

RESEARCH ARTICLE



Comparing BERT Against Traditional Machine Learning Models in Text Classification

Eduardo C. Garrido-Merchan^{1*}, Roberto Gozalo-Brizuela¹ and Santiago Gonzalez-Carvajal²

¹*Quantitative Methods Department, Comillas Pontifical University, Spain*

²*Artificial Intelligence Department, Polytechnic University of Madrid, Spain*

Abstract: The Bidirectional Encoder Representations from Transformers (BERT) model has arisen as a popular state-of-the-art model in recent years. It is able to cope with natural language processing (NLP) tasks such as supervised text classification without human supervision. Its flexibility to cope with any corpus delivering great results has made this approach very popular in academia and industry, although other approaches have been used before successfully. We first present BERT and a review on classical NLP approaches. Then, we empirically test with a suite of different scenarios the behavior of BERT against traditional term frequency – inverse document frequency vocabulary fed to machine learning models. The purpose of this work is adding empirical evidence to support the use of BERT as a default on NLP tasks. Experiments show the superiority of BERT and its independence of features of the NLP problem such as the language of the text adding empirical evidence to use BERT as a default technique in NLP problems.

Keywords: BERT, natural language processing, machine learning, comparison

1. Introduction

Natural language processing (NLP) methodologies have flourished, and lots of papers solving different tasks of the field, such as text classification (Aggarwal & Zhai, 2012), named entity recognition, or summarization (Puente et al., 2013), have been published. We can differentiate, mainly, between two types of approaches to NLP problems: firstly, linguistic approaches (Cambria & White, 2014) that generally use different features of the text that the experts on the domain consider that are relevant have been extensively used. Those features could be combinations of words or n-grams (Stamatatos, 2011), grammatical categories, unambiguous meanings of words, words appearing in a particular position, categories of words, and much more. These features could be built manually for a specific problem or can be retrieved by using different linguistic resources (Besançon et al., 2010) such as ontologies.

On the other hand, machine learning (ML) (Manning & Schütze, 1999) and deep learning-based approaches (Otter et al., 2020) classically have analyzed annotated corpora of texts inferring which features of the text, typically in a bag of words fashion (Zhang et al., 2010) or by n-grams, are relevant for the

classification automatically. Both approaches have their pros and cons; concretely, linguistic approaches have great precision, but their recall is low as the context where the features are useful is not as big as the one processed by ML models. Although the precision of classical NLP systems was, until recently, generally better as the one delivered by ML (Garrido-Merchán, 2015). Nevertheless, recently, thanks to the rise of computation, ML text classification dominates in scenarios where huge sizes of texts are processed.

Generally, linguistic approaches consist in applying a series of rules, which are designed by linguistic experts (Khurana et al., 2023). An example of linguistic approach can be found at Hutto & Gilbert (2014). The advantage of these type of approaches over ML-based approaches is that they do not need large amounts of data. Regarding ML-based approaches, they usually have a statistical base. We can find many examples of these type of approaches: Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2018), transformers (Vaswani et al., 2017), etc.

Another issue with traditional NLP approaches is multilingualism (Bikel & Zitouni, 2012). We can design rules for a given language, but sentence structure, and even the alphabet, may change from one language to another, resulting in the need to design new rules. Some approaches such as the universal networking language (UNL) standard (Uchida & Zhu, 2001) try to circumvent this issue, but the multilingual resource is hard to build and requires experts on the platform. Another problem with UNL approaches and related ones would be, given a specific

*Corresponding author: Eduardo C. Garrido-Merchan, Quantitative Methods Department, Comillas Pontifical University, Spain. Email: ecgarrido@comillas.edu

language, the different forms of expression, i.e., the way we write in, for example, Twitter, is very different from the way we write a more formal document, such as a research paper (Farzindar & Inkpen, 2015).

BERT is a NLP model that was designed to pre-train deep bidirectional representations from unlabeled text and, after that, be fine-tuned using labeled text for different NLP tasks (Devlin et al., 2018). That way, with BERT model, we can create state-of-the-art models for many different NLP tasks (Devlin et al., 2018). We can see the results obtained by BERT in different NLP tasks at Devlin et al. (2018).

In this work, we compare BERT model (Devlin et al., 2018) with a traditional ML NLP approach that trains ML models in features retrieved by the term frequency – inverse document frequency (TF-IDF) (Zhang et al., 2005) algorithm as a representative of these traditional approaches (Trstenjak et al., 2014). With this technique, we avoid the construction of a linguistic resource that needs expert supervision, simulating it with the punctuation retrieved for any term by the TF-IDF technique. We lose precision by doing this operation but gain recall.

We have carried out four different experiments about text classification. In all of them, we have used two different classifiers: BERT and a traditional classifier created in the way that we have just explained. In this work, we start by presenting some related work, then we describe the models we have used in our experiments, after that, we describe the experiments we have carried out and show the obtained results and, finally, we present the conclusions drawn from the work and some future lines of work.

2. Literature Review

In this section, we summarize the main comparisons against advanced models such as the BERT transformer and classical NLP. Recently, BERT has achieved state-of-the-art results in a broad range of NLP tasks (Devlin et al., 2018), so the question that is discussed is whether classical NLP techniques are still useful in comparison to the outstanding behavior of BERT and related models.

It is interesting to study how does the BERT model represent the steps of the traditional NLP pipeline (Tenney et al., 2019) in order to make a fair comparison. The main conclusion of this paper is that their work shows that the model adapts to the classical NLP pipeline dynamically, revising lower-level decisions on the basis of disambiguating information from higher-level representations. In other words, we can think of BERT as a generalization of the traditional NLP pipeline, hence being more dynamic.

An argument that defends classical ML NLP approaches is that the BERT approach needs huge amounts of texts to deliver proper results. An interesting work, Usherwood and Smit (2019) who focus on a pure empirical comparison of BERT and ULMFiT (Rother & Rettberg, 2018) w.r.t traditional NLP approaches in low-shot classification tasks where we only have \$100\$–\$1000\$ labeled examples per class, shows how BERT, representing the best of deep transfer learning, is the best-performing approach, outperforming top classical ML models, thanks to the use of transfer learning (Devlin et al., 2018). In our work, we are going to test this hypothesis under different problems that also involve texts in different languages.

A common critique of classical NLP practitioners is that the BERT model and ML methodologies can be fooled easily, committing errors that may be severe in certain applications and that can be easily solved by symbolic approaches. Following this reasoning, in this work (Jin et al., 2019) the authors present the

TextFooler baseline that generates adversarial text in order to fool BERT's classification (Jin et al., 2019). We wonder if these experiments are representative of common scenarios and hypothesize that, although it is true that some texts may fool BERT, they are not representatives of common problems. In order to test this hypothesis, we are going to measure the results given by BERT in different languages. If BERT fail in these problems, then these adversaries may be common. Although if BERT outperforms classical approaches under standard circumstances, then we can state that these adversarial attacks may not be common.

3. The BERT Model and the Traditional ML NLP Methodology

Having reviewed related work, we will now introduce the traditional NLP approaches that we are comparing with BERT and then the details of the BERT model.

3.1. Term frequency – Inverse document frequency

A classical way to deal with a supervised learning NLP task is to build a bag-of-words model with the most weighted words given by the TF-IDF algorithm.

Assuming there are N documents in the collection, and that term t_i occurs in n_i of these documents. Then, inverse document frequency can be computed as:

$$idf(t_i) = \log \frac{N}{n_i}. \quad (1)$$

Actually, the original measure was an integer approximation to this formula, and the logarithm was base 2. However, Aggarwal and Zhai (2012) are the most commonly cited form of IDF. For more information, we refer the reader to the original source (Robertson, 2004).

On the other hand, given a term t_i , we denote by tf_i the frequency of the term t_i in the document under consideration (Robertson, 2004).

Finally, TF-IDF is defined for a given term t_i in a given document as follows:

$$tfidf(t_i) = tf_i \cdot idf(t_i).$$

In our experiments, regarding the standard NLP algorithms, we will be using TF-IDF to build a vocabulary for a ML model. Further details are introduced in Section 4.

3.2. Bidirectional encoder representations from transformers

We now explain what we consider to be the state-of-the-art technique on NLP. Regarding the BERT model, there are two steps in its framework: *pre-training* and *fine-tuning* (Devlin et al., 2018). During pre-training, the model is trained on unlabeled large corpus. For fine-tuning, the model is initialized with the pre-trained parameters, and all the parameters are fine-tuned using labeled data for specific tasks.

BERT's model architecture is a multi-layer bidirectional transformer encoder (Devlin et al., 2018) based on the original implementation described in Vaswani et al. (2017).

This kind of encoder is composed of a stack of $N = 6$ identical layers. Each of these layers has two sub-layers. The first one is a multi-head self-attention mechanism, and the second one is a simple position-wise fully connected feedforward network.

It employs a residual connection (He et al., 2016) around both sub-layers, followed by a layer normalization, that is, the output of each sub-layer is $LayerNorm(x + Sublayer(x))$, where $Sublayer(x)$ is the function implemented by the sub-layer (Vaswani et al., 2017).

In relation to multi-head self-attention, first, we need to define scaled dotproduct attention. It is defined as follows:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V,$$

where Q is the matrix of queries, K is the matrix of keys, V is the matrix of values, and d_k is the dimension of the Q and K matrices. Now, we can define the multi-head attention as

$$MultiHead(Q, K, V) = Concat(head_1, \dots, head_n)W^O,$$

where $head_i = Attention(QW_i^Q, KW_i^K, VW_i^V)$. Multi-head attention consists of projecting the queries, keys, and values h times with different, learned linear projections to d_k , d_k , and d_v (dimension of the values matrix), respectively. Then, on each of these projected versions of the queries, keys, and values, we perform the attention function in parallel, yielding in d_v -dimensional output values. Finally, these are concatenated and projected, resulting in the final values (Vaswani et al., 2017). Self-attention means that all of the keys, values, and queries come from the same place.

BERT represents a single sentence or a pair of sentences (e.g., the pair (*question, answer*)) as a sequence of tokens according to the following features: BERT uses WordPiece embeddings (Wu et al., 2016). The first token of the sequence is “[CLS]”. When there is a pair of sentence, in the sequence, they are separated by the “[SEP]” token. And, an embedding is added to every token indicating whether it belongs to the first or the second sentence. For a given token, its input representation is constructed by summing the corresponding token, position, and segment embeddings (Devlin et al., 2018).

Pre-training is divided into *masked LM* and *next sentence prediction*. The first one consists in masking some percentage of the input tokens at random (using the “[MASK]” token) and then predict those masked tokens. The second one consists in, given two sentences A and B, 50% of the time B is the actual next sentence that follows A (labeled as IsNext), and 50% of the time B is a random sentence from the corpus (labeled as NotNext) (Devlin et al., 2018).

Fine-tuning is straightforward since the self-attention mechanism in the transformer allows BERT to model many downstream tasks. For each task, we simply plug in the specific inputs and outputs into BERT and fine-tune all the parameters (Devlin et al., 2018).

A random sampling method was employed for the study. The sample for the study consisted of 22 first year in-service postgraduate science teachers from one of the Colleges of Education in Bhutan. Thirteen male (59%) teachers and 9 female (40.9%) teachers participated in the study. The sample comprised of 7 biology teachers (31.8%), 10 chemistry teachers (45.5%), and 5 physics teachers (22.7%). Since the participation for this study was purely on a voluntary basis, only 22 out of 39 in-service postgraduate science teachers took part in the study. The overall response rate was recorded at 56%.

4. Experiments

In order to compare BERT model with respect to the traditional ML NLP methodology, we have designed four experiments that are described throughout the section.

In these experiments, we will be using TfidfVectorizer from sklearn Python 3 module. After using TF-IDF to preprocess the text, we will be using Predictor from auto ml module (in the third and fourth experiments) and H2OAutoML from h2o module (in the second experiment), to find the best model to fit the data. In the first experiment, we will, instead, show how much work needs to be done in order to get close to the results obtained, with no effort, using BERT model. For this purpose, we will be using many sklearn models and study their results in depth.

Regarding BERT’s implementation, we have used the pre-trained BERT model from ktrain Python 3 module. This model expects the following directory structure: a directory which must contain two subdirectories: **train** and **test**. Each one of them, in turn, must contain one subdirectory per class (named after the name of the class they represent). And, finally, each class directory must contain the “.txt” files (their name is irrelevant) with the texts that belong to the class they represent.

4.1. IMBD experiment

In the first experiment, we have downloaded the IMDB dataset from the following website. It contains 50,000 movie reviews (25,000 to train the model and 25,000 to test it) to perform sentiment analysis, a popular supervised learning text classification task. The dataset is classified into two different classes: positive and negative movie reviews.

We have compared the behavior of a pre-trained default BERT model w.r.t different popular ML models such as Support Vector Classifier (SVC) or logistic regression that use a vocabulary extracted from a TF-IDF model obtaining the following results (Table 1).

Table 1
Accuracy retrieved by the different methodologies in the IMDB experiment over the validation set

Model	Accuracy
BERT	0.9387
Voting classifier	0.9007
Logistic regression	0.8949
Linear SVC	0.8989
Multinomial NB	0.8771
Ridge classifier	0.8990
Passive aggressive classifier	0.8931

As we can see, BERT outperforms the rest of the models. It is noteworthy that obtaining these results with the traditional approaches has been far more complicated than obtaining this result with BERT.

4.2. RealOrNot tweets experiment

Our second experiment deals with the RealOrNot tweets written in English. We have downloaded the dataset from the following website. The task to solve here is pure binary text classification. It contains tweets classified into two different

classes: tweets about a real disaster and tweets which are not about a real disaster.

We have just used the *tweet* and *class* columns. We have also used the re Python 3 module to preprocess the tweets (#anything -> hashtag, @anyone -> entity, etc.). After that, we have generated the directory structure that we need to use BERT model (using 75% data to train and 25% data to validate). The obtained results are summarized in Table 2.

Table 2
RealOrNot experiment results

Model	Accuracy	Kaggle score
BERT	0.8361	0.83640
H2OAutoML	0.7875	0.77607

Finally, we have classified the data from the Kaggle competition with BERT. We have scored **0.83640**. We can see this result here (Santiago González). Regarding the traditional approaches, the best classifier from the h2o module has turned out to be the H2OStackedEnsembleEstimator: Stacked Ensemble with model key StackedEnsemble BestOfFamily AutoML 20200221 120302. And, its score in the competition has been **0.77607**.

4.3. Portuguese news experiment

Description. Having seen that BERT has outperformed an AutoML technique and other classical ML models using a vocabulary built from a traditional NLP technique such as TF-IDF in the English language, we choose to change the language to see if the BERT model also behaves well. We have downloaded the Portuguese news dataset from the following website. It contains articles from the news classified into nine different classes: ambiente, equilibriosaude, sobretudo, educacao, ciencia, tec, turismo, empreendedorsocial, and comida.

We have just used the *article text* and *class* columns. We have generated the directory structure that we need to use BERT model (using 75% data to train and 25% data to validate obtaining the following results (Table 3).

Table 3
Portuguese news experiment results

Model	Accuracy	Kaggle score
BERT	0.9093	0.91196
Predictor (auto ml)	0.8480	0.85047

Finally, we have classified the data for the Kaggle competition scoring a **0.91196** accuracy. We can see this result here (Santiago González). Regarding the traditional methods, the best classifier has turned out to be a GradientBoostingClassifier. And, the score in the competition of this model has been **0.85047**.

4.4. Chinese hotel reviews experiment

Description. Our last experiment involves a completely different language, Peninsular Chinese simplified characters zh-CN, where we hypothesize that, given that the way of expressing this language is through different symbols that are not separated by spaces BERT may not output a good result. The experiment is a sentiment analysis problem involving Chinese hotel reviews. We have downloaded the dataset from the

following website. It contains hotel reviews classified into two different classes: positive hotel reviews and negative hotel reviews.

In this experiment, we have used 85% of the data to train the model and 15% of the data to validate it. Results are given in Table 4.

Table 4
Chinese hotel reviews results

Model	Accuracy
BERT	0.9381
Predictor (auto ml)	0.7399

We can observe how, independently of the language and its characteristics, BERT behavior outperforms classical NLP approach.

Finally, we have tried to do some predictions with BERT using Google Translator. For example, we have tried to predict a class for: 这家酒店的风景和服务都非常糟糕, which means “the view and service of this hotel are very bad.” The predicted class for this hotel review has been neg, which is correct.

Regarding the traditional approaches, the best model has turned out to be a GradientBoostingClassifier. But in this case, the model has been pretty bad, since the probability for both classes is very close. In this experiment, the importance of transfer learning has become apparent, since the dataset was pretty small compared to the ones used in the previous experiments.

5. Conclusions and Further Work

In this work, we have introduced the BERT model and the classical NLP strategy where a ML model is trained using the features retrieved with TF-IDF and hypothesize about the behavior of BERT w.r.t these techniques in search of a default technique to tackle NLP tasks. We have introduced four different NLP scenarios where we have shown how BERT has outperformed the traditional NLP approach, adding empirical evidence of its superiority in average NLP problems w.r.t. classical methodologies. It is also noteworthy the importance of transfer learning. We have been able to obtain this results, thanks to pre-training. Transfer learning has become more apparent in experiment 4.4 (which has the smallest dataset among all the experiments). We are nevertheless aware of the limitations of the BERT model. Although it seems that it is a good default for NLP tasks, its results can be improved. In order to do so, we would like to research in a hyperparameter auto-tuned BERT model for any new NLP task with Bayesian optimization. We would like to use that autotuned BERT to enable classification of language messages for robots (Garrido-Merchán & Molina, 2020; Garrido-Merchán et al., 2020) showing consciousness correlated behaviors.

Recommendations

The finding revealed that the lack of training for both teachers and students was the main factor that prevented them from using educational technology tools in teaching and learning ecology. Therefore, training on educational technology for both teachers and students is recommended. Since educational technology tools have arose excitement and curiosity among students, they recommended other module tutors to use educational technology tools as well. Educational technology tools integrated in the module will be further replicated by students teacher during teaching practice or as a full-fledge teacher. Therefore, tutors were

recommended to use a variety of educational technology tools in learning, teaching, and an assessment.

Acknowledgements

The authors gratefully acknowledge the use of the facilities of Centro de Computación Científica (CCC) at Universidad Autónoma de Madrid. The authors also acknowledge financial support from Spanish Plan Nacional I+D+i, grants TIN2016-76406-P and TEC2016-81900-REDT.

Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

References

- Aggarwal, C. C., & Zhai, C. (2012). *Mining text data*. Germany: Springer.
- Besaçon, R., De Chalendar, G., Ferret, O., Gara, F., Mesnard, O., Laïb, M., & Semmar, N. (2010). LIMA: A multilingual framework for linguistic analysis and linguistic resources development and evaluation. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, 3697–3704.
- Bikel, D., & Zitouni, I. (2012). *Multilingual natural language processing applications: From theory to practice*. USA: IBM Press.
- Cambria, E., & White, B. (2014). Jumping NLP curves: A review of natural language processing research. *IEEE Computational Intelligence Magazine*, 9(2), 48–57. <https://doi.org/10.1109/MCI.2014.2307227>
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv PrePrint:1810.04805*. <https://doi.org/10.48550/arXiv.1810.04805>
- Farzindar, A., & Inkpen, D. (2015). Natural language processing for social media. *Synthesis Lectures on Human Language Technologies*, 8(2), 1–166. <https://doi.org/10.48550/arXiv.1810.04805>
- Garrido-Merchán, E., & Molina, M. (2020). A machine consciousness architecture based on deep learning and Gaussian processes. In *Hybrid Artificial Intelligent Systems: 15th International Conference*, 15, 350–361. https://doi.org/10.1007/978-3-030-61705-9_29
- Garrido-Merchán, E. C., Molina, M., & Mendoza, F. M. (2020). An artificial consciousness model and its relations with philosophy of mind. *arXiv PrePrint:2011.14475*. <https://doi.org/10.48550/arXiv.2011.14475>
- Garrido-Merchán, E. (2015). *Context-adaptive supervised expansion of polarized lexicons*. Master's Thesis, Comillas Pontifical University.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- Hutto, C., & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1), 216–225. <https://doi.org/10.1609/icwsm.v8i1.14550>
- Jin, D., Jin, Z., Zhou, J. T., & Szolovits, P. (2019). Is BERT really robust? Natural language attack on text classification and entailment. *arXiv PrePrint:1907.11932*, 2.
- Khurana, D., Koli, A., Khatter, K., & Singh, S. (2023). Natural language processing: State of the art, current trends and challenges. *Multimedia Tools and Applications*, 82(3), 3713–3744. <https://doi.org/10.1007/s11042-022-13428-4>
- Manning, C., & Schütze, H. (1999). *Foundations of statistical natural language processing*. USA: MIT Press.
- Otter, D. W., Medina, J. R., & Kalita, J. K. (2020). A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2), 604–624. <https://doi.org/10.1109/TNNLS.2020.2979670>
- Puente, C., Olivas, J. A., Garrido, E., & Seisdedos, R. (2013). Creating a natural language summary from a compressed causal graph. In *2013 Joint IFSA World Congress and NAFIPS Annual Meeting*, 513–518. <https://doi.org/10.1109/IFSA-NAFIPS.2013.6608453>
- Robertson, S. (2004). Understanding inverse document frequency: On theoretical arguments for IDF. *Journal of Documentation*, 60(5), 503–520. <https://doi.org/10.1108/00220410410560582>
- Rother, K., & Rettberg, A. (2018). ULMFiT at germeval-2018: A deep neural language model for the classification of hate speech in German tweets. In *14th Conference on Natural Language Processing*, 113–119.
- Stamatatos, E. (2011). Plagiarism detection using stopword n-grams. *Journal of the American Society for Information Science and Technology*, 62(12), 2512–2527. <https://doi.org/10.1002/asi.21630>
- Tenney, I., Das, D., & Pavlick, E. (2019). BERT rediscovered the classical NLP pipeline. *arXiv preprint:1905.05950*. <https://doi.org/10.48550/arXiv.1905.05950>
- Trstenjak, B., Mikac, S., & Donko, D. (2014). KNN with TF-IDF based framework for text categorization. *Procedia Engineering*, 69, 1356–1364. <https://doi.org/10.1016/j.proeng.2014.03.129>
- Uchida, H., & Zhu, M. (2001). The universal networking language beyond machine translation. In *International Symposium on Language in Cyberspace*, 26–27.
- Usherwood, P., & Smit, S. (2019). Low-shot classification: A comparison of classical and deep transfer machine learning approaches. *arXiv PrePrint:1907.07543*. <https://doi.org/10.48550/arXiv.1907.07543>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., . . . , & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., . . . , & Dean, J. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint:1609.08144*. <https://doi.org/10.48550/arXiv.1609.08144>
- Zhang, Y. Z., Gong, L., & Wang, Y. C. (2005). An improved TF-IDF approach for text classification. *Journal of Zhejiang University-Science A*, 6, 49–55.
- Zhang, Y., Jin, R., & Zhou, Z. H. (2010). Understanding bag-of-words model: A statistical framework. *International Journal of Machine Learning and Cybernetics*, 1, 43–52. <https://doi.org/10.1007/s13042-010-0001-0>

How to Cite: Garrido-Merchan, E. C., Gozalo-Brizuela, R., & Gonzalez-Carvajal, S. (2023). Comparing BERT Against Traditional Machine Learning Models in Text Classification. *Journal of Computational and Cognitive Engineering*, 2(4), 352–356, <https://doi.org/10.47852/bonviewJCCE3202838>