

# Spam filtering by quantitative profiles

M. Grendár<sup>a</sup>, J. Škutová<sup>a</sup>, V. Špitalský<sup>a</sup>

<sup>a</sup>*Slovanet a.s., Záhradnícka 151, 821 08 Bratislava, Slovakia*

---

## Abstract

Instead of the “bag-of-words” representation, in the quantitative profile approach to spam filtering and email categorization, an email is represented by an  $m$ -dimensional vector of numbers, with  $m$  fixed in advance. Inspired by Sroufe et al. [Sroufe, P., Phithakkitnukoon, S., Dantu, R., and Cangussu, J. (2010). Email shape analysis. In *LNCS*, 5935, pp. 18-29] two instances of quantitative profiles are considered: line profile and character profile. Performance of these profiles is studied on the TREC 2007, CEAS 2008 and a private corpuses. At low computational costs, the two quantitative profiles achieve performance that is at least comparable to that of heuristic rules and naive Bayes.

*Keywords:* email categorization, spam filtering, quantitative profile, character profile, line profile, Random Forest

---

## 1. Introduction

Spam is an unsolicited email message. From the receiver’s perspective, spam is an annoyance and thus it is necessary to block its delivery, for instance, by filtering it out.

Traditional approach to spam filtering and email categorization that is based on heuristic rules, naive Bayes filtering and/or text-mining suffers from several deficiencies. Among shortcomings of the traditional approach there are high computational costs, language dependence, necessity to update the heuristic rules, high number of rules, and vulnerability.

Instead of the “bag-of-words” representation, employed in text-mining and naive Bayes filtering, in the quantitative profile ( $QP$ ) approach that we propose, an email is represented by an  $m$ -dimensional vector of numbers with  $m$  fixed in advance. Inspired by Sroufe, Phithakkitnukoon, Dantu, and Cangussu [12], two instances of  $QPs$  are considered: line profile ( $LP$ ) and character profile ( $CP$ ). Informally put, the line profile of an email is a vector of lengths of the first  $m$  lines. The character profile is a histogram of characters. Of course, many other  $QPs$  are conceivable.

The main advantages of the two considered quantitative profiles are *i*) sound performance, *ii*) simple computability, *iii*) language-independence, *iv*) robustness to outlying emails, *v*) high scalability and *vi*) low vulnerability.

The considered two instances of the quantitative profile approach perform comparably to the naive Bayes filtering and heuristics-based approaches and perform very well also in a multi-language, non-English communication. Furthermore, the satisfactory performance is attained by means of a small set of easily computable quantitative features. A performance study was done on the TREC 2007, CEAS 2008 and a private corpuses. To demonstrate the power of the considered  $QPs$ , the profiles are obtained from raw emails, without any preprocessing. Consequently, we intentionally ignore the structure of emails and character encoding.

On the two  $QPs$ , the Random Forest algorithm substantially outperforms other classification algorithms (SVM, LDA/QDA, logistic regression). Thanks to the Random Forest, the  $QP$  approach gains robustness to emails with extreme-valued profiles as well as high scalability.

---

*Email addresses:* [marian.grendar@slovanet.net](mailto:marian.grendar@slovanet.net) (M. Grendár), [jana.skutova@slovanet.net](mailto:jana.skutova@slovanet.net) (J. Škutová), [vladimir.spitalsky@slovanet.net](mailto:vladimir.spitalsky@slovanet.net) (V. Špitalský)

In our view, classification and email categorization by *LP* (or *CP*) should have low vulnerability. For, the lines the lengths of which differentiates between spam and ham, change from corpus to corpus. For instance, in the CEAS 2008 corpus, the most important for deciding between spam and ham are the lengths of the (10, 17, 15, 14, 16)-th lines, whilst in the TREC 2007 corpus they are the (5, 13, 6, 15, 7)-th lines. As the training corpus is usually not available to a spammer, it should be not easy to evade the *LP* (or *CP*) filter.

As a by-product of a performance study of *CP* and *LP*, we note that in the TREC 2007 and CEAS 2008 corpuses the number of header lines is capable of discriminating between spam and ham, at a rate that is, in our view, too high. Consequently, the corpuses lead to overly optimistic performance of spam filtering methods.

The paper is organized as follows. In the next section we formally introduce the *QPs* mentioned above. Then, in Section 3, we describe the three email corpuses used for assessment of the *QPs*' performance. In Section 4 we describe measures used for performance evaluation. The results are summarized in Section 5. In the concluding section some directions for future research are briefly discussed. All the computations were performed with R [7]. To make the results reproducible, a supplementary material including the source code was prepared, cf. [4].

## 2. Quantitative profiles

The *quantitative profile (QP)* of an email is an  $m$ -dimensional vector of real numbers that represents the email. The dimension  $m$  of the profile is set in advance, and it is the same for all emails. In this paper we consider two particular *QPs* – the line profile and the character profile. These profiles can be introduced by means of a simple probabilistic model.

An email is represented as a realization of a vector random variable, that is generated by a hierarchical data generating process. The length  $n$  of an email is an integer-valued random variable, with the probability distribution  $F_n$ . Given the length, the email is represented by a random vector  $X_1^n = (X_1, \dots, X_n)$  from the probability distribution  $F_{X_1^n | n}$  with the support in  $\mathcal{A}^n$ , where  $\mathcal{A} = \{a_1, \dots, a_m\}$  is a finite set (alphabet) of size  $m = |\mathcal{A}|$ .

Then, the *character profile (CP)* of an email is an  $m$ -dimensional random vector  $CP = (CP_1, \dots, CP_m)$ , where

$$CP_j = \sum_{i=1}^n I_{\{X_i=a_j\}}, \quad j = 1, \dots, m,$$

and  $I$  is the indicator function.

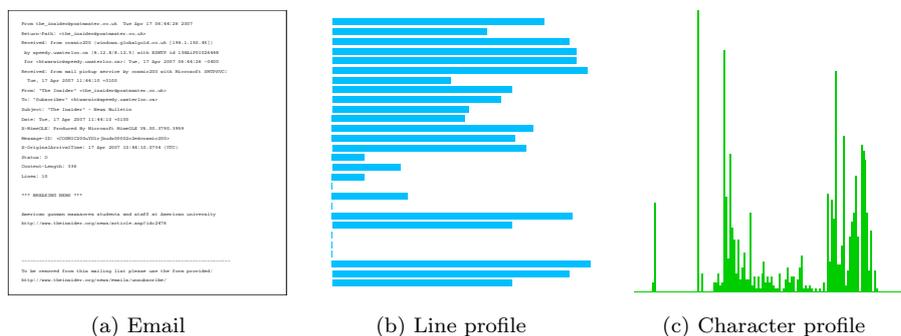


Figure 1: Graphical representation of the line and character profiles of an email

In order to introduce the other *QP*, it is necessary to select a special character (or a subset of characters) from the alphabet. Let  $k$  be the number of occurrences of the special character in an email and let  $T_j$

( $j = 1, \dots, k$ ) be the index of the  $j$ -th occurrence of the special character; put  $T_0 = 0$ . Then the *binary profile* ( $BP$ ) of an email is defined as a  $\tilde{k}$ -dimensional random vector  $BP = (BP_1, \dots, BP_{\tilde{k}})$ , where

$$BP_j = T_j - T_{j-1} - 1, \quad j = 1, \dots, \tilde{k}.$$

There,  $\tilde{k}$  is the maximum allowable number of the occurrences of the special character and it is set in advance. Hence, if  $k > \tilde{k}$ , the rest of the email is ignored. And if an email has  $k < \tilde{k}$ ,  $BP_j = 0$  for  $j > k$ .

In this work, an email is understood as a stream of bytes without any preprocessing, so that  $\mathcal{A}$  is taken to be the ASCII character set. The end-of-line is taken for the special character. Consequently, the binary profile becomes the *line profile* ( $LP$ ) since  $BP_j$  is the length of the  $j$ -th line of an email in bytes. We fix the maximal number  $\tilde{k}$  of considered lines to 100.

### 3. Data sets

To assess the performance of the two quantitative profiles, we consider three email corpuses: the publicly available TREC 2007 and CEAS 2008 corpuses, and a private corpus.

The TREC 2007 corpus [3] comprises over 75 000 emails (25 220 hams and 50 199 spams), of which approximately 67% is spam, and the rest are ham emails. For the training phase we used the first 50 000 emails. The rest forms the test set. The ratio of spam to ham in the training and test sets is approximately 2:1.

The other publicly available corpus, used for the performance analysis of  $QPs$ , is CEAS 2008 [2]. It consists of 137 705 emails (27 126 hams and 110 579 spams). The corpus was hand labeled [9]. To form the training set, we used the first 90 000 emails. The test set comprises the remaining emails. The ratio of spam to ham in the training and test sets is approximately 4:1.

Performance of spam filtering algorithms is typically assessed on English language corpuses, such as the above mentioned TREC and CEAS. When applied to non-English emails, their performance may be different. Due to the support from a Slovak internet services provider, we enjoyed the opportunity of access to a private corpus created in 2010, that comprises mainly non-English emails. Structure of the corpus is summarized in Table 1. The training set we used consists of 11 050 emails and the test set consists of 12 200 emails.

Table 1: Composition of the private corpus

corpus	s.ham	advert	notify	spam	total
train	6837	1611	1225	1377	11 050
test	3650	1409	5758	1383	12 200

The private corpus was hand labeled. Emails were placed into one of the two groups: ham and spam. Moreover, ham was divided into advert, solicited ham (denoted s.ham) and notify. Based on the language of the major part of an email, the emails were placed into one of the four language groups: Slovak/Czech, English, German, and other. The corpus comprises 64% of the Slovak/Czech ham in the training set and 38% in the test set.

### 4. Classification algorithm and performance measures

Quantitative profiles serve as an input to a classification algorithm. In this work we use the Random Forest classifier, introduced by Breiman [1] and ported to R by Liaw and Wiener [6], with the default settings. We have employed also LDA, logistic regression, LASSO and SVM [5], but these methods performed much worse.

To evaluate *QPs*' performance, we calculate the false positive rate  $fpr = FP/(TN + FP)$  and the false negative rate  $fnr = FN/(TP + FN)$ , where *TP* (*TN*) stands for the number of true positive (true negative) emails, i.e. the correctly recognized spam (ham) emails; and *FP* (*FN*) stands for the number of false positive (false negative) emails, i.e. the incorrectly recognized ham (spam) emails, respectively. We also present the receiver operating characteristic (*ROC*) curve, i.e. the graph of the true positive rate vs. the false positive rate, obtained as functions of the decision threshold. The area *AUC* under the *ROC* curve is also reported.

## 5. Results

In this section we summarize performance of the basic *QPs* on the three corpuses mentioned above.

### 5.1. Comparison of quantitative profiles with SpamAssassin and Bogofilter

For the sake of comparison, we report also results for SpamAssassin (*SA*) [11], version 3.3.1, off-line and without the Bayes filter, and Bogofilter (*BF*) [8], version 1.2.2 with the default configuration. On the private corpus, the output from *SA* was processed by the Random Forest classification algorithm, as it attains much better performance than *SA* with the default weights. In addition to much better performance, it allows for email categorization, which is impossible with the default *SA*. On the public corpuses, however, the default *SA* classification performs better. Bogofilter allows a binary classification only. *BF* was learnt in the batch mode.

Table 2: *fnr* (%) at fixed *fpr* = 0.5% or *fpr* = 1%

filter	private		TREC07		CEAS08	
	at 0.5%	at 1%	at 0.5%	at 1%	at 0.5%	at 1%
<i>CP</i>	14.39	11.64	2.53	0.49	4.38	4.25
<i>LP</i>	21.33	20.10	0.30	0.13	0.39	0.27
<i>SA</i>	12.68	10.10	35.87	30.51	76.14	69.92
<i>BF</i>	13.05	7.38	0.40	0.06	0.47	0.36

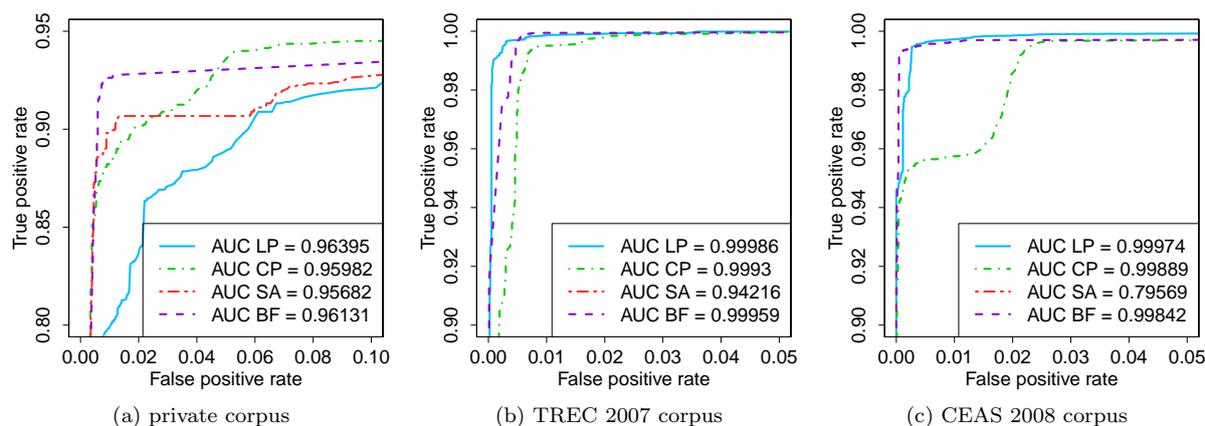


Figure 2: *ROC* curves

On the private corpus, for  $fpr = 0.5\%$  the best performance is attained by *SA* and *BF*, and *CP* performs only slightly worse, cf. Table 2 and Figure 2. *LP* is slightly less effective, and it attains  $fnr$  around 21% at  $fpr = 0.5\%$ .

With the exception of *SA*, on the public corpuses all the studied filters attain much better performance than on the private corpus, see also Section 5.5. The two *QPs* perform much better than *SA* as Table 2 as well as Figures 2b and 2c indicate. The line profile attains better performance (smaller  $fnr$ ) than *BF* at the 0.5% level of  $fpr$ .

### 5.2. Email categorization for the private corpus

On the private corpus, performance of *CP* and especially *LP* in spam filtering (i.e. binary categorization) is worse than that of *SA* and *BF*. However, in categorization of emails into one of the four categories, both *CP* and *LP* perform much better in the most interesting category of solicited ham. Misclassification table for *CP* and *SA* is in Table 3.

Table 3: *CP* and *SA* confusion tables for categorization, private corpus

count	<i>CP</i>				<i>SA</i>			
	advert	s.ham	notify	spam	advert	s.ham	notify	spam
advert	530	837	10	32	515	830	17	47
ham	26	3597	27	0	95	3290	248	17
notify	20	237	5499	2	24	235	5498	1
spam	37	221	12	1113	44	161	14	1164

### 5.3. Comparison with email shape analysis

Sroufe et al. [12] suggest to filter spam by means of its shape, which the authors define (using our terminology) as a smoothed line profile of email body, where smoothing is performed by the kernel smoother. Sroufe et al. also report the total error of 30%, based on a preliminary study on the TREC corpus. Further, the authors find the performance very good, 'considering that no content or context was even referenced', cf. [12], p. 26. The line profile, that is inspired by the email shape analysis, attains on the TREC corpus  $fpr = 4.23\%$  a  $fnr = 17.00\%$ , when the threshold is not optimized and email headers are intentionally not taken into account, cf. Section 5.5. This gives the total error around 12.29% and indicates that smoothing is unnecessary.

### 5.4. Reduction of the feature space

It is also important to know to what extent the dimension of the *QP* feature space can be reduced without substantive reduction of the classifier's performance. To this end the top 20 and the top 50 features were considered, where the ranking of features was provided by the Random Forest's measure of the mean decrease of accuracy. The study was done on the private corpus and solely the email body was considered.

In the binary classification the top 20 features of the line profile attain essentially the same performance as the entire line profile of the length 100. In the case of *CP*, to attain the full-set performance, the top 50 features out of 256 are needed. The same holds for *SA*. However, in the case of the email categorization it is not possible to reduce the dimension of *SA* features without substantive decrease in accuracy.

### 5.5. Why are TREC and CEAS misleading?

All the considered email filters except of *SA* attain much better performance on the public corpuses than on the private one; cf. Table 2. In search for explanation we have noted that in the public corpuses, unlike to the private corpus, the number of header lines contains information that is substantive for spam filtering. Figure 3 depicts the distribution of the number of lines in header, for spam and ham, in the TREC corpus.

Once the email header is not taken into account, and solely the email body is processed, performance of *LP* and *CP* worsens and becomes comparable to that on the private corpus; cf. Table 4. In Table 4, *LPH* (*CPH*) denotes the line (character) profile of email header and *LPB* (*CPB*) denotes the line (character) profile of email body, respectively.

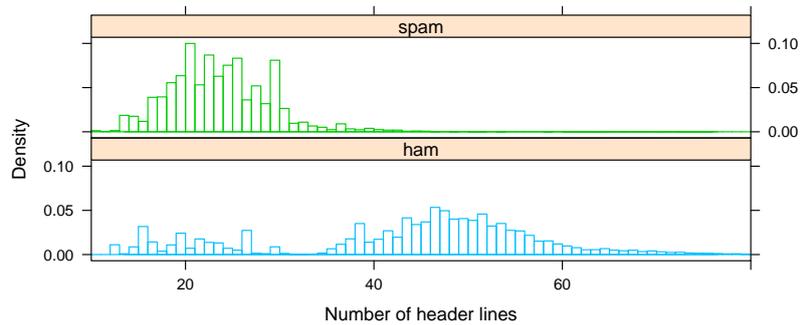


Figure 3: Distribution of emails with respect to the number of header lines, for spam and ham, in the TREC 2007 corpus

Table 4: *fnr* (%) at fixed *fpr* = 0.5% or *fpr* = 1%

filter	private		TREC07		CEAS08	
	at 0.5%	at 1%	at 0.5%	at 1%	at 0.5%	at 1%
<i>CPH</i>	16.51	13.71	0.19	0.05	0.78	0.30
<i>LPH</i>	18.06	15.11	1.78	0.12	0.68	0.36
<i>CPB</i>	14.24	11.92	15.05	5.47	4.70	4.51
<i>LPB</i>	21.26	18.58	45.01	43.67	7.07	6.35

The decline of performance of *CP* and *LP* caused by exclusion of email headers supports the hypothesis that in the TREC 2007 and CEAS 2008 corpuses the profiles of headers carry a substantive information for discriminating between spam and ham.

### 5.6. Summary of the performance study

The empirical study implies that the simple and easily obtainable line and character profiles attain at least comparable performance as the optimally tuned SpamAssassin, which is based on hundreds of fixed rules, and the performance of character profiles is close to that of Bogofilter. Particularly, on the public corpuses *LP* is better than *BF* and *SA*. On the private corpus *CP* attains comparable performance as *BF* and *SA*, and *LP* is slightly worse.

## 6. Conclusions

Motivated by Sroufe et al. [12], we have proposed the quantitative profile approach to email classification. In this report we explored two quantitative profiles, the line profile and the character profile. The profiles are obtained from raw emails, without any preprocessing. The computational costs of the two profiles are minimal. Performance of the profiles was studied on the TREC 2007, CEAS 2008 corpuses and a private, multi-lingual corpus. The two quantitative profiles attained at least comparable performance as the optimally tuned SpamAssassin and the batch-mode learnt Bogofilter. Besides the good performance, the two quantitative profiles are language independent and the resulting filter is robust to outlying emails, highly scalable and has low vulnerability.

As a by-product, we have noted that the number of header lines in the TREC 2007 and CEAS 2008 corpuses contain rather strong information on the email class. The corpuses thus lead to overly optimistic performance of spam filters.

In the near future we plan to explore quantitative profiles based on size and structure of emails, on the symbolic dynamics, and another instances of the binary profile. Also, the profiles are worth employing in a semi-supervised email categorization.

## 7. Acknowledgement

Stimulating feedback from Ján Gallo and Stanislav Záriš is gratefully acknowledged. This paper was prepared as a part of the project “SPAMIA”, MŠ SR 3709/2010-11, supported by the Ministry of Education, Science, Research and Sport of the Slovak Republic, under the heading of the state budget support for research and development.

## References

- [1] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32.
- [2] Cormack, G. V., and Lynam, T. R. (2008). CEAS 2008 corpus.  
<http://plg.uwaterloo.ca/~gvcormac/ceascorpus>
- [3] Cormack, G. V., and Lynam, T. R. (2007). TREC 2007 corpus.  
<http://plg.uwaterloo.ca/~gvcormac/treccorpus07/about.html>
- [4] Grendár, M., Škutová, J., and Špitalský, V. (2011). Supplement to “Spam filtering by quantitative profiles”.  
<http://www.savbb.sk/~grendar/spam/SupplementToQuantitativeProfiles.pdf>
- [5] Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*, 2-nd ed., Springer, New York.
- [6] Liaw, A., and Wiener, M. (2002). Classification and regression by randomForest. *R News*, 2(3), 18-22.
- [7] R Development Core Team (2010). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.  
<http://www.R-project.org>
- [8] Raymond, E. S., Relson, D., Andree, M., and Louis, G. 2004. Bogofilter.  
<http://Bogofilter.sourceforge.net>
- [9] Segal, R., Bratko, A., and Cormack, G. (2008). CEAS 2008 Spam Filter Challenge, conference talk,  
<http://www.ceas.cc/2008/challenge/results.pdf>
- [10] Sing, T., Sander, O., Beerenwinkel, N., and Lengauer, T. (2009). ROCR: Visualizing the performance of scoring classifiers. R package version 1.0-4.  
<http://CRAN.R-project.org/package=ROCR>
- [11] SpamAssassin.  
<http://spamassassin.apache.org>
- [12] Sroufe, P., Phithakkitnukoon, S., Dantu, R., and Cangussu, J. (2010). Email shape analysis. In *Distributed Computing and Networking*, Lecture Notes in Computer Science, K. Kant et al. (eds), 5935/2010, pp. 18-29.