Webis at TREC 2019: Decision Track

Extended Abstract

Alexander Bondarenko[1,*] Maik Fröbe[1,*] Vaibhav Kasturia[1,*] Michael Völske[2,*]
Benno Stein[2] Matthias Hagen[1]

[1]Martin-Luther-Universität Halle-Wittenberg
first.last@informatik.uni-halle.de

[2]Bauhaus-Universität Weimar
first.last@uni-weimar.de

ABSTRACT

This paper gives an overview of the Webis group's participation in the TREC 2019 Decision Track. Our idea is to axiomatically re-rank the top-k results of BM25F for those topics that seem to be argumentative. For the re-ranking, we use five axioms that capture signals of argumentativeness and information credibility.

1 INTRODUCTION

Our approach to the scenario of the TREC 2019 Decision Track focuses on two requirements that search results in a decision support scenario should fulfill: (1) the results should express a preference or state pros and cons in the area to be decided, and (2) the results should do so in a well-founded manner. These requirements are particularly important for "high-stakes" decisions such as this year's Decision Track's medical topics. We address the first requirement by trying to favor results that are *more argumentative*, and the second by trying to favor results that are *more credible*.

All our runs follow the same basic process of first retrieving an initial top-k ranking using BM25F, and then potentially re-ranking these results using an axiomatic result re-ranking framework [6]. In particular, the re-ranking is addressing those queries that seem to ask for argumentative results. We thus combine three axioms focusing on the argumentativeness of the initial top-k search results—already employed in our last year's Common Core runs [3]—with two new axioms targeting search result credibility. With the judgments for our runs, we want to explore different weighting schemes for aggregating the axioms' ranking preferences.

2 WEBIS DECISION TRACK RUNS

We describe our approach to identify argumentative queries and introduce the axioms and the weighting schemes used to re-rank the top-k BM25F results of argumentative queries.

2.1 Identifying Argumentative Queries

We manually inspected the titles of the topics used in this year's Decision Track and concluded that all of them could be labeled as argumentative queries (i.e., relevant results should contain some form of argumentation). Moreover, since all topics are medical topics, information credibility is of particular concern. Thus, our proposed approach employs an axiomatic re-ranking which tries to favor argumentative and credible results. This differs from our last year's TREC Common Core Track runs where we did not consider the information credibility aspect.

2.2 Obtaining Search Results

For each topic, we submit the title as a query to an Elasticsearch index that contains four fields per document: (1) title, (2) main content, (3) meta description, and (4) keywords extracted from the meta description. All fields are extracted during the indexing pipeline of Elastic ChatNoir [2]. We use the default BM25F similarity of Elasticsearch and assign equal weight to all four fields, but manipulate this initial ranking with respect to credibility as follows.

By manual assessment of the OpenDNS categories,[1] we identified the following as credible or at least related to the Decision Track guidelines: : `Research/Reference`, `Educational Institutions`, `Government`, `Health and Fitness`, `News/Media`, and `Non-Profits`. In the original BM25F ranking, we thus move to the top the documents that belong to any of those "credible" categories (the OpenDNS service assigns some categories to each URL) while keeping the relative order of these "credible" documents according to their scores. In effect, results with URLs not mapped to any of the "credible" OpenDNS categories are ranked below the ones from "credible" hosts independent of their original BM25F score.

2.3 Argumentative Unit Detection

Our idea is to include information about argumentativeness and credibility in a re-ranking pipeline. To this end we try to identify argumentative units in result documents by using the two sequence tagging models TARGER [4] (available as an API) and MARGOT [8] (available as a library). These models take a raw text as input and return the text tagged with the information where argument premises and claims start and end. The taggers' detections could be either combined, intersected, or just the detections of an individual tagger could be used.

2.4 Identifying Medical Entities

To identify argumentative units mentioning more "credible" sources of medical information, we want to identify "medical sources" in the form of named entities related to medicine. As the basis of a simple rule-based system, we create a list of keywords commonly used in the names of medical institutions (e.g., `hospital`, `institute`, `research center`) and medical professions / degrees (e.g., `Dr.`, `MBBS`, `PhD`). In addition to medical institutions and professions, we also add

*These four authors contributed to the paper equally.
TREC 2019, November 13–15, 2019, Gaithersburg, Md.

[1]https://community.opendns.com/domaintagging/

to the list the names of news channels and newspapers found on Wikipedia since the track guidelines state that information coming from news sources should be considered as credible. Any match of the medical keywords or a newspaper / news channel will thus be viewed as a mention of a credible "medical entity". The underlying idea of using credibility in the re-ranking then is to favor documents from the initial BM25F results that contain argumentative units mentioning medical entities as sources.

2.5 Re-Ranking Axioms

We use five axioms to re-rank the top-50 baseline results (BM25F + re-ranking according to "credible" OpenDNS domains). For every pair of documents in the top-50 results, the re-ranking computes a ranking preference based on a collaborative voting of the axioms.

The basic idea of axiomatic thinking in information retrieval [1] is to identify axioms (i.e., constraints) that good retrieval models should fulfill. Employing an axiomatic re-ranking approach [6], our last year's TREC Common Core Track runs showed some promising improvements for some argumentative queries [3]. The axioms compare argumentative units (cf. Section 2.3) in the initial baseline results. We extend this axiom set by two credibility axioms and directly embed them in the axiomatic re-ranking pipeline.

2.5.1 Axioms Capturing Argumentativeness.

Axiom ArgUC (Argumentative Units Count). The general idea of the ArgUC axiom is to favor documents that contain more argumentative units.

Formalization. Let q be an argumentative query, d_1 and d_2 be two retrieved documents, $count_{\text{Arg}}(d)$ be the number of argumentative units in document d, and let $\approx_{10\%}$ indicate "equality" up to a 10% difference. If $length(d_1) \approx_{10\%} length(d_2)$ and if $count_{\text{Arg}}(d_1) > count_{\text{Arg}}(d_2)$, then $d_1 >_{\text{ArgUC}} d_2$.

Axiom QTArg (Query Term Occurrence in Argumentative Units). Retrieved documents usually consist of argumentative and non-argumentative units or text passages. The general idea of the QTArg axiom is to favor documents where the query terms appear closer to (or better: in) argumentative units.

Formalization. Let $q = \{t\}$ be an argumentative single-term query, d_1 and d_2 be two retrieved documents, and let Arg_d be the set of argumentative units of a document d. If $length(d_1) \approx_{10\%} length(d_2)$ and if $t \in A$ for some $A \in Arg_{d_1}$ but $q \notin A'$ for all $A' \in Arg_{d_2}$, then $d_1 >_{\text{QTArg}} d_2$.

Axiom QTPArg (Query Term Position in Argumentative Units). Following the general observation that in relevant documents the query terms occur closer to the beginning [10, 12], the QTPArg axiom will favor documents where the first appearance of a query term in an argumentative unit is closer to the beginning of the document.

Formalization. Let $q = \{t\}$ be an argumentative single-term query, d_1 and d_2 be two retrieved documents, and let the first position in an argumentative unit of a document d where the term t appears be denoted by $1^{st} position(t, Arg_d)$. If $length(d_1) \approx_{10\%} length(d_2)$ and if $1^{st} position(t, Arg_{d_1}) < 1^{st} position(t, Arg_{d_2})$, then $d_1 >_{\text{QTPArg}} d_2$.

2.5.2 Axioms Capturing Credibility.

Axiom MEArg (Medical Entities in Argumentative Units). Following the heuristic that the source of an argument matters, the MEArg axiom will favor documents in which argumentative units contain "medical entities" (cf. Section 2.4).

Formalization. Let q be an argumentative query, d_1 and d_2 be two retrieved documents, and let the number of identified medical entities in the argumentative units of a document d be denoted by $numberME(Arg_d)$. If $length(d_1) \approx_{10\%} length(d_2)$ and if $numberME(Arg_{d_1}) > numberME(Arg_{d_2})$, then $d_1 >_{\text{MEArg}} d_2$.

Axiom aSLDoc (Average Sentence Length in Document). Following the observations that bad readability is one component of bad information quality and thus reduced credibility [5] and that a simple measure for readability is the average sentence length [5], the aSLDoc axiom will favor documents with average sentence length between 12–20 words since this is viewed as a simplistic indicator of good readability [9, 11].

Formalization. Let q be an argumentative query, d_1 and d_2 be two retrieved documents, and let $sentLength(d)$ be the average sentence length (in words) of a document d. If $length(d_1) \approx_{10\%} length(d_2)$ and if $12 \leq sentLength(d_1) \leq 20$ and $sentLength(d_2) < 12$ or $sentLength(d_2) > 20$, then $d_1 >_{\text{aSLDoc}} d_2$.

2.6 Axiom Weights

In addition to the five argumentativeness and credibility axioms, we also employ an axiom ORIG [6]. The ORIG axiom simply returns the preferences corresponding to the baseline retrieval system's ranking—BM25F + OpenDNS re-ranking in our case (which is different to our previous year's baseline [3]). The six different axioms (including ORIG) are weighted to linearly combine the respective preference matrices analog to the original axiomatic re-ranking pipeline [6]. The weight of an axiom then directly influences its impact on the document re-ranking—the more weight one axiom has compared to the others, and the more often its precondition is fulfilled for document pairs, the more impact it has on the final re-ranking.

3 WEBIS RUNS

Our runs were simply created to get as many votes as possible for the baseline retrieval system's top-50 results (BM25F + OpenDNS re-ranking). The goal is to compare different weighting schemes in the re-ranking which could even be schemes not imagined at the time of taking part in the Decision Track. To this end, every run tried to include in its top-10 and top-20 results some results from the baseline's top-50 that were not already included in another run. Based on the judgments obtained for the baseline system's top-50 results, we wanted to examine the following three axiom weighting schemes.

3.1 Axiom Weighting Schemes

The six different axioms (including ORIG) are weighted to linearly combine the respective preference matrices analog to the original axiomatic re-ranking pipeline, where the weight then influences the axiom's impact on the document re-ranking. A weighting scheme may be combined with the argumentative units detected by TARGER or MARGOT [4, 8]. We apply the following three weighting schemes.

Equal Weights (EW). All axioms get the same weight. This way, any agreement of the preferences of a pair of the new axioms may overrule the ORIG axiom preference when the no other axiom "supports" the ORIG preference.

Majority Voting (MV). The axioms are assigned weights such that document pairs are re-ranked iff the majority of the new axioms (at least 3 out of 5 axioms) agree to overrule the ORIG preference.

Total Agreement (TA). The axioms are assigned weights such that document pairs are re-ranked only when all the new axioms agree. It is not necessary for all axioms to have the same weight, although all of them have to be in agreement to overrule the ORIG axiom.

4 EVALUATION

We compare the different axiomatic re-rankings to the baseline ranking (BM25F+OpenDNS ranking) using three different metrics: nDCG@3, Precision@1 and NLRE@all (NLRE [7] combines aspects of relevance and credibility).

Table 1 shows the results for each topic and an average for all topics—grouped by weighting scheme and argumentative unit tagger. The axiomatic re-ranking with any weighting scheme and any tagger somewhat improves the NLRE scores (slight exception: topic 27). Argument unit tagging with MARGOT yields slightly better NLRE scores than TARGER.

The nDCG and precision metrics are compared at lower depths to emphasize the importance of the top-most results in web search scenarios and since precision@1 also corresponds to the case of building a medical question answering system. Using the TA weighting scheme (total agreement) with both argument taggers shows a slight improvement in precision. In contrast, nDCG@3 substantially decreases after axiomatic re-ranking.

The TA weighting scheme (total agreement) seems to better suited than the MV (majority voting) and the EW (equal weights) schemes. A reason could be that TA (total agreement) re-ranking decisions come with a higher "confidence" to change the original document ordering, while less agreement is needed in the other schemes. Some more thorough investigations and weighting fine tuning might help to improve the ranking quality further.

5 CONCLUSION

We have used an axiomatic re-ranking based on five axioms that aim to capture basic ideas about a search result's argumentativeness and credibility for queries that address decision making in a medical context. Our underlying hypothesis is that such queries will benefit from search results coming from trustworthy sources that argumentatively compare pros and cons of different options. The evaluation of our re-ranking approach using the judgments gathered for the baseline retrieval systems' top results (splitted over our runs) show that our current set of axioms is not really sufficient to improve standard scores of nDCG@3 or Precision@1 over a BM25F baseline but that some improvements in terms of NLRE are possible.

Interesting steps for future work include: (1) query expansion with synonyms and aliases of the query terms, (2) better medical entity identification via a more sophisticated medical entity recognizer, (3) improved weighting schemes based on learning axiom weights, (4) better detection of argumentative queries by deeper investigation of the relevance judgments, and (5) formulation of more sophisticated ideas to capture argumentativeness and credibility of documents.

ACKNOWLEDGMENTS

This work has been partially supported by the DFG through the project "ACQuA: Answering Comparative Questions with Arguments" (grant HA 5851/2-1) as part of the priority program "RATIO: Robust Argumentation Machines" (SPP 1999).

REFERENCES

[1] Enrique Amigó, Hui Fang, Stefano Mizzaro, and ChengXiang Zhai. 2017. Axiomatic Thinking for Information Retrieval: And Related Tasks. In *Proceedings of SIGIR 2017*. 1419–1420.
[2] Janek Bevendorff, Benno Stein, Matthias Hagen, and Martin Potthast. 2018. Elastic ChatNoir: Search Engine for the ClueWeb and the Common Crawl. In *Proceedings of ECIR 2018*. 820–824.
[3] Alexander Bondarenko, Michael Völske, Alexander Panchenko, Chris Biemann, Benno Stein, and Matthias Hagen. 2018. Webis at TREC 2018: Common Core Track. In *Proceedings of TREC 2018*.
[4] Artem Chernodub, Oleksiy Oliynyk, Philipp Heidenreich, Alexander Bondarenko, Matthias Hagen, Chris Biemann, and Alexander Panchenko. 2019. TARGER: Neural Argument Mining at Your Fingertips. In *Proceedings of ACL 2019 (Demos)*. 195–200.
[5] Gunther Eysenbach, John Powell, Oliver Kuss, and Eun-Ryoung Sa. 2002. Empirical Studies Assessing the Quality of Health Information for Consumers on the World Wide Web: A Systematic Review. *JAMA* 287, 20 (2002), 2691–2700.
[6] Matthias Hagen, Michael Völske, Steve Göring, and Benno Stein. 2016. Axiomatic Result Re-Ranking. In *Proceedings of CIKM 2016*. 721–730.
[7] Christina Lioma, Jakob Grue Simonsen, and Birger Larsen. 2017. Evaluation Measures for Relevance and Credibility in Ranked Lists. In *Proceedings of ICTIR 2017*. 91–98.
[8] Marco Lippi and Paolo Torroni. 2016. MARGOT: A Web Server for Argumentation Mining. *Expert Syst. Appl.* 65 (2016), 292–303.
[9] Mike Markel. 2010. Technical Communication. *Bedford/St Martin's* (2010).
[10] Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to Match Using Local and Distributed Representations of Text for Web Search. In *Proceedings of WWW 2017*. 1291–1299.
[11] Casi Newell. 2014. Editing Tip: Sentence Length. *AJE Scholar (online post)* https://www.aje.com/en/arc/editing-tip-sentence-length/
[12] Adam D. Troy and Guo-Qiang Zhang. 2007. Enhancing Relevance Scoring with Chronological Term Rank. In *Proceedings of SIGIR 2007*. 599–606.

Table 1: Evaluation results on the individual topics and averaged for BM25F+OpenDNS filtering as the baseline retrieval system (B) and our re-ranking approaches with different weighting schemes: total agreement (TA), majority voting (MV) and equal weights (EW) using the two different argument taggers MARGOT and TARGER.

	nDCG@3							P@1							NLRE@all						
		MARGOT			TARGER				MARGOT			TARGER				MARGOT			TARGER		
ID	B	TA	MV	EW	TA	MV	EW	B	TA	MV	EW	TA	MV	EW	B	TA	MV	EW	TA	MV	EW
1	0.38	0.30	0.70	0.62	0.30	0.30	0.15	0	0	1	1	0	0	0	0.60	0.73	0.73	0.73	0.73	0.72	0.72
2	0.85	0.23	0.73	0.23	0.35	0.15	0.12	1	1	1	1	1	0	0	0.46	0.64	0.64	0.65	0.64	0.64	0.65
3	0.77	0.00	0.00	0.00	0.00	0.47	0.00	1	0	0	0	0	1	0	0.75	0.87	0.87	0.87	0.87	0.88	0.87
4	0.38	0.38	0.38	0.50	0.50	0.50	0.38	1	1	1	1	1	1	1	0.60	0.68	0.68	0.68	0.68	0.68	0.68
5	0.53	0.47	0.00	0.00	0.00	0.23	0.00	0	1	0	0	0	0	0	0.52	0.69	0.69	0.66	0.68	0.68	0.67
6	0.57	0.83	0.66	0.00	0.83	0.70	0.57	1	1	1	0	1	1	1	0.55	0.76	0.76	0.74	0.76	0.76	0.75
7	0.13	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.61	0.67	0.66	0.66	0.67	0.67	0.66
8	0.00	0.20	0.20	0.16	0.20	0.20	0.16	0	0	0	0	0	0	0	0.83	0.84	0.84	0.83	0.84	0.84	0.84
9	0.38	0.30	0.23	0.70	0.23	0.23	0.47	0	0	0	1	0	0	1	0.75	0.82	0.82	0.82	0.82	0.82	0.82
10	0.23	0.00	0.00	0.12	0.00	0.00	0.00	1	0	0	0	0	0	0	0.72	0.80	0.80	0.80	0.80	0.80	0.80
11	0.48	0.68	0.16	0.00	0.68	0.68	0.20	1	1	0	0	1	1	0	0.43	0.58	0.57	0.57	0.58	0.58	0.57
12	0.41	0.85	0.85	0.23	0.85	0.77	0.00	0	1	1	0	1	1	0	0.64	0.75	0.75	0.75	0.75	0.75	0.75
13	0.32	0.00	0.00	0.00	0.16	0.16	0.00	1	0	0	0	0	0	0	0.58	0.74	0.74	0.75	0.75	0.75	0.75
14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.65	0.74	0.74	0.74	0.74	0.74	0.74
15	0.50	0.53	0.00	0.15	0.53	0.53	0.00	1	0	0	0	0	0	0	0.71	0.77	0.77	0.77	0.77	0.77	0.77
16	0.27	0.23	0.23	0.00	0.23	0.23	0.00	0	1	1	0	1	1	0	0.57	0.65	0.65	0.65	0.65	0.65	0.65
17	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.68	0.77	0.77	0.77	0.77	0.77	0.77
18	0.70	0.77	0.77	0.77	0.77	1.00	0.53	1	1	1	1	1	1	0	0.61	0.80	0.80	0.80	0.80	0.80	0.78
19	0.00	0.77	0.77	0.00	0.77	0.77	0.00	0	1	1	0	1	1	0	0.49	0.64	0.64	0.63	0.64	0.64	0.63
20	0.41	0.38	0.15	0.12	0.38	0.38	0.00	0	1	0	0	1	1	0	0.60	0.67	0.66	0.66	0.67	0.67	0.67
21	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.58	0.77	0.77	0.77	0.77	0.77	0.77
22	0.38	0.15	0.15	0.23	0.15	0.15	0.23	1	0	0	1	0	0	1	0.60	0.75	0.75	0.75	0.75	0.75	0.75
23	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.78	0.86	0.86	0.86	0.86	0.86	0.85
24	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.59	0.73	0.73	0.73	0.73	0.73	0.73
25	0.00	0.23	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.64	0.76	0.77	0.76	0.76	0.77	0.76
26	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.77	0.85	0.85	0.85	0.86	0.86	0.86
27	0.68	0.68	0.68	0.36	0.52	0.48	0.00	1	1	1	0	1	1	0	0.86	0.85	0.85	0.85	0.85	0.85	0.85
28	0.65	0.47	0.47	0.50	0.47	0.47	0.38	1	1	1	1	1	1	0	0.77	0.89	0.89	0.88	0.89	0.89	0.88
29	0.35	0.70	0.70	0.15	0.70	0.47	0.50	1	1	1	0	1	1	1	0.76	0.77	0.77	0.77	0.78	0.78	0.78
30	0.00	0.00	0.00	0.15	0.00	0.00	0.12	0	0	0	0	0	0	0	0.59	0.70	0.70	0.69	0.70	0.70	0.69
31	0.00	0.12	0.12	0.00	0.12	0.00	0.12	0	0	0	0	0	0	0	0.80	0.86	0.86	0.86	0.86	0.87	0.86
32	1.00	0.47	0.30	1.00	0.47	0.47	0.53	1	1	0	1	1	1	0	0.56	0.65	0.65	0.65	0.65	0.65	0.65
33	0.12	0.00	0.00	0.00	0.00	0.00	0.00	0	0	0	0	0	0	0	0.56	0.68	0.68	0.67	0.68	0.68	0.67
34	0.15	0.23	0.23	0.00	0.23	0.23	0.00	0	1	1	0	1	1	0	0.50	0.73	0.73	0.73	0.73	0.73	0.74
35	0.77	0.70	0.70	0.47	0.70	0.77	0.00	1	1	1	1	1	1	0	0.72	0.83	0.83	0.82	0.83	0.83	0.82
36	0.35	0.23	0.23	0.00	0.23	0.23	0.23	1	1	1	0	1	1	0	0.73	0.77	0.76	0.77	0.77	0.76	0.75
37	0.77	0.47	0.30	0.00	0.47	0.47	0.00	1	1	0	0	1	1	0	0.69	0.85	0.86	0.85	0.86	0.86	0.85
38	1.00	0.47	0.47	0.00	0.47	0.47	0.47	1	1	1	0	1	1	1	0.49	0.70	0.70	0.69	0.70	0.70	0.70
39	1.00	1.00	0.53	0.00	1.00	0.00	0.00	1	1	0	0	1	0	0	0.57	0.73	0.72	0.73	0.73	0.72	0.73
40	0.00	0.23	0.15	0.15	0.23	0.38	0.23	0	1	0	0	1	0	1	0.65	0.74	0.74	0.75	0.74	0.74	0.74
41	0.85	0.73	0.47	0.53	0.73	0.73	0.41	1	1	1	0	1	1	0	0.52	0.61	0.60	0.58	0.61	0.61	0.58
42	0.73	0.12	0.12	0.00	0.12	0.27	0.00	1	0	0	0	0	0	0	0.75	0.85	0.85	0.85	0.85	0.85	0.85
43	0.00	0.47	0.47	0.23	0.47	0.53	0.47	0	1	1	1	1	1	1	0.81	0.90	0.89	0.89	0.89	0.90	0.89
44	0.70	0.77	0.77	0.00	0.77	0.77	0.00	1	1	1	0	1	1	0	0.48	0.79	0.79	0.78	0.79	0.79	0.78
45	0.88	0.12	0.12	0.00	0.12	0.12	0.00	1	0	0	0	0	0	0	0.75	0.91	0.91	0.91	0.91	0.91	0.91
46	0.32	0.00	0.00	0.00	0.00	0.00	0.00	1	0	0	0	0	0	0	0.31	0.70	0.71	0.70	0.70	0.70	0.68
47	1.00	0.59	0.47	0.00	0.59	0.47	0.00	1	1	1	0	1	1	0	0.59	0.74	0.75	0.74	0.74	0.75	0.75
48	0.27	0.23	0.12	0.00	0.23	0.15	0.00	0	1	0	0	1	0	0	0.71	0.84	0.84	0.84	0.84	0.84	0.84
49	0.00	0.38	0.38	0.38	0.38	0.50	0.50	0	1	1	1	1	1	1	0.60	0.68	0.68	0.68	0.68	0.68	0.68
50	0.00	0.15	0.15	0.00	0.15	0.15	0.00	0	0	0	0	0	0	0	0.81	0.91	0.90	0.91	0.91	0.91	0.91
AVG	0.39	0.33	0.28	0.16	0.32	0.31	0.14	0.50	0.54	0.40	0.22	0.52	0.46	0.18	0.64	0.76	0.76	0.73	0.73	0.72	0.72