



Forward and Backward Feature Selection for Query Performance Prediction

Sébastien Dejean, Radu Tudor Ionescu, Josiane Mothe, Md Zia Ullah

► To cite this version:

Sébastien Dejean, Radu Tudor Ionescu, Josiane Mothe, Md Zia Ullah. Forward and Backward Feature Selection for Query Performance Prediction. 35th Annual ACM Symposium on Applied Computing (SAC 2020), Mar 2020, Brno, Czech Republic. pp.690-697, <10.1145/3341105.3373904>. <hal-02942304>

HAL Id: hal-02942304

<https://hal.science/hal-02942304v1>

Submitted on 17 Sep 2020

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.



HAL Authorization



Open Archive Toulouse Archive Ouverte

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is an author's version published in:
<http://oatao.univ-toulouse.fr/26302>

Official URL
<https://doi.org/10.1145/3341105.3373904>

To cite this version: Déjean, Sebastien and Ionescu, Radu Tudor and Mothe, Josiane and Ullah, Md Zia *Forward and Backward Feature Selection for Query Performance Prediction*. (2020) In: 35th Annual ACM Symposium on Applied Computing (SAC 2020), 30 March 2020 - 3 April 2020 (Brno, Czech Republic).

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

Forward and Backward Feature Selection for Query Performance Prediction

Sébastien Déjean

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

Josiane Mothe

IRIT UMR5505 CNRS, INSPE, Univ. de Toulouse
Toulouse, France
Josiane.Mothe@irit.fr

Radu Tudor Ionescu

University of Bucharest
Bucharest, Romania
raducu.ionescu@gmail.com

Md Zia Ullah

IRIT UMR5505 CNRS
Toulouse, France
mdzia.ullah@irit.fr

ABSTRACT

The goal of query performance prediction (QPP) is to automatically estimate the effectiveness of a search result for any given query, without relevance judgements. Post-retrieval features have been shown to be more effective for this task while being more expensive to compute than pre-retrieval features. Combining multiple post-retrieval features is even more effective, but state-of-the-art QPP methods are impossible to interpret because of the black-box nature of the employed machine learning models. However, interpretation is useful for understanding the predictive model and providing more answers about its behavior. Moreover, combining many post-retrieval features is not applicable to real-world cases, since the query running time is of utter importance. In this paper, we investigate a new framework for feature selection in which the trained model explains well the prediction. We introduce a step-wise (forward and backward) model selection approach where different subsets of query features are used to fit different models from which the system selects the best one. We evaluate our approach on four TREC collections using standard QPP features. We also develop two QPP features to address the issue of query-drift in the query feedback setting. We found that: (1) our model based on a limited number of selected features is as good as more complex models for QPP and better than non-selective models; (2) our model is more efficient than complex models during inference time since it requires fewer features; (3) the predictive model is readable and understandable; and (4) one of our new QPP features is consistently selected across different collections, proving its usefulness.

KEYWORDS

Information retrieval, Query performance prediction, QPP, feature selection, predictive model, selective model, linear regression, AIC criterion.

<https://doi.org/10.1145/3341105.3373904>

1 INTRODUCTION

In information retrieval (IR), query performance prediction (QPP) aims at automatically predicting the effectiveness of a system for a given query, without relevance judgments. QPP is useful to inform an IR system whether a query is difficult or not, allowing the system to process it differently. For example, in case of a difficult query, the system could either apply a specific automatic query reformulation or engage in an interactive session with the user in order to provide a better answer [8].

Query performance prediction uses query features that are extracted prior to running the query through the system (pre-retrieval) and/or from the initially-retrieved documents (post-retrieval). Intuitively, a good QPP feature should significantly correlate with the actual effectiveness of the IR system. Post-retrieval QPP features have been found to be more effective than pre-retrieval features, although they are much more expensive to calculate, as they need the IR system to run the query in order to make the prediction. While the first studies on QPP used single features [12, 13, 17, 24, 32], a more recent path is to combine various query features [17, 26, 30, 39, 45]. While combining multiple post-retrieval features improves accuracy, the method becomes applicable in real-world scenarios only if the number of features is limited to just a few, due to the increased computational time required for obtaining these features. A state-of-the-art method for combining QPP features, that, however, does not take into account these critical issues, is Raiber et al.'s [26]. Their system uses a pairwise learning-to-rank model that combines several existing QPP features. It uses a large number of features, deeming it unlikely to be implemented in real-world systems. Moreover, their method results in a non-interpretable model, due to the employed machine learning (ML) method.

Model interpretability refers to fairness, accountability, and transparency in machine learning [1, 6, 22], either for compulsory reasons (e.g. in the banking domain, the decision on accepting/rejecting mortgage) or because the end users want to understand the decisions taken by the ML model [2, 43]. Although fairness and transparency are not yet considered as requirements by search engine users, these features could become more popular with the growing awareness of the influence of search engines on social media users' opinion through the information these engines recommend to the users using sophisticated ML algorithms based on past queries (e.g. influence in political pools, fake news diffusion or unwanted ads). Moreover, public authorities may also require transparency in the near future for users' rights defense and privacy purposes [34]. Linear models (e.g. linear regression, SVM with the linear kernel) ensure this transparency, although some recent studies also explain how deep networks make decisions [7, 31].

Nevertheless, to our best knowledge, there is no previous effort to build interpretable models for QPP. In addition to the previously mentioned advantages, an interpretable QPP model would have another huge advantage over non-interpretable ones, considering our lack of understanding of query difficulty. Gaining additional insights from an interpretable model about the difficulty of a query would allow us to propose means for the system to overcome this difficulty.

In this paper, we propose a QPP approach that combines various features, yet, results in an interpretable and transparent model, so that we know the influence of each feature on the prediction. As a matter of fact, interpretability and transparency should not prevail in detriment of effectiveness. Hence, the model interpretability vs. effectiveness trade-off challenge corresponds to our first research question:

RQ1: *Can we design an interpretable and transparent model for QPP that is as effective as complex state-of-the-art black-box models?*

One could argue that the prediction performance improves as the model considers more and more features. However, this statement holds only up to some point [38], due to the curse of dimensionality. The curse of dimensionality is particularly problematic when few training examples are available, which happens for QPP evaluated on international reference collections (the only ones available for academic research). A smaller number of features reduces the model uncertainty and improves performance because fewer parameters have to be estimated in the model. Moreover, using more features increases the processing time of the model to the point where it becomes less applicable in real-world scenarios. These considerations are mentioned in present guidelines for IR practitioners [15, 36], being crucial for QPP due to the reliance on post-retrieval features which ensure effectiveness. On the other hand, feature selection poses the challenge of finding the appropriate criteria or strategies to select features in an optimal way. We tackle this problem in our second research question:

RQ2: *How selective can a white-box model be, without degrading prediction performance as compared to a non-selective one?*

To solve our two research questions, we develop a new feature selection model for QPP, whose main advantage compared to other related feature selection and QPP models is that it is parameter-free, making it applicable without tuning.

More precisely, our proposed framework is based on an iterative model selection procedure founded on linear regression, one of the most popular yet readable ML approaches. Linear regression is also known for its simplicity, following the Ockham's razor problem-solving principle that essentially states that "simpler solutions are more likely to be correct than complex ones" [4, 27]. However, rather than calculating the importance of each feature in one shot as linear regression does, we implement an iterative process which, at each step, adds a new feature to test or removes the least performing feature. Moreover, our approach uses a model selection criterion and is able to consider a large set of candidate features. While this approach has been used in machine learning [10], to our knowledge, it has never been considered in QPP.

In terms of performance prediction, we found that our model is consistently better than non-selective linear regression. We also compare our model with the penalized regression model called LASSO [17], which selects features by shrinking some feature coefficients to zero. The results reveal that our proposed model outperforms LASSO in most of the cases. Moreover, compared to LASSO, our method is parameter free and keeps fewer features, thus being less costly to use in the real world. Finally, we compare our model to that of Raiber *et al.* [26], which is the most recent approach that combines features and uses the same evaluation setting as ours. We found no statistically significant difference between our model and that of Raiber *et al.* [26], while our model is simpler, interpretable, and uses fewer features. We also investigate the inference times required by our model versus Raiber's *et al.* [26] model, during prediction of the query performance. The time evaluation shows that our model requires less time than Raiber's *et al.* [26]. We also found that our method consistently selects a specific group of features across collections for different folds and trials; one of these features is the *QFTERM* proposed in this work.

The remainder of this paper is structured as follows: Section 2 includes the related work. Section 3 presents our framework of step-wise model selection for query performance prediction. Section 4 presents the data collections, evaluation metrics, and experimental settings used for the evaluation part. Section 5 reports the evaluation of the proposed framework and the answers to our research questions. Finally, Section 6 concludes this work and presents some future directions.

2 RELATED WORK

The core objective of our paper is to define an optimal readable model that combines query features, selects the most important ones, and can explain the predicted values as opposed to black-boxes. Work related to our paper is about (a) query performance predictors and (b) methods to combine predictors.

Query performance predictors. Query performance prediction aims at automatically estimating the performance of a query without relevance judgment [41]. *Pre-retrieval predictors* were defined first, and they can be calculated prior to any search for the given query. Examples of pre-retrieval predictors are the Inverse Document Frequency (IDF) [35] or SynSet (the average number of query term senses) [25]. Further pre-retrieval QPPs have been defined in the literature, including the *CLARITY score* [12], the

query complexity [24], and the query scope [19]. However, the post-retrieval features have been shown to be more effective [18, 23, 32].

Post-retrieval predictors require to search through the documents to compute their scores and thus to predict the query difficulty. For example, Diaz [14] found that “low correlation between scores of topically close documents often implies a poor retrieval performance” and suggested a spatial analysis of retrieval scores for QPP. Indeed, several QPPs from the literature rely on document scores. Examples of post-retrieval predictors are: the agreement between the entire query results and the results obtained when using sub-queries [41], Query Feedback (QF) [45], Weighted Information Gain (WIG) [45], CLARITY [12], Normalized Query Commitment (NQC) [33], and score-distribution models [13]. Roitman *et al.* proposed an enhanced QPP estimator based on calibrating the retrieved document scores through learning document-level features [29]. Zamani *et al.* [42] proposed a NeuralQPP method based on integrating the retrieval scores, the term distribution, and the continuous representation of the top-retrieved documents by training a neural network with multiple weak supervision signals.

Combining query features. Several previous studies attempted to combine multiple query features or predictors. Bashir [3] employed a genetic algorithm to combine multiple pre-retrieval features and showed that it is more effective than using any single predictor. However, a straightforward way to combine features to predict a target value is by linear regression, and most of the related works combining features that way [17, 45].

Zhou and Croft [45] combined WIG and QF post-retrieval predictors in a linear way, showing that the combination improves performance. Shtok *et al.* [32] proposed a framework based on statistical decision theory to estimate the utility of a document ranking for QPP, considering four predictors (WIG, QF, CLARITY, NQC). They reached to the same conclusion as [45], namely that WIG and QF are worth combining. Collins-Thompson *et al.* [11] used a regression tree for QPP by combining features based on divergences between language or topic model representations, such as simplified clarity [20], query drift [40], clarity [12], expansion drift [44], and expansion clarity [11].

Hauff [16] used the absolute shrinkage and selection operator (LASSO) penalization when combining pre-retrieval features using linear regression. The LASSO penalization in linear regression aims at making the model sparse by removing features that roughly correspond to the smallest coefficients of the model. Even if LASSO exhibits proficiency in selecting the most important features, it relies on a parameter that has to be tuned optimally, for instance, using a cross-validation approach that can be time-consuming. Therefore, we rather opted for a stepwise approach with a straightforward implementation since it is parameter free.

Another closely related work is that of Raiber *et al.* [26]. They proposed a pairwise learning-to-rank model, that combines several existing pre-and post-retrieval QPP features through a two-stage training [21] process. In the first stage, the SVM-rank-based training combines several variants of individual post-retrieval features (e.g. NQC [33]) calculated for different hyper-parameter values and for several QPP features. In the second stage, another SVM-rank-based training combines all the QPPs from the first stage, while weighting them according to the weights learned in the first step. The main

reason why the two stages are needed is the (large) number of features (and/or feature variants) compared to the relatively small number of training examples. Although this framework [26] shows convincing performance, (a) it is computationally expensive at inference time since all the features are used in the final model, for both training and inference. Moreover, there are as many SVM-rank-based training procedures as the number of QPP features (first stage) plus an additional SVM training in the second stage and (b) the method is not parameter free, specifically in its adaption to sparse SVM. Since it is computationally expensive to extract the many QPP features required by Raiber’s framework [26] in the second stage, we rather develop a selective model which requires only a few features.

3 STEPWISE MODEL SELECTION FOR QPP

In this section, we describe our novel framework for selecting features to be used in the query performance predictive model. It employs an iterative process which relies on model selection theory in the context of linear regression and aims at combining various query features into a readable model. Not all the features are equally important and our model aims at optimizing the feature selection. Moreover, we use an iterative algorithm in order to select the best predictive model. While our model belongs to the group of models that have a solid mathematical background, we think it is worth providing the basics of linear regression for readers unfamiliar to ML, since our model is based on an adaption of it.

Linear model as a basis for predictor combination. Our model is founded on the theory of linear models [28]. A linear model links a response variable y to several predicting variables $x^j, j = 1, \dots, p$. In our context, x^j refers to a query feature and y refers to a performance measure representing the ground-truth effectiveness, that our model aims at predicting. We can model the performance measure according to query features and express it as the following linear model:

$$y_i = \beta_1 \cdot x_i^1 + \beta_2 \cdot x_i^2 + \dots + \beta_p \cdot x_i^p + \beta_0 + \varepsilon_i, \forall i \in \{1, \dots, n\}, \quad (1)$$

where i is the index of a query. A standard assumption for linear models is that $\varepsilon_i \sim \mathcal{N}(\mu, \sigma^2)$, expressing that the residuals contain only noise. Equivalently, the linear model can be expressed using vector and matrix notations:

$$Y = X\beta' + \varepsilon, \quad (2)$$

where Y and ε are n -dimensional vectors, β is a $(p+1)$ -dimensional row vector of weights, β' is the transposed (column) vector, and X is a $n \times (p+1)$ matrix containing n training examples. Our choice toward a linear model for QPP is driven by better interpretability and by the theoretical background we can rely on.

Parameter estimation. In the linear model, the unknown parameters β_j can be estimated using maximum likelihood. In statistics, the likelihood function expresses the way the parameters to be estimated are associated to the data actually observed. Maximizing the likelihood function consists in finding the values of the parameters that plausibly describe the observations. Using the previously defined setting, the likelihood function is given by:

$$L(\beta, \sigma^2) = (2\pi\sigma^2)^{-n/2} \exp \left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - X_i\beta)^2 \right). \quad (3)$$

Maximizing $L(\beta, \sigma^2)$ is equivalent to maximizing $\log L(\beta, \sigma^2)$ which is easier to handle and maximize. We thus want to maximize:

$$\log L(\beta, \sigma^2) = -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log(\sigma^2) - \frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - X_i \beta)^2. \quad (4)$$

The maximum (log-)likelihood is reached when the partial derivative according to each parameter is zero. This leads to the following estimators for β and σ^2 :

$$\begin{cases} \hat{\beta} = (X'X)^{-1}X'Y \\ \hat{\sigma}^2 = \frac{1}{n} (Y - X\hat{\beta})'(Y - X\hat{\beta}) \end{cases} \quad (5)$$

Once the parameters are estimated, the model can infer the fitted values for the performance measure y . Statistical testing (or equivalently confidence intervals) can be used to assess the significance of the parameters. The null hypothesis relies on the nullity of the parameters, i.e. on the uselessness of the associated features to predict the effectiveness of the system. It is interpreted through one p-value associated with each parameter. The p-value can be viewed as the probability to make an error when rejecting the null hypothesis. In other words, it is the probability of considering that the feature is not relevant to predict the performance measure, while the opposite is true. Based on the p-value, a feature could be excluded but that feature might be important if combined with others. Thus, we consider combining features rather than accepting or rejecting individual features.

Another crucial issue in linear modeling is variable selection, especially when dealing with a relatively large number of predicting variables. To address this issue, one has to go beyond elementary indicators, that mechanically increase with the number of variables. We chose to focus on the Akaike Information Criterion (AIC) [10] for assessing the goodness of fit of a linear model.

Akaike Information Criterion. This criterion is defined from the log-likelihood and uses a penalty to limit the number of parameters in the model. The function to minimize is defined as follows:

$$AIC = -2 \log L(\beta, \sigma^2) + 2 \cdot k \quad (6)$$

where k is the number of retained predictors.

AIC aims at selecting the most important features of the linear model. If a feature is kept, then a parameter is estimated to assess its influence in the linear model. As mentioned above, the parameters in a linear model can be estimated through the maximization of the (log-)likelihood. AIC is based on the opposite of the log-likelihood, thus requiring minimization. However, the penalty term added in AIC ($2 \cdot k$) depends on the number of parameters k to be estimated in the model (the same as the number of features included to predict the performance): the higher the number of parameters, the higher the penalty. Therefore, using AIC will ensure that the model does not use "too many" features and will keep only the most significant ones. Moreover, AIC can be used as a stopping rule for stepwise algorithms for model selection, as shown below.

An iterative stepwise selection algorithm. Our proposed framework is based on a model selection approach. This means different models are fitted with different subsets of features and the system selects the best model. This selection process is iterative.

More precisely, we employ a stepwise algorithm which mixes two strategies: forward and backward. Basically, the forward strategy starts from the model with no predictors and adds at each step the feature with the smallest p-value, thus possibly, the most useful because its coefficient in the linear model can be considered as significantly different from zero with a very low risk (quantified by the p-value) to be wrong. A stopping rule is based on a threshold for the p-value. On the other hand, the backward strategy starts from the complete model with all available features and, at each step, removes the feature with the highest p-value. In this case, the p-value can be interpreted as the probability to make an error if we consider that the coefficient of the feature is not null. The stepwise strategy combines the forward and backward strategies by attempting to remove a feature (applying backward) each time another one is added in the model (applying forward). This strategy is improved using AIC as a criterion instead of considering a threshold on the p-value. That is what we use in the following. Although this approach has been used in other ML tasks, it has never been used for QPP. We believe that AIC is worth investigating because of the cost of using multiple post-retrieval prediction features.

Our stepwise algorithm is an automatic model specification based on the AIC criterion. Starting from the complete model, the stepwise algorithm aims at decreasing the AIC at each step, using one of the two possible operations: (a) Remove one variable (obviously, this is the only option at the first step when starting from the complete model); (b) Add one variable removed in an earlier step.

The algorithm stops when the AIC criterion cannot be further decreased by removing or adding a variable. We note that our iterative feature selection algorithm is employed only at train time and it does not affect inference time.

Figure 1 illustrates the selection process when eight variables are used. The initial model consists of 8 variables. In Step 1, we cannot add any variable, the only possibility is to remove one. Eight models are built consisting each of 7 variables. Let us assume that the model without V_6 got the lowest AIC. The model without V_6 is the starting point for Step 2. In Step 2, we can either add V_6 or remove one of the other 7 variables. We thus test these 8 possible models and keep the one with the lowest AIC. Let us assume that removing V_3 is the best. We now have a model with 6 variables where V_3 and V_6 do not belong to. This model is used to start Step 3. We can either add one of these 2 variables or remove the third one (6 possibilities). Again, we test all the models. Let us assume that removing V_5 leads to the smallest AIC; we would keep the model with 5 variables for the next step.

4 DATA COLLECTIONS AND EVALUATION

Data collections. We considered four standard TREC collections from the ad-hoc task as follows: Robust, GOV2, WT10G, and ClueWeb12-B13. For Robust, there are approximately 500K newspaper articles. WT10G is composed of 1.6 million web/blog page documents. GOV2 includes 25 million web pages and ClueWeb12-B13 subset includes 50 million web pages. Table 1 summarizes a few features about the collections used for evaluation. The four TREC test collections also include topics. The "standard" format of a TREC topic statement comprises a topic ID, a title, a description, and a narrative. In our experiments, a query is composed of the topic

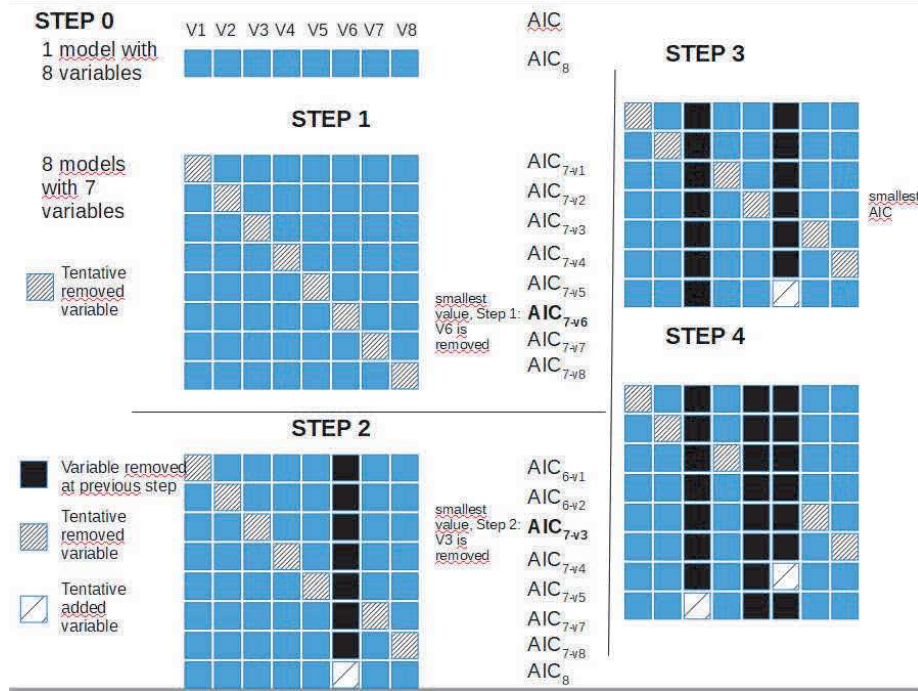


Figure 1: The four first steps of AIC stepwise model selection when starting with 8 variables. At each step, the best model is kept, either from removing a variable or adding one.

Table 1: Details of the collections used in the experiments.

Corpus	#Docs	Queries (Title only)
Robust	528,918	301-450, 601-700
WT10G	1,692,096	451-550
GOV2	25,205,179	701-850
Clueweb12B (CW12B)	52,343,021	201-300

title that contains two or three words representing the keywords a user could have used as a search query. Finally, the collections provide *qrels* (i.e. judged documents, relevant or non-relevant, for each query), which are used by the evaluation program *trec_eval*¹ in order to calculate the effectiveness of the IR system.

Query performance predictors. Several post-retrieval features have been proposed in the literature as QPP features and we reuse the main ones in this paper. We also propose two new post-retrieval features named *QFTERM* and *QFJSD* as variants of *QFDOC* [45]. Our proposed QPP features, as well as the state-of-the-art ones, are described as follows:

- *QFDOC* [45]: estimates the query feedback as the percentage of overlap at some rank between the returned document lists for the original query and the expanded query induced from the initially retrieved documents. It measures the query-drift.
- *QFTERM* (ours): we argue that the overlap at the document level, as computed by *QFDOC*, is too strict to estimate the discrepancy. We thus propose to relax this phenomenon at the term level, computing the percentage of overlap between the list of terms available in the top-retrieved documents for the original query and the term list for the expanded query. The higher the percentage of term overlaps,

the higher is the chance that the top-retrieved documents cover many relevant documents, since the expanded query is not too drifted away from the original query.

- *QFJSD* (ours): instead of computing the percentage of overlap at the document level between the top-retrieved documents for the original and the expanded queries, as *QFDOC* does, we rather measure the query feedback based on the similarity of term statistics between the two document lists, considering that a higher similarity value should correspond to a lower query-drift and a higher query performance. To estimate the similarity of term statistics between the two lists, we first build language models from the top-retrieved documents for the original query and the expanded query, respectively. Then, we apply the Jensen-Shannon divergence between the two language models to estimate how similar they are.
- *CLARITY* [12]: estimates the relative entropy between the relevance language models of the top retrieved documents and the corpus.
- *WIG* [45]: corresponds to the divergence between the mean of the top-retrieved document scores and the mean of the entire set of document scores.
- *NQC* [33]: is based on the standard deviation of the retrieved document scores.
- *UQC* [33]: is a variant of the *NQC* predictor, based on the standard deviation of the retrieved document scores without normalization.
- *SW1* [26]: is the ratio between the number of stop and non-stop words in each document, averaged over the top-retrieved documents for a query.

¹http://trec.nist.gov/trec_eval/

These QPP features can be estimated for different numbers of n -top-ranked feedback documents where n is a hyper-parameter. In this work, we consider 6 values of $n = \{10, 50, 100, 200, 500, \text{ and } 1000\}$. Moreover, to compute $QFDOC$, $QFTERM$, and $QFJSD$, we need to know the cutoff rank (termed $QFcut$) at which the percentage of overlap is computed for each hyper-parameter n . According to common practice [26, 29, 33], we define $QFcut = \min(50, n)$, i.e. the percentage of overlap is calculated for at most 50 documents.

Evaluation metrics. We use the Pearson and Spearman correlations between the predicted effectiveness value and the ground-truth effectiveness, as in previous works [5, 9, 17, 29, 30, 42]. To measure the ground-truth effectiveness of the system, we use AP (average precision) and NDCG (normalized discounted cumulative gain) since they are commonly adopted in related works [9, 17, 26, 29, 42].

Experimental settings. As a common practice in QPP evaluation [26, 30, 42], we randomly split the queries into two equally-sized sets and conduct two-fold cross-validation. We repeat these steps for 30 times and report the average results. Statistically significant differences of prediction performance are estimated using two-tailed paired t-test with Bonferroni correction ($p < 0.05$) computed over the 30 splits. Similar to previous works [12, 26, 32, 45], we chose the Language Modeling with Dirichlet smoothing and $\mu = 1000$ without query expansion (as implemented in Lemur Indri platform, using default parameters) to retrieve n documents for each query and to calculate the performance of the IR system (and thus, determine the results to be predicted in terms of AP or NDCG).

5 RESULTS AND DISCUSSIONS

Trade-off between sparsity and effectiveness. To answer our two research questions, we first study the correlation between the predicted and the ground-truth effectiveness (Table 2). When the Pearson correlation coefficient is employed, we measured the correlation between the predicted value and the actual value (the reference system is the language model with $\mu = 1000$). When the Spearman correlation coefficient is used, we measured the correlation between the ranks of the queries obtained when ordered by the predicted effectiveness and the actual effectiveness.

In Table 2, the first row is obtained by using the linear model (LM) with all the features in a single step (1S), a baseline that achieves readability, but not sparsity. The models listed on the subsequent rows use a two-stage approach as in [26], where the second step is either SVM-rank [26] (second row) or one of the models that ensure interpretability, as follows: LM (third row) refers to the linear model; LASSO (fourth row) is the LASSO selection [37] that ensure sparsity. Finally, the last row (AIC FS) corresponds to our proposed method of feature selection (FS) using the AIC criterion.

From the results displayed in Table 2, we observe that Raiber et al.’s [26] two-stage framework outperforms the one-stage linear regression baseline. Indeed, rows 2 to 5 indicate significant increases in correlation compared to the first baseline (see Δ in the table), irrespective of the model, the collection or the correlation measure being used. This result was expected considering the state-of-the-art results, but was worth checking².

More interestingly, we notice that, in Table 2, the three models we implemented (LM, LASSO and AIC FS) are generally (a) close to one another in terms of results, (b) without significant differences with respect to Raiber et al. [26], apart from a few cases (see \downarrow and \uparrow in the table). These results show that it is possible to use an interpretable model without decreasing effectiveness. From these three models, only LASSO and AIC FS ensure sparsity; we thus focus next on the results obtained by these models.

LASSO and AIC FS perform almost the same apart from a few cases where LASSO is slightly better (ROBUST AP and GOV2) or, conversely, where AIC is slightly better (WT10G). However, one important observation is that LASSO has a shrinkage parameter λ that needs to be tuned. In our experiments, we fitted it using 10-fold cross-validation, thus using the same data for both parameter fitting and model training, giving a clear advantage to LASSO in our experiments. In a preliminary set of experiments, we found that the transfer learning of λ (learning the parameter on a collection and using that value for another collection), that would make a fair comparison between LASSO and AIC FS, did not work at all for LASSO. Hence, λ has to be tuned separately for each collection, which is a clear drawback of LASSO compared to our parameter-free selection model based on the AIC criterion.

While one could have hypothesized that using all the features (the 1st and 3rd rows in Table 2) would have outperformed all the selective methods, this is not the case. This result is likely due to the curse of dimensionality [38]. This is an important result, as one crucial advantage of selective methods is to avoid calculating some features at inference time, that are not kept in the trained model. This also follows the Ockham’s principle “the law of brevity”, that is “more things should not be used than are necessary.” In the case of AIC, not only a limited number of features have to be computed, but the results are also better than using all features.

Sparsity and time complexity of the resulting models. As eight predictors are not that many, one may argue that this is a reasonable number which does not require feature selection. However, each feature can take valuable time in order to be computed, e.g. WIG requires 3.8 seconds on average for each query from the TREC-ROBUST collection on a machine having 8GB of RAM and processing on a single core. Thus, sparsity is an important issue when the number of predictors is large and/or costly to compute. Moreover, a smaller number of predictors leads to a simpler model and easier interpretation [46].

To investigate the sparsity of our model, we computed the number of features selected in the second stage for the different models. We start from eight QPP features in the second stage, two folds, and 30 trials. Since different numbers of features may be selected by a model across trials, we compute the average number of selected features. Table 3 reports the average number of features selected by Raiber et al. [26] and AIC FS models. Remarkably, our method, AIC, selects much fewer features than Raiber et al. [26] across all collections. The average number of features selected by AIC ranges from 2 to 5, while maintaining a similar performance to SVM-rank [26], that uses 7 to 8 features. This makes our method more applicable in real-world systems.

²The linear model with a single step (1S LM) does not work well apart from the Robust collection.

Table 2: Correlation of predicted and ground-truth effectiveness for scores with Pearson (r) and for ranks with Spearman (ρ). The first line uses a one step process. The others use a two steps process, the first step being SVM-rank based. Δ indicates statistically significant improvement over the Linear model (1S LM). \uparrow (resp. \downarrow) indicates statistically significant increase (resp. decrease) over the model of Raiber et al. [26], according to a paired t-test ($p < 0.05$) with a Bonferroni correction.

Method	ROBUST				WT10G				GOV2				CW12B			
	AP		NDCG		AP		NDCG		AP		NDCG		AP		NDCG	
	r	ρ	r	ρ	r	ρ	r	ρ	r	ρ	r	ρ	r	ρ	r	ρ
1S LM (One stage linear model)	.39	.41	.41	.43	.01	.01	.04	.05	.11	.23	.13	.25	.02	.01	-.00	-.02
SVM of Raiber et al. [26]	.48 Δ	.50 Δ	.49 Δ	.49 Δ	.27 Δ	.22 Δ	.27 Δ	.29 Δ	.45 Δ	.46 Δ	.45 Δ	.46 Δ	.41 Δ	.36 Δ	.37 Δ	.36 Δ
LM (Linear model)	.46 Δ	.47 Δ	.48 Δ	.48 Δ	.29 Δ	.29 Δ	.30 Δ	.36 Δ	.47 Δ	.48 Δ	.44 Δ	.47 Δ	.40 Δ	.37 Δ	.37 Δ	.34 Δ
LASSO	.48 Δ	.48 Δ	.49 Δ	.49 Δ	.21 Δ	.20 Δ	.23 Δ	.27 Δ	.48 Δ	.50 Δ	.45 Δ	.48 Δ	.38 Δ	.35 Δ	.32 Δ	.30 Δ
AIC FS (Feature selection)	.46 Δ	.47 Δ	.49 Δ	.49 Δ	.24 Δ	.24 Δ	.25 Δ	.29 Δ	.45 Δ	.47 Δ	.43 Δ	.45 Δ	.38 Δ	.35 Δ	.33 Δ	.30 Δ

Table 3: Average number of features selected in the models and average inference time required to predict the query performance using the same configurations as Table 2. We use a machine with an AMD Opteron 6262HE 1.6 GHz CPU, 8GB of RAM, a single thread.

	Method	ROBUST	WT10G	GOV2	CW12B
AP	Raiber’s SVM	8 (.042s)	6 (.036s)	8 (.033s)	6 (.023s)
	AIC FS	3 (.019s)	3 (.025s)	4 (.023s)	3 (.015s)
NDCG	Raiber’s SVM	8 (.041s)	7 (.037s)	8 (.017s)	7 (.020s)
	AIC FS	3 (.026s)	4 (.029s)	4 (.014s)	3 (.013s)

In Table 3, we also report the inference times required by our model versus Raiber’s et al. [26] model, respectively, during prediction of the query performance. The time evaluation shows that AIC FS takes less time than Raiber’s et al. [26] SVM to predict the query performance, since AIC FS requires fewer QPP features.

Most important features and interpretability. For a deeper understanding, we analyze the features selected by our AIC model, presenting the usefulness of each feature in Figure 2 (similar results are obtained for NDCG). We compute the number of times a feature was selected (SF) by our model, for two folds in 30 trials. Then, we compute the percentage of selecting that feature as $\frac{SF}{TF} \times 100$, where TF denotes the number of times a feature could be selected, i.e. the number of folds times the number of trials. This is done for all individual features per collection.

One interesting finding is that *QFTERM*, one of the feature we proposed in this paper, is consistently selected as an important feature. *WIG* is also consistently important. *UQC* and *SW1* are more collection dependant, the first one being more important for ROBUST, WT10G, and GOV2, while the second one being more important for GOV2 and CW12B.

The proposed AIC FS method is easily interpretable and enables one to understand the trained model and to know the impact of each query features in QPP. To illustrate this, we include below a model obtained on GOV2 collection:

$$QPP = .136 \cdot SW1 + .126 \cdot QFTERM + .101 \cdot UQC + .091 \cdot WIG + .034 \cdot CLARITY + .184 \quad (7)$$

We can see that the weight of the two most important features (*SW1* and *QFTERM*) are 30% higher than the two following ones (*UQC* and *WIG*); *CLARITY* is much less important.

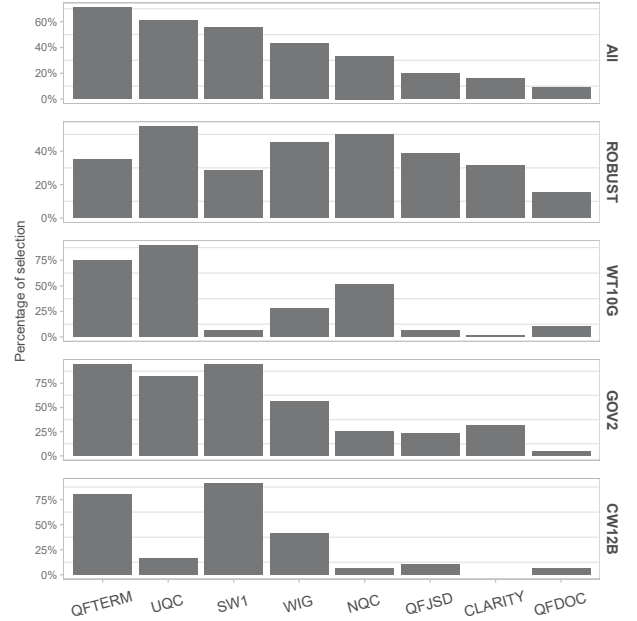


Figure 2: Distribution of features selected by our AIC FS model across collections for AP. Each percentage is computed as the number of times a feature is selected divided by the total number of times a feature could be selected.

6 CONCLUSION

In this paper, we promoted the use of model selection for combining query performance predictors. Our aim was twofold: (i) we wanted to contribute to shifting the effort from complex predictive models to simpler models and (ii) we wanted to develop a sparse and easy to interpret model. Indeed, we have shown that the trade-off between simplicity and effectiveness is in favour of our model. We showed that our selective framework achieves similar results in terms of correlation measures, while having the great advantage of using a limited number of features, and thus, being much cheaper for implementation in real-world systems. Our predictive approach based on the AIC selection strategy provides the best trade-off between the prediction accuracy and the number of features to be

computed. During inference, the features selected by our framework are the only ones which need to be calculated. That gives a key advantage to our framework compared to the literature. Moreover, our model is interpretable, which is very important at this stage of query difficulty research. These results open the path to a better understanding of system failures by analyzing the model deeper.

To answer this challenge, in future work, we will analyze the influence of each of the selected features individually as well as their cost/effectiveness trade-off. Because the current results reveal that our new version of QFDOC feature, namely QFTERM, is the most important QPP feature in the model across collections and performance metrics, it will be worth to continue defining new QPP features to try to further improve query difficulty prediction.

Acknowledgement. This work has been partially funded by the European Union's Horizon 2020 H2020-SU-SEC-2018 under the Grant Agreement n°833115 (PREVISION project).

REFERENCES

- [1] Philip Adler, Casey Falk, Sorelle A Friedler, Tionney Nix, Gabriel Rybeck, Carlos Scheidegger, Brandon Smith, and Suresh Venkatasubramanian. 2018. Auditing black-box models for indirect influence. *Knowledge and Information Systems* 54, 1 (2018), 95–122.
- [2] Athanasios Andreou, Giridhari Venkatadri, Oana Goga, Krishna Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Investigating ad transparency mechanisms in social media: A case study of Facebook's explanations. In *Network and Distributed System Security Symposium (NDSS)*.
- [3] Shariq Bashir. 2014. Combining pre-retrieval query quality predictors using genetic programming. *Applied Intelligence* 40, 3 (2014), 525–535.
- [4] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. 1987. Occam's razor. *Information processing letters* 24, 6 (1987), 377–380.
- [5] David Carmel and Elad Yom-Tov. 2010. Estimating the query difficulty for information retrieval. *Synthesis Lectures on Information Concepts, Retrieval, and Services* 2, 1 (2010), 1–89.
- [6] Carlos Castillo. 2019. Fairness and Transparency in Ranking. *SIGIR Forum* 52, 2 (Jan. 2019), 64–71.
- [7] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N. Balasubramanian. 2018. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 839–847.
- [8] Adrian-Gabriel Chifu, Sébastien Déjean, Stefano Mizzaro, and Josiane Mothe. 2017. Human-Based Query Difficulty Prediction. In *European Conference on Information Retrieval, ECIR*. Springer, 343–356.
- [9] Adrian-Gabriel Chifu, Léa Laporte, Josiane Mothe, and Md Zia Ullah. 2018. Query Performance Prediction Focused on Summarized Letor Features. In *International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. ACM, 1177–1180.
- [10] Gerda Claeskens, Nils Lid Hjort, et al. 2008. *Model selection and model averaging*. Vol. 330. Cambridge University Press Cambridge.
- [11] Kevyn Collins-Thompson and Paul N Bennett. 2010. Predicting query performance via classification. In *European Conf. on Information Retrieval*. 140–152.
- [12] Steve Cronen-Townsend, Yun Zhou, and W Bruce Croft. 2002. Predicting query performance. In *International ACM SIGIR conference on Research and development in information retrieval*. 299–306.
- [13] Ronan Cummins. 2014. Document score distribution models for query performance inference and prediction. *ACM Transactions on Information Systems (TOIS)* 32, 1 (2014), 2.
- [14] Fernando Diaz. 2007. Performance prediction using spatial autocorrelation. In *International ACM SIGIR conference on Research and development in information retrieval*. 583–590.
- [15] Isabelle Guyon and André Elisseeff. 2003. An Introduction to Variable and Feature Selection. *Journal of Machine Learning Research* 3 (March 2003), 1157–1182.
- [16] Claudia Hauff. 2010. Predicting the effectiveness of queries and retrieval systems. *SIGIR Forum* 44, 1 (2010), 88. <https://doi.org/10.1145/1842890.1842906>
- [17] Claudia Hauff, Leif Azzopardi, and Djoerd Hiemstra. 2009. The combination and evaluation of query performance prediction methods. In *European Conference on IR Research on Advances in Information Retrieval, ECIR*. Springer-Verlag, 301–312.
- [18] Claudia Hauff, Djoerd Hiemstra, and Francisca de Jong. 2008. A survey of pre-retrieval query performance predictors. In *ACM conference on Information and knowledge management*. 1419–1420.
- [19] Ben He and Iadh Ounis. 2004. Inferring query performance using pre-retrieval predictors. In *String processing and information retrieval*. Springer, 229–248.
- [20] Ben He and Iadh Ounis. 2006. Query performance prediction. *Information Systems* 31, 7 (2006), 585–594.
- [21] Thorsten Joachims. 2006. Training Linear SVMs in Linear Time. In *International ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '06)*. 217–226.
- [22] Bruno Lepri, Nuria Oliver, Emmanuel Letouze, Alex Pentland, and Patrick Vinck. 2018. Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology* 31, 4 (2018), 611–627.
- [23] Serge Molina, Josiane Mothe, Dorian Roques, Ludovic Tanguy, and Md Zia Ullah. 2017. IRIT-QFR: IRIT Query Feature Resource. In *Int. Conference and Labs of the Evaluation Forum, CLEF*. Springer, 69–81.
- [24] Josiane Mothe and Ludovic Tanguy. 2005. Linguistic features to predict query difficulty. In *ACM Conference on research and Development in Information Retrieval, SIGIR, Predicting query difficulty-methods and applications workshop*. 7–10.
- [25] Josiane Mothe and Ludovic Tanguy. 2007. Linguistic Analysis of Users' Queries: Towards an adaptive Information Retrieval System. In *IEEE Conference on Signal-Image Technologies and Internet-Based System*. 77–84.
- [26] Fiana Raiber and Oren Kurland. 2014. Query-performance Prediction: Setting the Expectations Straight. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 13–22.
- [27] Carl Edward Rasmussen and Zoubin Ghahramani. 2001. Occam's razor. In *Advances in neural information processing systems*. 294–300.
- [28] Alvin C Rencher and G Bruce Schaalje. 2008. *Linear models in statistics*. John Wiley & Sons.
- [29] Haggai Roitman. 2017. An Enhanced Approach to Query Performance Prediction Using Reference Lists. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 869–872.
- [30] Haggai Roitman. 2018. An Extended Query Performance Prediction Framework Utilizing Passage-Level Information. In *International ACM SIGIR Conference on Theory of Information Retrieval*. 35–42.
- [31] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *IEEE International Conference on Computer Vision (ICCV)*. 618–626.
- [32] Anna Shtok, Oren Kurland, and David Carmel. 2010. Using statistical decision theory and relevance models for query-performance prediction. In *Int. ACM SIGIR conference on Research and development in information retrieval*. 259–266.
- [33] Anna Shtok, Oren Kurland, David Carmel, Fiana Raiber, and Gad Markovits. 2012. Predicting query performance by query-drift estimation. *ACM Transactions on Information Systems (TOIS)* 30, 2 (2012), 11.
- [34] Xuemeng Song, Xiang Wang, Liqiang Nie, Xiangnan He, Zhumin Chen, and Wei Liu. 2018. A Personal Privacy Preserving Framework: I Let You Know Who Can See What. In *International ACM SIGIR Conference on Research & Development in Information Retrieval*. 295–304.
- [35] Karen Spärck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation* 28, 1 (1972), 11–21.
- [36] J.H. Stapleton. 2009. *Linear Statistical Models*.
- [37] Robert Tibshirani. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* (1996), 267–288.
- [38] G. V. Trunk. 1979. A Problem of Dimensionality: A Simple Example. *IEEE Trans. Pattern Anal. Mach. Intell.* 1, 3 (March 1979), 306–307.
- [39] Eduardo Vicente-López, Luis M de Campos, Juan M Fernández-Luna, and Juan F Huete. 2018. Predicting IR personalization performance using pre-retrieval query predictors. *J. of Intelligent Information Systems* 51, 3 (2018), 597–620.
- [40] Mattan Winaver, Oren Kurland, and Carmel Domshlak. 2007. Towards robust query expansion: model selection in the language modeling framework. In *International ACM SIGIR conference on Research and development in information retrieval*. 729–730.
- [41] Elad Yom-Tov, Shai Fine, David Carmel, and Adam Darlow. 2005. Learning to estimate query difficulty. In *International ACM SIGIR conference on Research and development in information retrieval*. 512–519.
- [42] Hamed Zamani, W. Bruce Croft, and J. Shane Culpepper. 2018. Neural Query Performance Prediction Using Weak Supervision from Multiple Signals. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 105–114.
- [43] Sihai Dave Zhao, Giovanni Parmigiani, Curtis Huttenhower, and Levi Waldron. 2014. Más-o-menos: a simple sign averaging method for discrimination in genomic data analysis. *Bioinformatics* 30, 21 (2014), 3062–3069.
- [44] Yun Zhou and W. Bruce Croft. 2006. Ranking Robustness: A Novel Framework to Predict Query Performance. In *Int ACM Conference on Information and Knowledge Management (CIKM '06)*. ACM, 567–574.
- [45] Yun Zhou and W. Bruce Croft. 2007. Query Performance Prediction in Web Search Environments. In *International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '07)*. ACM, 543–550.
- [46] Hui Zou and Trevor Hastie. 2005. Regularization and variable selection via the elastic net. *J. of the Royal Statistical Society: Series B (Statistical Methodology)* 67, 2 (2005), 301–320.