



HAL
open science

Can we predict QPP? An approach based on multivariate outliers

Adrian-Gabriel Chifu, Sébastien Déjean, Moncef Garouani, Josiane Mothe,
Diégo Ortiz, Md Zia Ullah

► **To cite this version:**

Adrian-Gabriel Chifu, Sébastien Déjean, Moncef Garouani, Josiane Mothe, Diégo Ortiz, et al.. Can we predict QPP? An approach based on multivariate outliers. 46th European Conference on Information Retrieval (ECIR 2024), Mar 2024, Glasgow, United Kingdom. pp.458–467, 10.1007/978-3-031-56063-7_38 . hal-04465059

HAL Id: hal-04465059

<https://hal.science/hal-04465059>

Submitted on 19 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Can we predict QPP? An approach based on multivariate outliers

Adrian-Gabriel Chifu¹[0000-0003-4680-5528], Sébastien Déjean²[0000-0001-9610-5306], Moncef Garouani⁴[0000-0003-2528-441X], Josiane Mothe^{3,5}[0000-0001-9273-2193], Diégo Ortiz⁵[0009-0002-7272-5644], and Md Zia Ullah⁶[0000-0002-4022-7344]

¹ Aix Marseille Université, Université de Toulon, CNRS, LIS, Marseille, France
`adrian.chifu@univ-amu.fr`

² IMT, UMR5219 CNRS, UPS, Univ. de Toulouse, Toulouse, France
`sebastien.dejean@math.univ-toulouse.fr`

³ IRIT, UMR5505 CNRS, Université de Toulouse, INSPE, UT2J, Toulouse, France
`josiane.mothe@irit.fr`

⁴ IRIT, UMR 5505 CNRS, Université Toulouse Capitole, UT1, Toulouse, France
`moncef.garouani@irit.fr`

⁵ IRIT, UMR 5505, Toulouse, France `diego.ortiz@irit.fr`

⁶ Edinburgh Napier University, Edinburgh, UK `m.ullah@napier.ac.uk`

Abstract. Query performance prediction (QPP) aims to forecast the effectiveness of a search engine across a range of queries and documents. While state-of-the-art predictors offer a certain level of precision, their accuracy is not flawless. Prior research has recognized the challenges inherent in QPP but often lacks a thorough qualitative analysis. In this paper, we delve into QPP by examining the factors that influence the predictability of query performance accuracy. We propose the working hypothesis that while some queries are readily predictable, others present significant challenges. By focusing on outliers, we aim to identify the queries that are particularly challenging to predict. To this end, we employ multivariate outlier detection method. Our results demonstrate the effectiveness of this approach in identifying queries on which QPP do not perform well, yielding less reliable predictions. Moreover, we provide evidence that excluding these hard-to-predict queries from the analysis significantly enhances the overall accuracy of QPP.

Keywords: Information Retrieval · Query performance prediction · QPP · Post-retrieval features · Multivariate outlier detection

1 Introduction

A search engine aims to process and answer any user query by retrieving relevant documents. However, the performance of a particular search engine can vary substantially depending on the specific queries it encounters [20, 15].

Query performance prediction (QPP) addresses the crucial task of predicting how effective a system will be on a given query [37, 7, 35, 8]. This problem is of

paramount importance for two primary reasons. Firstly, when a query is expected to be difficult, it may be necessary to adopt a specialized, albeit potentially costly, approach. Conversely, when a query is predicted to be easy, a simpler and more cost-effective method can be employed [11, 2, 14].

QPP accuracy is typically assessed by measuring the correlation between the predicted and the actual performance values [35, 31, 19, 12]. This approach is sound as long as the underlying assumptions and conditions for applying correlation measurements are respected [3].

Systems vary in their approach to processing a given query. This variability encompasses processes such as automatic query reformulation, the choice of the weighting/matching function (BM25 [32], a Language model [36], or an LLM-based model [16, 26, 9, 38], for example), and the application of document re-ranking models [25]. Consequently, different systems will perform differently on the same queries [14]. QPP is thus considered with regard to the system it predicts the performance for. This implies that a QPP predictor should adapt its behavior to suit a particular system. This could be a reason why post-retrieval QPPs, which use the results of the search, tend to be more accurate than pre-retrieval ones, which rely solely on the query and the set of documents [21, 22].

We found for example that *LemurTF_IDF*, a post-retrieval QPP corresponding to a Letor feature [6, 8] has a higher correlation with the actual Average precision (AP) obtained with the LGD weighting function [10] than with the JS weighting function [1] (resp. 0.522 and 0.504 Pearson correlation). On the other hand, two predictors will behave differently on a particular system (Fig. 1).




Fig. 1. QPPs behave differently on the same system and set of queries (TREC78). Predicted AP (X-axis) and actual AP (Y-axis) obtained with LGD weighting function and their Pearson correlation. Left side In_expC2 ($\rho=0.484$) - Right side LemurTF_IDF QPP ($\rho=0.552$).

In Figure 1, we can see that the correlation is relatively weak, which is accurately reflected by the 0.484 (resp. 0.552) Pearson correlation values. Current QPP models lack accuracy. Single features, even post-retrieval ones could not demonstrate high correlation with actual performance [21, 7, 18]; even the com-

bination of predictors, which is out of the scope of this paper, has not been very successful [13, 37, 31, 8, 27, 33].

In this paper, we aim to analyze performance prediction deeper. Past studies showed that predictors do not work well when considering all queries. We hypothesize that prediction accuracy can significantly vary among queries, and our objective is to identify the queries for which the predictor fails to estimate the effectiveness. In other words, we want to predict the predictability of performance and address the following research question: Is it possible to predict the queries for which a QPP can provide an accurate effectiveness estimation for a given system? To put it differently, can we concentrate on those queries that are more likely to yield accurate predictions and perhaps automatically disregard queries that are more challenging to predict?

Some queries can be considered outliers (*abnormally* easy or difficult); similarly, predictions may have abnormal values. We hypothesize that the queries with abnormal prediction values are difficult to predict or get unreliable predictions. To tackle this problem, we consider multivariate outlier detection as a means to identify these hard-to-predict queries. Since a given QPP may behave differently to estimate the effectiveness of a system, the idea is to consider multiple QPPs in the query identification phase.

We consider several effectiveness measures and several benchmark collections. We show that our hypothesis could pave the way for a new research direction on QPP on accuracy predictability.

2 Identification of difficult to predict queries by TRC

Here, we hypothesize that some predictions are outliers because the queries are difficult to predict. We aim to identify those queries for which we anticipate the QPP will not be accurate.

Johnson defines an outlier as an observation in a data set that appears to be inconsistent with the remainder of that set of data [24]. Univariate methods consider each variable independently, so only observations that appear odd for that variable are detected, while in the case of multivariate outlier detection, the interactions among different variables are compared. Multivariate outliers are a combination of unusual scores on several variables [30], and the idea is to detect the observations that are located relatively far from the center of the data distribution [5]. To identify outliers, we thus consider here multivariate outliers and identify the queries for which predictions are abnormal for different QPPs.

Mahalanobis distance is a common criterion for multivariate outlier detection. Applied to QPP, the Mahalanobis distance for a given query q_i , from a set of n queries $Q = \{q_1, q_2, \dots, q_n\}$, can be defined as follows:

$$D_M(q_i, Q) = \sqrt{(q_i - \bar{q})^T V_n^{-1} (q_i - \bar{q})}, \quad (1)$$

where the query vector q_i is composed of m ($n \gg m$) predictor values $q_i = (p_1^i, p_2^i, \dots, p_m^i)$, \bar{q} is the mean vector of the queries, and V_n^{-1} is the inverse of the covariance matrix of the queries. The superscript T denotes the transpose

of the vector. We used the Transformed Rank Correlations (TRC) for multivariate outlier detection function implemented in R in this paper [4]⁷. Instead of working with the raw data, which might be skewed or have a non-normal distribution, the TRC method works with the ranks of the data. This involves sorting the data from smallest to largest and using their position in this sorted order as their value. It then calculates the Spearman rank correlation or a similar non-parametric correlation measure between all pairs of variables. This helps in understanding the underlying structure and relationships in the data, which is crucial for identifying outliers in a multivariate context. The TRC method calculates a distance measure for each observation in the multivariate space. This distance is typically based on rank correlations and reflects how far away each observation is from the central trend of the data. It then determines a threshold for these distances (default value is set to the 0.95 quantile of F-distribution). Observations with a distance greater than this threshold are flagged as outliers.

3 Data

In this study, we use TREC adhoc collections: TREC78 with 528K documents and 100 queries (351 – 450) and WT10G which consists of 1.692M documents and 100 queries (451 – 550). We use Average Precision (AP@1000) [34] and normalized Discounted Cumulative Gain (nDCG@20) [23] as measures to evaluate the search engine performance, which are common in adhoc retrieval. Like previous studies on QPP [13, 37, 7, 35, 31, 8], we use Pearson correlation as a measure to evaluate the accuracy of the QPP, in addition to plots as recommended in statistics. Pearson correlation measures the linear correlation between the predicted effectiveness of queries and their actual effectiveness. It indicates how well the QPP method predictions align with the actual query performances. Although it may not be a perfect measure [28], we kept it in this study. Alternative measures will be used in future work. We will consider other measures such as Mean Squared Error (MSE), Mean Absolute Error (MAE) or its normalized variant. The former measures the average of the squares of the errors between predicted and actual query effectiveness. The latter measures the average of the absolute differences between predicted and actual values. It is less sensitive to outliers compared to MSE. These two measures could help in understanding the magnitude of the prediction errors.

We consider a series of systems that treat the same queries over the same set of documents. In this study, we construct systems using Terrier [29]. These systems differ on several factors: the scoring function employed and the variant of the query reformulation module utilized, if any. We report the results on the best system according to the considered collection and performance measures, the best in terms of the average effectiveness over the set of queries. We also used a reference system based on LM Dirichlet [36].

In this study, we consider the first four Letor features which have been shown as among the most accurate for single feature QPP [6, 8]. Letor features have been

⁷ <https://cran.r-project.org/web/packages/modi/modi.pdf>

initially used for retrieved document re-ranking [6]. Letor features are associated with each (query, retrieved document) pair ⁸. To obtain a single value for each query for a given letor feature, the values are aggregated over the documents. We used maximum as the aggregation function which has been shown as the most accurate for QPP [8]. The four features we kept are LemurTF_IDF, In_expC2, InB2, and InL2, aggregated using the maximum function. We also consider four state-of-the-art QPP: Normalized query commitment (NQC) [35], Unnormalized query commitment (UCQ) [35], QF (query feedback) [37] and WIG [37].

These features will be used individually to predict system performance. Although methods that combine features have been shown more accurate, here the objective is to conduct a first study on how to detect queries that are difficult to predict. We though opted for simple prediction models (single feature); we will analyze the proposed method for more complex models. Complex models that involve several features need to be trained. Here since we are not training a model, we do not need to split the data into train and test.

4 Results




Fig. 2. Outlier queries are different according to the QPP. Queries detected as outliers by our method are in red. The four plots correspond to the four QPPs. The X-axis is not meaningful. Y-axis corresponds to the predicted values by the considered Letor QPP. QQP values are calculated for the TREC78 collection.

Outlier queries are different according to the considered QPP, which shows the importance of using multivariate outlier detection (Fig. 2). For example, query 403 is a clear outlier for InL2 QPP (top right of the right-side sub-figure)

⁸ using Terrier (see <http://terrier.org/docs/v5.1/learning.html>)

Table 1. Pearson correlation is consistently better when queries our method detects as difficult to predict are removed. Correlation between actual effectiveness (either AP or ndcg) on TREC78 and WT10G. We report the number of outlier queries detected, the correlation when outliers are removed using univariate and multivariate methods and when all the queries are considered for the 4 Letor features and 4 other features. The reference system rows indicate the effectiveness on average for the considered measure and set of queries; this is the system that obtained the highest effectiveness measure over a set of system configurations we tried. A unilateral test of significance for the difference between the correlations with “No Outliers” and “All” were performed using the cocor R package [17]. A * after a correlation indicates that the p-value of the test is lower than 0.05.

Collection	WT10G	WT10G	TREC78	TREC78	Collection	WT10G	WT10G	TREC78	TREC78
Measure	NDCG	AP	NDCG	AP	Measure	NDCG	AP	NDCG	AP
Reference system	0.444	0.236	0.524	0.238	Ref system	0.4528	0.187	0.5235	0.2538
Outliers	16	16	19	18	Outliers	24	24	42	42
LemurTF_IDF - Univariate	0.337	0.342	0.544	0.658	NQC - No Outliers	0.330	0.310*	0.246	0.227
LemurTF_IDF - Outliers only	0.206	0.292	0.095	0.350	NQC - All	0.097	0.051	0.219	0.174
LemurTF_IDF - No Outliers	0.438	0.468	0.601*	0.700*	UQC - No Outliers	0.350	0.351	0.382	0.340
LemurTF_IDF - All	0.365	0.393	0.381	0.522	UQC - All	0.206	0.151	0.333	0.283
In_expC2 - Univariate	0.423	0.368	0.607	0.631	WIG - No Outliers	-0.015	-0.007	0.059	-0.020
In_expC2 - No Outliers	0.391	0.350	0.607*	0.635	WIG - All	0.077	0.041	0.002	-0.103
In_expC2 - All	0.425	0.371	0.418	0.484	QF - No Outliers	0.357	0.369	0.493	0.438
InB2 - Univariate	0.329	0.286	0.542	0.536	QF - All	0.283	0.281	0.469	0.403
InB2 - No Outliers	0.286	0.214	0.530	0.543					
InB2 - All	0.336	0.274	0.372	0.416					
InL2 - Univariate	0.264	0.341	0.380	0.426					
InL2 - No Outliers	0.258	0.347	0.458	0.491					
InL2 - All	0.340	0.353	0.398	0.446					

5 Conclusion

Studies on QPP generally focus on prediction accuracy [13, 37, 7, 35, 31, 8, 18], but seldom on the difficulty of that prediction. QPP is clearly a difficult task, since current predictors are not very accurate. In this study we show that difficult to predict queries can be detected. We used multivariate outlier detection for that; it has the advantage to consider different QPPs to detect the queries for which the prediction may not be accurate. We also show that removing these automatically detected queries, we have a higher accuracy of the predictor. That means that we know that the predictor is not accurate for some queries that we can identify. This result pave the way to a new research direction: the prediction of the accuracy of the prediction. Some predictor may be accurate for certain queries and not for others. In future work, we will re-examine the results on more benchmark collections and reference retrieval systems. In this study we considered a single feature at a time to make the prediction. In future work, we will consider trained models that combine several features. That will also allow us to consider methods that allow models to abstain from making a prediction when they are not sufficiently confident.

References

1. Amati, G.: Frequentist and bayesian approach to information retrieval. In: European Conference on Information Retrieval. pp. 13–24. Springer (2006)

2. Amati, G., Carpineto, C., Romano, G.: Query difficulty, robustness, and selective application of query expansion. In: McDonald, S., Tait, J. (eds.) *Advances in Information Retrieval*. pp. 127–137. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
3. Anscombe, F.J.: Graphs in statistical analysis. *The American Statistician* **27**(1), 17–21 (1973)
4. Béguin, C., Hulliger, B.: Multivariate outlier detection in incomplete survey data: the epidemic algorithm and transformed rank correlations. *Journal of the Royal Statistical Society Series A: Statistics in Society* **167**(2), 275–294 (2004)
5. Ben-Gal, I.: Outlier detection. *Data mining and knowledge discovery handbook* pp. 131–146 (2005)
6. Cao, Z., Qin, T., Liu, T.Y., Tsai, M.F., Li, H.: Learning to rank: from pairwise approach to listwise approach. In: *ICML*. pp. 129–136 (2007)
7. Carmel, D., Yom-Tov, E.: *Estimating the query difficulty for information retrieval*. Morgan & Claypool Publishers (2010)
8. Chifu, A.G., Laporte, L., Mothe, J., Ullah, M.Z.: Query performance prediction focused on summarized letor features. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. pp. 1177–1180 (2018)
9. Clark, K., Luong, M.T., Le, Q.V., Manning, C.D.: Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555* (2020)
10. Clinchant, S., Gaussier, E.: Information-based models for ad hoc ir. In: *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. pp. 234–241 (2010)
11. Cronen-Townsend, S., Zhou, Y., Croft, W.B.: Predicting query performance. In: *International ACM SIGIR conference on Research and development in information retrieval*. pp. 299–306 (2002)
12. Datta, S., MacAvaney, S., Ganguly, D., Greene, D.: A ‘pointwise-query, listwise-document’ based query performance prediction approach. In: *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. pp. 2148–2153 (2022)
13. De, R., Grivolla, J., Jourlin, P., de Mori, R.: Automatic classification of queries by expected retrieval performance. In: *ACM SIGIR 2005 Workshop on Predicting Query Difficulty—Methods and Applications* (2005) (2005)
14. Déjean, S., Mothe, J., Ullah, M.Z.: Studying the variability of system setting effectiveness by data analytics and visualization. In: *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 10th International Conference of the CLEF Association, CLEF 2019, Lugano, Switzerland, September 9–12, 2019, Proceedings 10*. pp. 62–74. Springer (2019)
15. Deveaud, R., Mothe, J., Ullah, M.Z., Nie, J.Y.: Learning to adaptively rank document retrieval system configurations. *ACM Transactions on Information Systems (TOIS)* **37**(1), 3 (2018). <https://doi.org/10.1145/3231937>
16. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
17. Diedenhofen, B., Musch, J.: cocor: A comprehensive solution for the statistical comparison of correlations. *PLoS ONE* **10**(4) (2005)
18. Faggioli, G., Formal, T., Marchesin, S., Clinchant, S., Ferro, N., Piwowarski, B.: Query performance prediction for neural ir: Are we there yet? In: *European Conference on Information Retrieval*. pp. 232–248. Springer (2023)

19. Faggioli, G., Zendel, O., Culpepper, J.S., Ferro, N., Scholer, F.: smare: a new paradigm to evaluate and understand query performance prediction methods. *Information Retrieval Journal* **25**(2), 94–122 (2022)
20. Harman, D., Buckley, C.: Overview of the reliable information access workshop. *Information Retrieval* **12**(6), 615 (2009)
21. Hauff, C., Hiemstra, D., de Jong, F.: A survey of pre-retrieval query performance predictors. In: *ACM CIKM*. pp. 1419–1420 (2008)
22. Jafarzadeh, P., Ensan, F.: A semantic approach to post-retrieval query performance prediction. *Information Processing & Management* **59**(1), 102746 (2022). <https://doi.org/https://doi.org/10.1016/j.ipm.2021.102746>, <https://www.sciencedirect.com/science/article/pii/S0306457321002284>
23. Järvelin, K., Kekäläinen, J.: Ir evaluation methods for retrieving highly relevant documents. *ACM SIGIR Forum* **51**(2), 243–250 (2017)
24. Johnson, R.A., Wichern, D.W., et al.: *Applied multivariate statistical analysis*. Prentice hall Upper Saddle River, NJ (2002)
25. Liu, T.Y., et al.: Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval* **3**(3), 225–331 (2009)
26. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019)
27. Mizzaro, S., Mothe, J., Roitero, K., Ullah, M.Z.: Query performance prediction and effectiveness evaluation without relevance judgments: Two sides of the same coin. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. pp. 1233–1236 (2018)
28. Mothe, J.: On correlation to evaluate QPP. In: *Proceedings, C.W. (ed.) QPP++ 2023: Query Performance Prediction and Its Evaluation in New Tasks Workshop co-located with 45th European Conference on Information Retrieval (ECIR)*. vol. 3366, p. 29 – 36 (2023), <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85152226741&partnerID=40&md5=a5dea754c0149bc72f573d7f89457568>
29. Ounis, I., Amati, G., Plachouras, V., He, B., Macdonald, C., Johnson, D.: Terrier information retrieval platform. In: *Advances in Information Retrieval: 27th European Conference on IR Research, ECIR 2005, Santiago de Compostela, Spain, March 21-23, 2005. Proceedings 27*. pp. 517–519. Springer (2005)
30. Peña, D., Prieto, F.J.: Multivariate outlier detection and robust covariance matrix estimation. *Technometrics* **43**(3), 286–310 (2001)
31. Raiber, F., Kurland, O.: Query-performance prediction: Setting the expectations straight. In: *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*. pp. 13–22. SIGIR '14, ACM, New York, NY, USA (2014). <https://doi.org/10.1145/2600428.2609581>, <http://doi.acm.org/10.1145/2600428.2609581>
32. Robertson, S., Zaragoza, H., et al.: The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval* **3**(4), 333–389 (2009)
33. Roy, D., Ganguly, D., Mitra, M., Jones, G.J.: Estimating gaussian mixture models in the local neighbourhood of embedded word vectors for query performance prediction. *Information Processing & Management* **56**(3), 1026–1045 (2019). <https://doi.org/https://doi.org/10.1016/j.ipm.2018.10.009>, <https://www.sciencedirect.com/science/article/pii/S0306457318302437>
34. Sakai, T.: On the reliability of information retrieval metrics based on graded relevance. *Information processing & management* **43**(2), 531–548 (2007)

35. Shtok, A., Kurland, O., Carmel, D., Raiber, F., Markovits, G.: Predicting query performance by query-drift estimation. *ACM Transactions on Information Systems (TOIS)* **30**(2), 11 (2012)
36. Zhai, C., Lafferty, J.: A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems (TOIS)* **22**(2), 179–214 (2004)
37. Zhou, Y., Croft, W.B.: Query performance prediction in web search environments. In: *ACM SIGIR*. pp. 543–550 (2007)
38. Zhu, Y., Yuan, H., Wang, S., Liu, J., Liu, W., Deng, C., Dou, Z., Wen, J.R.: Large language models for information retrieval: A survey. *arXiv preprint arXiv:2308.07107* (2023)