

2019

## TEST: A terminology extraction system for technology related terms

Murhaf Hossari

*Trinity College Dublin, Ireland, murhaf.hossari@adaptcentre.ie*

Soumyabrata Dev

*Trinity College Dublin, Ireland, soumyabrata.dev@adaptcentre.ie*

John Kelleher

*Technological University Dublin, john.kelleher@tudublin.ie*

Follow this and additional works at: <https://arrow.tudublin.ie/adaptcon>



Part of the [Computer Sciences Commons](#)

---

### Recommended Citation

Murhaf Hossari, Soumyabrata Dev, and John D. Kelleher. 2019. TEST: A Terminology Extraction System for Technology Related Terms. In Proceedings of the 2019 11th International Conference on Computer and Automation Engineering (ICCAE 2019). Association for Computing Machinery, New York, NY, USA, 78–81. DOI:10.1145/3313991.3314006

This Conference Paper is brought to you for free and open access by the Adapt Research Centre at ARROW@TU Dublin. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@TU Dublin. For more information, please contact [arrow.admin@tudublin.ie](mailto:arrow.admin@tudublin.ie), [aisling.coyne@tudublin.ie](mailto:aisling.coyne@tudublin.ie), [vera.kilshaw@tudublin.ie](mailto:vera.kilshaw@tudublin.ie).

Funder: Science Foundation Ireland

# TEST: A Terminology Extraction System for Technology Related Terms

Murhaf Hossari  
ADAPT Centre, Trinity College  
Dublin  
Dublin, Ireland  
murhaf.hossari@adaptcentre.ie

Soumyabrata Dev  
ADAPT Centre, Trinity College  
Dublin  
Dublin, Ireland  
soumyabrata.dev@adaptcentre.ie

John D. Kelleher  
ADAPT Centre, Technological  
University Dublin  
Dublin, Ireland  
john.d.kelleher@dit.ie

## ABSTRACT

Tracking developments in the highly dynamic data-technology landscape are vital to keeping up with novel technologies and tools, in the various areas of Artificial Intelligence (AI). However, It is difficult to keep track of all the relevant technology keywords. In this paper, we propose a novel system that addresses this problem. This tool is used to automatically detect the existence of new technologies and tools in text, and extract terms used to describe these new technologies. The extracted new terms can be logged as new AI technologies as they are found on-the-fly in the web. It can be subsequently classified into the relevant semantic labels and AI domains. Our proposed tool is based on a two-stage cascading model – the first stage classifies if the sentence contains a technology term or not; and the second stage identifies the technology keyword in the sentence. We obtain a competitive accuracy for both tasks of sentence classification and text identification.

## CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Applied computing** → **Document management and text processing**; • **Information systems** → *Information retrieval*;

## KEYWORDS

text classification, term extraction, TEST, natural language processing.

## ACM Reference Format:

Murhaf Hossari, Soumyabrata Dev, and John D. Kelleher. 2019. TEST: A Terminology Extraction System for Technology Related Terms. In *Proceedings of 11th International Conference on Computer and Automation Engineering (ICCAE 2019)*. Perth, Australia, 4 pages. <https://doi.org/10.1145/3313991.3314006>

## 1 INTRODUCTION

With the ubiquity of online resources and the massive growth of news and blog articles, it has been extremely difficult for the end users to keep abreast of all the latest advances. Nowadays, data science and artificial intelligence have been applied to a multitude

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ICCAE 2019, February 23-25, 2019, Perth, WN, Australia

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6287-0/19/02...\$15.00

<https://doi.org/10.1145/3313991.3314006>

of applications [6], ranging from e-mail spam filtering, advertising and marketing industry [4, 7], and social media. Particularly, in these realm of online articles and blog articles, the growth in the amount of information is almost in an exponential nature. Therefore, it is important for the researchers to develop a system that can automatically parse the millions of documents, and identify the key technological terms. Such system will greatly reduce the manual work, and save several man-hours.


In this paper, we develop a machine-learning based solution, that is capable of the discovery and extraction of information related to AI technologies. We call the system **TEST**, that stands for **Terminology Extraction System for Technology-related terms**. This will greatly help the users to build the knowledge about these technologies, and also predict the trends and sentiments around them by observing the textual data found in tech news articles and blogs.

To achieve this, we combined AI techniques such as supervised and unsupervised learning on our text datasets. We use the concepts drawn from natural language processing, such as word embeddings, and conventional machine learning techniques to propose a state-of-the-art text extraction system. Our proposed ensemble method achieved a competitive performance for both sentence classification and term classification stages. The system will help the users to get insights about AI technologies in the market, and how they interact or relate to each other. This would have a great impact on decision making, and would reduce the manual review needed to peruse daily news articles to be able to analyse the top trends. This solution will make the analysis task more focused and clear.

The extracted data and insights can be also used and augmented to build more visually helpful insights. This can eventually be used to create short article snippets, that would summarise the news articles and present structured and meaningful insights about the data science technologies, in a clear and concise fashion.

## 2 RELATED WORK

In the literature, there are plethora of application areas for text classification. It is extensively used in news filtering and spam detection [8], document categorisation and text summarisation [2]. A good survey of text classification algorithms can be found in [1]. Most of the traditional techniques use decision trees, rule-based classifiers, bayesian classifiers etc. Recently, neural-network based deep networks are also used for the purpose of text summarisation and text classification. However, they do not propose a working system for these purposes – we bridge this gap in the literature by proposing a novel technology extraction system.



**Figure 1: Illustration of TEST system that automatically identifies the tech keywords from a sentence. The sentence containing technology term are highlighted in green colour. The extracted words are also colour-coded in terms of their semantics. Technology keywords are marked in yellow, and organisation names are marked in blue colour. (Text snippet extracted from theregister.co.uk.)**

Figure 1 illustrates our proposed TEST system. As an illustration, we copy verbatim a text paragraph from theregister.co.uk. Our TEST system parses each sentence in the paragraph. It identifies whether a sentence in the paragraph contains a technology-related term. In the next stage of the cascading model, the TEST system automatically identifies the keywords in the sentence. In our example, the keywords *TensorFlow* and *PyTorch* are identified as the technology terms, while *Google* can be identified as an organisation name and linked at a later stage to the technology *TensorFlow* as having an ownership relation.

The rest of the paper is organised as follows. Section 3 discusses the method, in which we prepare the dataset from internet articles and tech blogs. In Section 4, we discuss the two-stage cascading model of our proposed TEST system. We discuss the objective evaluation of our proposed system in Section 5. Finally, Section 6 concludes the paper, and discusses the potential future work.

### 3 DATASET PREPARATION

In order to train a machine learning model that is capable of extracting technology terms from text, we need a labelled dataset where the technology terms are properly annotated and tagged. To the best of our knowledge, there are no datasets available in the literature that classifies whether a sentence contains a technological keyword, and also annotates the technological keyword in a sentence. Therefore, we created a new dataset of example texts, that mostly comes from technology news articles and blogs. We crawled approximately 500 articles from various online sources, that are varied in size. We created a dataset of 150 thousands paragraphs. Together with this dataset, we also manually populated a list of technology keywords comprising single and multi-word technology terms. For example, the list contains single technology terms such as *Cortana* and, multiple-word technology terms as *Apache Hive* and *Google Cloud Natural Language API*.

Using this manually collected list of technology terms, we attempted to automatically annotate the crawled text articles. We used string matching, and we annotated every token in the text as being part of technology term  $T$  or other  $O$ . The result of this process was 3310000 tokens. The resulting dataset was highly imbalanced, due to the fact that technology terms are less frequent

than the rest of the vocabulary. In order to reduce the imbalance, we propose a cascading model where the tokens are grouped into bigger blocks (for example sentences), and then the whole block is labelled based on the fact that it contains a technology term or not. More details of this cascading model is described in the subsequent section. We used this dataset of 3.3 million tokens for our subsequent experiments.

### 4 CASCADING METHOD

In this section, we propose a cascading method that is useful in our scenario of unbalanced dataset. The first stage of the cascading method classifies whether a sentence contains a technology term or not. In case the sentence contains a technology term, the second stage of the cascading method is used, which extracts the technology keyword from the sentence. Such cascading method greatly helps in ignoring the cases when the sentence does not contain a technology term. The cascading method looks at sentence level classification as a first step, and then looks at term extraction. For this purpose, we group tokens in larger groups/collections for example sentences, and we tag each sentence as ‘contains a technology term’ or ‘doesn’t contain a technology term’.

We run the automatic annotation using the previously mentioned pattern of string matching. The resulting dataset contains 243336 sentences, out of which only about 10 thousands sentences were positive examples (i.e. contain technology terms). This resulting dataset is highly unbalanced, as there are more instances of sentences containing a technology term, as compared to otherwise. We use a random downsampling technique to remove the impact of imbalance nature, and generate a balanced dataset. The positive samples are the minority cases, whereas the negative samples are larger in number and constitute the majority cases. In order to create a balanced dataset, we consider all the 10 thousand positive examples, and then perform a random selection of 10 thousand negative examples from the remaining sentences. The resultant dataset is balanced in nature, containing equal number of positive- and negative- samples. We use this balanced dataset for our subsequent experiments and analysis.

### 4.1 Text Classification

In our proposed system, we use Facebook’s fastText [5] implementation to train our own word representations using all the textual corpus we have. We used the skipgram method in order to learn the word embeddings which is based on the concepts drawn from deep learning for natural language processing, and particularly word2vec method. It involves representing words as n-dimensional vectors. Subsequently, we used these word vectors to represent the sentences. We train our model using the corpus of 243336 sentences using the mentioned approach. After training, each word in the sentence is represented using a 300 dimensional vector. We average the vector representations of all the words in a sentence, across their corresponding elements to compute the vector representation of the entire sentence. Using this approach, we represent each sentence in the corpus with a 300 dimensional vector also.

Finally, we use a softmax function  $f$  over the binary labels – T and 0 to estimate the final label of the sentence. We assume that the total number of sentences in the corpus is  $N$ . Our objective [5] in this task of text classification is to minimise the following objective function:

$$-\frac{1}{N} \sum_{n=1}^N y_n \log(f(BAx_n)), \tag{1}$$

where  $A, B$  are weight matrices,  $x_n$  in the feature vector of  $n$ -th sentence, and  $y_n$  is the corresponding binary label.

### 4.2 Term Extraction

In the second stage of the cascading model, we are interested in extracting the tech keyword from the sentence. Similar to the text classification stage, the term extraction stage also considers the same dataset of 10 thousand positive examples, wherein each example/sentence contains a technology term. Using this dataset of 10 thousand positive samples, we expanded into a more detailed dataset comprising the individual tokens of the sentences. We tagged each token in the 10 thousand sentences individually, and labelled each token with a binary label – token is part of a technology term labelled as T, and token is not part of a technology term labelled as 0. This newly created dataset is used for the second layer of the cascading model, which deals with term extraction.

We use the labelled tokens in the sentences to train the term extraction model using Stanford Named Entity Recognition (NER) tool [3]. The NER tool uses well-engineered natural language processing feature descriptors, and represents each sentences of the corpus into discriminative features. We train a Conditional Random Field (CRF) sequence model, using the generated features. The output of the CRF model labels each token in a sentence as T or 0. Hence, we can identify the technological term(s) in a single sentence as a sequence of tokens with T labels.

We use the following features to train the CRF model:

- Current Word
- Previous Word
- Next Word
- Current Word Character all n-grams
- Current POS Tag
- Surrounding POS Tag Sequence

- Current Word Shape
- Surrounding Word Shape Sequence
- Presence of Word in Left Window (window size 4)
- Presence of Word in Right Window (window size 4)

The CRF model for NER is a well-established model, in order to estimate the probability of a hidden state <sup>1</sup>, with some *a priori* given observations [3]. We define the transition probabilities between two adjacent states as  $\phi_c$ . The term  $\phi_c$  is often described as the *clique potential*. Therefore, we can define the probability  $p$  of a chain of state sequences as:

$$p = \prod_{i=1}^N \phi_i(s_{i-1}, s_i), \tag{2}$$

where  $\phi_i$  is the clique potential for position  $i$ , with respect to the transition between states  $s_{i-1}$  and  $s_i$ .

In summary, our TEST system can be schematically described as Fig. 2. The TEST system is based on two-stage cascading model – text classification and term extraction.

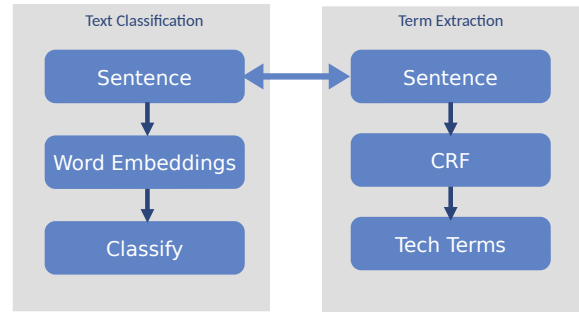


Figure 2: Detailed system diagram of TEST.

## 5 RESULTS AND EVALUATION

We evaluated the performance of our proposed system based on the performance of each of the two stages of the cascading method. As discussed earlier, the first stage deals with the sentence classification, whereas the second stage deals with the term extraction.

### 5.1 Subjective Evaluation

In Fig. 3, we present a few examples of how our system performs at run-time. In these examples, our proposed TEST system examines the text in certain technology articles found on the web. The system returns a list of sentences with potential technology terms, and another list of technology terms (can be multi-word technology term). In Fig. 3(a), our proposed system identifies that the following sentence *The idea initial spark for Portal... production team* contains a technology word. It also identifies the technology terms such as *Portal* and *Portal Plus* as indicated in yellow. However, in some cases the system misses some technology terms, such as *Facebook Phone* or *Amazon Echo*. Our system in these cases gets confused when organisation names make part of the technology term. Hence, it displays an incorrect prediction. Similar observations can be found in Fig. 3(b) and Fig. 3(c).

<sup>1</sup>In the areas of natural language processing, a state is defined as one of the possible events, that constitutes the stochastic model of the CRF.

The idea initial spark for Portal was struck a year and a half ago, leading Facebook to build out the Building 8 product team. In the intervening months, speculation has ramped up that the company was building either a Facebook Phone or Amazon Echo competitor. The latter, of course, was much closer to the truth (a least for now), though in some ways, Portal and Portal Plus are their own beasts entirely.

(a) Sample illustration shows how sentence level and term level detection can catch different technology terms.

- **Motion Auto focus:** A camera mode that allows you to tap a subject once and track it while it moves.
- **Lens Suggestions:** A new mode for Google Lens.
- **Titan M:** A new security chip with a cool name that Google touts for providing enterprise-level security.
- **Wireless charging:** Either a big deal to you or it's not.

(b) Sample illustration shows that the sentence level and term level detection can overlap in a few cases.

Seva, which is currently in Beta, certainly isn't the first company to try to solve this issue. Jain believes that with a modern application of AI and machine learning and single sign-on, Seva can provide a much more user-centric approach than past solutions simply because the technology wasn't there yet.

The way they do this is by looking across the different information types. Today they support a range of products including Gmail, Google Calendar, Google Drive,, Box, Dropbox, Slack and JIRA, Confluence. Jain says they will be adding additional services over time.

(c) Sample illustration demonstrates that the system can detect multiple sentences and technology terms.

**Figure 3: Illustration showing the results in executing the TEST system against technology articles. The green colour marks the sentences with potential technology terms which were detected by the system. The yellow colour marks the technology terms that are detected by the system. Finally, the red colour marks the terms that are missed by the system.**

## 5.2 Objective Evaluation

In order to provide an objective evaluation of our system, we report the F-score of the two stages. We use F-score measure to evaluate how good our system performs in a real-world dataset.

Suppose,  $TP$ ,  $FP$ ,  $TN$  and  $FN$  represent the true positive, false positive, true negative and false negative samples of the task. The F-score is defined as:

$$\text{F-Score} = \frac{(2 \times \text{Precision} \times \text{Recall})}{(\text{Precision} + \text{Recall})}, \quad (3)$$

where precision and recall are respectively defined as  $\text{Precision} = \frac{TP}{TP+FP}$  and  $\text{Recall} = \frac{TP}{TP+FN}$ .

In the first stage comprising the sentence classification part, we split our dataset of 20 thousand sentences into 3 parts – training set

70%, validation set 15% and test set 15%. We maintained a balanced distribution between the positive and negative examples while splitting. In the second stage comprising the term extraction from a sentence, we used a similar split. We ensured that every token in the data gets a positive (technology) or a negative (non technology) tag. We also used the same assessment metric of F-score to evaluate how well this stage performs.

Table 1 summaries the performance of our proposed system. We observe that the proposed TEST system has a competitive score in both its stages. We obtain a F-score of 0.93 and 0.96 respectively in both the stages.

Stage	Activity	F-score
I	Sentence classification	0.93
II	Term extraction	0.96

**Table 1: Performance evaluation of our proposed system.**

It is important to benchmark these results with other similar technology text extraction systems. However, because of the lack of the existence of such similar systems, we could not benchmark our performance with other similar systems.

## 6 CONCLUSION

In this paper, we propose an end-to-end system called TEST that automatically extracts the technology terms from text. Our system is trained on a large corpus of technology-related news articles and blogs, and has a competitive performance in detecting and extracting tech terms from a sentence. In the future, we plan to use the sentiment around these technologies in order to evaluate and predict the impact of those technologies and tools in any specific AI area. In addition to word embeddings, we also plan to borrow techniques from convolutional neural networks (CNN), Recurrent Neural Networks (RNNs), Long Short Term Memory Units (LSTMs) etc., to propose an ensemble-based TEST system.

## ACKNOWLEDGMENTS

The ADAPT Centre for Digital Content Technology is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

## REFERENCES

- [1] C. C. Aggarwal and C. Zhai. 2012. A survey of text classification algorithms. In *Mining text data*. Springer, 163–222.
- [2] S. Chakrabarti, B. Dom, R. Agrawal, and P. Raghavan. 1997. Using taxonomy, discriminants, and signatures for navigating in text databases. In *VLDB*, Vol. 97. 446–455.
- [3] J. R. Finkel, T. Grenager, and C. Manning. 2005. Incorporating non-local information into information extraction systems by Gibbs sampling. In *Proceedings of the 43rd annual meeting on association for computational linguistics*. Association for Computational Linguistics, 363–370.
- [4] M. Hossari, S. Dev, M. Nicholson, K. McCabe, A. Nautiyal, C. Conran, J. Tang, X. Wei, and F. Pitié. 2018. ADNet: A Deep Network for Detecting Adverts. In *Proc. Irish Conference on Artificial Intelligence and Cognitive Science (AICS 2018)*.
- [5] A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759* (2016).
- [6] J. D. Kelleher and B. Tierney. 2018. *Data Science*. The MIT Press.
- [7] A. Nautiyal, K. McCabe, M. Hossari, S. Dev, M. Nicholson, C. Conran, D. McKibben, J. Tang, X. Wei, and F. Pitié. 2018. An Advert Creation System for Next-Gen Publicity. In *Proc. European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*.
- [8] M. Sahami, S. Dumais, D. Heckerman, and E. Horvitz. 1998. A Bayesian approach to filtering junk e-mail. In *Learning for Text Categorization: Papers from the 1998 workshop*, Vol. 62. Madison, Wisconsin, 98–105.