

Comment Ranking Diversification in Forum Discussions

Curtis G. Northcutt, Kimberly A. Leon, Naichun Chen
 Massachusetts Institute of Technology
 Cambridge, MA, USA
 {cgn, kimleon, naichun}@mit.edu

ABSTRACT

Viewing consumption of discussion forums with hundreds or more comments depends on ranking because most users only view top-ranked comments. When comments are ranked by an ordered score (e.g. number of replies or up-votes) without adjusting for semantic similarity of near-ranked comments, top-ranked comments are more likely to emphasize the majority opinion and incur redundancy. In this paper, we propose a top K comment diversification re-ranking model using Maximal Marginal Relevance (MMR) and evaluate its impact in three categories: (1) semantic diversity, (2) inclusion of the semantics of lower-ranked comments, and (3) redundancy, within the context of a HarvardX course discussion forum. We conducted a double-blind, small-scale evaluation experiment requiring subjects to select between the top 5 comments of a diversified ranking and a baseline ranking ordered by score. For three subjects, across 100 trials, subjects selected the diversified (75% score, 25% diversification) ranking as significantly (1) more diverse, (2) more inclusive, and (3) less redundant. Within each category, inter-rater reliability showed moderate consistency, with typical Cohen-Kappa scores near 0.2. Our findings suggest that our model improves (1) diversification, (2) inclusion, and (3) redundancy, among top K ranked comments in discussion forums of online courses.

INTRODUCTION

Text ranking systems (e.g. Facebook post comments, Amazon product reviews, Reddit forums) are ubiquitous, yet many face a common problem. When posts (e.g. reviews or comments) are ranked primarily by text content and rating (e.g. like/unlike, \uparrow/\downarrow , $+/-$, number of replies, etc.), similar posts tend to receive similar scores. Moreover, higher ranking posts tend to exclusively represent the majority opinion, since there are more users in the majority group to up-vote posts sharing their sentiment. For large forums with thousands of posts, viewers may only be exposed to the majority opinion when they only view top-ranked posts. If the ground truth semantics of each comment were known a priori, comment scores could be

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 the author(s) 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.

L@S 2017, April 20 - 21, 2017, Cambridge, MA, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4450-0/17/04...\$15.00

DOI: <http://dx.doi.org/10.1145/3051457.3054016>

normalized by the number of comments with similar semantics, avoiding this problem. Unfortunately this is not the case. Instead, there are a multitude of techniques to approximate semantic similarity [12, 7, 13].

We consider the comment ranking diversity problem in the context of an online edX course, *Harvardx Christianity Through Its Scriptures*, where increased visibility of the diversity of comments across thousands of learners may aid in debunking misconceptions held by the majority of forum respondents. edX forums are organized hierarchically into topics > comments > replies (an example topic is depicted in Figure 1). Our focus is the ranking of comments and we use the number of replies as the score for each comment, although by default, edX comments are ranked chronologically.

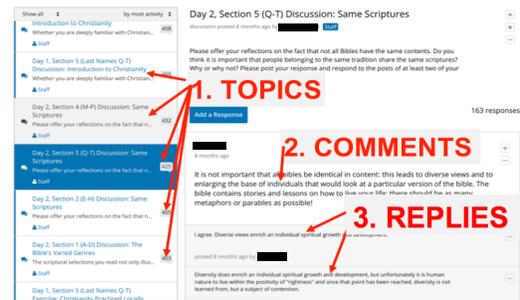


Figure 1. An example topic used to illustrate the organization of an edX discussion forum. edX forums are organized hierarchically into topics > comments > replies. Our focus is the ranking of comments.

In this paper, we develop an algorithm for forum comment ranking diversification using maximal marginal relevance (MMR) to linearly interpolate between the original ranking score (relevance) and the similarity of an item to higher-ranked items (diversity) estimated using PCA + TFIDF, and evaluate our model using a blind experiment requiring subjects to compare our diversified ranking to a baseline relevance ranking.

BACKGROUND AND RELATED WORK

The crux of diversification is a well-trained comment embedding model that accurately captures the semantic similarity between two documents. Text embedding is a well-studied problem at the word-level [12] and document-level [10]. In this section, we consider increasingly complex methods for comment similarity, followed by methods for ranking documents and how it relates to diversification.

One of the simplest document embedding representations is TFIDF [9] which uses a "bag of words" (nBOW) counts model, normalized by word count per document frequency. Although

TFIDF works well on some tasks [1], it ignores word ordering and suffers a performance loss for longer documents. TFIDF performs well when combined with matrix decomposition methods like PCA or LSA. More sophisticated approaches such as word2vec [12], LDA [2], and Gated CNN [11] offer classification accuracy improvements, but are task-specific. These models are compared in Table 1. A state-of-the-art (2016) LSTM similarity model uses a Siamese recurrent architecture to combine the word2vec embeddings of all words in a document, and trains using a Manhattan loss on the output of the two LSTMs [13]. Although this method would likely offer improvements, simpler models were sufficient for our task.

Model	Method	Scaling Sensitivity
TFIDF	F	False
PCA + TFIDF	M+S	True
LSA + TFIDF	M+S	True
NMF + TFIDF	M+S	True
LDA + TFIDF	T	False
Word2Vec + TFIDF	W+S	False
Word2Vec + nBOW	W+S	False
Gated CNN + TFIDF	W+S	False

Table 1. A comparison of the comment embedding models evaluated in this study. Method symbols are abbreviated as: T=Topic, M=Matrix Factorization, W=Local Window, F=Frequency, S=Semantic

The task of forum comment ranking can be thought of as a search task, where common methods like PageRank [14] and RankSVM [6]) are used to identify the most relevant document for a given query. In our case, relevance is determined a priori by comment score, and instead our focus is diversification of this ranking. Diversification has been successfully applied to the task of online shopping [4], with the task of reducing abandonment in shopping queries by providing a diversified selection of options. In this paper, we elect a more general approach, MMR [3], which we describe more in Section 3.2.

TECHNICAL APPROACH

Our methodology consists of four ordered components: (1) Automated generation of gold data, (2) Evaluation of comment embedding models, (3) Implementing diversification in comment rankings, and (4) Measuring efficacy of diversification. We describe these components in the following sections.

Dataset

edX forums are organized hierarchically by topic > comments > replies as shown in Figure 1. We consider diversification at the comments level (within a single topic). In the context of this study, we focus on the comment rankings for topics in the forum discussions of an edX course, *HarvardX: HDS3221.2x Christianity Through Its Scriptures*, obtained via web-scraping. Comment scores were set equal to the number of replies for each comment. Forum text was tokenized with stop-words removed and over 100,000 comments were analyzed.

Automated Gold Data Generation

We used a novel method to generate large gold datasets, without human labeling, by sampling comments across highly differing topics and generating a pairwise cosine similarity matrix for these comments. This matrix contains binary labels, a (1) if comments were taken from the same topic (Gold 1 pairs) or (0) if comments were taken from different topics

(Gold 0 pairs). For exclusive sets of topics, we generated both train and test gold datasets to evaluate our selection of different comment embedding models discussed in 4.1.

Maximal-Marginal Relevance (MMR)

MMR is an iterative algorithm, at each step selecting the comment which maximizes a modified score (Equation 1).

$$\hat{s} := \lambda \cdot s - (1 - \lambda) \cdot c \quad (1)$$

A single parameter λ adjusts the trade-off between the original comment score, s , and its maximum cosine similarity among all comments that have already been added to the new ranking, c , to produce the updated score, s' . For example, $\lambda = 1$ ranks entirely by score and $\lambda = 0$ selects maximally diverse comments irrespective of score. In this study, we evaluate two settings of the parameter, $\lambda = 0.75$ and $\lambda = 0.25$ in comparison with a baseline where $\lambda = 1$.

Comment Embedding Model Selection

Diversification with MMR hinges on a comment embedding model that accurately captures the semantic similarity between two comments. Eight models were evaluated (Table 1).

Two evaluation metrics were used to compare these models. (1) The median quantile difference defined as the difference in average cosine similarity percentile rank (quantile) of Gold 1 pairs minus that of Gold 0 pairs. We recommend this metric as it is unbiased and captures relative ranking. (2) The accuracy of logistic regression using a given model's pairwise comment cosine similarity matrix as input and the gold binary labels as output. Our two metrics consistently ranked all models.

Using the best performing model for these two metrics, comment similarity was computed using cosine similarity [8]. In our case, the best model was PCA + TFIDF comment embeddings, as seen in Table 2 in the Results section.

MMR Evaluation Experiment

Since the comment order for the course we experimented on is chronological, we used ordering by score (number of replies, $\lambda = 1$) for our baseline ranking. We conducted a small-scale re-ranking evaluation experiment requiring subjects to choose among two unidentified ordered lists of comments: (1) the top 5 comments of our diversified ranking and (2) the top 5 comments of a baseline ranking ordered only by score, their true identities unknown. Three subjects evaluated 100 trials. The Cohen-Kappa score [5] was used to measure inter-rater reliability. For each trial, subjects were presented with three items (an example trial is shown in Figure 2):

1. The forum's topic question
2. Two lists, A and B. One of these lists is the top five comments ordered by score (baseline). The other is the top five diversified (re-ranked) comments
3. A random comment C from this forum not included in (2) where C's probability of being chosen was proportional to number of replies (higher rank = more likely to be chosen).

Both the order in which lists A and B were shown to subjects and trial order were randomized to ensure the true labels for

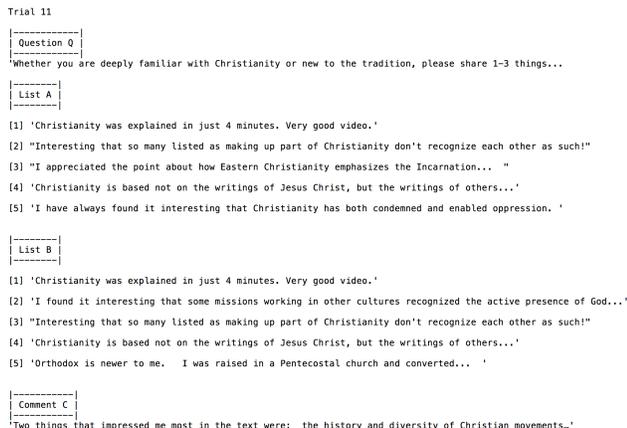


Figure 2. A trial in the MMR evaluation experiment.

list A and B were unrecoverable within and across subjects. For each double-blind trial, each subject answered 3 questions:

- Inclusion Experiment:** Which list, A or B, has a comment that resembles the semantics of comment C?
- Diversity Experiment:** Which list, A or B, best captures a diverse set of all potential answers to this question Q?
- Redundancy Experiment:** Which list, A or B, contains more redundant comments?

If our comment embedding model accurately captures pairwise semantic similarity, we would expect the diversified ranking to be chosen more often for "inclusion" and "diversity", and less often for "redundancy".

Among the 100 trials for each subject, 75 trials used $\lambda = 0.25$ (ranked more by diversity) and 25 trials used $\lambda = 0.75$ (ranked more by score). More trials were taken for $\lambda = 0.25$ to offset increased stochasticity when selecting low-scored (but diverse) comments. Neglecting comment score increases variation in ranking. Additional trials mitigated increased variance.

RESULTS AND DISCUSSION

This section is divided into two parts. Since diversification relies on accurate semantic similarity scores, in Section 4.1 we evaluate comment embedding models on our gold dataset. Then, in Section 4.2, we evaluate our model in a double-blind subject experiment comparing our diversified ranking against a baseline ranking ordered by score.

Model	Median Quantile Difference	Logistic Regression Accuracy
TFIDF	0.338009	0.841026
PCA + TFIDF	0.433901	0.866738
LSA + TFIDF	0.431401	0.867211
NMF + TFIDF	0.416625	0.860729
LDA + TFIDF	0.128716	0.815278
Word2Vec + TFIDF	0.204944	0.815278
Word2Vec + nBOW	0.166570	0.815278
Gated CNN + TFIDF	0.115584	0.786467

Table 2. Comparison of various comment embedding methods. Median quantile difference computes the difference in average cosine similarity rank (percentile) of Gold 1 pairs - Gold 0 pairs. Logistic regression predicts the accuracy of the gold labels trained using each model's pairwise cosine similarity matrix as input.

Comment Embedding Models

For our task, word-level comment embedding methods (word2vec, Gated CNN, LDA) performed worse than a simple TFIDF vector representation alone, with a classical application of dimensionality reduction using PCA achieving highest accuracy on our gold dataset. Table 2 captures the performance of different embedding models on our gold test set, for both median quantile difference and logistic regression accuracy. In the rest of this section, we discuss potential reasons for this.

Comparing the use of the TFIDF embedding to the use of PCA and LSA affirms that there is benefit to employing dense embeddings. More unexpectedly, word2vec and Gated CNN, when combined with TFIDF, did not perform as well as TFIDF. A likely suspect is that our word2vec model was trained on the Google News corpus, which is a semantically different and much broader corpus than learner comments in an online course. As a result, word embeddings related to the course content were compressed into a smaller space relative to the broader embeddings of the model.

Given that comments were on average 78 words in length, and "bag of words" ignores ordering and contextual information, it is less surprising that PCA and LSA outperformed nBOW and TFIDF models. As PCA offered a marginal performance improvement over LSA, PCA + TFIDF was chosen as our final comment embedding model.

MMR Evaluation

Table 3 lists the results of the blind evaluation experiment. The fraction of subject responses selecting the diversified (MMR) ranking is depicted in Figure 3. The MMR ranking with $\lambda = 0.75$ (ranked more by score) outperformed the baseline in every experiment (experiments are described in 3.3), while rankings with $\lambda = 0.25$ (ranked more by diversity) did not perform significantly better or worse than the baseline.

λ	Experiment	Trials	Baseline Trials	MMR Trials
0.25	inclusion	225	0.52	0.48
	diverse	225	0.46	0.54
	redundant	225	0.51	0.49
0.75	inclusion	75	0.32	0.68
	diverse	75	0.37	0.63
	redundant	75	0.61	0.39

Table 3. Depicts the aggregated subject counts of the blind evaluation experiment. For each (λ , experiment) group, the number of times either list was chosen is tallied. The two rightmost columns capture the normalized counts. The baseline ranking is generated with MMR and $\lambda = 1$ (ranked only by score).

For moderate diversification ($\lambda = 0.75$), the MMR ranking was chosen significantly more often than the baseline ranking for both diversity and inclusion experiments, and significantly less often than the baseline for the redundancy experiment, suggesting our model mitigates redundancy and majority biases in the top K comments. However, for extreme diversification ($\lambda = 0.25$) the fraction of responses choosing the MMR ranking was nearly 0.5 (completely random when compared with the baseline ranking) across all three experiment groups. The cause is likely two fold. Firstly, ranking correlates with relevance, therefore, replacing more high-ranking comments

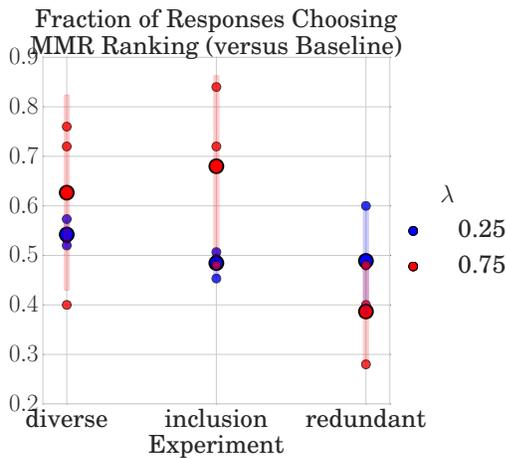


Figure 3. Depicts the fraction of trials choosing the diversified (MMR) ranking for each λ , experiment pair. Higher values for the "diverse" and "inclusion" experiments and lower values for the "redundant" experiment suggest MMR's efficacy, with $\lambda = 0.75$ outperforming $\lambda = 0.25$. The large, encircled points depict the means of each λ , experiment pair and the translucent bars depict the standard error of each mean. The smaller points depict individual rater scores.

with diverse, but lower-ranked (and less relevant) comments, may negatively impact all three experiments. Secondly, lower-ranked comments may be off-topic, lower quality, or harder to parse, leading to a simulated random choice.

Reliability and Agreement Among Test Subjects

Since only three subjects were included in our experiment, each evaluating 100 trials, we consider the inter-rater reliability among the three subjects to validate the consistency in our findings. Table 4 lists the Cohen's Kappa score for all pairs of subjects, for each experiment group. Although a small number of pairs were inconsistent, most showed moderate consistency.

		other1	other2
diversity	subject 1	-0.010918	0.273901
	subject 2	0.179319	0.273901
	subject 3	-0.010918	0.179319
inclusion	subject 1	0.033912	0.146825
	subject 2	0.185336	0.146825
	subject 3	0.033912	0.185336
redundancy	subject 1	-0.025851	0.135560
	subject 2	0.211045	0.135560
	subject 3	-0.025851	0.211045

Table 4. Cohen's Kappa pairwise inter-rater reliability scores.

CONCLUSION

Discussion forums play a vital role in online courses, yet due to the large scale of MOOCs, comment rankings often suffer from majority biases and redundancy. The primary contributions of this paper are (1) design and evaluation of a top K comment diversification re-ranking algorithm and (2) experimental evidence suggesting a significant increase in diversity and inclusion and decrease in redundancy when our algorithm is used to rank comments versus a baseline relevance ranking. We encourage MOOC and forum platforms to consider the importance of ranking diversification on learning and user experience, and hope our findings inspire future consideration.

ACKNOWLEDGEMENTS

We graciously thank Y-Lan Boureau (Facebook AI Research) for her mentorship and suggestion to use MMR and embeddings. We are also very grateful to Regina Barzilay (MIT) for her guidance in model selection and framework.

REFERENCES

1. Akiko Aizawa. 2003. An information-theoretic perspective of tf-idf measures. *Information Processing & Management* 39, 1 (2003), 45–65.
2. David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3, Jan (2003), 993–1022.
3. Jaime Carbonell and Jade Goldstein. 1998. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st annual intl. ACM SIGIR conf. on Research and development in information retrieval*. ACM, 335–336.
4. Olivier Chapelle, Shihao Ji, Ciya Liao, Emre Velipasaoglu, Larry Lai, and Su-Lin Wu. 2011. Intent-based diversification of web search results: metrics and algorithms. *Info. retrieval* 14, 6 (2011), 572–592.
5. J. Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educ. and Psych. Measurement* 20, 1 (1960), 37.
6. Yajuan Duan, Long Jiang, Tao Qin, Ming Zhou, and Heung-Yeung Shum. 2010. An empirical study on learning to rank of tweets. In *Proceedings of the 23rd International Conference on Computational Linguistics*. Association for Computational Linguistics, 295–303.
7. Susan T Dumais, George W Furnas, Thomas K Landauer, Scott Deerwester, and Richard Harshman. 1988. Using latent semantic analysis to improve access to textual information. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 281–285.
8. Anna Huang. 2008. Similarity measures for text document clustering. In *Proceedings of the sixth new zealand computer science research student conference (NZCSRSC2008)*, Christchurch, New Zealand. 49–56.
9. Thorsten Joachims. 1996. *A Probabilistic Analysis of the Rocchio Algorithm with TFIDF for Text Categorization*. Technical Report. DTIC Document.
10. Quoc V Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *ICML*, Vol. 14. 1188–1196.
11. Tao Lei, Hrishikesh Joshi, Regina Barzilay, Tommi Jaakkola, Katerina Tymoshenko, Alessandro Moschitti, and Lluís Marquez. 2016. *Semi-supervised question retrieval with gated convolutions*. Association for Computational Linguistics (ACL), 1279–1289.
12. Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Adv. in neural information processing systems*. 3111–3119.
13. Jonas Mueller and Aditya Thyagarajan. 2016. Siamese Recurrent Architectures for Learning Sentence Similarity. In *Thirtieth AAAI Conference on Artificial Intelligence*.
14. Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank citation ranking: bringing order to the web. (1999).