling**; *Neural network*; *Data quality*; • **Information systems** → First order logic.

KEYWORDS

sentiment analysis, inconsistency, data quality, Markov logic network, logical inference

ACM Reference Format:

Wissam Maamar Kouadri, Salima Benbernou, Mourad Ouziri, Themis Palpanas, and Iheb Ben Amor. 2020. SentiQ: A Probabilistic Logic Approach to Enhance Sentiment Analysis Tool Quality. In *Proceedings of Wisdom '20: Workshop on Issues of Sentiment Discovery and Opinion Mining (Wisdom '20)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

With the proliferation of social media, people are increasingly sharing their sentiments and opinions online about products, services, individuals, and entities, which has spurred a growing interest in sentiment analysis tools in various domains [8, 11, 13, 14, 17, 25, 27]. The customer opinion, if yielded correctly, is crucial for the decision-making of any organization. Thus, numerous studies [3, 13, 16, 23, 24] try to automate sentiment extraction from a massive volume of data by identifying the polarity of documents, i.e., positive, negative, or neutral.

Nevertheless, sentiment analysis of social media data is still a challenging task [10] due to the complexity and variety of natural language through which the same idea can be expressed and interpreted using different text. Many research work have adopted the consensus that semantically equivalent documents should have the same polarity [3, 5, 12, 22, 26, 28]. For instance [5] have attributed the same polarity labels to the semantically equivalent couples (event/effect) while [12] have augmented their sentiment dataset using paraphrases and assign the original document's polarity to the generated paraphrases.

However, we found that these tools do not detect this similarity and assign different polarity labels to semantically equivalent documents; hence, considering *in-tool* inconsistency where the sentiment analysis tool attribute different polarities to the semantically equivalent documents and *inter-tool inconsistency* where different sentiment analysis tools attribute different polarities to the same documents that have a single polarity. This inconsistency can be translated by the fact that at least one tool has given an incorrect polarity. Consequently, returning an incorrect polarity in the query can be misleading, and leads to poor business decision.

Few works have used inconsistencies to improve systems' accuracy, such as [20], that considers various labeling functions and minimizes the inter-tool inconsistency between them based on different factors: correlation, primary accuracy, and labelling abstinence. However, in [20], we resolve the inconsistency statistically, and ignore the semantic dimension that could enhance the results' quality. The work in [5] has proposed to create a corpus of (event/effect) pairs for sentiment analysis by minimizing the sentiment distance between semantically equivalent (event/effect) pairs. In our work, we study the effect of solving the two types of inconsistency on accuracy. We focus more on the improvement that we can obtain by

resolving in-tool inconsistency between the documents i.e., resolving inconsistency such that all semantically equivalent documents get the same polarity label and resolving both inconsistencies. To the best of our knowledge, the only work studying polarity inconsistency does this at word-level [9], by checking the polarity consistency for sentiment words inside and across dictionaries.

Our work is the first to study the effect of resolving the polarity inconsistency on accuracy for in-tool inconsistency, *and* inter-tool inconsistency on document data. We seek to converge to the golden truth by resolving in-tool and inter-tool inconsistencies. Each document has a unique polarity, by resolving in-tool and inter-tool inconsistency, we minimize the gap of incorrect labels and converge to the gold truth. Such a method can be applied on any classification task in natural language processing.

Contributions. In summary, we make the following contributions:

- We study the impact of inconsistency on the accuracy of the sentiment analysis tools.
- We propose SentiQ, an approach that resolves both polarity inconsistencies: in-tool and inter-tool. The approach we are proposing is based on our earlier work to handle the inconsistency in big data [2] on one side and on the probabilistic logic framework, Markov Logic Network, on the other side.
- We present preliminary experimental results using news headlines datasets [4] and the sentiment treebank dataset [24]. When compared to the majority voting to resolve inter-tool inconsistencies, our framework leads to the efficiency of using the semantic dimension in optimizing the accuracy by resolving both in-tool and inter-tool inconsistencies.
- Following the lessons learned from our experimental evaluation, we discuss promising future research directions, including the semantic dimension's use in different integration problems, such as truth inference in crowd-sourcing and accuracy optimization of different classification problems.

Paper Outline. In the remainder of the paper, we present in section 2 a motivation through a real example. In section 3, we provide some preliminaries used in our work. In sections 4 and 5, we discuss the SentiQ model based on Markov Network logic (MLN) while in section 6, we present our experiments and discussions.

2 MOTIVATING EXAMPLE

We consider the following real life example collected from twitter and that represents statements about Trump's restrictions on Chinese technology such that $D = \{d_1, \dots, d_9\}$ and:

- d_1 : Chinese technological investment is the next target in Trump's crackdown.
- d_2 : Chinese technological investment in the US is the next target in Trump's crackdown.
- d_3 : China urges end to United States crackdown on Huawei.
- d_4 : China slams United States over unreasonable crackdown on Huawei.
- d_5 : China urges the US to stop its unjustifiable crackdown on Huawei.
- d_6 : Trump softens stance on China technology crackdown.
- d_7 : Donald trump softens threat of new curbs on Chinese investment in American firms.
- d_8 : Trump drops new restrictions on China investment.

A_i	Id	P_{tb}	P_{sw}	P_o	P_h
A_1	d_1	Neutral	Negative	Neutral	Negative
	d_2	Negative	Negative	Neutral	Negative
A_2	d_3	Negative	Positive	Neutral	Negative
	d_4	Negative	Negative	Neutral	Negative
	d_5	Negative	Negative	Neutral	Negative
A_3	d_6	Neutral	Positive	Neutral	Positive
	d_7	Negative	Negative	Negative	Positive
	d_8	Negative	Positive	Neutral	Positive
	d_9	Neutral	Negative	Neutral	Positive

Table 1: Predicted polarity on dataset D by different tools

- d_9 : Donald Trump softens tone on Chinese investments.

We call each element of this dataset D a *Document*. We notice that D can be clustered on subsets of semantically equivalent documents. For instance, d_1 and d_2 are semantically equivalent as they both express the idea that the US is restricting Chinese technological investments. We denote this set by A_1 and we write: $A_1 = \{d_1, d_2\}$ and $A_2 = \{d_3, d_4, d_5\}$, which express that the Chinese government demands the US to stop the crackdown on Huawei, and $A_3 = \{d_6, \dots, d_9\}$ which conveys the idea that Trump reduces restrictions on Chinese investments. We have: $D = A_1 \cup A_2 \cup A_3$. We analyse D using three sentiment analysis tools: Stanford Sentiment Treebank [24], Sentiwordnet [1] and Vader [13]. In the rest of this paper, we refer to the results of these tools using the polarity functions: P_{tb} , P_{sw} , P_o ; we use P_h to refer to the ground truth. Table 1 summarizes the results of the analysis.

We know that each document has a single polarity, so each precise tool should find this polarity, and a difference in prediction results is a sign that at least one tool is erroneous on this document. We also know that semantically equivalent documents should have the same polarity. However, in this real-life example, we observe different tools attributing different polarities for the same document (e.g., only P_{tb} attributes the correct polarity to d_3 in A_2), which represent an inter-tool inconsistency. Also, the same tool attributes different polarities for semantically equivalent documents (for e.g., P_{tb} considers d_6 as Neutral and d_7 as Negative) which represent an in-tool inconsistency. A trivial method to resolve those inconsistencies is to use majority voting, inside the cluster of documents, or between functions. However, when applying the majority voting baseline on this example, we found that the polarity is *Negative* in A_2 which represents the correct polarity of the cluster while we found that the polarity is *Negative* in A_3 , which is not a correct polarity in this case. Because with simple majority voting, we got only a local vision of the polarity function, and we ignore its behavior on the rest of the data.

3 PRELIMINARIES

Definition 3.1. (Sentiment Analysis)

Sentiment Analysis is the process of extracting a polarity $\pi \in \{+, -, 0\}$ from a document d_i . With + for Positive polarity, - for Negative polarity and 0 for Neutral polarity. In this paper, we refer to polarity functions as P_{tk} s.t. $P_{tk} : D \rightarrow \pi$. We refer to the set of all functions as Π s.t. $\Pi = \{P_{t_1}, \dots, P_{t_n}\}$

Definition 3.2. (Polarity Consistency)

Cluster: cluster is a set of semantically equivalent documents:

for a cluster $A_l = \{d_1, \dots, d_n\}$ we have $\forall d_i, d_j \in A_l, d_i \xleftrightarrow{s} d_j$.

Sentiment Quality: we define the polarity consistency of a given cluster A_l as the two following rules:

In-tool Consistency means that semantically equivalent documents should get the same polarity, s.t.:

$$\forall d_i, d_j \in A, P_* \in \Pi P_*(d_i) = P_*(d_j) \quad (1)$$

Inter-tool Consistency means that all polarity functions should give the same polarity to the same document:

$$\forall d_i \in A, P_{t_k}, P_{t'_k} \in \Pi P_{t_k}(d_i) = P_{t'_k}(d_i) \quad (2)$$

Definition 3.3. (Markov Logic Network (MLN))

In this section, We recall Markov logic network (MLN) model [6, 21] which is a general framework for joining logical and Probability. **MLN** is defined as a set of weighted first-order logic (FOL) formula with free variables $L = \{(l_1, w_1), \dots, (l_n, w_n)\}$, with $w_i \in \mathbb{R} \cup \infty$ and l_i an FOL constraint. With a set of constants $C = \{c_1, \dots, c_m\}$, it constitutes the Markov network $M_{L,C}$. The $M_{L,C}$ contains one node for each predicate grounding that its value is 1 if the grounding is true and 0 otherwise. Each formula of L is represented by a feature node that its value is 1 if the formula l_i grounding is true and 0 otherwise. The syntax of the formulas that we adopted in this paper is the FOL syntax.

World x over a domain C is a set of possible grounding of MLN constraints over C .

Hard Constraints are constraints with infinite weight $w_i = \infty$. A world x that violates these constraints is impossible.

Soft Constraints are constraints with a finite weight ($w_i \in \mathbb{R}$) that can be violated.

World's Probability is the probability distribution of possible worlds x in $M_{L,C}$ given by

$$Pr(X = x) = \frac{1}{Z} \exp\left(\sum_i w_i, n_i(x)\right)$$

, where $n_i(x)$ is the number of the true grounding of F_i in x and Z is a normalization factor.

Grounding. We define grounding as the operation of replacing predicate variables by constants from C .

4 SENTIQ: AN MLN BASED MODEL FOR INCONSISTENCY

The polarity inconsistency is a complex problem due to the tool and document natures and the relations between them. This problem can be solved using semantics to model the relations between tools, documents, and statistic dimension to optimize both the inconsistency and the accuracy of the system —this why we chose *MLN* to model the resulted inconsistent system. We present the details of our semantic model in this section.

4.1 Semantic Model's Components

Our semantic model is a knowledge-base $KB = \langle R, F \rangle$, where

- (1) R is a *set of rules* (FOL formulas) defining the vocabulary of our application which consists of concepts (sets of individuals) and relations between them.
- (2) F is a *set of facts* representing the instances of the concepts or individuals defined in R .

We represent each document by the concept *Document*, each polarity function in the system by its symbol and the polarity that it attributes to the *Document*. For instance, $P_{tb+}(d_1)$, P_{tb0} , and P_{tb-} represent respectively the polarities (+, 0, -) attributed to the *Document*(d_1) by the polarity function P_{tb} . Each *Document* is *Positive*, *Negative*, or *Neutral*. This is represented respectively by the concepts *IsPositive*, *IsNegative*, and *IsNeutral*. We also have the relation *sameAs* as a semantic similarity between documents in the input dataset clusters. For instance, *sameAs*(d_1, d_2) indicates that the documents *Document*(d_1) and *Document*(d_2) are semantically equivalent.

4.2 Rule modeling for inconsistency

We define two types of rules from R in our framework, *Inference rules* and *Inconsistency rules*:

Inference rules IR

The inference rules allow deriving the implicit instances. They model the quality of the polarity at in-tool and inter-tool levels. They are soft rules that add an uncertainty layer to different polarity functions based on the inconsistency of tools.

• **In-tool consistency rules.** This set of rules models the fact that all the documents of the cluster should have the same polarity. They are defined as follows (for the sake of clarity we omitted the predicate *Document*(d_i) in all logical rules) :

$$IR1 : sameAs(?d_i, ?d_j) \wedge IsPositive(?d_j) \rightarrow IsPositive(?d_i)$$

$$IR2 : sameAs(?d_i, ?d_j) \wedge IsPositive(?d_i) \rightarrow IsPositive(?d_j)$$

The rule *IR1* denotes that if two documents d_i and d_j are semantically equivalent (expressed with *sameAs* relation), they got the same polarity, which translates the in-tool consistency defined in equation 1. The *sameAs* relation is transitive, symmetric, and reflexive. We express the symmetry by duplicating the rule for both documents of the relation (rules *IR1* and *IR2* instead of only one rule). For instance, when applying the rule *IR1* on the relation *sameAs*(d_1, d_2) and the instances *IsNeutral*(d_1) and *IsNegative*(d_2), we infer the new instance *IsNeutral*(d_2). The instance *IsNegative*(d_1) is inferred when applying the rule *IR2*. The transitivity is handled in the instantiating step (algorithm 1) and we ignore the reflexivity of the relation because it does not infer additional knowledge. Note that *IR1* and *IR2* are examples of rules. The set of rules is presented in Algorithm 2.

• **Inter-tool consistency rules.** These rules model the inter-tool consistency described in equation 2 by assuming that each function gives the correct polarity to the document. For example, given the instances $P_{tb-}(d_2)$ the rule *IR* infers *IsNegative*(d_2). For each tool in the system, we create the following rules by replacing P_{t_k*} with the polarity function of the tool.

$$IR3 : P_{t_k+}(?d_i) \rightarrow IsPositive(?d_i) \quad IR4 : P_{t_k-}(?d_i) \rightarrow IsNegative(?d_i)$$

$$IR5 : P_{t_k0}(?d_i) \rightarrow IsNeutral(?d_i)$$

Those rules are soft rules that allow us to represent inconsistencies in the system and attribute a ranking to the rules that we use in the in-tools uncertainty calculation. The idea behind this modeling is that if the inter-tool consistency is respected, all tools will attribute the same polarity to this document; otherwise, the document will have different polarities (contradicted polarities). To represent this contradiction, we define, next, inconsistency rules.

Inconsistency rules ICR

They are considered as hard rules that represent the disjunction between polarities since each document has a unique polarity.

$$ICR1 : IsPositive(?d_i) \rightarrow \neg IsNegative(?d_i) \wedge \neg IsNeutral(?d_i)$$

$$ICR2 : IsNegative(?d_i) \rightarrow \neg IsPositive(?d_i) \wedge \neg IsNeutral(?d_i)$$

$$ICR2 : IsNeutral(?d_i) \rightarrow \neg IsPositive(?d_i) \wedge \neg IsNegative(?d_i)$$

These rules generate negative instances that create inconsistencies used in learning inference rules weights.

For instance, consider the following instances $P_{tb-}(d_3)$ and $P_{sw+}(d_3)$ from the motivating example. By applying the inter-tool consistency inference rules, we infer: $IsNegative(d_3)$ and $IsPositive(d_3)$. However, F appears consistent even it contains polarity inconsistencies. We get the inconsistency once applying the inconsistency rules. We get: $\neg IsPositive(d_3)$, $\neg IsNeutral(d_3)$, $\neg IsNegative(d_3)$ that represent an apparent inconsistency in F .

4.3 MLN based model for inconsistency resolution

As depicted in Figure 1, the proposed inconsistency resolution process follows four main phases:

- Inference of implicit knowledge: The system infers all implicit knowledge needed for the inconsistencies before applying the learning procedure of *MLN*.
- Detection of inconsistencies: Having the explicit and implicit knowledge, we apply the inconsistency rules to discover the inconsistency.
- Inference of correct polarity: Using the saturated fact set F and R , the system learns first the weights of $M_{R,F}$, and use them to infer the correct polarities.
- Resolve the in-tool inconsistencies : Since we are in an uncertain environment, we can still have some in-tool inconsistencies after the previous phase, that we resolve by applying a weighted majority voting.

The phases will be detailed in the next section.

5 SENTIQ:THE INCONSISTENCY RESOLUTION

In this section we discuss the reasoning process to solve the inconsistencies and improve the accuracy.

5.1 Facts generation

Our data are first saved in a relational database, where each table represents a concept, and the table content represents the concept's domain. For that, instantiating our data follows the steps of Algorithm 1.

Each function and its polarity is represented by a table. The content of the table is the document ID that got this polarity by the

function. The instantiating process converts the content of the database to logic predicates that we use in our reasoning. The purpose of this algorithm is to fill in the set F with the prior knowledge needed in the reasoning. Our prior knowledge is the documents, polarities attributed by the functions to documents, and the semantic similarity between documents represented by the *SameAs* predicate. We note that we do not consider the ground truth. We adopt an unsupervised approach because inconsistency resolution is useful when we do not know the valid prediction from the invalid ones.

Algorithm 1 Instantiating

Input : Database with prior knowledge

Output : F :Set of generated Facts (polarities and same as)

```

1: procedure INSTANTIATING
2:   //Step1: Add all Polarities attributed to documents
3:   for each  $P_{t_k} \in Functions$  :
4:     for each  $d_i \in P_{t_k}^+$  :  $F.add(P_{t_k}^+(d_i))$ 
5:     for each  $d_i \in P_{t_k}^-$  :  $F.add(P_{t_k}^-(d_i))$ 
6:     for each  $d_i \in P_{t_k}^0$  :  $F.add(P_{t_k}^0(d_i))$ 
7:   //Step2: Add sameAs relations
8:    $clusters = groupeByClusterId(D)$ 
9:   for each  $cluster \in clusters$  :
10:    for  $i \in \{0, \dots, len(cluster)\}$  :
11:      for  $j \in \{i+1, \dots, len(cluster)\}$  :
12:        if  $SameAs(d_i, d_j) \notin F$  :
13:           $F.add(SameAs(d_i, d_j))$ 
14:   return  $F$ 
```

5.2 Implicit knowledge inference Algorithm

In *MLN*, the learning is done only on the available knowledge in F . For this, we infer all implicit knowledge in the system before applying the learning procedure. The inference procedure is presented in Algorithm 2. This inference phase is crucial for an integrated learning since most polarity knowledge are implicit. For instance, consider the two documents d_3 and d_4 from the motivating example. We have $P_{sw+}(d_3)$ and $P_{sw-}(d_4)$, by inferring documents polarities using inter-tool consistency rules *IR3* and *IR4*, we get $IsPositive(d_3)$ and $IsNegative(d_4)$. When applying the in-tool consistency rules on the previous concepts and the relation $sameAs(d_4, d_3)$ (*IR1* and *IR2*), we infer the new polarities $IsNegative(d_3)$ and $IsPositive(d_4)$.

We ensure that we inferred all implicit knowledge by redoing the inference until no new knowledge are inferred. Such process is called inference by saturation.

5.3 Inconsistency inference Algorithm

After inferring all implicit knowledge in the set F , we apply the inconsistency rules *ICR* that allow to explicitly define the inconsistencies as it is presented in Algorithm 3. We apply this rules on a saturated knowledge base because most inconsistencies are implicit. For instance, if we apply the inconsistency rules directly after inferring the polarities $IsPositive(d_3)$ and $IsNegative(d_4)$, we get $\neg IsNegative(d_3)$, and $\neg IsPositive(d_3)$. However, when applying the in-tool consistency rules on the previous concepts and relation $sameAs(d_4, d_3)$ (saturation process), we obtain $IsNegative(d_3)$

Algorithm 2 Implicit Knowledge Inference Algorithm

Input : F with prior knowledge
Output : Saturated F

```

1: procedure INFERENCE
2:   //step1: Infer Polarities by applying
3:   //inter-tool consistency rules
4:   functions =  $D.getFunctions()$ 
5:   for each  $P_{t^+}(d_i) \in functions$  :
6:     if  $IsPositive(d_i) \notin F$  :  $F.add(IsPositive(d_i))$ 
7:   for each  $P_{t^-}(d_i) \in functions$  :
8:     if  $IsNegative(d_i) \notin F$  :  $F.add(IsNegative(d_i))$ 
9:   for each  $P_{t^0}(d_i) \in functions$  :
10:    if  $IsNeutral(d_i) \notin F$  :  $F.add(IsNeutral(d_i))$ 
11:   //step2: Infer Polarities by applying
12:   //in-tool consistency rules
13:   sameAsRelations =  $getSameAs(F)$ 
14:   repeat:
15:     for each  $SameAs \in sameAsRelations$  :
16:       if  $IsPositive(d_i) \in F \wedge IsPositive(d_j) \notin F$  :
17:          $F.add(IsPositive(d_j))$ 
18:       if  $IsPositive(d_j) \in F \wedge IsPositive(d_i) \notin F$  :
19:          $F.add(IsPositive(d_i))$ 
20:       if  $IsNegative(d_i) \in F \wedge IsNegative(d_j) \notin F$  :
21:          $F.add(IsNegative(d_j))$ 
22:       if  $IsNegative(d_j) \in F \wedge IsNegative(d_i) \notin F$  :
23:          $F.add(IsNegative(d_i))$ 
24:       if  $IsNeutral(d_i) \in F \wedge IsNeutral(d_j) \notin F$  :
25:          $F.add(IsNeutral(d_j))$ 
26:       if  $IsNeutral(d_j) \in F \wedge IsNeutral(d_i) \notin F$  :
27:          $F.add(IsNeutral(d_i))$ 
28:   until: no new inferred instance
29:   return:  $F$ 

```

and $IsPositive(d_4)$, and when applying the inconsistency rules on this instances, we get $\neg IsPositive(d_3)$ and $\neg IsNegative(d_4)$ which represents an implicit inconsistency in the fact set F . Therefore, applying the inconsistency rules on F after the saturation process is an important step in our reasoning procedure, because it shows all inconsistencies even the implicit ones.

5.4 MLN reasoning to reduce inconsistency

Here we discuss how to reduce the inconsistencies discovered in the previous phase, by applying the *MLN* approach. The reasoning process will first learn the rules' weights of R and after will use them to infer the correct polarities.

Grounding. The grounding algorithm enumerates all possible assignments of formulas to its free variables. (the set of possible worlds). We used the grounding algorithm described in [19] because it speeds up the inference process. We adopted the closed world assumption; hence we consider all groundings that are not present in the Fact set as false.

Learning. To learn the rules' weights, we use the discriminative training described in [18]. The training consists of optimizing the conditional log-likelihood given by:

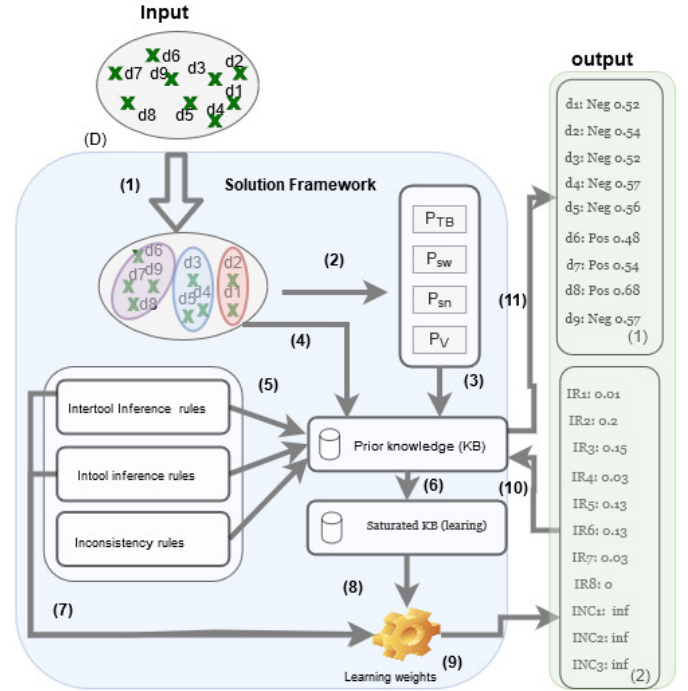


Figure 1: SentiQ overview

$$\log -P(Y = y|X = x) = \log Z_x - \sum_i w_i n_i(x, y)$$

where X represents priors (saturated inconsistent fact set), Y the set of queries (in our case: $IsPositive(d)$, $IsNegative(d)$, $IsNeutral(d)$), Z_x the normalization factor over the set of worlds, and $n_i(x, y)$ the number of the correct groundings of the formula l_i (the inference rules) in the worlds where Y holds.

We used in the optimization the Diagonal Newton discriminative method described in [18] that calculates the *Hessian* of the negative conditional log-likelihood given by:

$$\frac{\partial}{\partial w_i \partial w_j} - \log P(Y = y|X = x) = E_w[n_i n_j] - E_w[n_i] E_w[n_j]$$

With E_w the expectation. We call the inference procedure MC-SAT to estimate the number of satisfied (correct) formulas (n_i , n_j).

We can see that we consider the rules independently in the learning process. We calculate the number of each formula's correct grounding separately in the world; hence we do not take into consideration the implicit knowledge, which justifies the inference of all implicit knowledge and inconsistencies before learning.

Inference. The inference in *MLN* [21] contains two steps, grounding step, where we sample all possible worlds based on the priors and construct a large weighted Sat formula used in satisfiability calculation, and search step to find the best weight assignment to this Sat formula. In our work, we used the marginal inference algorithm that estimates the atoms' probability and returns the query answer with a probability score representing the confidence. It uses the MC-Sat algorithm, which combines satisfiability verification with MCMC by calling in each step the SampleSat algorithm that is a combination of

Simulated Annealing and WalkSat. Note that the walkSat algorithm selects in each iteration an unsatisfiable clause, selects an atom from the clause, and flip its truth value to satisfy the clause.

5.5 MLN-based Reasoning to Enhance Accuracy

Majority voting could be a solution to the inconsistency problem. However, this trivial method takes into consideration only the voting subset (cluster) and ignores information about the voters (polarity functions) from the other voting subsets (other clusters), which may hurt the accuracy.

To enhance the quality in terms of accuracy of the inconsistency issue resolution, the process in SentiQ follows two steps:

Step1. We use *MLN* to model the different inconsistencies and select the most appropriate polarity of the set (phase 1 to phase 3 of the process). We illustrate in Figure 1 the global workflow of our system. As input, we have an unlabeled dataset D (1) that we cluster to group the semantically equivalent documents in clusters. Then, (2) we extract the polarities from the documents using different polarity functions (P_{tb} , P_{sw} , P_v). After that, (4) we construct our knowledge base KB by creating first the fact set F (Algorithm 1). (5) We infer all implicit knowledge by applying inference rules (IR) on the *Fact set* until saturation (Algorithm 2). Then we apply inconsistency rules (ICR) to generate different inconsistencies between polarities (Algorithm 3). (7) We learn the weights of inference rules. (8) The output of the learning procedure is a set of weighted inference rules that we apply on the *prior knowledge* to infer the most appropriate polarities for documents.

Running the motivating example in this system shows an improvement in both the consistency and accuracy (accuracy of 88% and a low inconsistency).

Step2. (phase 4 of the process) As we still have inconsistencies from the previous step, we propose to resolve those remaining inconsistencies by using weighted majority voting with as weights the polarities probability, which leads to an accuracy of 100% on the motivating example.

Algorithm 3 Discover inconsistencies

Input : Saturated F

Output : Inconsistent F

```

1: procedure INCONSISTENCY INFERENCE
2:   //step1: get all polarities from the F
3:   //and apply inconsistency rules
4:   polarities = F.getPolarities()
5:   for each Polarity  $\in$  polarities :
6:     if Polarity == IsPositive( $d_i$ ) :
7:       F.add( $\neg$ IsNegative( $d_i$ ))
8:       F.add( $\neg$ IsNeutral( $d_i$ ))
9:     if Polarity == IsNegative( $d_i$ ) :
10:      F.add( $\neg$ IsPositive( $d_i$ ))
11:      F.add( $\neg$ IsNeutral( $d_i$ ))
12:     if Polarity == INeutral( $d_i$ ) :
13:       F.add( $\neg$ IsPositive( $d_i$ ))
14:       F.add( $\neg$ IsNegative( $d_i$ ))

```




Figure 2: Accuracy optimization on stanford tree-bank




Figure 3: Accuracy optimization on news headlines

6 EXPERIMENTAL EVALUATION

Tools. In our experiments, we use five representative sentiment analysis tools, a convolutional neural network with word embedding as P_{cnn_txt} [16], a convolutional neural network with character embedding as P_{char_cnn} [7], [13] as P_v , [1] as P_{sw} and [24] as P_{tb} . We chose these tools because of their performance and their association with different methods' categories, so they have different behaviors within inconsistency.

Dataset. We studied the effect of inconsistency resolution on accuracy using two publicly available datasets for sentiment analysis: News headlines dataset [4] and the test data of the sentiment tree-bank dataset [24] (sst). To consider the in-tool inconsistency, and for experimental purposes, we augmented the datasets with paraphrases using a generative adversarial network (GAN) [15].

For each document in the dataset, we generated three paraphrased documents with the same polarity as the original one. These datasets allow us to study the effect of resolving in-tool and inter-tool inconsistency on accuracy. Note that in our future work, we use a clustering method on the data to create our clusters.

Statistics about the datasets are presented in Table 2

Statistics	# elements	# Positive	# Neutral	# Negative
News_heads	1583	891	37	655
SST	3505	1481	640	1384

Table 2: Statistics on datasets.

Tools	Original Acc				M _V in – tool			
	size = 25	100	500	1500	size = 25	100	500	1500
P_{char_cnn}	0.62	0.55	0.50	0.506	0.69	0.59	0.50	0.52
P_{cnn_txt}	0.5	0.48	0.52	0.505	0.54	0.54	0.57	0.55
P_{sw}	0.34	0.35	0.34	0.33	0.38	0.39	0.35	0.33
P_{tb}	0.38	0.38	0.38	0.40	0.46	0.41	0.42	0.44
P_v	0.5	0.35	0.34	0.033	0.38	0.39	0.35	0.33
inter_tool MV	0.5	0.47	0.51	0.506	0.42	0.54	0.52	0.515
SentiQ	0.76	0.70	0.60	0.56	N/A	N/A	N/A	N/A

Table 3: Accuracy of tools before/after inconsistency resolution. The best performance for each dataset size is marked in bold.

Experiments. To evaluate the efficiency of resolving inconsistencies using SentiQ on the accuracy of the system, we compare it to the Majority Voting (MV) baseline. We use MV to resolve the in-tool inconsistency, inter-tool inconsistency, and both inconsistencies; then, we calculate the accuracy on the dataset after resolving contradictions. The majority voting for in-tool inconsistency resolution consists of calculating the most repeated polarity in the cluster and attributes it to all cluster documents :

$P_{t_k}(A) = \operatorname{argmax}_{\{+,0,-\}} \{ \sum_{d_i \in A} \mathbb{1}_{(P_{t_k}(d_i)=+)} , \sum_{d_i \in A} \mathbb{1}_{(P_{t_k}(d_i)=0)} , \sum_{d_i \in A} \mathbb{1}_{(P_{t_k}(d_i)=-)} \}$. Inter-tool inconsistency resolution using majority voting consists of attributing to the document the polarity attributed by most tools:

$P_*(d_i) = \operatorname{argmax}_{\{+,0,-\}} \{ \sum_{P_{t_k} \in \Pi} \mathbb{1}_{(P_{t_k}(d_i)=+)} , \sum_{P_{t_k} \in \Pi} \mathbb{1}_{(P_{t_k}(d_i)=0)} , \sum_{P_{t_k} \in \Pi} \mathbb{1}_{(P_{t_k}(d_i)=-)} \}$. Resolving both inconsistencies with MV consists of considering in the cluster all polarities given by polarity functions and attributing to each document the most repeated polarity.

Accuracy Optimization with SentiQ. To evaluate the accuracy improvement obtained by SentiQ, we run SentiQ on the two datasets News headlines and SST. The Figures 2,3 present the accuracy of resolving inconsistencies using SentiQ on the two datasets SST and news headlines respectively with the two queries $IsNegative(d)$ and $IsPositive(d)$ and the polarity functions P_{char_cnn} , P_{text_cnn} , P_{sw} and P_v .

We observe an accuracy improvement of 0.629 and 0.56 on the two datasets SST and the news headlines, respectively. These preliminary results prove the efficiency of resolving both in-tool inconsistency and inter-tool inconsistency using SentiQ to improve the accuracy. To analyze the performances and limits of SentiQ, we compare it in the next section to the MV baseline in the presence of variable-sized datasets.

Accuracy optimization and dataset size. The results are presented in the Table 3.

We evaluate the accuracy optimization of polarity functions on samples of different sizes (25, 100, 500 and 1500) from the news headlines dataset using SentiQ and MV to resolve in-tool inconsistencies, inter-tool inconsistencies, and both of them. "Original Acc" represents the original accuracy of the polarity function on this dataset, while "MV in-tool" represents the accuracy on different samples after resolving in-tool inconsistency using MV. "Inter-tool MV" represents the overall accuracy of the system after solving inter-tool inconsistencies, and the last line of the table represents the

accuracy obtained after inferring the polarity of the whole system using our SentiQ.

Results. We observe that resolving in-tool inconsistency increases the accuracy of tools in most of the cases. The only case where we have accuracy degradation corresponds to the tool P_v , where accuracy changes from $acc = 0.5$ to $acc = 0.38$ after resolving inconsistencies. When analyzing the data of this case, we found that most of this tool's predictions were *Neutral* instead of the data ground truth. As a result, majority voting falsified the results of the correctly predicted instances.

Resolving inter-tool inconsistency using majority voting decreases effectiveness in the case of tools that are incompatible in terms of accuracy (i.e., having widely different accuracy scores). Like the case of the two samples of size=25 and size=100 of the Table 3, where the weak nature of P_v , P_{sw} , and P_{tb} on the datasets has influenced the performance of the voting system (accuracy decreased from 0.62 to 0.5 on the dataset of size 25 and from 0.55 to 0.47 on the dataset of size 100). SentiQ addresses this problem, because it weighs different tools based on the inconsistencies on the whole dataset. SentiQ provides an accuracy improvement of 0.76 on the first dataset, 0.70 on the second, and 0.60 on the third dataset, outperforming majority voting.

This leads to other research problems, especially that of scalability, since we could not run experiments with a larger dataset, due to the high inference complexity of the Markov solver. Therefore, we need a more efficient Markov logic solver adapted to analyze large scale social media data.

We also observe that the MLN solver deletes some rules from the model (by attributing a negative, or a 0 weight), which can penalize the inference. The final results of the system could be boosted by adding business rules that can improve the polarity inference. This approach can be applied to various problems such as truth inference in crowd-sourcing, and other classification problems. We proved that resolving both in-tool and inter-tool inconsistency outperforms using only inter-tool inconsistencies.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we presented an MLN-based approach to solve inconsistencies and improve classification accuracy. Our results show the efficiency of including semantics to resolve in-tool inconsistency. The initial results of SentiQ are promising and confirm that resolving in-tool inconsistency boosts accuracy. However, to test SentiQ efficiency in resolving inconsistencies and improving the accuracy

of social media data, we need MLN solvers that can scale with the data size. Finally, we plan to investigate the use of domain expert rules for improving the polarity inference of SentiQ.

8 ACKNOWLEDGEMENT:

This work has been supported by the ANRT French program and IMBA consulting.

REFERENCES

- [1] Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining.. In *Lrec*. 2200–2204.
- [2] Salima Benbernou and Mourad Ouziri. 2017. Enhancing data quality by cleaning inconsistent big RDF data. In *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 74–79.
- [3] Erik Cambria, Soujanya Poria, Devamanyu Hazarika, and Kenneth Kwok. 2018. SenticNet 5: discovering conceptual primitives for sentiment analysis by means of context embeddings. In *Proceedings of AAAI*.
- [4] Keith Cortis, André Freitas, Tobias Daudert, Manuela Huerlimann, Manel Zarrouk, Siegfried Handschuh, and Brian Davis. 2017. Semeval-2017 task 5: Fine-grained sentiment analysis on financial microblogs and news. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. 519–535.
- [5] Haibo Ding and Ellen Riloff. 2018. Weakly supervised induction of affective events by optimizing semantic consistency. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [6] Pedro M. Domingos and Daniel Lowd. 2019. Unifying logical and statistical AI with Markov logic. *Commun. ACM* 62, 7 (2019), 74–83. <https://doi.org/10.1145/3241978>
- [7] Cicero Dos Santos and Maira Gatti. 2014. Deep convolutional neural networks for sentiment analysis of short texts. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*. 69–78.
- [8] Mauro Dragoni and Giulio Petrucci. 2018. A fuzzy-based strategy for multi-domain sentiment analysis. *International Journal of Approximate Reasoning* 93 (2018), 59–73.
- [9] Eduard C. Dragut, Hong Wang, A. Prasad Sistla, Clement T. Yu, and Weiyi Meng. 2015. Polarity Consistency Checking for Domain Independent Sentiment Dictionaries. *IEEE Trans. Knowl. Data Eng.* 27, 3 (2015), 838–851.
- [10] DI Hernández Farias and Paolo Rosso. 2017. Irony, sarcasm, and sentiment analysis. In *Sentiment Analysis in Social Networks*. Elsevier, 113–128.
- [11] Ronen Feldman. 2013. Techniques and applications for sentiment analysis. *Commun. ACM* 56, 4 (2013), 82–89.
- [12] Guohong Fu, Yu He, Jiaying Song, and Chaoyue Wang. 2014. Improving Chinese sentence polarity classification via opinion paraphrasing. In *Proceedings of The Third CIPS-SIGHAN Joint Conference on Chinese Language Processing*. 35–42.
- [13] CJ Hutto Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International Conference on Weblogs and Social Media (ICWSM-14)*. Available at (20/04/16) <http://comp.social.gatech.edu/papers/icwsm14.vader.hutto.pdf>.
- [14] Stephan Greene and Philip Resnik. 2009. More than words: Syntactic packaging and implicit sentiment. In *Proceedings of human language technologies: The 2009 annual conference of the north american chapter of the association for computational linguistics*. Association for Computational Linguistics, 503–511.
- [15] Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial Example Generation with Syntactically Controlled Paraphrase Networks. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. 1875–1885.
- [16] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882* (2014).
- [17] Efthymios Kouloumpis, Theresa Wilson, and Johanna D Moore. 2011. Twitter sentiment analysis: The good the bad and the omg! *Icwsml* 11, 538–541 (2011), 164.
- [18] Daniel Lowd and Pedro Domingos. 2007. Efficient weight learning for Markov logic networks. In *European conference on principles of data mining and knowledge discovery*. Springer, 200–211.
- [19] Feng Niu, Christopher Ré, AnHai Doan, and Jude Shavlik. 2011. Tuffy: Scaling up statistical inference in markov logic networks using an rdbms. *arXiv preprint arXiv:1104.3216* (2011).
- [20] Alexander Ratner, Stephen H. Bach, Henry R. Ehrenberg, Jason A. Fries, Sen Wu, and Christopher Ré. 2020. Snorkel: rapid training data creation with weak supervision. *Vldb J.* 29, 2 (2020), 709–730. <https://doi.org/10.1007/s00778-019-00552-1>
- [21] Matthew Richardson and Pedro M. Domingos. 2006. Markov logic networks. *Mach. Learn.* 62, 1-2 (2006), 107–136. <https://doi.org/10.1007/s10994-006-5833-1>
- [22] Julian Risch and Ralf Krestel. 2018. Aggression identification using deep learning and data augmentation. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*. 150–158.
- [23] Aliaksei Severyn and Alessandro Moschitti. 2015. Twitter sentiment analysis with deep convolutional neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 959–962.
- [24] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*. 1631–1642.
- [25] Mikalai Tsytarau and Themis Palpanas. 2016. Managing Diverse Sentiments at Large Scale. *IEEE Trans. Knowl. Data Eng.* 28, 11 (2016), 3028–3040.
- [26] Soroush Vosoughi, Prashanth Vijayaraghavan, and Deb Roy. 2016. Tweet2vec: Learning tweet embeddings using character-level cnn-lstm encoder-decoder. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 1041–1044.
- [27] Yequan Wang, Aixin Sun, Jialong Han, Ying Liu, and Xiaoyan Zhu. 2018. Sentiment Analysis by Capsules. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*. 1165–1174.
- [28] Jason W Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196* (2019).