

Document Difficulty Framework for Semi-Automatic Text Classification

Miguel Martinez-Alvarez¹, Alejandro Bellogin², and Thomas Roelleke¹

¹ Queen Mary, University of London
{miguel, thor}@eeecs.qmul.ac.uk
² Centrum Wiskunde & Informatica (CWI)
alejandro.bellogin@cwi.nl

Abstract. Text Classification systems are able to deal with large datasets, spending less time and human cost compared with manual classification. This is achieved, however, in expense of loss in quality. Semi-Automatic Text Classification (SATC) aims to achieve high quality with minimum human effort by ranking the documents according to their estimated certainty of being correctly classified. This paper introduces the Document Difficulty Framework (DDF), a unification of different strategies to estimate the document certainty, and its application to SATC. DDF exploits the scores and thresholds computed by any given classifier. Different metrics are obtained by changing the parameters of the three levels the framework is based upon: how to measure the confidence for each document-class (evidence), which classes to observe (class) and how to aggregate this knowledge (aggregation). Experiments show that DDF metrics consistently achieve high error reduction with large portions of the collection being automatically classified. Furthermore, DDF outperforms all the reported SATC methods in the literature.

1 Introduction and Motivation

Automatic Text Classification (TC) provide much faster and cheaper classification than human experts. However, even though there have been large improvements in the last decades, human experts achieve higher quality. Since the introduction of automatic classifiers, two alternative options can be applied. Firstly, a full-automatic classification system is applied, where every document is classified according to the decisions made by the classifier. Secondly, a completely manual classification is performed, where human experts classify each document. The main drawbacks of the latter option are its huge cost and potential unfeasibility for large collections. On the other hand, the quality achieved by the manual approach will be higher. Full automatic TC is preferred if large datasets are used (i.e. webpage classification) or when lower quality is not as important as the human effort required. On the other hand, manual classification is the best option for systems that require high-quality and have medium size such as law or medical data. In general, the time saved by using an automatic system is leveraged with the possible quality loss with respect to the manual classification.

This research focuses on an intermediate solution using Semi-Automatic Text Classification (SATC) [8, 1]. The main goal is to achieve high quality with minimum human

effort or, more specifically, to use human experts only for the documents that the automatic system is more likely to misclassify. Therefore, maximising the quality, while minimising the cost. Given a set of documents to be labelled and a specific classifier, a SATC algorithm needs to rank the documents according to the likelihood of their classification decisions to be correct. The ranking allows experts to inspect documents iteratively, starting with the most uncertain ones, until a specific point, where the rest of the documents are automatically assigned. In addition, this strategy can be applied with variable resources over time (i.e. less human experts could be available).

This paper introduces the **Document Difficulty Framework (DDF)**, a family of document certainty algorithms, and its application to SATC. DDF exploits the document-class confidence scores computed by a classifier and the class thresholds given by any class-based thresholding strategy to calculate the certainty of each document. This implies that the class scores for all documents have to be computed. The framework defines an array of different metrics, depending on three different dimensions: how the document-class evidence is computed (evidence), which classes will be considered (class), and how to aggregate a document-based certainty (aggregation).

The remainder of this paper is organised as follows: Section 2 presents the background and related research. Section 3 introduces DDF, and analyses its different variations. Experiments are explained and analysed in Sections 4. Finally, Section 5 concludes the paper and presents the future work.

2 Background and Related Research

2.1 Multi-Label Text Classification

In multi-label TC (ML-TC), each document can belong to multiple classes. As a result, a classification process assigns a boolean value to each pair $(d_j, c_i) \in D \times C$, where D is a set of documents and C is a set of predefined categories [9]. To achieve this goal, most classifiers use a two-step procedure. Firstly, the classifier produces a score for each pair (d_j, c_i) , and then a thresholding strategy decides, for each of the scores, if that value implies that the document belongs to the class (T) or not (F). Given a classifier σ and a threshold function δ , both processes can be mathematically denoted as: $\sigma \in \Sigma : D \times C \rightarrow [0, \infty]$, $\delta : \Sigma \times D \times C \rightarrow \{T, F\}$.

There are different types of thresholding strategies, based on documents or topics. For instance, using RCut [12], the R classes with higher score for a specific document are selected. On the other hand, SCut computes the threshold per class which maximises its quality (i.e. measured using F_1) [12]. If the number of training examples for a category is small, SCut can compute a very high or very low threshold from a global point of view. To address these drawbacks, SCutFBR.1 and SCutFBR.0 were introduced [12]. These metrics modify the behaviour of SCut in the case that the quality obtained for a class is lower than a value (fbr). The former strategy uses the highest ranked document as threshold value, while SCutFBR.0 assigns an infinite value. As a result, SCutFBR.0 does not assign any document in any class with lower quality than fbr .

2.2 Semi Automatic Text Classification (SATC)

SATC assumes that neither manual, nor full automatic classification is the optimum solution. This situation appears when full automatic classification achieves lower than required quality, and a full manual classification is either too expensive or unfeasible due to lack of resources. The foundation of SATC is that if we are able to separate the documents with high probability to be correctly classified, and the ones that are probably wrong, the latter can be inspected by human experts while the former will be automatically classified. As a result, the resources (the human experts) are optimised, while the quality remains high. To solve this task, SATC methods rank the documents to be classified according to their uncertainty. SATC assumes that the documents with higher certainty are probably better classified, whereas the documents with higher uncertainty are incorrectly classified. Therefore, the quality is maximised if the human annotators inspect the documents starting from the ones with higher uncertainty. The possibility of combining human and automatic classification has been suggested before [13, 9, 6]. However, only two approaches have been proposed: Document Difficulty and Utility-Theoretic Ranking. Both of them have been proven to be well-suited for SATC. Nonetheless, their performance have never been compared in the literature.

Document Difficulty [8] uses the classification scores and thresholds as evidence to compute the document certainty, where the labels with greater certainty are those with larger relative difference with respect to their threshold. The aggregation of label confidences is performed by averaging the confidences for those classes the document will be labelled in. The reason for the name similarity is that DDF extends and generalises the principles we explained in [8], mainly exploiting the classification scores and the threshold values within a classifier-independent framework. However, while only one method was proposed in that work, DDF represents a family of certainty metrics, where the previous metric appears naturally as an special case. In addition, the evaluation is also different. Our previous work evaluated the quality of the subset of automatically classified documents, instead of the whole collection. This research analyses the quality of the full test set, including both manual and automatic classified documents subsets.

The Utility-Theoretic Ranking (UT) method [1] optimises the global quality of the system, exploiting the potential benefit of manually inspecting each document, using the confidence scores of a classifier, and the gain in terms of quality that could be achieved, if that label is actually correct. The main conceptual difference with our approach is that UT exploits the collection information, whereas DDF focuses on each document independently. Furthermore, DDF exploits threshold information, and class filtering for the aggregated document certainty.

Similar to SATC, Active Learning (AL) ranks documents according to their benefit in the learning process, selecting which unlabelled documents should be manually labelled and being included as training examples. However, SATC focuses on the classification step, while AL operates in the training phase, selecting the documents from which the classifier can learn the most. Extensive research has been done related to single-label AL [7, 10]. However, very limited research has tried to address the same problem in a multi-label environment [4, 11].

2.3 Semi Automatic Text Classification Evaluation

SATC is evaluated using traditional classification quality measures such as micro-averaged- F_1 , once the human and the automatic decisions have been combined. This approach, introduced by Berardi *et al.* [1], provides quality values for different proportions of the collection being automatically classified, where the most uncertain documents are manually classified. In addition, the goal of SATC is not only to compute the quality but to analyse how it varies depending on the number of documents considered. Therefore, quality variations with respect to the full automatic quality with the same classifier are also computed. The main issue is that the relative quality increase is fully dependent on the base quality, when all the documents are automatically classified. For instance, in some cases, a 100% quality increase is impossible (i.e. full automatic classification achieving 95% quality), while more than 100% is possible for others, making a comparison over different classifiers impossible. Berardi *et al.* [1] addressed these challenges and introduced two alternatives based on the error reduction with respect to the full automatic system, instead of its quality increase. Error Reduction at rank (ER) measures the error reduction with a specific number of documents being automatically classified, where $E_p(n)$ models the error (defined as 1-quality) achieved by a classifier p with n documents being manually classified,

$$ER_p(n) = \frac{E_p(0) - E_p(n)}{E_p(0)} \quad (1)$$

Normalised Error Reduction at rank (NER) subtracts the error reduction at rank n achieved by a random ranker ($\frac{n}{|Te|}$, where $|Te|$ is the size of the documents to be classified) from ER in order to obtain more meaningful quality values,

$$NER_p(n) = ER_p(n) - \frac{n}{|Te|} \quad (2)$$

A third metric was also proposed by the same authors to include the specific position of each document into the evaluation: Expected Normalised Error Reduction (ENER) exploits the probability of a human expert inspecting n documents ($P_s(n)$),

$$ENER_p = \sum_{n=1}^{|Te|} P_s(n) \cdot NER_p(n) \quad (3)$$

$P_s(n)$ can follow different probability distributions, Berardi *et al.* [1] suggested the definition shown below, where p models the probability of the next document to be inspected,

$$P_s(n) = \begin{cases} p^{n-1} \cdot (1 - p) & \text{if } n \in \{1, \dots, |Te|\} \\ p^{n-1} & \text{if } n = |Te| \end{cases} \quad (4)$$

The value of p can be defined as a function of the expected ratio (ξ) of documents being manually classified. $p = \frac{1}{\xi \cdot |Te|}$. Therefore, p is computed for different expected ratios of manually classified documents. Extended information about this evaluation can be found in Berardi *et al.* [1].

Table 1. DDF Levels. c_i represents a class, d a document, and s the classifier’s score for each document-class pair. $t(c_i)$ is the threshold for c_i , and $q(c_i)$ is the estimated quality for c_i .

Evidence ϵ ; given $d, c_i, t(c_i)$		Class γ ; given $d, c_i, t(c_i), \epsilon$		Aggregation α ; given $\gamma(\epsilon, d, \cdot)$	
S	$s(d, c_i)$	A	$\epsilon(d, c_i)$	M	$\max_{c_i \in C} \gamma(\epsilon, d, c_i)$
D	$\ln(1 + s(d, c_i) - t(c_i))$	P	$\begin{cases} \epsilon(d, c_i) & \text{if } s(d, c_i) \geq t(c_i) \\ 0 & \text{otherwise} \end{cases}$	A	$\text{avg}_{c_i \in C} \gamma(\epsilon, d, c_i)$
R	$\ln(1 + \frac{ s(d, c_i) - t(c_i) }{t(c_i)})$			W	$\text{avg}_{c_i \in C} q(c_i) \cdot \gamma(\epsilon, d, c_i)$

3 Document Difficulty Framework for SATC

The Document Difficulty Framework (DDF) is a family of document certainty metrics within the context of TC. DDF extends and generalises the principles we explained in previous research [8], exploiting the classification scores and the threshold values within a classifier-independent framework to compute the document certainty. This computation is divided into three different levels, inspired by the comparison of multi-label AL metrics by Esuli *et al.* [4]: evidence, class and aggregation. The **evidence** level computes the confidence value for each document and category, using their classification score and the class threshold. The **class** level specifies which classes are to be considered in the final aggregation. The **aggregation** level combines the filtered confidence levels, producing a document-based certainty.

DDF is based on the composition of the three transformation functions, one for each level, where ϵ represents the evidence, γ the class, and α the aggregation level,

$$\text{certainty}(d) = \alpha(\{\gamma(\epsilon(d, \cdot))\}) \quad (5)$$

A method within the framework consists in a specific strategy for each level. Table 1 summarises the different candidates analysed herein for each DDF levels. Each method is represented as the concatenation of three letters, representing the strategy followed in each one of the levels. For instance, following Table 1, the difficulty measured by the DPA variation is defined as follows,

$$\text{certainty}(d) = \text{avg}_{c_i \in C: s(d, c_i) \geq t(c_i)} \ln(1 + |s(d, c_i) - t(c_i)|) \quad (6)$$

This method considers absolute distance for evidence function (D), only for positive classes (P), and computing the final aggregation as the average topic confidence (A).

3.1 Evidence Level

The evidence level is responsible for computing a confidence value for a document-topic pair. It is modelled as a function $\epsilon \in \mathcal{E} : D \times C \rightarrow [0, \infty]$, where D denotes the set of documents, and C the classes. Three candidates are considered for this level: **score** (S), absolute **difference** (D) and **relative** difference (R).

The first strategy (*S*) follows the same principle as relevance sampling [7], exploiting the score obtained by the classifier. It assumes that the higher the value the more relevant the score. Therefore, classes with higher scores are the ones with more certainty.

The second method (*D*) exploits the score and the threshold, assuming that larger distances imply lower uncertainty and higher chance that the document is correctly classified. This assumption is similar to uncertainty sampling [7]. A logarithmic function is applied to limit the effect of very large differences.

The last method (*R*) applies the same principles as the difference approach. However, it uses a relative difference instead of the absolute value. The rationale is that the absolute distances can be misleading. For instance, a distance of 0.2 would be much more important if the threshold is 0.05 than if it is 0.6.

3.2 Class Level

The class level behaves as a filter, selecting whether to exploit the certainty of a specific label, and hence, if it will be available at the next aggregation step or not. It is defined as a function $\gamma \in \Gamma : \mathcal{E} \times D \times C \rightarrow [0, \infty]$, where a composition with an element $\epsilon \in \mathcal{E}$ would be applied. Two candidates are considered for this level: **all** (*A*) and **positive** (*P*).

The first strategy (*A*) consists on not applying any filtering, hence considering all the confidence scores for a specific document.

The second method (*P*) selects the classes for which the classification score is higher or equal than the threshold. These are the classes which will be assigned to the document if automatic classification is applied. This strategy aims to focus the difficulty computation on the positive labels. In TC, the positive labels are more representative than the negative ones due to the fact that the number of positive classes for a specific document is usually much smaller than the number of negative ones. For example, the average number of classes per document in Reuters-21578 is 1.24, while the number of classes is 90. This approach assumes that if all classes are observed, the document certainties are somehow diluted because most of the documents will obtain a high confidence that do not belong to a large subset of the classes.

3.3 Aggregation Level

When multi-label data is used, a certainty value per document has to be provided, since the ranking of the labels can not be used to select nor to rank documents. For example, even if 90% of the most certain labels are selected, it is impossible to decide which documents should be automatically classified. For this reason, the filtered evidence per class should be combined into a single certainty metric for each document. This level is defined as a function $\alpha \in \mathcal{A} : \{\Gamma\} \times D \rightarrow [0, \infty]$. Typically, it will be applied to the set of possible functions $\gamma \in \Gamma$, one for each class $c_i \in C$. This is equivalent, taken a document d and an evidence level function ϵ as inputs, to the set $\Gamma(\gamma, \epsilon, d) = \{\gamma(\epsilon, d, c_i) : c_i \in C\}$. It should be noted that a one-to-one relation exists between the set of classes used in the definition of Γ and Γ itself, and thus, this notation could be

simplified as in Table 1. Three candidates are considered for this level: **maximum** (M), **average** (A) and **weighted** (W).

The first method (M) selects the most certain class for each document. The goal is to rank higher documents with at least one class correctly classified. This is specially important for collections with a low number of classes per document.

The second method (A) averages the confidence values for the filtered classes, providing a general estimation of how certain the class labels are.

The third method (W) uses an averaged weighted linear combination (WLC), based on the quality estimation per class. Classes with low expected quality are weighted less, because even if their assignment seems certain, it is likely to be a misclassification. The estimated quality values are obtained in the cross-validation phase.

4 Results and Discussion

4.1 Experimental Set-Up

The quality of the certainty algorithms for SATC is evaluated using ER and ENER [1]. ER is plotted with different percentages of the collection being manually classified, while the ENER metric provides values to directly compare the quality achieved by DDF methods with other state of the art approaches. Two traditional TC collections (`Reuters-21578` and `20-newsgroups`) are used:

`20-newsgroups` is a collection of approximately 20,000 newsgroup documents and 20 classes, with almost uniform distribution of documents over classes (Obtained from <http://people.csail.mit.edu/jrennie/20NewsGroups/>). The split for the collection is based on time as it is suggested. Cross-posting emails have not been considered. This collection has been selected to observe the behaviour of DDF with a single label collection.

`Reuters-21578` contains structured information about newswire articles that can be assigned to several classes. Two variations of the “ModApte” split are used. `Reuters-21578` uses only documents that belong to classes with at least one training and one test document. As a result, there are 7770/3019 documents for training and testing, observing 90 classes with a highly skewed distribution over classes (same as Yang *et al.* [12]). On the other hand, `Reuters-21578-115` uses documents belonging to classes with at least one training or testing document. This configuration has 9603/3299 documents for train and test respectively, and 115 classes. `Reuters-21578-115` allows a direct comparison with the results presented by Berardi *et al.* [1].

Three different families of classifiers have been used, namely Naive Bayes (Weka [5] implementation), SVM using LibSVM [3], and our own implementation of k -NN. Documents are represented using *ltc* [2] and χ^2 is used for feature selection: 5,000 features for NB and k -NN for `20-newsgroups`, and 3,000 for the others. SVM uses 10,000 features (based on full automatic classification experiments). Stemming and stop-words removal have been applied. `SCutFBR1` [12] thresholding strategy is applied with an *fbr* value of 0.3 and a 5-fold cross-validation process, optimising micro-average- F_1 . The SVM scores are obtained by running LibSVM with the option (-b 1).

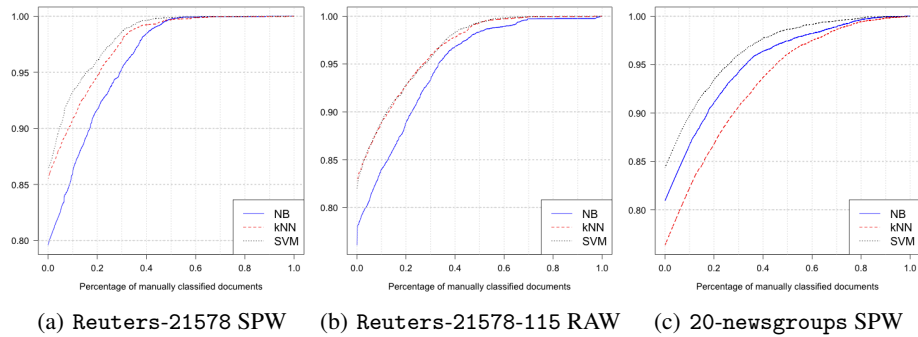


Fig. 1. Micro-averaged F_1 evaluation for different ratios of manually classified documents.




Fig. 2. Error Reduction for different ratios of manually classified documents.

4.2 Results

Figure 1 shows the absolute quality (micro-averaged F_1) depending on the percentage of manually classified documents. Due to clarity, only the best DDF metric per collection is shown. It illustrates how the best DDF metrics achieve high quality levels, while manually assigning a small subset of the collections. For instance, for Reuters-21578, micro-average F_1 of more than 95% can be achieved with as few as 20% of the documents manually classified. Furthermore, they also show that perfect quality is achieved with approximately 50 and 60% of documents manually classified for Reuters-21578 and Reuters-21578-115 respectively. 20-newsgroups appears to be a more challenging collection for SATC. Perfect quality is only achieved only after 80% of documents are inspected by experts. The main reason for this seems to be the uniformed distribution of documents over classes and the high similarity between some of the classes. The best performing model in all cases is SVM, while k -NN is the second best algorithm, despite the fact that it performs poorly when applied to 20-newsgroups.

Table 2. 20-newsgroups ENER wrt. the expected ratio of manually classified docs (ξ). Best results per model in bold, best overall is underlined. Improvement (%) wrt. RPA between brackets.

	NB			kNN			SVM		
	0.05	0.1	0.2	0.05	0.1	0.2	0.05	0.1	0.2
RPA	.097	.164	.230	.073	.127	.185	.121	.194	.251
SAA	-.014 (-114)	-.021 (-113)	-.030 (-113)	.038 (-48)	.058 (-55)	.074 (-60)	.044 (-64)	.046 (-77)	.037 (-85)
SAM	.090 (-6)	.151 (-8)	.211 (-9)	.070 (-5)	.116 (-9)	.161 (-13)	.119 (-1)	.195 (1)	.265 (6)
SAW	.012 (-88)	.030 (-82)	.061 (-74)	.041 (-45)	.062 (-51)	.081 (-56)	.012 (-90)	.033 (-83)	.069 (-72)
SPA	.097 (0)	.164 (0)	.228 (-1)	.073 (-1)	.123 (-3)	.175 (-6)	.120 (-0)	.197 (2)	.267 (6)
SPM	.096 (-0)	.160 (-3)	.218 (-5)	.072 (-2)	.119 (-6)	.165 (-11)	.120 (-1)	.196 (1)	.266 (6)
SPW	.097 (1)	.166 (1)	.231 (0)	.073 (-1)	.125 (-2)	.180 (-3)	.121 (1)	.198 (2)	.268 (7)
DAA	.094 (-3)	.158 (-4)	.219 (-5)	.069 (-6)	.116 (-9)	.170 (-8)	.119 (-2)	.194 (0)	.260 (4)
DAM	.050 (-48)	.077 (-53)	.112 (-51)	.048 (-34)	.081 (-37)	.119 (-35)	.079 (-35)	.136 (-30)	.200 (-20)
DAW	.091 (-6)	.155 (-6)	.216 (-6)	.064 (-12)	.107 (-16)	.156 (-16)	.117 (-3)	.193 (-1)	.259 (3)
DPA	.097 (1)	.166 (1)	.234 (2)	.073 (0)	.127 (0)	.186 (0)	.121 (0)	.197 (2)	.262 (4)
DPM	.096 (-0)	.163 (-1)	.226 (-2)	.073 (-1)	.124 (-3)	.174 (-6)	.121 (-0)	.196 (1)	.261 (4)
DPW	.097 (1)	.167 (2)	.235 (2)	.073 (0)	.128 (1)	.188 (1)	.121 (0)	.198 (2)	.265 (6)
RAA	.089 (-8)	.153 (-7)	.215 (-6)	.068 (-8)	.116 (-9)	.171 (-7)	.119 (-1)	.195 (1)	.260 (4)
RAM	.066 (-31)	.113 (-31)	.172 (-25)	.053 (-28)	.095 (-25)	.145 (-22)	.092 (-24)	.156 (-20)	.216 (-14)
RAW	.086 (-11)	.150 (-8)	.213 (-7)	.064 (-13)	.107 (-16)	.159 (-14)	.118 (-3)	.193 (-0)	.260 (3)
RPM	.096 (-1)	.160 (-3)	.221 (-4)	.073 (-0)	.124 (-2)	.176 (-5)	.120 (-0)	.193 (-0)	.250 (-0)
RPW	.097 (0)	.165 (1)	.232 (1)	.073 (0)	.128 (0)	.187 (1)	.121 (0)	.195 (1)	.256 (2)

ER evaluation is shown in Figure 2. Although the best model from this perspective is SVM, it illustrates almost overlapped curves for all different classifiers, specially for Reuters-21578-115. This result strongly supports the generalisation of DDF metrics.

Tables 2-4 allow to directly compare different SATC methods. They show the ENER quality evaluation for all the DDF candidates, with different percentages of the documents expected to be manually classified (ξ). The best performing models presented in the literature are chosen as baselines: The baselines for Reuters-21578-115 are the results reported by Berardi *et al.* [1] for their Utility-Theoretic method (*UT*). The method we introduced in [8], modelled as RPA within the DDF framework, is used as baseline for the other two collections. *UT* quality for 20-newsgroups and Reuters-21578 was not reported by Berardi *et al.* [1]. In all cases, the performance of DDF is higher when SVM is used, instead of NB or *k*-NN. For Reuters-21578-115, one DDF metric (SVM with a RAW configuration) outperforms both baselines (*UT* and RPA), as well as any other DDF candidates. The improvements are as high as 14 and 50% with respect to *UT* and RPA, respectively. In addition, DPA and DPW also outperform *UT* when $\xi = 0.2$. All collections confirm the quality of DDF, with several candidates outperforming our previously proposed method (RPA) [8].

For 20-newsgroups, there is almost not difference between the performance of candidates applying average aggregation and those applying weighted aggregation (i.e. DPA vs DPW). The main reason for this is that the classes quality are very similar. Furthermore, although no one of the best candidates includes the aggregation based on the maximum confidence (surprisingly being a single-label collection), this strategy achieves high quality (i.e. SPM is virtually as good as the best candidate for SVM).

Reuters-21578 and Reuters-21578-115 are analysed together as their main difference is the existence of documents without any correct class for the latter. The first

Table 3. Reuters-21578 ENER wrt. the expected ratio of manually classified docs (ξ). Best results per model in bold, best overall is underlined. Increment (%) wrt. RPA between brackets.

	NB			kNN			SVM		
	0.05	0.1	0.2	0.05	0.1	0.2	0.05	0.1	0.2
RPA	.101	.178	.245	.139	.233	.317	.156	.250	.321
SAA	.027 (-73)	.032 (-82)	.026 (-89)	.091 (-35)	.155 (-33)	.226 (-29)	.050 (-68)	.052 (-79)	.042 (-87)
SAM	.089 (-12)	.162 (-9)	.244 (-0)	.102 (-27)	.176 (-25)	.251 (-21)	.171 (10)	.269 (8)	.349 (9)
SAW	.140 (38)	.210 (19)	.283 (16)	.095 (-32)	.162 (-30)	.235 (-26)	.157 (0)	.246 (-1)	.321 (0)
SPA	.103 (2)	.184 (4)	.267 (9)	.135 (-3)	.221 (-5)	.304 (-4)	.193 (24)	.297 (19)	.376 (17)
SPM	.094 (-8)	.166 (-6)	.248 (1)	.120 (-14)	.189 (-19)	.260 (-18)	.171 (9)	.270 (8)	.351 (9)
SPW	.104 (2)	.185 (4)	.269 (10)	.137 (-2)	.225 (-3)	.309 (-2)	.192 (23)	.297 (19)	.376 (17)
DAA	.105 (4)	.183 (3)	.265 (8)	.163 (17)	.234 (1)	.300 (-5)	.175 (12)	.262 (5)	.337 (5)
DAM	.055 (-45)	.093 (-48)	.130 (-47)	.064 (-54)	.122 (-48)	.196 (-38)	.041 (-74)	.074 (-71)	.135 (-58)
DAW	.081 (-20)	.157 (-11)	.242 (-1)	.167 (20)	.240 (3)	.308 (-3)	.157 (0)	.248 (-1)	.326 (2)
DPA	.105 (4)	.192 (8)	.277 (13)	.143 (2)	.238 (2)	.323 (2)	.187 (19)	.293 (17)	.372 (16)
DPM	.096 (-5)	.174 (-2)	.257 (5)	.126 (-10)	.197 (-16)	.266 (-16)	.164 (5)	.262 (5)	.342 (7)
DPW	.105 (4)	.193 (8)	.280 (14)	.144 (3)	.241 (4)	.329 (4)	.187 (20)	.294 (18)	.374 (16)
RAA	.115 (14)	.189 (6)	.267 (9)	.138 (-1)	.229 (-2)	.310 (-2)	.171 (9)	.260 (4)	.334 (4)
RAM	.082 (-20)	.135 (-24)	.178 (-27)	.064 (-54)	.131 (-44)	.209 (-34)	.119 (-24)	.192 (-23)	.250 (-22)
RAW	.126 (24)	.205 (16)	.285 (17)	.158 (13)	.245 (5)	.323 (2)	.211 (35)	.295 (18)	.361 (13)
RPM	.093 (-8)	.150 (-16)	.193 (-21)	.126 (-10)	.195 (-16)	.258 (-18)	.136 (-13)	.203 (-19)	.254 (-21)
RPW	.106 (5)	.190 (7)	.267 (9)	.142 (2)	.239 (3)	.325 (3)	.170 (9)	.270 (8)	.344 (7)

Table 4. Reuters-21578-115ENER wrt. the expected ratio of manually classified docs (ξ). Best results per model in bold, best overall is underlined. Improvement (%) wrt. Utility-Theoretic (UT) Ranking between brackets (Berardi *et al* [1]).

	NB			kNN			SVM		
	0.05	0.1	0.2	0.05	0.1	0.2	0.05	0.1	0.2
UT	.145	.221	.285	.145	.221	.285	.145	.221	.285
SAA	.015 (-90)	.019 (-91)	.014 (-95)	.060 (-58)	.109 (-51)	.172 (-40)	.032 (-78)	.035 (-84)	.032 (-89)
SAM	.111 (-23)	.165 (-25)	.230 (-19)	.058 (-60)	.113 (-49)	.186 (-35)	.127 (-12)	.206 (-7)	.282 (-1)
SAW	.133 (-9)	.187 (-16)	.248 (-13)	.065 (-55)	.116 (-48)	.180 (-37)	.127 (-12)	.200 (-10)	.267 (-6)
SPA	.068 (-53)	.142 (-36)	.224 (-21)	.063 (-56)	.132 (-40)	.216 (-24)	.103 (-29)	.196 (-11)	.283 (-1)
SPM	.054 (-63)	.119 (-46)	.200 (-30)	.056 (-61)	.112 (-49)	.185 (-35)	.092 (-37)	.179 (-19)	.265 (-7)
SPW	.067 (-53)	.142 (-36)	.224 (-21)	.064 (-56)	.134 (-39)	.220 (-23)	.102 (-30)	.196 (-11)	.283 (-1)
DAA	.097 (-33)	.153 (-31)	.221 (-23)	.139 (-4)	.199 (-10)	.260 (-9)	.114 (-21)	.194 (-12)	.274 (-4)
DAM	.032 (-78)	.047 (-79)	.063 (-78)	.070 (-52)	.121 (-45)	.184 (-35)	.055 (-62)	.100 (-55)	.170 (-40)
DAW	.039 (-73)	.101 (-54)	.181 (-36)	.141 (-3)	.202 (-8)	.265 (-7)	.089 (-39)	.173 (-22)	.260 (-9)
DPA	.069 (-52)	.147 (-33)	.230 (-19)	.067 (-54)	.144 (-35)	.234 (-18)	.101 (-30)	.200 (-9)	.292 (2)
DPM	.055 (-62)	.121 (-45)	.201 (-29)	.059 (-59)	.120 (-46)	.195 (-31)	.091 (-37)	.182 (-18)	.271 (-5)
DPW	.069 (-52)	.148 (-33)	.232 (-19)	.067 (-54)	.146 (-34)	.237 (-17)	.102 (-30)	.201 (-9)	.293 (3)
RAA	.136 (-6)	.186 (-16)	.243 (-15)	.121 (-16)	.199 (-10)	.273 (-4)	.129 (-11)	.201 (-9)	.269 (-5)
RAM	.077 (-47)	.113 (-49)	.143 (-50)	.062 (-57)	.125 (-43)	.199 (-30)	.104 (-28)	.170 (-23)	.235 (-18)
RAW	.147 (2)	.204 (-8)	.264 (-7)	.133 (-8)	.209 (-6)	.280 (-2)	.167 (15)	.240 (9)	.306 (7)
RPA	.055 (-62)	.124 (-44)	.195 (-32)	.065 (-55)	.141 (-36)	.231 (-19)	.087 (-40)	.178 (-19)	.266 (-7)
RPM	.041 (-72)	.081 (-64)	.118 (-58)	.060 (-59)	.120 (-46)	.194 (-32)	.073 (-49)	.148 (-33)	.221 (-22)
RPW	.067 (-54)	.143 (-36)	.220 (-23)	.066 (-54)	.144 (-35)	.236 (-17)	.097 (-33)	.193 (-13)	.282 (-1)

Table 5. Average ENER evaluation for DDF patterns and $\xi = 0.1$

Collection	Model	S**	D**	R**	*A*	*P*	**A	**M	**W
20newsgroups	NB	.108	.148	.151	.107	.164	.131	.137	.139
	kNN	.100	.114	.116	.095	.125	.111	.110	.109
	SVM	.144	.186	.188	.149	.196	.170	.179	.168
Reuters21578.115	NB	.129	.119	.142	.131	.130	.129	.108	.154
	kNN	.119	.155	.156	.155	.132	.154	.118	.158
	SVM	.169	.175	.188	.169	.186	.167	.164	.201
Reuters21578	NB	.157	.165	.174	.152	.179	.160	.147	.190
	kNN	.188	.212	.212	.188	.220	.218	.168	.225
	SVM	.238	.239	.245	.211	.271	.236	.211	.275

observation, in terms of performance, is that models considering positive classes lose their competitiveness against selecting all classes, for Reuters-21578-115. The reason is that this strategy was conceived for collections where test documents have at least one correct class, as documents with no classes are assigned a large uncertainty. This also explains the poor performance of our previous method [8], with decreases of more than 50% ENER, both with respect to the best DDF metric and UT. Another interesting difference is that, while there is a clear winner for Reuters-21578-115 (RAW), there is none for Reuters-21578. Furthermore, the qualities achieved by the best model in Reuters-21578 are significantly higher than those for Reuters-21578-115. This means that the addition of documents without correct classes makes the SATC problem more complex to solve, or at least that DDF metrics are less suited for this type of datasets. Results also show that SAA (and SAW for 20-newsgroups because of the similar qualities per class) is only suited for classifiers that do not normalise the scores per document. SAA performs as a random ranker for this type of classifiers which include the versions of NB and SVM presented on this paper. If the classification scores are normalised per document ($\sum_{c_i \in C} s(d, c_i) = 1$), SAA produces the same difficulty, independently of the document. Other very poor metric is DAM, because the highest confidence based on difference is usually based on a very low (or even zero) score. Therefore, the certainty computation will be uniquely based on this information.

Table 5 summarises the average quality for candidates sharing two strategies, with $\xi = 0.1$ (arbitrarily chosen). For instance, S** averages the error reduction for every variation using positive labels (this is, it encapsulates information about SAA, SAM, SAW, SPA, SPM, and SPW). This analysis provides information about which strategies are better for each level in different conditions, and it helps to understand some of the previously reported results from a general perspective. For the evidence level, the best strategy is the relative difference, independently of classifier and collection. This result confirms our previous assumptions made for the RPA method [8]. The class level illustrates that the selection of positive classes achieves good quality, as long as the assumption that all the documents have at least one correct class is correct. Otherwise, all classes should be considered. The aggregation level shows that the exploitation of quality estimation outperforms the other strategies for Reuters-21578 and Reuters-21578-115. All strategies perform similarly for 20-newsgroups.

5 Conclusions and Future Work

SATC represents a largely unexplored task within TC which is critical in environments where high quality classification is needed, but resources are limited. Its main goal is to achieve high quality with minimum human effort, minimising the potential cost. This research introduces DDF, a document certainty framework based on classification scores and class thresholds, and its application to SATC.

DDF generalises several methods by abstracting three different levels, specifying how to manipulate the scores and thresholds to obtain a document certainty measure. Results show that DDF metrics achieve virtually perfect classification with as low as 50% of documents being classified. SVM is the best classifier for DDF and RAW is its best overall variation, with the exception of *20-newsgroups*, where SPW and DPW are the best ones. DDF outperforms all the previously proposed methods in the literature for SATC. The strategy analysis shows that the best models should include a relative difference of scores, and the exploitation of estimated class quality. In addition, observing only the positive classes for a document achieves better quality, but only if all documents belong to at least one class.

Future work will provide deeper analysis of the combination between difference strategies, as well as their behaviour for different ratios of the collection. The combination of DDF and the Utility-Theoretic Ranking method, and the combination of DDF values as features for a meta-ranker are also interesting lines of research.

References

1. G. Berardi, A. Esuli, and F. Sebastiani. A utility-theoretic ranking method for semi-automated text classification. In *SIGIR*, 2012.
2. C. Buckley, G. Salton, and J. Allan. The effect of adding relevance information in a relevance feedback environment. In *SIGIR*, 1999.
3. C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2011.
4. A. Esuli and F. Sebastiani. Active learning strategies for multi-label text classification. In *ECIR*, 2009.
5. M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: an update. *SIGKDD Explor. Newsl.*, 2009.
6. L. S. Larkey and W. B. Croft. Combining classifiers in text categorization. In *SIGIR*, 1996.
7. D. D. Lewis and W. A. Gale. A sequential algorithm for training text classifiers. In *SIGIR*, 1994.
8. M. Martinez-Alvarez, S. Yahyaei, and T. Roelleke. Semi-automatic document classification: Exploiting document difficulty. In *ECIR*, 2012.
9. F. Sebastiani. Machine learning in automated text categorization. *ACM Comput. Surv.*, 2002.
10. S. Tong and D. Koller. Support vector machine active learning with applications to text classification. *J. Mach. Learn. Res.*, 2002.
11. B. Yang, J.-T. Sun, T. Wang, and Z. Chen. Effective multi-label active learning for text classification. In *SIGKDD*, 2009.
12. Y. Yang. A study on thresholding strategies for text categorization. In *SIGIR*, 2001.
13. Y. Yang and X. Liu. A re-examination of text categorization methods. In *SIGIR*, 1999.