Durham Research Online

Deposited in DRO:

06 December 2021

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Atapour-Abarghouei, Amir and Bonner, Stephen and McGough, Andrew Stephen (2021) 'Rank over Class: The Untapped Potential of Ranking in Natural Language Processing.', 2021 IEEE International Conference on Big Data (IEEE BigData 2021) <https://bigdataieee.org/BigData2021/>.

Further information on publisher's website:

<https://bigdataieee.org/BigData2021/>

Publisher's copyright statement:

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Rank over Class: The Untapped Potential of Ranking in Natural Language Processing

Amir Atapour-Abarghouei
Department of Computer Science
Durham University
Durham, UK

amir.atapour-abarghouei@durham.ac.uk

Stephen Bonner
School of Computing
Newcastle University
Newcastle, UK

stephen.bonner3@newcastle.ac.uk

Andrew Stephen McGough
School of Computing
Newcastle University
Newcastle, UK

stephen.mcgough@newcastle.ac.uk

Abstract—Text classification has long been a staple within Natural Language Processing (NLP) with applications spanning across diverse areas such as sentiment analysis, recommender systems and spam detection. With such a powerful solution, it is often tempting to use it as the go-to tool for all NLP problems since when you are holding a hammer, everything looks like a nail. However, we argue here that many tasks which are currently addressed using classification are in fact being shoehorned into a classification mould and that if we instead address them as a ranking problem, we not only improve the model, but we achieve better performance. We propose a novel end-to-end ranking approach consisting of a Transformer network responsible for producing representations for a pair of text sequences, which are in turn passed into a context aggregating network outputting ranking scores used to determine an ordering to the sequences based on some notion of relevance. We perform numerous experiments on publicly-available datasets and investigate the applications of ranking in problems often solved using classification. In an experiment on a heavily-skewed sentiment analysis dataset, converting ranking results to classification labels yields an approximately 22% improvement over state-of-the-art text classification, demonstrating the efficacy of text ranking over text classification in certain scenarios.

Index Terms—Sentiment Analysis, Text Ranking, Text Classification, Natural Language Processing, Deep Learning

I. INTRODUCTION

Recent advances in machine learning have led to significant strides in various active areas of research, with many applications already integrated into our daily lives. However, as natural languages are the primary method of human communication, learning-based Natural Language Processing (NLP) is receiving an ever-increasing level of attention within both academia and industry. Amongst the various applications of NLP, text classification [1], [2] has arguably blazed the trail due to the simplicity of its definition and its numerous use cases. From online content tagging [3] to sentiment analysis [4], text classification has always been at the forefront of natural language processing.

With the emergence of deep learning, various approaches have addressed text classification via feed-forward networks using bag-of-word inputs [5], recurrent neural networks that consider structural elements of the text [6] and convolutional neural networks capable of detecting position-invariant patterns in the text [7]. Transformers [8], however, have arguably been the greatest advancement in NLP with significant

improvements enabled by large-scale pre-trained language models, taking advantage of deeper architectures and larger corpora of text used for training to enable a stronger and better representation learning process.

However, despite the significant advances in text classification, the field is still fraught with challenges. For instance, text classification can be highly subjective due to the presence of textual ambiguity, biased annotations and unknown classes. Examples of this are widely seen in the numerous publicly-available movie review datasets [9], [10], [11], commonly used as benchmarks for text classification. As an example, imagine the following sentence from a movie review:

Example 1: *Movie Review - while plagued with a plodding mess of a narrative, the sincere performance of the main character salvages the clichéd dialogue and provides some escapism from the distorted perspective of the protagonist.*

Any observer would have trouble labelling this review as positive or negative. In the same vein, the performance of a machine learning model trained on subjectively annotated reviews would solely depend on the presence of similar words, patterns and structures of the text in other less ambiguous and more concrete data points in the dataset. This level of subjectivity and ambiguity essentially makes a fair and accurate classification of the sentiment of such passages very difficult using learning-based models.

Another significant issue in text classification stems from significant imbalances commonly found in existing datasets. If opinions on a specific topic are extracted from social media to be used as training data, most would lie on the extreme ends of the spectrum as individuals with extreme beliefs are more likely to voice their opinions publicly, thus creating a skew in the dataset towards more extreme opinions [12]. In other scenarios, if subjective passages are labelled by multiple annotators with an averaging method determining the final label, most of the passages will inevitably fall near the centre of the labelling interval due to the subjectivity of the annotation [13]. Such imbalances can impede accurate classification and necessitate significant algorithmic intervention.

There are numerous other similar shortcomings in text classification, depending on the scope of the problem, the nature of the classes and the features available in the dataset.

Here, we argue that many such issues can be resolved by reformulating certain text classification problems as *text ranking*. It is important to note that we are in no way suggesting that ranking should completely replace classification. On the contrary, classification remains an integral solution to a vast portion of the problems in computational linguistics. Our objective here is to investigate the possibility of reformulating certain key problems as ranking to address the common challenges hindering classification solutions.

Ranking is defined as a derivation of ordering over a list of items that maximises the utility of the entire list [14]. It is a significant component of many information retrieval systems with applications including web search, recommender systems, document summarisation and question answering [15].

Ranking, by nature, is different from classification and regression in that a classifier or a regressor attempts to assign a specific class label or value to an individual data point, while the objective of a ranking approach is to optimally sort a list, such that, for some notion of *relevance*, the items within the list with higher relevance scores appear earlier in the list.

Note that while we primarily discuss the efficacy of ranking compared to classification in certain problems, the same conclusions can be made about regression solutions. The reason why classification is chosen as the main point of comparison is that the existing literature is dominated by text classification solutions to sentiment analysis [16], [17], [18], [19] due to the availability of discrete class labels in benchmark datasets [9], [10], [11] and the difficulty of attaining accurate continuous regression labels for subjective text data.

In this paper, we argue that for certain problems, ranking is a more appropriate formulation than classification (or regression). Take the movie review from *Example 1*. While it might be difficult for a human or a neural network to accurately classify this as a positive or a negative review or associate it with an absolute sentiment value or class, it is significantly easier to rank this review with respect to other reviews as *more positive* or *more negative*, especially if they are placed in some known shared *context*, for example all the reviews given for the same movie. In this paper, we explore the capabilities of text ranking and demonstrate its efficacy using experiments conducted over publicly-available datasets.

In short, the main contributions of this work are as follows:

- With the end-to-end training of a Transformer, a context aggregating dense network and a pair-wise ranking loss, text passages are accurately ranked based on a given ranking label or score (Section III-A).
- The potential of text ranking is demonstrated using various datasets [11], [20] with the learning process *pivoting* around the different contexts the passages can be placed in, *e.g.* the writer of a given passage or the topic of the passage, and the effectiveness of this contextual ranking is explored (Sections IV-A and IV-B).
- When ranking results are artificially converted to class labels, our approach is capable of outperforming state-of-the-art text classification models on heavily-skewed text classification datasets [21] (Section IV-C).

To enable reproducibility, an implementation of our ranking approach is publicly available¹.

II. RELATED WORK

Having first appeared in the literature in the 1940s [22], ranking gained prominence as the foundation of modern search engines towards the end of the last millennium [23]. Given a query q and a collection P of passages p that match the query, the goal is to rank all the passages p in P according to some notion of relevance to q so that the *best* results appear earlier in the ranked list. Note that while we, in this work, focus on passage ranking, the same concepts can equally be applied to any data modality.

Though traditionally solved via boolean, vector space and probabilistic models [24], [25], ranking is now commonly performed using learning-based approaches [15], taking advantage of labelled data and some parametrised function to map feature vectors extracted from list items to real values used as ranking scores. This function is subsequently used to sort the items. Ranking approaches can be point-wise, pair-wise or list-wise.

Point-wise approaches [26] utilise a classifier or a regressor trained to predict the relevance score of a passage, which is subsequently used to sort the items, with respect to a given query. Pair-wise techniques [27], on the other hand, consider a pair of items at each pass with the overall objective being to minimise the number of inversions, where the items are in the wrong order relative to the ground truth, in the ranked list. List-wise approaches [28] attempt to solve for the optimal ranking for the entire list all at once. This is often done either via a specific loss function design, based on the unique properties of the items [29], or by directly optimising certain information retrieval metrics [30].

In this paper, we propose a pair-wise ranking approach, capable of outperforming classification in certain natural language processing problems. Unlike point-wise approaches, a pair-wise technique, such as the one proposed here, receives more context as it directly compares list items. List-wise models, despite considering the explicit global context and removing the need for post-processing, are more difficult to optimise, slower to converge and suffer from training instability issues, which is why a pair-wise ranking approach is implemented for this study.

III. PROPOSED APPROACH

Consisting of two sub-networks trained end-to-end, our overall approach receives two passages as its input and outputs two scalar values which can be used to determine whether the passages have been input in the correct order. See Figure 1.

The first sub-network (left) is a Transformer that produces a representation of the input passages. The second sub-network (right) is a context aggregating multilayer perceptron, which receives the latent vectors representing the passages (produced by the Transformer network) and outputs ranking scores, used to determine the ordering of the two input passages.

¹<https://github.com/atapour/rank-over-class>

Fig. 1: The overall model approximates the ranking function, f . Input passages, x_1 and x_2 , are first passed through a Transformer network. A multilayer perceptron then takes the concatenated Transformer outputs ($z'_1 \parallel z'_2$) and produces the ranking scores, $f(x_1)$ and $f(x_2)$, used to rank the inputs.

The proposed approach performs contextually sound comparisons between passage pairs to ensure meaningful ranking results. Details of the approach are discussed in the following.

A. Context is King

In text classification, either it can be assumed that the entirety of the training dataset falls within the same context or the context is completely ignored. Context, in this setting, refers to a unifying element which the passages can be grouped by and provides a coherent background for understanding what those passages represent. To get a clearer picture of what is meant by context, let's refer to the following examples:

Example 2: a human observer is given three random passages and asked to assess their quality. One passage is an excerpt from a technical report, one a poem and the other from a history textbook.

Example 3: a human observer is given three random passages and asked to assess their quality. All three passages contain answers from different individuals to the same question.

In *Example 2*, the passages are from different sources, follow different linguistic patterns and address different types of information. As such, they lack a shared context. This makes objectively comparing their quality extremely difficult, even for a human expert. However, in *Example 3*, the passages are all answers to the same question and follow a shared context. This means the problem can be solved in a more meaningful and objective manner. In this vein, our approach first attempts to introduce the idea of *context*.

To bring the notion of *context* of the passages into the learning process, we first group the passages in our dataset based on some shared context, thereby pivoting the learning process around it. This provides a stronger background for the model's representation learning capabilities as the model, at any given point in time, is only being trained on passages from the same context. From now on, we will refer to this as the *contextual pivot point*. Subsequently, once the passages are grouped based on a contextual pivot point, all possible *combinations* of passages within the groups are extracted and used as training data for our pair-wise ranking model.

The overall pipeline of our approach is seen in Figure 1. A pair of input passages are passed through the layers of a headless Transformer network to get a latent vector containing the representation of each passage. In our experiments, BERT [31], GPT2 [32], RoBERTa [33] and ALBERT v2 [34] are used to obtain the feature vectors but any other Transformer model can similarly be used. As seen in Figure 1, the resulting feature vectors of the two passages are concatenated and used as the input to a four-layered multilayer perceptron (context-aggregating network). The multilayer perceptron assesses the relationship between the two inputs and regresses to two values (ranking scores), subsequently used to rank the passages.

Note that both the input passage pairs and the output score pairs are correspondingly-ordered and any change in the ordering of the input passages affects the output ranking scores. However, this does not mean that the results will be inconsistent when the input changes, but that the interpretation of the output ranking scores depends on the order of the input passages (similar to a comparison operator). Trained end to end, the entire model uses a ranking loss function (explained in Section III-B) and pair-wise ranking is enabled via the comparison of the output values.

B. Loss Function

While most ranking approaches traditionally utilise loss functions such as the sigmoid cross-entropy for binary relevance labels, pair-wise logistic loss or softmax cross-entropy [35], we make use of a margin-based ranking loss, which has been effectively used for representation learning in knowledge graph embedding models [36] by separating the positive samples from the negative samples within the dataset by a given margin. Using this loss function, we can take advantage of its strong representation learning capabilities to extract more robust features from the passages and better learn their compositional relationship during ranking. Formally, for any passage pair p_i and p_j , the ranking label, E , is determined as:

$$E(p_i, p_j) = \begin{cases} 1, & p_i \text{ ranked higher than } p_j \\ -1, & p_j \text{ ranked higher than } p_i \end{cases}, \quad (1)$$

Classification Model	(a) Stack Exchange			(b) Fine Foods		
	Accuracy	F ₁ Score	AUC	Accuracy	F ₁ Score	AUC
ALBERT [34]	0.272	0.278	0.618	0.722	0.617	0.810
RoBERTa [33]	0.281	0.291	0.635	0.705	0.590	0.809
GPT2 [32]	0.285	0.298	0.622	0.765	0.738	0.813
BERT [31]	0.292	0.309	0.648	0.768	0.721	0.828

TABLE I: (a) Quality assessment of Stack Exchange posts and (b) sentiment analysis of Fine Food Reviews.

where E is the ground truth ranking label. Consequently, for the set of all passage pairs $\Psi = \{(p_i, p_j); E(p_i, p_j)\}$, a pair is fed into our model, which approximates the desired ranking function f . The loss function is thus defined as:

$$\mathcal{L} = \sum_{(p_i, p_j) \in \Psi} \max(0, -E(p_i, p_j) \times (f(p_i) - f(p_j)) + \gamma), \quad (2)$$

where $f(p_i)$ and $f(p_j)$ are the ranking scores produced by the overall model, f , and γ is the margin enforced between $f(p_i)$ and $f(p_j)$. In our experiments, $\gamma = 2$ empirically yields the most favourable results. This loss measures ranking violations of passage pairs, allowing the network to learn discriminative features to enforce a clearer distinction between the passages.

C. Implementation Details

For all experiments, the smallest pre-trained versions of the Transformers provided by the HuggingFace library [37] are used. Text sequences are truncated if they exceed 128 tokens, except for experiments in Section IV-D, where sequences of 512 tokens are used. The multilayer perceptron uses Linear-BatchNorm-PReLU modules with a dropout of 0.2 for each layer during training. All implementations are done in PyTorch [38], with AdamW [39] providing the optimization ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-8$).

An advantage of our approach is that there are very few hyperparameters associated with it. The only hyperparameters that might have an effect on the overall results of the approach in the experiments are the margin value in the loss function (γ) and the optimisation parameters (*e.g.* learning rate). These hyperparameters are empirically selected through a grid search. For our hyperparameter searches, 10% of our Stack Exchange dataset (explained in Section IV-A) is randomly selected as the training dataset and 2% as the test set. The margin value γ is selected from values of $\{0, 0.1, 1, 2, 5, 10\}$ and the learning rate α is selected from $\{1e-6, 4e-6, 8e-6, 1e-5, 4e-5, 8e-5, 1e-4, 4e-4, 8e-4\}$. Our experiments indicate that the best results are obtained when the margin value is $\gamma = 2$ for all models and the learning rate is $\alpha = 4e-5$ for BERT and GPT2 and $\alpha = 4e-6$ for ALBERT and RoBERTa.

All the Transformer networks used in our work are trained from a fixed pre-trained state and the only portion of our overall models that actually starts from random initialisation is the four-layer context aggregating fully-connected network. This greatly reduces the effect of random initialisation on our

approach as the size of the context aggregating multilayer perceptron pales in comparison to the size of the Transformer networks.

All experiments are carried out using two NVIDIA Titan RTX GPUs in parallel with a combined memory of 48 GB on an Arch Linux system with a 3.30GHz 10-core Intel CPU and 64 GB of memory. Note that for large NLP models, such as those used in this work, this hardware has limited capabilities, which is why smaller Transformer models and datasets are used in the approach.

IV. EXPERIMENTAL RESULTS

To evaluate our approach, we perform extensive experiments using four publicly-available datasets, Stack Exchange [20], Fine Food Reviews [11], tweets about self-driving cars [21] and the MS MARCO passage ranking dataset [40]. The datasets are selected for their level of difficulty and prevalence in the sentiment analysis and text classification literature.

Note that the tweets about self-driving cars are downloaded from an anonymised publicly available dataset [21] and not extracted by the authors of this paper. In addition to the experiments conducted to assess the efficacy of our technique, we also perform ablation studies and demonstrate the importance of every component of the proposed approach.

A. Stack Exchange

A large Question and Answer Forum made up of numerous communities focusing on different subjects, Stack Exchange [20] offers a space where users can post questions about specific topics for which they can receive answers from other users. Any user can annotate whether an answer is useful or not by voting for it favourably (up-vote) or unfavourably (down-vote). The original asker can also mark one answer as the best. As these votes are expected to be based on quality, subjective though they might be, they provide an opportunity for a learning-based approach to assess the quality of each answer for a potential recommendation system.

A naïve way to assess post quality (in this work only the *answers* are considered for quality assessment in the training and testing datasets) is using a classification-based solution that predicts the quality of posts using classes corresponding to how well the post has done (or can do) in terms of up-votes. The ground truth quality scores can be generated based on the number of votes an answer has received normalised by the number of votes the original question has received.

	Approach	Ranking Metrics @ k			Pair Labels
		MRR	NDCG	MAP	Accuracy
(a) Stack Exchange	w/ ALBERT	0.656	0.772	0.662	0.854
	w/ RoBERTa	0.671	0.780	0.678	0.882
	w/ GPT2	0.760	0.806	0.764	0.895
	w/ BERT	0.762	0.814	0.768	0.910
(b) Review (User)	w/ ALBERT	0.970	0.982	0.976	0.985
	w/ RoBERTa	0.972	0.984	0.979	0.990
	w/ GPT2	0.975	0.986	0.982	0.991
	w/ BERT	0.977	0.986	0.984	0.990
(c) Review (Product)	w/ ALBERT	0.972	0.984	0.976	0.981
	w/ RoBERTa	0.970	0.982	0.972	0.985
	w/ GPT2	0.976	0.987	0.982	0.990
	w/ BERT	0.974	0.985	0.981	0.990

TABLE II: Results of passage ranking using the proposed approach with data from (a) Stack Exchange, (b) Food Reviews with users as the contextual pivot point and (c) Food Reviews with products as the pivot point.

This normalisation is important since questions with higher vote counts generally receive more engagement from the users. As such, the normalisation process removes the possibility of bad answers to more popular questions with a higher vote count overwhelming better answers to less popular questions. The quality scores are then categorised into 5 classes via a simple histogram.

The dataset is extracted from the communities of *Ask Ubuntu*, *Cryptography*, *Data Science*, *Network Engineering*, *Unix & Linux* and *Webmasters*. 250,000 posts are randomly selected for the training dataset and 50,000 for testing. State-of-the-art text classification models [31], [32], [33], [34] are trained for 3 epochs (to convergence) to classify the posts based on their quality scores.

As seen in Table I (a), all classification models fail to assess the quality of the posts beyond randomly guessing, evidenced by the low accuracy, F₁ Score and Area Under the ROC Curve (AUC). This is explained by the lack of context providing a coherent backdrop for the passages as they are associated with different users, questions and communities. Note that while in this paper, we focus on classification as a commonly used baseline, experiments with a regression-based models, predicting the quality score of the passages directly, demonstrated similarly bad performance, explained by the extreme sensitivity of regression models to data imbalances.

Using *Example 2* in Section III-A, we explained how even a human would have trouble assessing the text quality of unrelated passages without any shared context. The same concept applies to this dataset, where even a human observer would have difficulty comparing the quality of answers written by different people to different questions. This is why a contextual pivot point is necessary.

The Stack Exchange post quality annotations (user votes) are within the context of each question, *i.e.* users vote on post quality based on how well the post addresses the specific

Ranking Model	Pair Labels	
	Accuracy	F ₁ Score
Ranking w/ ALBERT	0.905	0.908
Ranking w/ RoBERTa	0.910	0.917
Ranking w/ GPT2	0.932	0.936
Ranking w/ BERT	0.945	0.950

TABLE III: Ranking results of the proposed approach over 1,633 comparisons of temporally-ordered passages from 50 users with at least 20 answers that have unique scores [20].

question. Therefore, measuring the quality of the posts within each question would be very easy as the desired task directly aligns with the ground truth data. As a result, we experiment with unique users as the contextual pivot point. There are 17,085 unique users within our dataset with an average of 20.1 answers per user.

In our experiments, we assess the quality of passages posted by individual users by ranking them. While evaluating the quality of posts for *each individual user* is possible via text classification, since classification cannot explicitly take context into account, it would entail training a *separate classification model for each individual user* (17,085 different classification models for the 17,085 users), which is intractable. Our ranking approach, however, is trained end to end over the entire dataset once and can provide an accurate measure of the quality for the posts across all questions, all users and all communities.

With the user chosen as the contextual pivot point, the passages are grouped based on unique users and all combinations of passages are extracted as training data ($\approx 8,000,000$ pairs). Similar to the classification setup, ground truth ranking labels (Eqn. 1) are based on the number of votes an answer has received normalised by the number of votes the original question has received. To enable evaluation with the appropriate ranking metrics, 100 random users with at least 100 answers that have unique quality scores are selected as the test set.

Table II (a) demonstrates better results are obtained with BERT and GPT2. Common ranking metrics, including Normalized Discounted Cumulative Gain (NDCG), Mean Average Precision (MAP) and Mean Reciprocal Rank (MRR) are used to evaluate the results with $k = 10$. Note that k denotes how many items are considered relevant in the ranked list - *i.e.* relevance is only considered in the top k items in the list for the ranking metrics.

We also measure the accuracy of the approach applied to test passages (with success referring to an accurate ranking of pairs - Eqn. 1). As seen in Table II (a), despite the fact that the ground truth labels are crowd-sourced and not objective, accurate ranking is achieved.

An interesting application of ranking the passages with the contextual pivot point being the users is the possibility of tracking the change in the skill level of users over time. To test this, we extract a test set of 50 random users with at least 20 answers that have unique ranking scores. Considering that

Fig. 2: Test accuracy of pair rankings on the Fine Foods reviews dataset using our model trained for 40,000 steps. When ranking is pivoted around users, the models reach convergence faster, due to the overpowering effect of user writing styles.

the model receives no information about the time of the posts, if the resulting ranking matches the order at which passages were posted, the approach is capable of tracking the skill level of users, which can be of value in downstream recommendation and forum moderation systems, with opportunities for recommending questions to users and ordering posts based on a user’s skill level, among others.

A numerical analysis of this is shown in Table III, which presents the results after evaluating the ranking of posts from 50 users. Since for this experiment, the accuracy of the comparison of the pairs is what matters, we present the accuracy metric over the 1,633 comparisons of temporally-ordered passage pairs. The promising results in Table III are indicative of the potential for better user behaviour and profile tracking, leading to more accurate recommendation and security systems.

We would like to also point out that while we demonstrate how classification is not a suitable solution for this problem (Table I (a)), a regression-based solution also suffers from the same issues. The quality scores for the Stack Overflow posts are segmented into 5 classes via a simple histogram. This leads to an unbalanced dataset as many of the posts are of roughly the same quality with a number of outliers with significantly different scores. These issues will remain the

same in a regression setup with most posts having the same score and the outliers skewing the possible range of scores the model is meant to learn, leading to potentially worse results than classification in general. This is why the literature on quality assessment/sentiment analysis predominantly focuses on classification solutions [41], [42], [43], [44].

Due to the fact that the majority of the posts are clumped together in terms of their score and the existence of outliers, a regression model completely fails to capture the underlying data distribution and will perform no better than classification. Our ranking model, however, does not suffer from issues stemming from data imbalance, better captures the patterns and the associations in the posts and their quality and produces promising results, as seen in Tables II (a) and III.

B. Fine Food Reviews

The Amazon Fine Food Reviews dataset [11] contains approximately 500,000 reviews of around 75,000 products from roughly 250,000 users. The objective is to predict the sentiment of the review based on scores (from 1 to 5) provided in the dataset. As opposed to the Stack Exchange data, the passages in this dataset are significantly easier to classify as they substantially follow the same context (product reviews).

Experimental results presented in Table I (b) demonstrate that classification methods are reasonably capable of classifying the reviews albeit with underwhelming results. 20% of the data is randomly selected as the test set and the text sequences are truncated to 128 tokens (words). The dataset is unbalanced with higher sentiment scores being more prevalent in the dataset. However, if the task of sentiment analysis is reformulated as ranking, significantly better performance can be expected.

For the purposes of our approach, the contextual pivot point for this dataset can be either the *user* (ranking the sentiment of the reviews from individual users for all the products they have reviewed) or the *product* (ranking the sentiment of the reviews of individual products from all users who have reviewed them). Here, we perform experiments based on both the users and the products as contextual pivot points.

Separate training and testing sets are created, with the product-based test set made up of 200 random products with reviews that have unique scores (only 5 unique scores exist) and the user-based test set made up of 200 random users with reviews that have unique scores. To evaluate ranking, the test reviews for each user or product are selected not to have the same score so the ranking of the ground truth list for each user/product can be easily compared against the ranking output of the approach. As our model always ranks a pair of passages, even if they are similar, only pairs with different ranking scores are passed into the model during training and in the test set, only passages with unique scores are used to evaluate the model.

Table II (b) shows the results of the user-based experiments and Table II (c) the results of the product-based experiments. As there are only 5 items for each user or product (5 sentiment classes), the ranking metrics are calculated with $k = 2$. *Accuracy* alludes to the correctness of pair rankings over all possible pairs in the test set. The metric values are higher than those of Section IV-A, as only five items exist in the list, but as seen in Table II, the results are extremely promising with the models achieving near perfect ranking in both experiments.

While highly accurate results are achieved across the board, an interesting observation is that it is easier for the model to learn the context and thus perform better ranking when the pairs are pivoted around users as opposed to products. Figure 2 demonstrates how all models consistently reach convergence faster when the contextual pivot point is the user. This is due to the powerful influence of the writing style of users with different reviews. While the reviews of all products share a clear context as evidenced by the results in Table II (c), user-based ranking is more easily learned (Figure 2).

C. Sentiment Analysis: Self-Driving Cars

This dataset [21] focuses on people’s opinions of autonomous driving and consists of 6,943 relevant tweets labelled from 1 to 5, with 1 indicating the most negative sentiment and 5 the most positive. This dataset is heavily skewed with over 61% of the data annotated as *neutral* (4,245 tweets with label 3) and less than 2% annotated as *very*

Classification Model	Evaluation Metrics		
	Accuracy	F ₁ Score	AUC
Liu et al. [33]	0.668	0.598	0.642
Lan et al. [34]	0.672	0.602	0.626
Radford et al. [32]	0.684	0.608	0.638
Devlin et al. [31]	0.688	0.620	0.634
Yang et al. [45]	0.692	0.626	0.642
Converted Ranking w/ ALBERT	0.872	–	–
Converted Ranking w/ RoBERTa	0.902	–	–
Converted Ranking w/ GPT2	0.898	–	–
Converted Ranking w/ BERT	0.916	–	–

TABLE IV: Accuracy results of state-of-the-art classification compared against our ranking approach converted to class labels on the Self-Driving Car Tweets dataset [21].

negative (110 tweets with label 1). This creates significant challenges for any classification approach and requires extreme measures to combat the imbalance in the dataset. Note that this datasets is not directly obtained from Twitter by the authors of this work and is instead downloaded as an anonymised publicly-available dataset [21].

Here, we evaluate the ability of our model to deal with such a challenging dataset. While all the tweets follow the same context (self-driving vehicles), the user that posted each tweet is ignored in this dataset and thus no contextual pivot point is available other than the shared subject of the tweets (self-driving vehicles). The dataset is split into an unbalanced training set and a balanced test set with the test data containing 100 tweets from each class (500 tweets in total). This leaves only 10 tweets from class 1 for training, further exacerbating the data imbalance problem.

Note that the only reason that the test set for this experiment is balanced is to enable a better analysis of the results of our comparators. Our ranking approach is not affected by any imbalances in the training and/or test datasets and is able to remain completely unchanged even if the test set was extremely unbalanced.

This is intuitively expected as when one ranks a series of data points based on some score, the distribution of the data and the score does not matter as the ranking technique only needs to predict an ordering of the data and not their class/value or score – each pair for comparison are handled independently and there is no longer a concept of balance to the data. For a classification or regression approach, however, distribution is of significant importance. We have opted for a balanced test set only to make the evaluation of the classification techniques easier.

As our approach is a ranking one, it is capable of producing an ordered list and does not predict absolute class labels. To enable a direct comparison with classification methods, here, we convert the resulting ranked list to a set of class labels, which can be trivial as the class labels represent sentiment scores from 5 to 1 (most positive to most negative) and our ranking approach is trained to rank based on sentiment scores

Ranking Approach	Dev
	MRR@10
Mitra and Craswell [46]	0.252
Mitra et al. [47]	0.254
Rosset et al. [48]	0.262
Nogueira et al. [49]	0.277
Qiao et al. [50]	0.311
Hofstatter et al.[51]	0.318
Our Approach w/ ALBERT	0.264
Our Approach w/ RoBERTa	0.267
Our Approach w/ GPT2	0.273
Our Approach w/ BERT	0.275

TABLE V: Comparison of contemporary ranking methods and our approach applied to MS MARCO. It is important to note that unlike many comparators, the proposed approach uses the smallest most basic version of the Transformers.

(most positive to most negative). This is accomplished by simply sorting the balanced test set of 500 tweets from most positive to most negative using our ranking approach, dividing the ranked list into 5 segments and assigning labels (5 to 1) to each segment from top to bottom.

It is important to note that the balance or imbalance of the test set has no bearing on the performance of our approach. For example, if only 5% of the test data belonged to class 5 (most positive) instead of the current 20%, it would make no difference to our approach as those would still be ranked at the top of the sorted list. Ranking does not suffer from such imbalance issues, which is why we argue it is more suited for solving such problems than classification solutions.

Metrics such as F_1 and AUC do not apply to a ranking approach even when the results are converted to class labels. However, accuracy provides an excellent metric for evaluating the performance of the state-of-the-art classification methods as well as our ranking technique. So while we report accuracy, F_1 and AUC for classification methods, our converted ranking is evaluated using accuracy as the primary metric (Table IV).

We would like to emphasise that this conversion of the ranking results to class labels is *not* necessary in the real world and is only done to evaluate our approach in direct comparison to classification. We argue that in certain problems, such as the one addressed in this section, ranking is a better solution and assigning an absolute class or value to a passage does not offer as much information as an ordered list of passages with a shared context provides.

In this experiment, all models are trained for 100,000 steps. Table IV shows how effective our approach is by achieving promising results despite the significant imbalance in the training dataset. Based on the F_1 and AUC scores of the state-of-the-art text classification methods [33], [34], [32], [31], [45], it is clear that the skew in the training dataset has severely affected the learning capabilities of the state-of-the-art classifiers.

On the other hand, ranking remains robust. In fact, the

converted results of our best ranking approach improve on the most effective text classifier [45] by about 22%.

D. MS MARCO

While the primary focus of this paper is not information retrieval, ranking is indelibly linked to information retrieval so much so that the metrics used to evaluate our results are predominantly from the information retrieval literature. In this vein, we also apply our approach to the publicly available benchmark dataset of MS MARCO [40].

We perform the passage ranking task within the benchmark using our proposed pipeline. The queries and the passages are concatenated into sequences of no more 512 tokens. Due to hardware restrictions, less than 1% of the entire available dataset is used for training and the smallest possible versions of the Transformer models are used.

Table V demonstrates the results of our approach compared to contemporary ranking approaches on the validation set of the MS MARCO dataset. Despite using smaller Transformer models and a fraction of the available training data, our proposed approach produces promising results and remains competitive with dedicated information retrieval techniques.

E. Ablation Studies

It is important to evaluate and demonstrate the contribution of every component of our approach. Consequently, we re-train our model with different components removed or replaced, with pair ranking accuracy being measured as the primary metric. For these experiments, 30% of the Stack Exchange dataset is selected for our ablation studies.

As an integral part of our pipeline is the Transformer model generating the feature vector representing the input sequences, the choice of the Transformer is a critical decision. We perform all our experiments with four commonly-used state-of-the-art Transformers. Table VI demonstrates that BERT produces superior results with GPT2 only remaining competitive despite its larger size. This is also supported by our other experiments with results presented in Tables II, III, IV and V.

Another important component of the proposed approach is the context aggregating multilayer perceptron (MLP) that receives the sequence representations of the two input passages and generates the scores, subsequently used to rank the passages. To evaluate its overall influence, the model is re-trained with a single linear layer replacing the entire multilayer perceptron that projects the sequence representation into a single scalar value that is passed into the loss function as the ranking score.

As seen in Table VI, while the approach still learns to rank the passages reasonably well without the context aggregating multilayer perceptron, the pair ranking accuracy significantly drops. This emphasises the importance of the MLP component to the overall predictive performance of our approach.

Within our proposed pipeline, we opt for the use of a single Transformer model to produce the latent vector representing both passages in the input pair. However, one could envisage using two separate Transformer models, each independently

Approach	Pair Labels
	Accuracy
Full Approach w/ ALBERT	0.861
Full Approach w/ RoBERTa	0.882
Full Approach w/ GPT2	0.899
Separate Transformers (BERT)	0.890
Approach w/o MLP (BERT)	0.815
Full Approach w/ BERT	0.902

TABLE VI: Pair-wise ranking accuracy results with different components of the proposed ranking approach removed or replaced, applied to 30% of the Stack Exchange dataset [20].

learning the representation of one of the passages. While this will almost double the number of learnable parameters and can potentially introduce training instabilities, it is possible that the higher parameter count will enhance the learning process.

However, as seen in Table VI, the pair ranking accuracy is reduced when separate Transformer networks not sharing weights are used for the passages. This is primarily due to the reduction in the number of training samples for each Transformer and the possible overfitting of each network. Additionally, by using the same network for all passages, the model will get a better sense of the entire dataset and can produce more robust representations.

V. LIMITATIONS AND FUTURE WORK

While ranking does not suffer from common issues that classification approaches suffer from, such as data imbalances, text subjectivity and lack of context, it does not suit every problem. Many text classification applications could be replaced by ranking with superior results, but there are certain situations where text classification remains the only option. For instance, a scenario wherein passages are meant to be categorised based on their topic cannot be addressed with ranking and requires a classification solution.

Moreover, text ranking cannot deal in absolutes and is only capable of relatively comparing items. For example, while ranking can provide an answer to whether a movie review is more or less positive than another, it cannot definitively say whether both the reviews are positive or negative. This issue could simply be addressed in real-world applications by identifying “sentinel” examples which are agreed, by consensus, to be positive / neutral / negative. Then using our ranking approach could indicate where an example is placed by comparing with the sentinels – for example if it is ranked above a positive sentinel then it is a positive example. In this way, during inference, other passages can then be assigned an absolute labels/scores relative to their position in the sorted list produced by our ranking model based on the previously known sentinel passages.

It is important to note, however, that the inability to accurately provide absolute sentiments is seen in sentiment

classification methods as well. Many classification solutions do not generalise to unseen data well and are incapable of absolute sentiment prediction if the vocabulary and linguistic structure of a new sentence has not been seen by the model within the training data.

Additionally, for our approach, we intend to address this in a future work with a double-headed model that simultaneously ranks and classifies the input passages. As the two heads can correct each other during training, better representation learning and thus more accurate results can be expected, as well as the ability to provide an absolute relevance and sentiment value for each passage.

Another avenue for future work would be to modify the loss function to cope with lists rather than pairs of passages. Whilst more difficult to optimise, list-wise ranking generally offers a better understanding of the context of the dataset, leading to more accurate results in addition to fewer post-processing requirements and thus better efficiency during inference.

VI. CONCLUSION

We have investigated the applications of ranking in natural language processing using a novel pair-wise ranking approach and different publicly-available datasets. Our text ranking approach makes use of state-of-the-art Transformer networks to generate a learned representation of a pair of text sequences. These representation vectors are subsequently used as the input to a context-aggregating multilayer perceptron, which combines the features representing the context and the content of the passages, assesses the relationship between the two input passages and regresses to two values denoting the ranking scores subsequently used to rank the passages simply based on the ordering of the predicted scores.

The entirety of the model is trained end to end using a margin-based ranking loss function. Experiments are carried out on four publicly available datasets, with all the evaluation demonstrating the effectiveness of ranking for potential recommendation, forum moderation and security applications. We also compare the results of our approach directly with state-of-the-art classification methods. A comparison of our ranking results converted to class labels with state-of-the-art classification results indicates an approximately 22% improvement, pointing to the efficacy of ranking over classification.

We kindly invite the readers to refer to the supplemental video: <https://youtu.be/Njjg3mWVE2g> for more information.

REFERENCES

- [1] K. Kowsari, K. Jafari Meimandi, M. Heidarysafa, S. Mendu, L. Barnes, and D. Brown, “Text classification algorithms: A survey,” *Information*, vol. 10, no. 4, p. 150, 2019.
- [2] S. Minaee, N. Kalchbrenner, E. Cambria, N. Nikzad, M. Chenaghlu, and J. Gao, “Deep learning based text classification: A comprehensive review,” *arXiv preprint arXiv:2004.03705*, 2020.
- [3] J. Salminen, V. Yoganathan, J. Corporan, B. J. Jansen, and S.-G. Jung, “Machine learning approach to auto-tagging online content for content marketing efficiency: A comparative analysis between methods and content type,” *Journal of Business Research*, vol. 101, pp. 203–217, 2019.
- [4] L. Zhang, S. Wang, and B. Liu, “Deep learning for sentiment analysis: A survey,” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 8, no. 4, p. e1253, 2018.

- [5] M. Iyyer, V. Manjunatha, J. Boyd-Graber, and H. Daumé III, "Deep unordered composition rivals syntactic methods for text classification," in *Int. Joint Conf. Natural Language Processing*, 2015, pp. 1681–1691.
- [6] S. Wan, Y. Lan, J. Guo, J. Xu, L. Pang, and X. Cheng, "A deep architecture for semantic matching with multiple positional sentence representations," in *AAAI Conf. Artificial Intelligence*, 2016.
- [7] J. Wang, Z. Wang, D. Zhang, and J. Yan, "Combining knowledge with deep convolutional neural networks for short text classification," in *Int. Joint Conf. Artificial Intelligence*, 2017, pp. 2915–2921.
- [8] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [9] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, "Learning word vectors for sentiment analysis," in *Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 2011, pp. 142–150.
- [10] B. Pang and L. Lee, "Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales," in *Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 2005, pp. 115–124.
- [11] J. J. McAuley and J. Leskovec, "From amateurs to connoisseurs: Modeling user expertise through online reviews," in *Int. Conf. World Wide Web*, 2013, pp. 897–908.
- [12] J. Finn, "Twitter users overrepresent extreme and fringe political ideas, not the mainstream," <https://screenrant.com/twitter-extreme-view-problem/>, 2020.
- [13] F. Rodrigues, F. Pereira, and B. Ribeiro, "Sequence labeling with multiple annotators," *Machine learning*, vol. 95, no. 2, pp. 165–181, 2014.
- [14] H. Schütze, C. D. Manning, and P. Raghavan, *Introduction to Information Retrieval*. Cambridge: Cambridge University Press Cambridge, 2008, vol. 39.
- [15] H. Li, "Learning to rank for information retrieval and natural language processing," *Synthesis Lectures on Human Language Technologies*, vol. 4, no. 1, pp. 1–113, 2011.
- [16] M. Zulqarnain, R. Ghazali, Y. M. M. Hassim, and M. Rehan, "A comparative review on deep learning models for text classification," *Indones. J. Electr. Eng. Comput. Sci.*, vol. 19, no. 1, pp. 325–335, 2020.
- [17] S. Minaee, N. Kalchbrenner, E. Cambria, N. Nikzad, M. Chenaghlu, and J. Gao, "Deep learning-based text classification: A comprehensive review," *ACM Computing Surveys*, vol. 54, no. 3, pp. 1–40, 2021.
- [18] A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov, "Bag of tricks for efficient text classification," *arXiv preprint arXiv:1607.01759*, 2016.
- [19] K. Mouthami, K. N. Devi, and V. M. Bhaskaran, "Sentiment analysis and classification based on textual reviews," in *Int. Conf. Information Communication and Embedded Systems*. IEEE, 2013, pp. 271–276.
- [20] I. Stack Exchange, *Stack Exchange Data Dump*, 2019, <https://archive.org/details/stackexchange>.
- [21] Twitter, *Twitter Sentiment Analysis: Self-Driving Cars*, 2017, <https://data.world/crowdfunder/sentiment-self-driving-cars>.
- [22] W. W. Leontief, "The structure of american economy, 1919-1929," Harvard University Press, Tech. Rep., 1941.
- [23] L. Page, S. Brin, R. Motwani, and T. Winograd, "The PageRank citation ranking: Bringing order to the web," Stanford InfoLab, Tech. Rep., 1999.
- [24] K. S. Jones, S. Walker, and S. E. Robertson, "A probabilistic model of information retrieval: Development and comparative experiments: Part 2," *Information Processing & Management*, vol. 36, no. 6, pp. 809–840, 2000.
- [25] R. Baeza-Yates and B. Ribeiro-Neto, *Modern Information Retrieval*. ACM Press New York, 1999, vol. 463.
- [26] F. C. Gey, "Inferring probability of relevance using the method of logistic regression," in *Int. Conf. Research and Development in Information Retrieval*, 1994, pp. 222–231.
- [27] C. J. Burges, "From RankNet to LambdaRank to LambdaMART: An overview," *Learning*, vol. 11, no. 23-581, p. 81, 2010.
- [28] X. Wang, C. Li, N. Golbandi, M. Bendersky, and M. Najork, "The LambdaLoss framework for ranking metric optimization," in *ACM Int. Conf. Information and Knowledge Management*, 2018, pp. 1313–1322.
- [29] Y. Lan, Y. Zhu, J. Guo, S. Niu, and X. Cheng, "Position-aware ListMLE: A sequential learning process for ranking," in *Conf. Uncertainty in Artificial Intelligence*, 2014, pp. 449–458.
- [30] M. Taylor, J. Guiver, S. Robertson, and T. Minka, "SoftRank: Optimizing non-smooth rank metrics," in *Int. Conf. Web Search and Data Mining*, 2008, pp. 77–86.
- [31] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Conf. North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, pp. 4171–4186.
- [32] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," *OpenAI Blog*, vol. 1, no. 8, p. 9, 2019.
- [33] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A robustly optimized BERT pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [34] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A lite BERT for self-supervised learning of language representations," *arXiv preprint arXiv:1909.11942*, 2019.
- [35] R. K. Pasumarthi, S. Bruch, X. Wang, C. Li, M. Bendersky, M. Najork, J. Pfeifer, N. Golbandi, R. Anil, and S. Wolf, "TF-Ranking: Scalable TensorFlow library for Learning-to-Rank," in *Int. Conf. Knowledge Discovery & Data Mining*, 2019, pp. 2970–2978.
- [36] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard, "Complex embeddings for simple link prediction," in *Int. Conf. Machine Learning*, 2016.
- [37] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Transformers: State-of-the-art natural language processing," *arXiv preprint arXiv:1910.03771*, 2019.
- [38] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, 2019, pp. 8026–8037.
- [39] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [40] P. Bajaj, D. Campos, N. Craswell, L. Deng, J. Gao, X. Liu, R. Majumder, A. McNamara, B. Mitra, T. Nguyen *et al.*, "MS MARCO: A Human Generated Machine Reading Comprehension dataset," *arXiv preprint arXiv:1611.09268*, 2016.
- [41] O. Ormandjieva, I. Hussain, and L. Kosseim, "Toward a text classification system for the quality assessment of software requirements written in natural language," in *Int. Workshop on Software Quality Assurance: in Conjunction with ESEC/FSE Joint Meeting*, 2007, pp. 39–45.
- [42] D. Ganguly, J. Leveling, and G. Jones, "Automatic prediction of aesthetics and interestingness of text passages," in *Int. Conf. Computational Linguistics: Technical Papers*. Dublin, Ireland: Association for Computational Linguistics, 2014, pp. 905–916.
- [43] T. U. Haque, N. N. Saber, and F. M. Shah, "Sentiment analysis on large scale amazon product reviews," in *IEEE Int. Conf. Innovative Research and Development*. IEEE, 2018, pp. 1–6.
- [44] R. Johnson and T. Zhang, "Deep pyramid convolutional neural networks for text categorization," in *Annual Meeting of the Association for Computational Linguistics*, 2017, pp. 562–570.
- [45] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "XLnet: Generalized autoregressive pretraining for language understanding," in *Advances in Neural Information Processing Systems*, 2019, pp. 5754–5764.
- [46] B. Mitra and N. Craswell, "An updated duet model for passage re-ranking," *arXiv preprint arXiv:1903.07666*, 2019.
- [47] B. Mitra, C. Rosset, D. Hawking, N. Craswell, F. Diaz, and E. Yilmaz, "Incorporating query term independence assumption for efficient retrieval and ranking using deep neural networks," *arXiv preprint arXiv:1907.03693*, 2019.
- [48] C. Rosset, B. Mitra, C. Xiong, N. Craswell, X. Song, and S. Tiwary, "An axiomatic approach to regularizing neural ranking models," in *Int. Conf. Research and Development in Information Retrieval*, 2019, pp. 981–984.
- [49] R. Nogueira, J. Lin, and A. Epistemic, "From doc2query to docTTTT-Tquery," *Online preprint*, 2019.
- [50] Y. Qiao, C. Xiong, Z. Liu, and Z. Liu, "Understanding the behaviors of BERT in ranking," *arXiv preprint arXiv:1904.07531*, 2019.
- [51] S. Hofstätter, N. Rekabsaz, C. Eickhoff, and A. Hanbury, "On the effect of low-frequency terms on neural-IR models," in *Int. Conf. Research and Development in Information Retrieval*, 2019, pp. 1137–1140.