

**This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.**

**Author(s):** Pandey, Gaurav; Ren, Zhaochun; Wang, Shuaiqiang; Veijalainen, Jari; Rijke, Maarten de

**Title:** Linear feature extraction for ranking

**Year:** 2018

**Version:** Accepted version (Final draft)

**Copyright:** © Springer Science+Business Media, LLC, part of Springer Nature 2018.

**Rights:** In Copyright

**Rights url:** <http://rightsstatements.org/page/InC/1.0/?language=en>

**Please cite the original version:**

Pandey, G., Ren, Z., Wang, S., Veijalainen, J., & Rijke, M. D. (2018). Linear feature extraction for ranking. *Information Retrieval*, 21(6), 481-506. <https://doi.org/10.1007/s10791-018-9330-5>

## Linear Feature Extraction for Ranking

Gaurav Pandey · Zhaochun Ren · Shuaiqiang Wang · Jari Veijalainen · Maarten de Rijke

Received: date / Accepted: date

**Abstract** We address the feature extraction problem for document ranking in information retrieval. We then propose LifeRank, a Linear feature extraction algorithm for Ranking. In LifeRank, we regard each document collection for ranking as a matrix, referred to as the *original matrix*. We optimize a *transformation matrix*, so that a new matrix (dataset) can be generated as the product of the original matrix and a transformation matrix. The transformation matrix projects high-dimensional document vectors into lower dimensions. Theoretically, there could be very large transformation matrices, each leading to a new generated matrix. In LifeRank, we produce a transformation matrix so that the generated new matrix can match the learning to rank problem. Extensive experiments on benchmark datasets show the performance gains of LifeRank in comparison with state-of-the-art feature selection algorithms.

**Keywords** Feature extraction · Dimension reduction · Learning to rank · Information retrieval

---

Gaurav Pandey  
University of Jyväskylä  
E-mail: gaurav.g.pandey@student.jyu.fi

Zhaochun Ren  
JD.com  
E-mail: renzhaochun@jd.com

Shuaiqiang Wang  
University of Jyväskylä  
E-mail: shuaiqiang.wang@jyu.fi

Jari Veijalainen  
University of Jyväskylä  
E-mail: jari.veijalainen@jyu.fi

Maarten de Rijke  
University of Amsterdam  
E-mail: derijke@uva.nl

## 1 Introduction

Document ranking is an essential component of information retrieval systems and web search engines. Recently, machine learning-based ranking techniques, referred to as “learning to rank,” have given rise to an active and growing research area, both in the information retrieval and machine learning communities [10, 15, 21, 39, 59]. A large number of learning to rank algorithms have been proposed, which incorporate more and more useful features, aiming to improve the performance of the ranking algorithms [33]. In a supervised setting, they first collect a set of training data, which includes a set of queries, each associated with a list of documents labeled by relevance degrees; with the training dataset, they train a ranking model that can order unseen documents according to their degree of relevance [20]. In this situation, dimension reduction inevitably becomes an important issue [16].

Firstly, dimension reduction can enhance the accuracy for many machine learning problems, including learning to rank. With dimension reduction techniques, a small set of more discriminative and less redundant features can be selected or generated for learning. Thus, better results could be achieved, as overfitting becomes less likely [38]. Also, the generalization ability of machine learning models could depend on the radius of training data points, which may decrease when the number of features decreases [5, 16, 56, 57].

Secondly, large number of features leads to high complexity in most learning to rank algorithms. Therefore, dimension reduction often leads to significant improvements in training and prediction efficiency, while maintaining, or having a limited negative impact on, accuracy. With accuracy being the primary metric, efficiency has also emerged as a crucial issue for evaluating learning to rank algorithms [10, 11, 55]. Training datasets and ranking features continue to expand, so as to obtain more accurate models. Furthermore, as a consequence of the dynamic character of the Web, ranking models need to be re-learned repeatedly, and the interval between re-learning procedures decreases sharply [33]. With dimension reduction techniques, fewer features are used, resulting in more efficient training and prediction.

Generally, there are two types of dimension reduction algorithms: feature *selection* and feature *extraction*. The former aims to select a subset of the original features for learning, while the latter attempts to generate a small set of new features from the original features [5, 35, 58]. Recently, feature selection for ranking has been investigated intensively [16, 17, 26, 29, 37, 40, 60]. To the best of our knowledge, the advantages of feature *extraction* have not yet been explored in learning to rank.

In this study, we address the *feature extraction problem for learning to rank*. In comparison with feature selection, the feature extraction problem has a much larger search space. For example, given  $n$  original features, feature selection selects a subset of features of size  $k$  (where  $k < n$ ) for learning. Here, for a particular value of  $k$ , the search space of the problem contains  $\binom{n}{k}$  possible solutions. The full search space that can include any number of features (i.e., all values of  $k$  in range 1 to  $n$ ), would lead to  $2^n - 1$  solutions. In comparison, for linear feature extraction, each extracted feature is a linear combination of original  $n$  features. Since the coefficient associated with each original feature can be any real number, the search space becomes infinite. The search space of non-linear feature extraction would be even larger, as it also includes

solutions involving non-linear combinations of features (e.g. polynomial combinations). Hence, with a larger search space, feature extraction has a greater possibility to achieve better performance than feature selection.

To address the problem of linear feature extraction for learning to rank, we propose LifeRank, a Linear feature extraction algorithm for Ranking. LifeRank regards each dataset for training, validation or test as a matrix, referred to as an *original matrix*, where each row vector represents a document with a set of features. With a given original matrix for training  $\mathbf{X}$ , LifeRank attempts to discover a transformation matrix  $\mathbf{T}$ , so that a new matrix (dataset)  $\mathbf{X}'$  can be generated as the product of the original matrix and a transformation matrix, i.e.,  $\mathbf{X}' = \mathbf{X}\mathbf{T}$ . Thus  $\mathbf{T}$  projects high-dimensional document vectors in  $\mathbf{X}$  into lower-dimensional ones in  $\mathbf{X}'$ . Theoretically, there could be a very large number of possible transformation matrices, each leading to a new generated matrix. LifeRank attempts to discover a transformation matrix to transform the original matrix (dataset) into a low-rank one for dimension reduction, on which learning to rank algorithms can achieve optimum results in comparison with other dimension-reduced matrices.

Our problem formulation is similar to principal component analysis (PCA) [23], and thus our algorithm LifeRank can be understood from the perspective of PCA. PCA is one of the most popular dimension reduction techniques in machine learning. When PCA is performed using singular valued decomposition (SVD) [28], the given matrix  $\mathbf{X}$  can be approximately decomposed into three low-rank matrices  $\mathbf{X} \approx \mathbf{P}\mathbf{\Sigma}\mathbf{Q}^T$ . Here,  $\mathbf{\Sigma}$  is composed of the singular values of  $\mathbf{X}$ ,  $\mathbf{P}$  and  $\mathbf{Q}$  are composed of the left and right singular vectors of  $\mathbf{X}$  respectively, and  $\mathbf{P}^T\mathbf{P} = \mathbf{Q}^T\mathbf{Q} = \mathbf{I}$  is equal to the identity matrix. Thus a new matrix  $\mathbf{X}' = \mathbf{P}\mathbf{\Sigma} \approx \mathbf{X}\mathbf{Q}$ . However, it should be noted that while PCA calculates  $\mathbf{X}'$  as an approximation of  $\mathbf{X}$ , in LifeRank  $\mathbf{X}$  is transformed to  $\mathbf{X}'$  using a transformation matrix.

In LifeRank, we formulate the learning to rank task by using a classical pairwise loss function. A pairwise loss function is used because such functions are fundamental, straightforward and intuitive for ranking. Besides, pairwise loss functions are consistent with the assumption that the labels of documents to rank lie in a rank-differentiable probability space [27], and they are upper bounds of measure-based ranking errors [12]. In the generated matrix, the column vectors represent the features. Since optimization over orthogonal features is beneficial to many machine learning problems [48, 49], we utilize the Lagrange multipliers method [1, 4] to impose orthonormality constraints on the column (feature) vectors of the transformed matrix, and then use gradient descent for optimization. With the transformation matrix  $\mathbf{T}$ , the training, validation and test datasets can be directly generated with matrix product.

Note that (1) LifeRank generalizes feature selection algorithms for the learning to rank task. Feature selection can be regarded as optimizing a transformation matrix  $\mathbf{T}$  so that the column vectors of  $\mathbf{T}$  meet the orthonormality constraints and each element in  $\mathbf{T}$  can only be either 0 or 1. (2) Although some deep learning-based ranking algorithms [47] also aim to generate a set of features for ranking, our problem is completely different: we try to construct our features based on some predesigned ranking features like term frequency (TF) and inverse document frequency (IDF), which have been comprehensively used in conventional learning to rank algorithms like Ranking SVM [9, 21] and RankBoost [15]. Deep learning-based algorithms, however, try to

build features based on word-level features in a corpus that differ substantially from conventional ranking features.

Our main contributions are as follows: (1) We address the feature extraction problem for learning to rank. Feature extraction is a category of comprehensively used dimension reduction techniques in many machine learning problems for performance gains in accuracy and efficiency, but to the best of our knowledge, feature extraction and its advantages have not been explored in learning to rank yet. (2) We propose LifeRank, a linear feature extraction algorithm that generates datasets to be utilized by the learning to rank task. (3) We perform extensive experiments on benchmark datasets and present the performance gains of LifeRank in comparison with the state-of-the-art feature selection algorithms.

The remainder of the paper is organized as follows. Section 2 reviews related work; Section 3 defines the feature extraction problem for ranking; Section 4 proposes LifeRank, a gradient descent-based algorithm. Section 5 introduces our experimental setup. Section 6 reports the experimental results, and Section 7 concludes the paper.

## 2 Related work

We discuss three types of related work: learning to rank, feature selection for ranking, and feature extraction for ranking.

### 2.1 Learning to rank for information retrieval

Learning to rank has received increased attention from both the machine learning and information retrieval community. While there is growing interest in *online* learning to rank [46] and in *counterfactual* learning to rank from online data [22], the bulk of the work on learning to rank concerns *offline* learning to rank, where explicit human annotations are used to label query, document pairs. Offline learning to rank is the focus of this paper. Given its effectiveness, many algorithms have been proposed, which mainly fall into three categories [11, 32]: pointwise, pairwise, and listwise.

Pointwise approaches, such as Pranking [14], McRank [31] and Subset Ranking [13], view each document in the training dataset as a learning instance, and utilize a classification or regression technique to predict the relevance categories or numerical/ordinal relevance scores for unlabeled data. Pairwise approaches, such as Ranking SVM [9, 21], RankBoost [15], RankNet [6], FRank [52], LambdaRank [7], and BoltzRank [54], regard a pair of documents as a learning instance, and try to learn a binary classifier that can predict the more relevant document to the given query from each pair of documents. Then the ranked lists of documents can be aggregated based on the pairwise preferences of the documents. Listwise approaches, such as ListNet [10], SVM-MAP [61], NDCGBoost [53], take the entire ranked list of documents as a learning instance, and attempt to construct a ranking model that can directly predict the full rankings of the documents. Recently, some hybrid algorithms have been proposed, such as FocusedRank [39], MixRank [8], targeting improvements in learning accuracy, efficiency, or both. More algorithms are surveyed in [11, 32, 33].

With the incorporation of more and more useful features for performance gains, dimension reduction inevitably becomes an important issue in the ranking problem [16]. With effective dimension reduction techniques, not only the efficiency of the algorithms could be improved, but also accuracy could be enhanced as a result of using more discriminative features with less redundancy and noise. Furthermore, the generalization of the ranking model can also be increased as a result of using fewer features [16].

## 2.2 Feature selection for ranking

Recently, considerable efforts have been made on feature selection for ranking. Geng et al [16] present GAS, one of the first attempts to incorporate the importance and similarity of features for ranking. In particular, it evaluates the importance of features with ranking metrics like MAP [3] and NDCG [19], and estimates the similarity between features using agreement between rankings, e.g., with Kendall  $\tau$  correlation coefficient [24]. Then it greedily selects a subset of features with maximum total importance scores and minimum total similarity scores. Metzler [34] proposes a greedy feature selection algorithm to be used within the Markov random field model for information retrieval. The model automatically generates models that are more effective than, or as effective as, models created by carefully selecting the features manually. Pan et al [40] investigate a boosted regression trees-based feature selection algorithm. It evaluates the importance of the features based on boosted trees. Then it selects features by maximizing the discounted importance of the features, where the importance of each feature is discounted by feature similarity. Yu et al [60] propose RankWrapper and RankSelect, two feature weighting and selection algorithms for learning to rank. They utilize ranking distances of nearest data points in order to identify the key features for ranking, demonstrating significant efficiency gains in comparison with GAS.

Gupta and Rosso [17] present a Kullback-Leibler (KL) divergence-based divergence metric, and select a subset of features for ranking based on features' expected divergence over the relevance classes and the importance of features. Lai et al [26] propose a joint convex optimization formulation for minimizing ranking errors while simultaneously conducting feature selection. This optimization formulation provides a flexible framework in which various importance measures and similarity measures of the features can easily be incorporated. Naini and Altingövde [37] adopt three greedy diversification strategies, maximal marginal relevance, MaxSum dispersion and modern portfolio theory, to the problem of feature selection for ranking. Laporte et al [29] propose a general framework for feature selection in learning to rank based on support vector machine (SVM); they investigate both classical convex regularizations (such as  $L1$  and weighted  $L1$ ) and non-convex regularization terms (such as log penalty, Minimax Concave Penalty (MCP) and  $Lp$  pseudo norm with  $p < 1$ ). Furthermore, they provided an accelerated proximal approach for solving the convex problems and a re-weighted  $L1$  scheme to address the non-convex regularizations.

All of these algorithms are meant to address *feature selection* for ranking. To the best of our knowledge, there is no work targeting *feature extraction* for ranking.

### 2.3 Feature extraction techniques

Feature extraction has been extensively used in various machine learning scenarios for performance gains in terms of accuracy and efficiency. Given its effectiveness, many approaches have been proposed, which are either linear or non-linear algorithms.

The main linear technique for feature extraction is principal component analysis (PCA) [23], which performs a linear mapping of high-dimensional data into a lower-dimensional space in such a way that the variance of the data in the low-dimensional representation is maximized. Canonical-correlation analysis (CCA) [18] is another popular linear feature extraction algorithm, which attempts to discover linear combinations of the original features that have maximal correlation with each other. In addition, several probabilistic algorithms, including probabilistic PCA [51], probabilistic CCA [2] and probabilistic partial CCA [36], have been proposed, where a set of latent variables are introduced for probabilistically interpreting these models.

Non-linear feature extraction algorithms can combine the original features to generate a set of features in a non-linear way. For example, the locally linear embedding (LLE) method [44] learns the global structure of non-linear manifolds to yield low-dimensional, neighborhood-preserving embeddings of high-dimensional inputs. Isomap [50] is capable of discovering the non-linear degrees of freedom that underly complex natural observations. It can efficiently compute a globally optimal solution and can be guaranteed to converge asymptotically to the true structure. Besides, some kernel techniques have been proposed to transform linear feature extraction algorithms into nonlinear ones. For example, kernel PCA [45] is a non-linear form of principal component analysis (PCA), which can efficiently compute principal components in high-dimensional feature spaces through the use of integral operator kernel functions.

Although feature extraction techniques have been extensively investigated and shown to demonstrate promising performance gains, to the best of our knowledge, they have not been explored yet in the context of the ranking problem.

## 3 Problem statement

### 3.1 Learning to rank for information retrieval

Let  $\mathcal{X}$  be a collection of documents, each represented by a vector of feature values. In information retrieval systems, given a query  $q$ , a list of documents from  $\mathcal{X}$  is returned as search results, where the documents are ranked according to their estimated relevance to  $q$ . Given a query  $q$ , the ground truth, i.e., relevance judgments of documents with respect to  $q$  (produced by human experts) is defined as a function  $rel : \mathcal{X} \rightarrow \mathbb{N}_0$ , where  $\mathbb{N}_0$  is the set of natural numbers (including 0).

Let  $f : \mathcal{X} \rightarrow \mathbb{R}$  be a ranking function assigning real valued relevance scores to documents. The goodness of ranking functions can be evaluated by a measure  $s$ , such as precision at  $n$  ( $P@n$ ), mean average precision ( $MAP$ ) [3], or normalized discounted cumulative gain ( $NDCG@n$ ) [19].

**Definition 1 (Learning to rank)** *Given a training dataset  $\mathcal{X}$  and an evaluation measure  $s$ , the problem of learning to rank is to learn a ranking function  $f$  from  $\mathcal{X}$  such that  $s(f)$  is maximized.*

### 3.2 Dimension reduction for ranking

In learning to rank, each dataset  $\mathcal{X}$  can be regarded as a document matrix  $\mathbf{X}_{m \times n}$  with  $m$  rows (documents) and  $n$  columns (features). In particular,  $\mathbf{x}_i$  is the  $i$ -th row of  $\mathbf{X}$ , and  $\mathbf{x}_i^\top$  is a  $n$ -dimensional (column) vector that represents a document with  $n$  features. Let  $g : \mathbb{R}^n \rightarrow \mathbb{R}^k$  ( $k \leq n$ ) be a mapping that projects an  $n$ -dimensional vector space into a  $k$ -dimensional space. Let  $L(\cdot)$  be the loss function for the learning to rank task. Our problem is to discover a mapping function  $g$  such that the obtained dataset  $\mathbf{X}' = g(\mathbf{X})$  minimizes the loss function.

**Definition 2 (Dimension reduction for ranking)** *Let  $\mathbf{X}_{m \times n}$  be a document matrix with  $m$  columns and  $n$  rows, where each column  $\mathbf{x}_i^\top$  is a  $n$ -dimensional vector, representing a document with  $n$  features. Let  $\mathcal{G}$  be the set of all possible mapping functions, where each element  $g : \mathbb{R}^n \rightarrow \mathbb{R}^k$  ( $k \leq n$ ) is used to project an  $n$ -dimensional vector space into a  $k$ -dimensional space. The dimension reduction for the learning to rank task tries to discover an optimum mapping function  $g^* \in \mathcal{G}$  such that:*

$$\arg \min_{g \in \mathcal{G}} L(g(\mathbf{X})), \quad (1)$$

where  $L(\cdot)$  is the loss function for the learning to rank task. Then the new dataset can be generated with  $g^*(\mathbf{X})$ .

In this paper, we consider linear feature extraction for learning to rank as it is the simplest and most straightforward feature extraction technique in machine learning. Here, each generated feature is a linear combination of the original features. It utilizes a transformation matrix  $\mathbf{T}$  to achieve the effectiveness of the mapping function, aiming to discover an optimal matrix  $\mathbf{T}$  such that the obtained dataset  $\mathbf{X}' = \mathbf{X}\mathbf{T}$  results in a minimal value of the loss function.

The problem can be understood from the perspective of PCA [23]. Using PCA, the given matrix  $\mathbf{X}$  can be approximately decomposed into three lower-rank matrices:

$$\mathbf{X} \approx \mathbf{P}\mathbf{\Sigma}\mathbf{Q}^\top, \quad (2)$$

where  $\mathbf{\Sigma}$  is composed of the singular values of  $\mathbf{X}$ ,  $\mathbf{P}$  and  $\mathbf{Q}$  are composed of the left and right singular vectors of  $\mathbf{X}$  respectively, and  $\mathbf{P}^\top\mathbf{P} = \mathbf{Q}^\top\mathbf{Q} = \mathbf{I}$  (the identity matrix). Thus, a new matrix  $\mathbf{X}'$  can be generated as follows:

$$\mathbf{X}' = \mathbf{P}\mathbf{\Sigma} \approx \mathbf{X}\mathbf{Q}. \quad (3)$$

The role of the transformation matrix  $\mathbf{T}$  in LifeRank is very similar to the right singular matrix  $\mathbf{Q}$  in PCA, where  $\mathbf{Q}$  maps the document vectors to another space spanned by the columns of  $\mathbf{Q}$  before transforming them through  $\mathbf{\Sigma}$  and going back through  $\mathbf{P}$ . Hence, in LifeRank we consider the orthonormality constraints of  $\mathbf{T}$  in our optimization process, i.e.,  $\mathbf{T}^\top\mathbf{T} = \mathbf{I}$ .

**Definition 3 (Constrained linear feature extraction for ranking)** Let  $\mathbf{X}_{m \times n}$  be a document matrix, where the transpose of each row, i.e.,  $\mathbf{x}_i^\top = \mathbf{d}_i$  is a  $n$ -dimensional vector, representing a document with  $n$  features. Linear feature extraction for ranking aims to optimize a transformation matrix  $\mathbf{T}_{n \times k}$  by solving the following optimization problem, so that a new document matrix  $\mathbf{X}'_{m \times k} = \mathbf{X}_{m \times n} \mathbf{T}_{n \times k}$  can be generated, where each document vector  $\mathbf{d}_i$  can be projected into  $k$ -dimensional vector  $\mathbf{d}'_i = \mathbf{T}^\top \mathbf{d}_i$ :

$$\arg \min_{\mathbf{T}} L(\mathbf{X}\mathbf{T}) \text{ such that } \mathbf{T}^\top \mathbf{T} = \mathbf{I}, \quad (4)$$

where  $L(\cdot)$  is the loss function for the learning to rank task.

Based on the optimized mapping function  $g$ , the new dataset can be generated by taking the product of the original matrix and the transformation matrix, i.e.,  $\mathbf{X}' = \mathbf{X}\mathbf{T}$ .

We have used the example of PCA to help us explain the mechanism of LifeRank. However, it should be noted that in PCA  $\mathbf{X}'$  is calculated as an approximation of  $\mathbf{X}$ , whereas in LifeRank we generate a transformed representation of the initial matrix, in order to achieve a better ranking performance. Hence, unlike PCA,  $\mathbf{X}'$  as computed in Definition 3 is not an approximation of  $\mathbf{X}$ , but a transformation.

## 4 The LifeRank algorithm

Given a high-dimensional dataset  $\mathcal{X}$ , LifeRank generates a new low-dimensional dataset  $\mathcal{X}'$  in two phases. In the first phase, LifeRank first preprocesses the training dataset  $\mathcal{X}$  into an original matrix  $\mathbf{X}$ . Then LifeRank optimizes the transformation matrix  $\mathbf{T}$  for  $\mathbf{X}$  according to the loss function in Equation 4. In the second phase, LifeRank generates low-dimensional training, validation and test matrices with the projection of  $\mathbf{T}$ . Then LifeRank constructs new datasets based on the low-dimensional data matrices.

### 4.1 Phase I: Generation of the transformation matrix

In this study, we utilize a classic pairwise learning to rank loss function to implement the function  $L(\cdot)$  in Definition 3. Pairwise loss functions are chosen because apart from being relatively simple and straightforward, they are also intuitive choices for ranking. Besides, with the assumption that the labels of documents to rank lie in a rank-differentiable probability space, pairwise loss functions are consistent [27] and provide upper bounds for measure-based ranking errors like NDCG [12]. Thus, minimizing a pairwise loss function will maximize the ranking measures [27].

First of all, the training dataset  $\mathcal{X}$  is preprocessed into an original matrix  $\mathbf{X}$  and other information  $\mathbf{I}_X$  consisting of identities of the documents and queries, relevance labels, etc. Let  $D = \{\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_m\}$  be the set of columns (document vectors) in the matrix  $\mathbf{X}_{m \times n}^\top$ . We regard each pair of vectors  $(\mathbf{d}_i, \mathbf{d}_j) \in D \times D$  as an instance, and the label  $y_{i,j} \in \{+1, -1\}$  indicates whether the relevance of the  $i$ -th document is higher or lower than the  $j$ -th document, corresponding to the given query. Let

$\{\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_k\}$  be the column vectors of  $\mathbf{T}$ . We try to discover a  $k$ -dimensional vector of weights  $\mathbf{w}$  such that:

$$\begin{aligned} \arg \min_{\mathbf{T}, \mathbf{w}, b} \quad & \sum_{\forall (\mathbf{d}_i, \mathbf{d}_j), i \neq j} \log \left( 1 + e^{-y_{i,j} (\mathbf{w}^\top \mathbf{T}^\top (\mathbf{d}_i - \mathbf{d}_j) + b)} \right) + \frac{\lambda}{2} \|\mathbf{w}\|^2 \\ \text{such that } \mathbf{t}_i^\top \mathbf{t}_j = & \begin{cases} 1, & i = j \\ 0, & i \neq j \end{cases}, \forall i, j = 1, 2, \dots, k, \end{aligned} \quad (5)$$

where the first part calculates the log loss of the ranking accuracy, the second part is the  $l_2$  norm of the parameters for regularization, and  $\lambda$  is the coefficient of the regularization term for trade-off.

We optimize the constrained loss function based on the Lagrange multipliers method [1, 4] in Equation 5. Let

$$\begin{aligned} \mathcal{L}(\mathbf{T}, \mathbf{w}, b, \mathbf{A}) = \quad & \sum_{\forall (\mathbf{d}_i, \mathbf{d}_j), i \neq j} \log \left( 1 + e^{-y_{i,j} (\mathbf{w}^\top \mathbf{T}^\top (\mathbf{d}_i - \mathbf{d}_j) + b)} \right) + \\ & \frac{\lambda}{2} \|\mathbf{w}\|^2 + \sum_{i,j=1,\dots,k \wedge i \neq j} \alpha_{i,j} \mathbf{t}_i^\top \mathbf{t}_j + \sum_{i=1}^k \alpha_{i,i} (1 - \mathbf{t}_i^\top \mathbf{t}_i), \end{aligned} \quad (6)$$

where  $\mathbf{A}$  is a matrix with  $k$  columns and  $k$  rows, and elements  $\alpha_{i,j}$ . Then, the optimum  $\mathbf{T}$ ,  $\mathbf{w}$  and  $b$  for minimizing  $\mathcal{L}$  are the exact results of Equation 5.

In Phase I, we utilize gradient descent to generate the training dataset and the transformation matrix. Initially, we assign all 1s to the vector  $\mathbf{w}_{k \times 1}$  so that all of the generated features in the ranking model have the same initial weight. We initialize the transformation matrix  $\mathbf{T}$  in a random manner, following work on matrix generalization problems like matrix factorization-based collaborative filtering [25]. After initialization, the weight vector  $\mathbf{w}$  and the factorized matrix can be updated iteratively with gradient descent until reaching convergence or the maximum number of iterations with the given learning rate. The gradients of the function  $\mathcal{L}$  with respect to the variables are calculated as follows:

$$\begin{aligned} \nabla_{\mathbf{w}} \mathcal{L} &= \sum_{\forall (\mathbf{d}_i, \mathbf{d}_j), i \neq j} \frac{-y_{i,j} \mathbf{T}^\top (\mathbf{d}_i - \mathbf{d}_j)}{1 + e^{y_{i,j} (\mathbf{w}^\top \mathbf{T}^\top (\mathbf{d}_i - \mathbf{d}_j) + b)}} + \lambda \mathbf{w} \\ \nabla_{\mathbf{t}_l} \mathcal{L} &= \sum_{\forall (\mathbf{d}_i, \mathbf{d}_j), i \neq j} \frac{-y_{i,j} w_l (\mathbf{d}_i - \mathbf{d}_j)}{1 + e^{y_{i,j} (\mathbf{w}^\top \mathbf{T}^\top (\mathbf{d}_i - \mathbf{d}_j) + b)}} + \\ & \quad \sum_{i \neq l} (\alpha_{l,i} + \alpha_{i,l}) \mathbf{t}_i - 2\alpha_{l,l} \mathbf{t}_l, \quad l = 1, \dots, k \\ \frac{\partial \mathcal{L}}{\partial b} &= \sum_{\forall (\mathbf{d}_i, \mathbf{d}_j), i \neq j} \frac{-y_{i,j}}{1 + e^{y_{i,j} (\mathbf{w}^\top \mathbf{T}^\top (\mathbf{d}_i - \mathbf{d}_j) + b)}} \\ \frac{\partial \mathcal{L}}{\partial \alpha_{i,j}} &= \begin{cases} \mathbf{t}_i^\top \mathbf{t}_j, & i \neq j \\ 1 - \mathbf{t}_i^\top \mathbf{t}_j, & i = j \end{cases} \quad \forall i, j = 1, \dots, k, \end{aligned} \quad (7)$$

where  $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_n$  are the column vectors of  $\mathbf{T}$ . Since gradient descent generally does not work with Lagrange multipliers, we use the basic differential multiplier method (BDMM) [41] for optimization, where the sign inversion for  $\alpha$  in Equation 8 makes the optimization stable. Given a learning rate  $\eta$ , the update formulas of the gradient descent method are:

$$\begin{aligned}
 \mathbf{w} &\leftarrow \mathbf{w} - \eta \nabla_{\mathbf{w}} \mathcal{L} \\
 \mathbf{t}_l &\leftarrow \mathbf{t}_l - \eta \nabla_{\mathbf{t}_l} \mathcal{L}, \text{ for } l = 1, \dots, k \\
 b &\leftarrow b - \eta \frac{\partial \mathcal{L}}{\partial b} \\
 \alpha_{i,j} &\leftarrow \alpha_{i,j} + \eta \frac{\partial \mathcal{L}}{\partial \alpha_{i,j}}, \text{ for } i, j = 1, \dots, k.
 \end{aligned} \tag{8}$$

#### 4.2 Phase II: Generation of low-rank datasets

In LifeRank, Phase II generates all of the datasets for learning to rank, including the training, validation and test datasets. According to Definition 3, for each original matrix  $\mathbf{X}$ , the generated matrix  $\mathbf{X}'$  can be obtained as a product of the original dataset  $\mathbf{X}$  and the transformation matrix  $\mathbf{T}$ , formally  $\mathbf{X}' = \mathbf{X}\mathbf{T}$ . Then, the new low-dimensional dataset  $\mathcal{X}'$  can be generated by integrating matrix  $\mathbf{X}'$  with other information  $\mathbf{I}_X$  that was filtered in the preprocessing step in Phase I.

#### 4.3 Pseudocode

The pseudocode of LifeRank as a dimension reduction algorithm for ranking is summarized in Algorithm 1. Given the number of generated features  $k$  and a set of standard learning to rank datasets, including a training dataset  $\mathcal{X}$ , a validation dataset  $\mathcal{V}$  and a test dataset  $\mathcal{E}$ , LifeRank tries to output new low-dimensional datasets  $\mathcal{X}'$ ,  $\mathcal{V}'$  and  $\mathcal{E}'$  for training, validation and test, respectively, for the learning to rank procedure.

Algorithm 1 implements the two phases of LifeRank: (I) Lines 1–8 generate the transformation matrix  $\mathbf{T}$  based on the original training dataset  $\mathcal{X}$ ; (II) Using  $\mathbf{T}$ , lines 9–10 generate the low-dimensional matrices for training  $\mathbf{X}'$ , validation  $\mathbf{V}'$  and test  $\mathbf{E}'$ . Then, line 11 constructs the low-dimensional training, validation and test datasets by directly integrating the low-rank matrices and their corresponding information filtered in the preprocessing steps in lines 1 and 9.

#### 4.4 Discussion

In this section, we reveal a connection between the feature selection for ranking problem and the linear feature extraction for ranking problem. In particular, from the perspective of linear transformations of matrices, the feature selection for ranking problem can be defined as in Definition 4.

**Algorithm 1: LifeRank: A Linear Feature Extraction Algorithm for Ranking**

**Input:** A training dataset  $\mathcal{X}$ , a validation dataset  $\mathcal{V}$ , a test dataset  $\mathcal{E}$ , the learning rate  $\eta$ , and the number of features  $k$  in the set of generated document.

**Output:** A generated training dataset  $\mathcal{X}'$ , validation dataset  $\mathcal{V}'$ , and test dataset  $\mathcal{E}'$ , each with  $k$  features.

```

// Phase I
1  $(\mathbf{X}, \mathbf{I}_X) \leftarrow \text{Preprocess}(\mathcal{X})$ ;
2  $\mathbf{T}, \mathbf{w}, \{\alpha_{i,j}\}_{i,j=1,\dots,k} \leftarrow \text{Initialize}(\mathbf{X}, k)$ ;
3 repeat
4    $\mathbf{w} \leftarrow \mathbf{w} - \eta \nabla_{\mathbf{w}} \mathcal{L}$ ;
5    $\mathbf{t}_l \leftarrow \mathbf{t}_l - \eta \nabla_{\mathbf{t}_l} \mathcal{L}$ , for  $l = 1, \dots, k$ ;
6    $b \leftarrow b - \eta \frac{\partial \mathcal{L}}{\partial b}$ ;
7    $\alpha_{i,j} \leftarrow \alpha_{i,j} + \eta \frac{\partial \mathcal{L}}{\partial \alpha_{i,j}}$ , for  $i, j = 1, \dots, k$ ;
8 until Reach convergence or the max iteration;

// Phase II
9  $(\mathbf{V}, \mathbf{I}_V), (\mathbf{E}, \mathbf{I}_E) \leftarrow \text{Preprocess}(\mathcal{V}, \mathcal{E})$ ;
10  $\mathbf{X}' \leftarrow \mathbf{X}\mathbf{T}$ ,  $\mathbf{V}' \leftarrow \mathbf{V}\mathbf{T}$ ,  $\mathbf{E}' \leftarrow \mathbf{E}\mathbf{T}$ ;
11  $\mathcal{X}', \mathcal{V}', \mathcal{E}' \leftarrow \text{GenerateDatasets}(\mathbf{X}', \mathbf{V}', \mathbf{E}', \mathbf{I}_X, \mathbf{I}_V, \mathbf{I}_E)$ ;

```

**Definition 4 (Feature selection for ranking)** Let  $\mathbf{X}_{m \times n}$  be a document matrix, where the transpose of each row  $\mathbf{x}_i^\top = \mathbf{d}_i$  is an  $n$ -dimensional vector, representing a document with  $n$  features. Feature selection for ranking aims to optimize a transformation matrix  $\mathbf{T}_{n \times k}$  by solving the following optimization problem, so that a new document matrix  $\mathbf{X}'_{m \times k} = \mathbf{X}_{m \times n} \mathbf{T}_{n \times k}$  can be generated, where each  $n$ -dimensional document vector  $\mathbf{d}_i$  can be projected into a  $k$ -dimensional vector  $\mathbf{d}'_i = \mathbf{T}^\top \mathbf{d}_i$ :

$$\arg \min_{\mathbf{T}} L(g(\mathbf{X}\mathbf{T})) \text{ such that } \begin{cases} \forall t_{i,j} \in \mathbf{T} : t_{i,j} \in \{0, 1\} \\ \mathbf{T}^\top \mathbf{T} = \mathbf{I}. \end{cases} \quad (9)$$

Based on the optimized mapping function  $g$ , the new low-rank matrix can be generated as a product of the original matrix and the transformation matrix, i.e.,  $\mathbf{X}' = \mathbf{X}\mathbf{T}$ .

The  $k$  columns of the transformation  $\mathbf{T}$  mentioned in Definition 4 present the  $k$  iterations of the feature selection processes. The constraints in Equation 9 guarantee that there is only one “1” in each column of the transformation matrix  $\mathbf{T}$  and the others are all “0,” indicating that each feature selection process only selects one feature. The second constraint  $\mathbf{T}^\top \mathbf{T} = \mathbf{I}$  guarantees that the position of the unique “1” in each column is different from other columns, which is the index of the selected feature in that step.

Since the elements in the transformation matrix  $\mathbf{T}$  can be any real numbers in Definition 3 while they are only either 0 or 1 in Definition 4, Definition 3 generalizes Definition 4, i.e., the problem of linear feature extraction for ranking generalizes the problem of feature selection for ranking. Because of this, linear feature extraction is expected to outperform or be at least as good as any feature selection technique. The linear feature extraction is expected to use more computational resources than feature selection, since former deals with the search space in real numbers and the

latter with binary case. However, this computational overhead is the tradeoff for the higher performance expected to be achieved by the extracted features, when utilized for learning to rank.

## 5 Experimental setup

### 5.1 Research questions

We list the research questions that guide the remainder of the paper.

- RQ1** What is the performance of LifeRank in generating low-dimensional datasets? Does LifeRank outperform state-of-the-art feature selection algorithms? (See §6.1)
- RQ2** Can the importance and redundancy of the features generated by LifeRank outperform those selected by feature selection algorithms? (See §6.2)
- RQ3** What is the effect of the orthonormality constraints of the transformation matrix in Equation 4? Does it help enhance the performance of ranking predictions? (See §6.3.)

### 5.2 Datasets

In this study, we use the MQ2007 and MQ2008 datasets from LETOR 4.0 [42] and OHSUMED from LETOR 3.0 [43] to evaluate our algorithm. The LETOR<sup>1</sup> datasets are commonly used benchmarks in learning to rank. LETOR 4.0 is the latest version, which was released in July 2009. It uses the Gov2 web page collection (~25M pages) and two query sets from the Million Query track of TREC 2007 and TREC 2008, which are referred to as MQ2007 and MQ2008. We use both MQ2007 and MQ2008 in our experiments. In MQ2007, there are about 1700 queries and about 70,000 query-document pairs, while MQ2008 has 800 queries and about 15,000 query-document pairs for training, validation and testing. In both datasets, each query-document pair has 46 features. We also use the OHSUMED dataset from LETOR 3.0, which was released in December 2008. OHSUMED is extracted from the online medical information database MEDLINE. It contains 106 queries and about 16,000 query-document pairs, where each query-document pair has 45 features.

In all the datasets that we use, relevance of documents with respect to queries is judged at three levels: 2 (definitely relevant), 1 (partially relevant), and 0 (not relevant). In our experiments, we use 5-fold cross validation. In each fold, 60% queries are used for training, 20% for validation and the remaining 20% for testing. The performance numbers reported are averaged over the five folds.

### 5.3 Baselines

LifeRank aims to generate low-dimensional datasets for ranking. In this paper, we utilize three baselines to evaluate the datasets generated by our algorithm:

<sup>1</sup> <http://research.microsoft.com/en-us/um/beijing/projects/letor/>

**Original datasets:** We firstly use the original LETOR datasets as our first baseline, on which no selection or generation has been performed.

**Datasets generated by GAS:** GAS [16] incorporates importance and similarity information of the features into ranking. It greedily selects a subset of features by maximizing the total importance scores meanwhile minimizing the total similarity scores.

**Datasets generated by FSMRank:** FSMRank [26] trains a feature selection model with machine learning, which can select a subset of features meanwhile minimizing the ranking errors.

We then run Linear Regression [30]-based learning to rank and RankSVM<sup>2</sup> [21] to determine how well these datasets can address the ranking problem. The former makes pointwise predictions on the relevance of the documents by linear regression, which is implemented in the RankLib learning to rank toolkit.<sup>3</sup> The latter predicts pairwise ranking relation between each pair of documents directly by support vector machine (SVM). These are classical pointwise and pairwise learning to rank algorithms, respectively, with which we can clearly demonstrate the effects of dimension reduction.

Since LifeRank uses a linear approach for feature extraction, it is expected to show effectiveness mainly for linear learning-to-rank methods. This is the reason why we have chosen SVMRank and Linear Regression for experimentation.

## 5.4 Evaluation measures

### 5.4.1 Measures for ranking

We use two standard ranking accuracy metrics to evaluate the rankings generated by learning to rank algorithms: mean average precision (*MAP*) [3] and normalized discount cumulative gain (*NDCG@n*) [19].

Statistical significance of observed differences between the performance of two runs is tested using a two-tailed paired t-test and is denoted using  $\blacktriangle$  (or  $\blacktriangledown$ ) for strong significance for  $\alpha = 0.01$ ; or  $\triangle$  (or  $\triangledown$ ) for weak significance for  $\alpha = 0.05$ .

### 5.4.2 Measures for features

We consider two metrics to evaluate the quality of the features: importance and redundancy.

The *importance* of each feature can be evaluated by the ranking performance when the feature is used as a ranking model to order the documents. In particular, we use *NDCG@5* for evaluation. Since for calculating these measures, for some features larger values correspond to higher ranks while for others smaller values lead to higher ranks, we utilize the strategy in GAS [16] for evaluation: We order the documents twice in ascending and descending manners respectively, and take the larger score as

<sup>2</sup> [https://www.cs.cornell.edu/people/tj/svm\\_light/svm\\_rank.html](https://www.cs.cornell.edu/people/tj/svm_light/svm_rank.html)

<sup>3</sup> <https://sourceforge.net/p/lemur/wiki/RankLib/>

the importance score of the features. Then we calculate the average importance of the features as the importance of the set of features  $F = \{f_1, f_2, \dots, f_k\}$ :

$$Imp(F) = \frac{1}{k} \sum_{i=1}^k \max \{eva(\mathcal{X}, f_i), eva(\mathcal{X}, -f_i)\},$$

where the function  $eva(\mathcal{X}, f_i)$  returns the evaluation results of the ranking model  $f_i$  on the dataset  $\mathcal{X}$ .

The *redundancy* of features can be defined as the average similarity between each pair of features. In practice, we regard each feature as a ranking model to order the documents, and then calculate the similarity between each pair of features as the average similarity of their document rankings associated to different queries. Let  $Q$  be the set of queries in the given dataset, each associated with a set of documents for ranking. The redundancy of the features  $F = \{f_1, f_2, \dots, f_k\}$  is calculated as follows:

$$Rdd(F) = \frac{2}{k(k-1)} \sum_{f_i, f_j \in F, i > j} \frac{1}{|Q|} \sum_{q \in Q} sim(\sigma_i^{(q)}, \sigma_j^{(q)}),$$

where  $\sigma_i^{(q)}$  is the ranking of the document associated to the query  $q$  when the feature  $f_i$  is used as the ranking model to order the documents. In this paper, we take the absolute value of Kendall's  $\tau$  correlation coefficient [24] as the similarity metric for rankings:

$$\tau(\sigma_i, \sigma_j) = \frac{N_c - N_d}{N_c + N_d}, \quad (10)$$

where  $N_c$  and  $N_d$  are the numbers of the concordant pairs and discordant pairs respectively between rankings  $\sigma_i$  and  $\sigma_j$ .

The range of  $\tau(\sigma_i, \sigma_j)$  is  $[-1, 1]$ , where the sign indicates that the correlation between  $\sigma_i$  and  $\sigma_j$  is either positive or negative, and the absolute value indicates the strength of the correlation. Since positive and negative values should not neutralize and we only consider the strength of the correlations, we take the absolute value of Kendall's  $\tau$  as the similarity metric in the definition of redundancy.

## 6 Experimental results

### 6.1 Performance on generated datasets

Tables 1, 2 and 3 list the results obtained in our experiments on the MQ2007, MQ2008 and OHSUMED datasets, respectively. They show the NDCG@1–10 and MAP scores for the RankSVM and Linear Regression learning to rank algorithms on 4 categories of datasets: the original datasets and 3 datasets generated by dimension reduction algorithms including GAS, FSMRank and our LifeRank. For each dimension reduction algorithm, we consider  $k$  generated features, with  $k = 5, 10, 15, 20$ . The results for the original dataset in the tables are independent of the value of  $k$ , but are repeated nevertheless for ease of comparison. The values in bold represent the best performance among GAS, FSMSVM and LifeRank.

**Table 1: Ranking performance on MQ2007 and selected/generated datasets.**

Statistical significance shown for LifeRank against Original, GAS and FSMSVM  
 Performance for Original is independent of value of k (corresponds to original dataset)

Performance for RankSVM											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
k = 5											
Original	0.4079	0.4007	0.4009	0.4030	0.4077	0.4129	0.4201	0.4275	0.4336 <sup>▽</sup>	0.4391 <sup>▽</sup>	0.4615 <sup>▽</sup>
GAS	0.3751	0.3807	0.3869 <sup>△</sup>	0.3946	0.3980	0.4068	0.4113	0.4177	0.4242	0.4290	<b>0.4563</b>
FSMSVM	0.3598 <sup>▲</sup>	0.3703 <sup>▲</sup>	0.3773 <sup>▲</sup>	0.3811 <sup>▲</sup>	0.3871 <sup>▲</sup>	0.3941 <sup>▲</sup>	0.3997 <sup>▲</sup>	0.4061 <sup>▲</sup>	0.4129 <sup>▲</sup>	0.4193 <sup>▲</sup>	0.4512
<b>LifeRank</b>	<b>0.3925</b>	<b>0.3925</b>	<b>0.3975</b>	<b>0.4009</b>	<b>0.4062</b>	<b>0.4118</b>	<b>0.4162</b>	<b>0.4209</b>	<b>0.4256</b>	<b>0.4312</b>	0.4539
k = 10											
Original	0.4079	0.4007	0.4009	0.4030 <sup>△</sup>	0.4077 <sup>△</sup>	0.4129 <sup>▲</sup>	0.4201 <sup>△</sup>	0.4275	0.4336	0.4391	0.4615
GAS	0.3897	0.3893 <sup>▲</sup>	0.3914 <sup>▲</sup>	0.3965 <sup>▲</sup>	0.4029 <sup>▲</sup>	0.4098 <sup>▲</sup>	0.4153 <sup>▲</sup>	0.4209 <sup>▲</sup>	0.4272 <sup>▲</sup>	0.4343 <sup>▲</sup>	0.4558 <sup>△</sup>
FSMSVM	0.3919	0.3917	0.3982 <sup>△</sup>	0.4027 <sup>△</sup>	0.4079 <sup>△</sup>	0.4134 <sup>▲</sup>	0.4187 <sup>▲</sup>	0.4239 <sup>▲</sup>	0.4290 <sup>▲</sup>	0.4349 <sup>▲</sup>	0.4593
<b>LifeRank</b>	<b>0.4037</b>	<b>0.4023</b>	<b>0.4089</b>	<b>0.4117</b>	<b>0.4161</b>	<b>0.4215</b>	<b>0.4264</b>	<b>0.4312</b>	<b>0.4370</b>	<b>0.4423</b>	<b>0.4634</b>
k = 15											
Original	0.4079	0.4007	0.4009	0.4030	0.4077	0.4129	0.4201	0.4275	0.4336	0.4391	0.4615
GAS	0.3942	0.3954	0.3998	0.4023	0.4063 <sup>▲</sup>	0.4126 <sup>△</sup>	0.4195	0.4256 <sup>△</sup>	0.4316 <sup>▲</sup>	0.4372 <sup>▲</sup>	0.4593 <sup>▲</sup>
FSMSVM	0.3905	0.3937 <sup>△</sup>	0.4005	0.4060	0.4106	0.4144	0.4201	0.4250 <sup>△</sup>	0.4309 <sup>▲</sup>	0.4365 <sup>▲</sup>	0.4589 <sup>△</sup>
<b>LifeRank</b>	<b>0.3972</b>	<b>0.4032</b>	<b>0.4061</b>	<b>0.4082</b>	<b>0.4141</b>	<b>0.4194</b>	<b>0.4242</b>	<b>0.4308</b>	<b>0.4376</b>	<b>0.4436</b>	<b>0.4635</b>
k = 20											
Original	0.4079	0.4007	0.4009	0.4030 <sup>△</sup>	0.4077 <sup>△</sup>	0.4129 <sup>▲</sup>	0.4201 <sup>△</sup>	0.4275	0.4336	0.4391	0.4615
GAS	0.4014	0.3967 <sup>▲</sup>	0.4007	0.4027 <sup>▲</sup>	0.4088 <sup>▲</sup>	0.4137 <sup>▲</sup>	0.4189 <sup>▲</sup>	0.4257 <sup>▲</sup>	0.4326 <sup>△</sup>	0.4394 <sup>△</sup>	0.4601 <sup>▲</sup>
FSMSVM	0.3882 <sup>▲</sup>	0.3929 <sup>▲</sup>	0.4004	0.4060	0.4096	0.4138 <sup>△</sup>	0.4203 <sup>△</sup>	0.4259	0.4310 <sup>△</sup>	0.4368 <sup>▲</sup>	0.4595 <sup>△</sup>
<b>LifeRank</b>	<b>0.4097</b>	<b>0.4077</b>	<b>0.4058</b>	<b>0.4106</b>	<b>0.4153</b>	<b>0.4213</b>	<b>0.4265</b>	<b>0.4311</b>	<b>0.4374</b>	<b>0.4438</b>	<b>0.4640</b>
Performance for Linear Regression											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
k = 5											
Original	0.3750	0.3854	0.3882	0.3926	0.3979	0.4034	0.4088	0.4154	0.4208	0.4277	0.4497
GAS	0.3712	0.3751	0.3797 <sup>△</sup>	0.3851 <sup>△</sup>	0.3894	0.3952	0.4011 <sup>△</sup>	0.4071 <sup>△</sup>	0.4136 <sup>△</sup>	0.4196 <sup>△</sup>	0.4462
FSMSVM	0.3554 <sup>▲</sup>	0.3634 <sup>▲</sup>	0.3673 <sup>▲</sup>	0.3741 <sup>▲</sup>	0.3780 <sup>▲</sup>	0.3872 <sup>▲</sup>	0.3929 <sup>▲</sup>	0.4002 <sup>▲</sup>	0.4076 <sup>▲</sup>	0.4145 <sup>▲</sup>	0.4489
<b>LifeRank</b>	<b>0.3852</b>	<b>0.3874</b>	<b>0.3908</b>	<b>0.3955</b>	<b>0.3962</b>	<b>0.4026</b>	<b>0.4089</b>	<b>0.4149</b>	<b>0.4213</b>	<b>0.4279</b>	<b>0.4507</b>
k = 10											
Original	0.3750	0.3854	0.3882	0.3926	0.3979	0.4034	0.4088	0.4154	0.4208	0.4277	0.4497
GAS	0.3886	0.3879	0.3881	0.3929	0.3980	0.4031	0.4090	0.4146	0.4198	0.4261	0.4491
FSMSVM	<b>0.3903</b>	<b>0.3951</b>	<b>0.3936</b>	0.3950	0.3991	0.4049	0.4111	0.4166	0.4223	0.4285	0.4494
<b>LifeRank</b>	0.3852	0.3920	0.3926	<b>0.3976</b>	<b>0.4026</b>	<b>0.4073</b>	<b>0.4146</b>	<b>0.4206</b>	<b>0.4270</b>	<b>0.4312</b>	<b>0.4507</b>
k = 15											
Original	0.3750	0.3854	0.3882	0.3926	0.3979	0.4034	0.4088 <sup>△</sup>	0.4154	0.4208 <sup>△</sup>	0.4277 <sup>△</sup>	0.4497
GAS	0.3767	0.3849	0.3913	0.3954	0.4000	0.4070	0.4133	0.4191	0.4250	0.4315	<b>0.4528</b>
FSMSVM	0.3800	0.3825	0.3855	0.3882 <sup>▲</sup>	0.3922 <sup>▲</sup>	0.3976 <sup>▲</sup>	0.4035 <sup>▲</sup>	0.4091 <sup>▲</sup>	0.4147 <sup>▲</sup>	0.4215 <sup>▲</sup>	0.4441 <sup>▲</sup>
<b>LifeRank</b>	<b>0.3842</b>	<b>0.3888</b>	<b>0.3943</b>	<b>0.3984</b>	<b>0.4019</b>	<b>0.4091</b>	<b>0.4161</b>	<b>0.4207</b>	<b>0.4274</b>	<b>0.4336</b>	0.4513
k = 20											
Original	0.3750	0.3854	0.3882	0.3926	0.3979	0.4034	0.4088	0.4154	0.4208	0.4277	0.4497
GAS	0.3783	<b>0.3897</b>	<b>0.3956</b>	<b>0.3990</b>	<b>0.4021</b>	<b>0.4079</b>	0.4120	0.4178	<b>0.4255</b>	<b>0.4313</b>	<b>0.4521</b>
FSMSVM	0.3742	0.3867	0.3916	0.3937	0.3970	0.4011	0.4062 <sup>△</sup>	0.4132	0.4184	0.4244 <sup>△</sup>	0.4483
<b>LifeRank</b>	<b>0.3828</b>	0.3894	0.3912	0.3948	0.4014	0.4068	<b>0.4121</b>	<b>0.4184</b>	0.4243	0.4306	0.4514

**Table 2: Ranking performance on MQ2008 and selected/generated datasets.**

Statistical significance shown for LifeRank against Original, GAS and FSMSVM  
 Performance for Original is independent of value of k (corresponds to original dataset)

Performance for RankSVM											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
k = 5											
Original	0.3712	0.3933 <sup>▲</sup>	0.4238 <sup>△</sup>	0.4485 <sup>△</sup>	0.4672 <sup>▲</sup>	0.4814 <sup>△</sup>	0.4875 <sup>△</sup>	0.4531 <sup>△</sup>	0.2234	0.2284	0.4714
GAS	0.3678	0.3983 <sup>△</sup>	0.4213 <sup>▲</sup>	0.4492 <sup>△</sup>	0.4665 <sup>▲</sup>	0.4800 <sup>▲</sup>	0.4862 <sup>▲</sup>	0.4522 <sup>▲</sup>	0.2207 <sup>▲</sup>	0.2246 <sup>▲</sup>	0.4714
FSMSVM	<b>0.3780</b>	0.4126	0.4384	0.4599	0.4761	0.4909	0.4968	0.4616	<b>0.2284</b>	0.2326	0.4776
<b>LifeRank</b>	0.3767	<b>0.4168</b>	<b>0.4389</b>	<b>0.4606</b>	<b>0.4806</b>	<b>0.4921</b>	<b>0.4976</b>	<b>0.4627</b>	0.2280	<b>0.2329</b>	<b>0.4788</b>
k = 10											
Original	0.3712	0.3933 <sup>▲</sup>	0.4238 <sup>△</sup>	0.4485 <sup>△</sup>	0.4672 <sup>▲</sup>	0.4814	0.4875 <sup>△</sup>	0.4531 <sup>▲</sup>	0.2234	0.2284 <sup>△</sup>	0.4714
GAS	0.3698	0.4015 <sup>△</sup>	0.4292	0.4565	0.4732	0.4862	0.4929	0.4551 <sup>△</sup>	0.2217 <sup>△</sup>	0.2266 <sup>▲</sup>	0.4776
FSMSVM	0.3759	0.4157	0.4371	0.4589	0.4781	<b>0.4925</b>	0.4962	0.4617	0.2276	0.2322	<b>0.4793</b>
<b>LifeRank</b>	<b>0.3763</b>	<b>0.4181</b>	<b>0.4384</b>	<b>0.4612</b>	<b>0.4796</b>	<b>0.4900</b>	<b>0.4972</b>	<b>0.4625</b>	<b>0.2281</b>	<b>0.2345</b>	0.4792
k = 15											
Original	0.3712	0.3933 <sup>▲</sup>	0.4238 <sup>▲</sup>	0.4485 <sup>▲</sup>	0.4672 <sup>▲</sup>	0.4814 <sup>▲</sup>	0.4875 <sup>▲</sup>	0.4531 <sup>▲</sup>	0.2234 <sup>▲</sup>	0.2284 <sup>▲</sup>	0.4714 <sup>△</sup>
GAS	0.3720	0.3984 <sup>△</sup>	0.4320	0.4533 <sup>△</sup>	0.4711 <sup>△</sup>	0.4851	0.4902 <sup>△</sup>	0.4543 <sup>▲</sup>	0.2223 <sup>▲</sup>	0.2279 <sup>▲</sup>	0.4742
FSMSVM	<b>0.3788</b>	<b>0.4165</b>	0.4351	0.4557	0.4761	0.4905	0.4967	0.4613	0.2265	0.2311	0.4788
<b>LifeRank</b>	0.3771	0.4140	<b>0.4379</b>	<b>0.4622</b>	<b>0.4804</b>	<b>0.4920</b>	<b>0.4972</b>	<b>0.4633</b>	<b>0.2310</b>	<b>0.2349</b>	<b>0.4792</b>
k = 20											
Original	0.3712	0.3933	0.4238 <sup>△</sup>	0.4485 <sup>▲</sup>	0.4672 <sup>△</sup>	0.4814 <sup>△</sup>	0.4875 <sup>▲</sup>	0.4531 <sup>▲</sup>	0.2234 <sup>△</sup>	0.2284 <sup>△</sup>	0.4714
GAS	0.3656	0.4027	0.4313	0.4527 <sup>△</sup>	0.4720	0.4850	0.4921	0.4566	0.2254	0.2298	0.4719
FSMSVM	<b>0.3737</b>	<b>0.4144</b>	0.4343	0.4588	0.4757	0.4902	0.4946	0.4590	0.2249	0.2302	<b>0.4754</b>
<b>LifeRank</b>	0.3712	0.4045	<b>0.4363</b>	<b>0.4619</b>	<b>0.4763</b>	<b>0.4903</b>	<b>0.4971</b>	<b>0.4617</b>	<b>0.2293</b>	<b>0.2338</b>	0.4751
Performance for Linear Regression											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
k = 5											
Original	0.3465	0.3617 <sup>▲</sup>	0.3961	0.4225	0.4407	0.4558 <sup>△</sup>	0.4684 <sup>△</sup>	0.4746 <sup>▲</sup>	0.4806 <sup>▲</sup>	0.4871 <sup>▲</sup>	0.4550 <sup>▲</sup>
GAS	0.3537	0.3691 <sup>△</sup>	0.3947	0.4184	0.4390	0.4584	0.4692	0.4776	0.4838	0.4892	0.4630
FSMSVM	0.3541	0.3755 <sup>△</sup>	0.4002	<b>0.4191</b>	0.4396 <sup>△</sup>	0.4563 <sup>▲</sup>	0.4681 <sup>△</sup>	0.4791 <sup>△</sup>	0.4842 <sup>△</sup>	0.4903 <sup>△</sup>	0.4593 <sup>△</sup>
<b>LifeRank</b>	<b>0.3691</b>	<b>0.3907</b>	<b>0.4089</b>	0.3955	<b>0.4491</b>	<b>0.4664</b>	<b>0.4772</b>	<b>0.4864</b>	<b>0.4924</b>	<b>0.4979</b>	<b>0.4685</b>
k = 10											
Original	0.3465 <sup>△</sup>	0.3617 <sup>▲</sup>	0.3961 <sup>△</sup>	0.4225	0.4407 <sup>△</sup>	0.4558 <sup>△</sup>	0.4684 <sup>▲</sup>	0.4746 <sup>▲</sup>	0.4806 <sup>▲</sup>	0.4871 <sup>▲</sup>	0.4550 <sup>▲</sup>
GAS	0.3605	0.3804	0.3990	<b>0.4282</b>	0.4490	0.4628	0.4742	0.4840	0.4895	0.4942	0.4669
FSMSVM	0.3512	0.3779	0.4017	0.4213	0.4442	0.4601	0.4702 <sup>△</sup>	0.4789 <sup>△</sup>	0.4847	0.4910	0.4622
<b>LifeRank</b>	<b>0.3698</b>	<b>0.3857</b>	<b>0.4085</b>	0.3976	<b>0.4501</b>	<b>0.4659</b>	<b>0.4791</b>	<b>0.4859</b>	<b>0.4908</b>	<b>0.4970</b>	<b>0.4686</b>
k = 15											
Original	0.3465	0.3617 <sup>▲</sup>	0.3961	0.4225	0.4407	0.4558 <sup>△</sup>	0.4684	0.4746 <sup>▲</sup>	0.4806 <sup>△</sup>	0.4871 <sup>△</sup>	0.4550 <sup>△</sup>
GAS	<b>0.3652</b>	0.3795	<b>0.4114</b>	<b>0.4302</b>	0.4497	0.4641	0.4757	0.4845	<b>0.4914</b>	<b>0.4964</b>	<b>0.4686</b>
FSMSVM	0.3631	0.3822	0.4053	0.4285	<b>0.4527</b>	<b>0.4669</b>	<b>0.4774</b>	0.4850	0.4901	0.4956	0.4647
<b>LifeRank</b>	0.3618	<b>0.3842</b>	0.4032	0.4296	0.4482	0.4657	0.4766	<b>0.4852</b>	0.4904	0.4955	0.4662
k = 20											
Original	0.3465 <sup>▲</sup>	0.3617 <sup>▲</sup>	0.3961	0.4225	0.4407 <sup>△</sup>	0.4558	0.4684	0.4746 <sup>△</sup>	0.4806 <sup>▲</sup>	0.4871 <sup>▲</sup>	0.4550 <sup>▲</sup>
GAS	0.3588	0.3801	<b>0.4100</b>	<b>0.4300</b>	0.4486	0.4650	0.4744	<b>0.4843</b>	<b>0.4906</b>	<b>0.4961</b>	0.4660
FSMSVM	0.3601 <sup>△</sup>	0.3803	0.4075	0.4266	0.4525	<b>0.4668</b>	<b>0.4772</b>	0.4833	0.4889	0.4945	0.4661
<b>LifeRank</b>	<b>0.3712</b>	<b>0.3820</b>	0.4018	0.3948	<b>0.4495</b>	0.4628	0.4748	0.4833	0.4897	0.4956	<b>0.4662</b>

**Table 3: Ranking performance on OHSUMED and selected/generated datasets.**

Statistical significance shown for LifeRank against Original, GAS and FSMSVM  
 Performance for Original is independent of value of k (corresponds to original dataset)

Performance for RankSVM											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
k = 5											
Original	0.5416	0.5076	0.4775	0.4565 <sup>△</sup>	0.4439 <sup>▲</sup>	0.4452	0.4433	0.4405	0.4338 <sup>▲</sup>	0.4300 <sup>▲</sup>	0.4345 <sup>▲</sup>
GAS	0.5332	0.4901	0.4739	0.4630 <sup>△</sup>	0.4578 <sup>▲</sup>	0.4503 <sup>△</sup>	0.4432 <sup>▲</sup>	0.4398 <sup>△</sup>	0.4374 <sup>▲</sup>	0.4340 <sup>▲</sup>	0.4642 <sup>▼</sup>
FSMSVM	0.5771 <sup>▽</sup>	0.4889	0.4772	0.4749	0.4694 <sup>△</sup>	0.4609	0.4622	0.4559	0.4529	0.4518	<b>0.4728</b>
<b>LifeRank</b>	0.5170	<b>0.5000</b>	<b>0.5015</b>	<b>0.4950</b>	<b>0.4905</b>	<b>0.4749</b>	<b>0.4701</b>	<b>0.4628</b>	<b>0.4684</b>	<b>0.4668</b>	0.4523
k = 10											
Original	0.5416	0.5076	0.4775	0.4565 <sup>▲</sup>	0.4439 <sup>▲</sup>	0.4452 <sup>▲</sup>	0.4433 <sup>▲</sup>	0.4405 <sup>▲</sup>	0.4338 <sup>▲</sup>	0.4300 <sup>▲</sup>	0.4345 <sup>▲</sup>
GAS	<b>0.5677</b>	<b>0.5390</b>	0.5088	0.4944	0.4873	0.4673	0.4652	0.4605	0.4549	0.4507 <sup>△</sup>	0.4466
FSMSVM	0.5296	0.4866 <sup>▲</sup>	0.4794 <sup>▲</sup>	0.4745 <sup>▲</sup>	0.4636 <sup>▲</sup>	0.4602 <sup>△</sup>	0.4558 <sup>▲</sup>	0.4531 <sup>△</sup>	0.4484 <sup>▲</sup>	0.4463 <sup>▲</sup>	0.4459
<b>LifeRank</b>	0.5518	0.5373	<b>0.5185</b>	<b>0.5053</b>	<b>0.4910</b>	<b>0.4811</b>	<b>0.4774</b>	<b>0.4699</b>	<b>0.4692</b>	<b>0.4663</b>	<b>0.4505</b>
k = 15											
Original	0.5416	0.5076	0.4775	0.4565 <sup>△</sup>	0.4439 <sup>▲</sup>	0.4452 <sup>△</sup>	0.4433	0.4405 <sup>△</sup>	0.4338 <sup>▲</sup>	0.4300 <sup>▲</sup>	0.4345 <sup>▲</sup>
GAS	0.5771	0.5068	0.4850	0.4713	0.4656	0.4552	0.4524	0.4494	0.4439 <sup>△</sup>	0.4419 <sup>▲</sup>	0.4402 <sup>▲</sup>
FSMSVM	<b>0.5834</b>	0.5317	0.5021	0.4856	0.4723	0.4660	0.4606	0.4557	0.4525	0.4500 <sup>△</sup>	0.4452 <sup>▲</sup>
<b>LifeRank</b>	0.5769	<b>0.5344</b>	<b>0.5065</b>	<b>0.4942</b>	<b>0.4835</b>	<b>0.4713</b>	<b>0.4672</b>	<b>0.4630</b>	<b>0.4632</b>	<b>0.4628</b>	<b>0.4536</b>
k = 20											
Original	0.5416	0.5076	0.4775 <sup>△</sup>	0.4565 <sup>▲</sup>	0.4439 <sup>▲</sup>	0.4452 <sup>▲</sup>	0.4433 <sup>▲</sup>	0.4405 <sup>▲</sup>	0.4338 <sup>▲</sup>	0.4300 <sup>▲</sup>	0.4345 <sup>▲</sup>
GAS	0.5519	0.5051	0.4838 <sup>△</sup>	0.4761 <sup>△</sup>	0.4710 <sup>▲</sup>	0.4520 <sup>▲</sup>	0.4473 <sup>▲</sup>	0.4463 <sup>▲</sup>	0.4401 <sup>▲</sup>	0.4405 <sup>▲</sup>	0.4387 <sup>▲</sup>
FSMSVM	0.5173 <sup>△</sup>	0.4848 <sup>△</sup>	0.4816 <sup>△</sup>	0.4766	0.4642 <sup>▲</sup>	0.4522 <sup>▲</sup>	0.4452 <sup>▲</sup>	0.4411 <sup>▲</sup>	0.4388 <sup>▲</sup>	0.4387 <sup>▲</sup>	0.4441 <sup>▲</sup>
<b>LifeRank</b>	<b>0.5805</b>	<b>0.5416</b>	<b>0.5204</b>	<b>0.5029</b>	<b>0.4976</b>	<b>0.4834</b>	<b>0.4765</b>	<b>0.4725</b>	<b>0.4669</b>	<b>0.4656</b>	<b>0.4519</b>
Performance for Linear Regression											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
k = 5											
Original	0.4830	0.4800	0.4749	0.4548	0.4483	0.4368	0.4389	0.4343	0.4311	0.4302	0.4333
GAS	0.4762	0.4491	0.4489	0.4466	0.4378	0.4286	0.4275	0.4209	0.4201	0.4202	0.4549
FSMSVM	<b>0.5271</b>	<b>0.4852</b>	<b>0.4686</b>	<b>0.4589</b>	<b>0.4513</b>	0.4403	0.4342	0.4335	0.4306	0.4292	<b>0.4655</b>
<b>LifeRank</b>	0.4941	0.4736	0.4554	0.4586	0.4485	<b>0.4459</b>	<b>0.4425</b>	<b>0.4402</b>	<b>0.4347</b>	<b>0.4341</b>	0.4599
k = 10											
Original	0.4830	0.4800	0.4749	0.4548	0.4483	0.4368	0.4389	0.4343	0.4311	0.4302	0.4333
GAS	0.5202	<b>0.4883</b>	<b>0.4680</b>	0.4597	0.4543	0.4460	0.4415	0.4354	0.4312	0.4275	0.4339
FSMSVM	0.5082	0.4556	0.4502	0.4395	0.4333	0.4245	0.4259	0.4223	0.4216	0.4180	0.4344
<b>LifeRank</b>	<b>0.5330</b>	0.4866	<b>0.4680</b>	<b>0.4681</b>	<b>0.4601</b>	<b>0.4494</b>	<b>0.4436</b>	<b>0.4423</b>	<b>0.4370</b>	<b>0.4366</b>	<b>0.4358</b>
k = 15											
Original	0.4830	0.4800	0.4749	0.4548	0.4483	0.4368	0.4389	0.4343	0.4311	0.4302	0.4333
GAS	0.5105	<b>0.4941</b>	0.4791	0.4654	0.4510	0.4412	0.4351	0.4331	0.4303	0.4301	0.4298 <sup>△</sup>
FSMSVM	0.4984	0.4611	0.4639	0.4567	0.4485	0.4400	0.4325	0.4321	0.4302	0.4282	<b>0.4405</b>
<b>LifeRank</b>	<b>0.5342</b>	0.4839	<b>0.5001</b>	<b>0.4781</b>	<b>0.4678</b>	<b>0.4529</b>	<b>0.4437</b>	<b>0.4412</b>	<b>0.4370</b>	<b>0.4353</b>	0.4399
k = 20											
Original	0.4830	0.4800	0.4749	0.4548	0.4483	0.4368	0.4389	0.4343	0.4311	0.4302	0.4333
GAS	0.5050	0.4943 <sup>△</sup>	0.4813 <sup>▲</sup>	<b>0.4706</b> <sup>▼</sup>	0.4546 <sup>▲</sup>	<b>0.4497</b> <sup>▽</sup>	0.4389 <sup>▲</sup>	0.4367 <sup>▲</sup>	0.4324 <sup>▲</sup>	0.4299 <sup>▲</sup>	0.4308 <sup>▲</sup>
FSMSVM	0.4984	0.4706	0.4574	0.4510	0.4489	0.4416	0.4370	0.4365	0.4341	0.4315	<b>0.4369</b>
<b>LifeRank</b>	<b>0.5264</b>	<b>0.5021</b>	<b>0.4882</b>	0.4686	<b>0.4571</b>	0.4467	<b>0.4439</b>	<b>0.4406</b>	<b>0.4359</b>	<b>0.4341</b>	<b>0.4369</b>

Overall, from the tables we can see that: (1) The performance of ranking algorithms can be maintained or slightly improved on the datasets generated by dimension reduction techniques. (2) The performance of the ranking algorithms on the datasets generated by LifeRank is higher than those generated by GAS and FSMRank in most cases. Let us now take a closer look.

### 6.1.1 Performance of RankSVM

For RankSVM, we can see that LifeRank clearly shows improvements over the original datasets for all the three benchmarks (MQ2007, MQ2008 and OHSUMED) in terms of NDCG@1–10 as well as MAP. The only exception is MQ2007 for  $k = 5$ , where the performance of LifeRank as well as the other generated datasets does not beat the original dataset. We can also see from the tables that LifeRank clearly outperforms other generated datasets (GAS and FSMSVM) on NDCG@1–10 for all the benchmarks and all values of  $k$ .

In terms of MAP, LifeRank outperforms the other generated datasets in most cases. The few exceptions include the case for MQ2007, when GAS has a higher MAP for  $k = 5$ . For MQ2008, FSMSVM attains slightly higher MAP score than LifeRank for  $k = 10$  and  $k = 20$ , but these differences are not significant. Also, for OHSUMED when  $k = 5$ , the MAP score attained by LifeRank is lower than FSMSVM and GAS, but it is still an improvement over the original dataset.

### 6.1.2 Performance of Linear Regression

Also in the case of Linear Regression, for all three benchmarks (MQ2007, MQ2008 and OHSUMED) LifeRank clearly shows improvements over the original datasets in terms of NDCG@1–10 as well as MAP. The only exception is MQ2007 for  $k = 5$ , where the original dataset performs better than LifeRank as well as the other generated datasets.

On NDCG@1–10, for MQ2007 LifeRank gives the best performance for all values of  $k$ , except for  $k = 20$ , where GAS gives the best performance. For MQ2008, LifeRank gives the best performances for  $k = 5$  and  $k = 10$  on NDCG@1–10. However, for  $k = 15$  and  $k = 20$ , there is mixed performance where all GAS, FSMSVM and LifeRank give best performances in certain cases. For, OHSUMED, LifeRank gives the best performance on NDCG@1–10 in most cases.

In terms of MAP, LifeRank gives the best performance for MQ2007 for  $k = 5$  and  $k = 10$ , whereas for  $k = 15$  and  $k = 20$  the best performance is given by GAS. For MQ2008, LifeRank outperforms others for all values of  $k$ , except for  $k = 15$  where the best performance is given by GAS. Moreover, for OHSUMED, FSMSVM outperforms the others for  $k = 5$  and  $k = 15$ , while LifeRank gives the best performance for  $k = 10$ . In case of  $k = 20$ , there is a tie between LifeRank and FSMSVM.

### 6.1.3 Statistical Significance Overview

In Tables 1, 2 and 3, markups are provided to denote the statistical significance between LifeRank and the following baselines: original dataset, GAS and FSMSVM. It should be noted that the original dataset is independent of the values of  $k$ , but is repeated in the table to indicate statistical significance between it and datasets generated by LifeRank for different values of  $k$ .

It can be observed from Table 1 that for MQ2007 in the case of RankSVM, there is strong to weak significance between LifeRank and the baselines in most cases across the metrics, while there is no significance shown against original for  $k = 15$ . Moreover, for  $k = 5$ , significance is shown against original and GAS in few cases. For Linear Regression, there is strong significance shown against FSMSVM for  $k = 5$  and  $k = 15$ , though there is not much significance shown for  $k = 10$  and  $k = 20$ . Also, weak significance is shown against GAS in few cases for  $k = 5$  and against original for  $k = 15$ .

Table 2 for MQ2008 shows no statistically significant differences between LifeRank and FSMSVM for RankSVM. There is weak to strong statistical significance for LifeRank against original dataset for most cases and against GAS mainly for  $k = 5, 10$  and  $15$ . For Linear Regression, LifeRank shows weak to strong statistical significance against original in most cases, GAS in no cases and FSMSVM in few cases. Moreover, Table 3 for OHSUMED shows statistical significance for RankSVM in many cases against the baselines, whereas there is statistical significance observed for Linear Regression for few cases. The comparative lack of statistical significance seen for MQ2008 and OHSUMED can most probably be attributed to the relatively small size of these datasets.

## 6.2 Quality of the generated features

Table 4 lists the quality scores of the features from four datasets: the original datasets and the three datasets generated by GAS, FSMRank and LifeRank, respectively, for  $k = 10$ .

**Table 4: Performance of the generated features.**

Datasets	MQ2007		MQ2008		OHSUMED	
	<i>Imp</i>	<i>Rdd</i>	<i>Imp</i>	<i>Rdd</i>	<i>Imp</i>	<i>Rdd</i>
Original	0.2671	0.4833	0.3297	0.5318	0.3763	0.5592
GAS	0.2643	<b>0.3242</b>	0.3235	<b>0.3308</b>	0.3603	<b>0.3904</b>
FSMRank	0.3005	0.4706	0.3723	0.5276	0.4170	0.5412
LifeRank	<b>0.3214</b>	0.4606	<b>0.4095</b>	0.5758	<b>0.4422</b>	0.8881

From the table we see that: (1) GAS can significantly reduce the redundancy of the features. The redundancy of the features selected by GAS is the lowest among the four datasets. However, the importance of the features selected by GAS is also lowest and even slightly lower than that of the original datasets. (2) FSMRank can improve

the importance of the features while reducing their redundancy, but the differences in terms redundancy are subtle. (3) LifeRank can sharply improve the importance of the features. The importance of the features generated by LifeRank is highest among the four datasets. Besides, the redundancy of the features can also be slightly reduced by LifeRank for MQ2007 but deteriorated for MQ2008 and OHSUMED. Worse redundancy for LifeRank in comparison with GAS and FSMSVM could be because of the reason that, while these baselines are feature selection methods, for LifeRank each extracted feature is a linear combination of the original features. Moreover, it can be observed that for the larger dataset MQ2007, redundancy for LifeRank is comparable to the baselines, and even better than FSMSVM. However, for smaller dataset MQ2008, the redundancy is worse than the baselines. For the smallest dataset OHSUMED, it is worse than the baselines by a greater difference.

### 6.3 Effect of the Orthonormality Constraints

To confirm that the orthonormality constraints used in LifeRank do indeed contribute to performance gains, we re-generated the datasets for the benchmarks MQ2007, MQ2008 and OHSUMED using LifeRank for  $k = 10$ , but this time without the incorporation of the constraints in its algorithm in Phase I (see Algorithm 1, line 1–8). Table 5 shows the comparison of performances of ranking algorithms, for datasets generated by LifeRank and LifeRank without orthonormality constraints (represented by LifeRank<sup>NO</sup>). Moreover, markups are presented in the table to denote to the statistical significance between LifeRank and LifeRank<sup>NO</sup>.

From the results in Table 5 we see that the datasets generated by LifeRank show significant improvements in performance over the datasets generated by LifeRank<sup>NO</sup> for both learning to rank algorithms, RankSVM and Linear Regression. Performance gains can be observed on all three benchmarks and across all performance measures (NDCG@1–10 and MAP). Hence, these results show that the usage of orthonormality constraints is beneficial in the LifeRank algorithm. Also, strong statistical significance between LifeRank and LifeRank<sup>NO</sup> can be observed for all three benchmarks for RankSVM as well as Linear Regression, across all performance measures, except for a small number of cases where weak or no statistical significance is seen.

## 7 Conclusion

In this paper, we have addressed the feature extraction problem for learning to rank, and have proposed LifeRank, a linear feature extraction algorithm for ranking. LifeRank regards each dataset for ranking as a matrix, referred to as the *original matrix*. We then optimize a *transformation matrix* by minimizing a classic pairwise learning to rank loss function, so that we can discover the optimal one that matches the ranking task. Then a new matrix (dataset) can be generated by the product of original matrix and transformation matrix. Extensive experiments on benchmark datasets show the performance gains of LifeRank in comparison with the state-of-the-art algorithms.

The performance of LifeRank has been evaluated for RankSVM and Linear Regression. In future work, its benefits for other learning to rank algorithms could be

**Table 5: Effect of orthonormality constraints on datasets for  $k = 10$ .**Statistical significance shown for LifeRank against LifeRank<sup>NO</sup>

Performance for RankSVM											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
MQ2007											
<b>LifeRank</b>	<b>0.4037<sup>▲</sup></b>	<b>0.4023<sup>▲</sup></b>	<b>0.4089<sup>▲</sup></b>	<b>0.4117<sup>▲</sup></b>	<b>0.4161<sup>▲</sup></b>	<b>0.4215<sup>▲</sup></b>	<b>0.4264<sup>▲</sup></b>	<b>0.4312<sup>▲</sup></b>	<b>0.4370<sup>▲</sup></b>	<b>0.4423<sup>▲</sup></b>	<b>0.4634<sup>▲</sup></b>
LifeRank <sup>NO</sup>	0.3808	0.3881	0.3923	0.3939	0.3996	0.4044	0.4081	0.4148	0.4198	0.4269	0.4528
MQ2008											
<b>LifeRank</b>	<b>0.3763</b>	<b>0.4181<sup>▲</sup></b>	<b>0.4384<sup>△</sup></b>	<b>0.4612<sup>▲</sup></b>	<b>0.4796<sup>▲</sup></b>	<b>0.4900<sup>△</sup></b>	<b>0.4972<sup>△</sup></b>	<b>0.4625<sup>△</sup></b>	<b>0.2281<sup>▲</sup></b>	<b>0.2345<sup>▲</sup></b>	<b>0.4792<sup>△</sup></b>
LifeRank <sup>NO</sup>	0.3618	0.3987	0.4264	0.4429	0.4657	0.4807	0.4887	0.4552	0.2199	0.2236	0.4704
OHSUMED											
<b>LifeRank</b>	0.5518	<b>0.5373<sup>△</sup></b>	<b>0.5185<sup>▲</sup></b>	<b>0.5053<sup>▲</sup></b>	<b>0.4910<sup>▲</sup></b>	<b>0.4811<sup>△</sup></b>	<b>0.4774<sup>▲</sup></b>	<b>0.4699<sup>▲</sup></b>	<b>0.4692<sup>▲</sup></b>	<b>0.4663<sup>▲</sup></b>	<b>0.4505<sup>▲</sup></b>
LifeRank <sup>NO</sup>	<b>0.5703</b>	0.4900	0.4707	0.4673	0.4595	0.4522	0.4450	0.4399	0.4360	0.4312	0.4404
Performance for Linear Regression											
NDCG	@1	@2	@3	@4	@5	@6	@7	@8	@9	@10	MAP
MQ2007											
<b>LifeRank</b>	<b>0.3852<sup>▲</sup></b>	<b>0.3920<sup>▲</sup></b>	<b>0.3926<sup>▲</sup></b>	<b>0.3976<sup>▲</sup></b>	<b>0.4026<sup>▲</sup></b>	<b>0.4073<sup>▲</sup></b>	<b>0.4146<sup>▲</sup></b>	<b>0.4206<sup>▲</sup></b>	<b>0.4270<sup>▲</sup></b>	<b>0.4312<sup>▲</sup></b>	<b>0.4507<sup>▲</sup></b>
LifeRank <sup>NO</sup>	0.3584	0.3691	0.3729	0.3764	0.3816	0.3862	0.3918	0.3980	0.4034	0.4084	0.4338
MQ2008											
<b>LifeRank</b>	<b>0.3698<sup>▲</sup></b>	<b>0.3857<sup>▲</sup></b>	<b>0.4085<sup>▲</sup></b>	<b>0.3976<sup>▲</sup></b>	<b>0.4501<sup>▲</sup></b>	<b>0.4659<sup>▲</sup></b>	<b>0.4791<sup>▲</sup></b>	<b>0.4859<sup>▲</sup></b>	<b>0.4908<sup>▲</sup></b>	<b>0.4970<sup>▲</sup></b>	<b>0.4686<sup>▲</sup></b>
LifeRank <sup>NO</sup>	0.3295	0.3519	0.3689	0.3934	0.4144	0.4337	0.4446	0.4550	0.4615	0.4682	0.4346
OHSUMED											
<b>LifeRank</b>	<b>0.5330</b>	<b>0.4866</b>	<b>0.4680<sup>△</sup></b>	<b>0.4681<sup>▲</sup></b>	<b>0.4601<sup>▲</sup></b>	<b>0.4494<sup>▲</sup></b>	<b>0.4436<sup>▲</sup></b>	<b>0.4423<sup>▲</sup></b>	<b>0.4370<sup>▲</sup></b>	<b>0.4366<sup>▲</sup></b>	<b>0.4358<sup>▲</sup></b>
LifeRank <sup>NO</sup>	0.4571	0.4383	0.4111	0.4014	0.3915	0.3844	0.3784	0.3710	0.3689	0.3632	0.3963

analysed. Moreover, nonlinear feature extraction techniques like some kernel tricks could be incorporated in LifeRank to further improve its performance. Besides, we plan to try more learning to rank loss functions like some state-of-the-art listwise loss functions for performance gains of our algorithm. In addition, we believe it would be interesting to establish theoretical results on dimension reduction for ranking, including feature extraction and feature selection-based algorithms, especially concerning retrieval performance.

**Acknowledgements** We would like to thank our anonymous reviewers for valuable comments and suggestions.

This research was supported by Ahold Delhaize, Amsterdam Data Science, the Bloomberg Research Grant program, the Dutch national program COMMIT, Elsevier, the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement nr 312827 (VOX-Pol), the Google Faculty Research Award program, the Microsoft Research Ph.D. program, the Netherlands Institute for Sound and Vision, the Netherlands Organization for Scientific Research (NWO) under project nrs CI-14-25, 652-002.001, 612.001.551, 652.001.003, and Yandex.

All content represents the opinion of the author(s), which is not necessarily shared or endorsed by their respective employers and/or sponsors.

## References

1. Arfken GB (2013) *Mathematical Methods for Physicists*. Academic Press
2. Bach FR, Jordan MI (2005) A probabilistic interpretation of canonical correlation analysis. Tech. Rep. 688, Department of Statistics, University of California, Berkeley
3. Baeza-Yates R, Ribeiro-Neto B (1999) *Modern Information Retrieval*. Addison Wesley
4. Bertsekas DP (1999) *Nonlinear Programming*. Athena Scientific
5. Blum AL, Langley P (1997) Selection of relevant features and examples in machine learning. *Artif Intell* 97(1):245–271
6. Burges C, Shaked T, Renshaw E, Lazier A, Deeds M, Hamilton N, Hullender G (2005) Learning to rank using gradient descent. In: *ICML*, pp 89–96
7. Burges CJ, Ragno R, Le QV (2007) Learning to rank with nonsmooth cost functions. In: *NIPS*, pp 193–200
8. Busa-Fekete R, Kégl B, Éltető T, Szarvas G (2013) Tune and mix: Learning to rank using ensembles of calibrated multi-class classifiers. *Machine Learning* 93(2-3):261–292
9. Cao Y, Xu J, Liu TY, Li H, Huang Y, Hon HW (2006) Adapting ranking SVM to document retrieval. In: *SIGIR*, pp 186–193
10. Cao Z, Qin T, Liu TY, Tsai MF, Li H (2007) Learning to rank: From pairwise approach to listwise approach. In: *ICML*, pp 129–136
11. Chapelle O, Chang Y, Liu TY (2011) Future directions in learning to rank. *Journal of Machine Learning Research* 14:91–100
12. Chen W, Liu TY, Lan Y, Ma ZM, Li H (2009) Ranking measures and loss functions in learning to rank. In: *NIPS*, pp 315–323
13. Cossock D, Zhang T (2008) Statistical analysis of Bayes optimal subset ranking. *IEEE Transactions on Information Theory* 54(11):5140–5154
14. Crammer K, Singer Y (2001) Pranking with ranking. In: *NIPS*, pp 641–647
15. Freund Y, Iyer R, Schapire RE, Singer Y (2003) An efficient boosting algorithm for combining preferences. *Journal of Machine Learning Research* 4(1):933–969
16. Geng X, Liu T, Qin T, Li H (2007) Feature selection for ranking. In: *SIGIR*, pp 407–414
17. Gupta P, Rosso P (2012) Expected divergence based feature selection for learning to rank. In: *COLING*, pp 431–440
18. Haroon DR, Szedmak S, Shawe-Taylor J (2004) Canonical correlation analysis: An overview with application to learning methods. *Neural Computation* 16(12):2639–2664
19. Järvelin K, Kekäläinen J (2002) Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems* 20(4):422–446
20. Joachims T, Li H, Liu TY, Zhai C (2007) Learning to rank for information retrieval (LR4IR 2007). *SIGIR Forum* 41(2):58–62
21. Joachims T, Finley T, Yu CNJ (2009) Cutting-plane training of structural SVMs. *Machine Learning* 77(1):27–59
22. Joachims T, Swaminathan A, de Rijke M (2018) Deep learning with logged bandit feedback. In: *ICLR 2018*

23. Jolliffe I (2002) *Principal Component Analysis*. Springer Series in Statistics, Springer
24. Kendall MG (1948) *Rank Correlation Methods*. C. Griffin
25. Koren Y, Bell R, Volinsky C (2009) Matrix factorization techniques for recommender systems. *Computer* 42(8):30–37
26. Lai H, Pan Y, Tang Y, Yu R (2013) FSMRank: Feature selection algorithm for learning to rank. *IEEE Transactions on Neural Networks and Learning Systems* 24(6):940–952
27. Lan Y, Guo J, Cheng X, Liu TY (2012) Statistical consistency of ranking methods in a rank-differentiable probability space. In: *NIPS*, pp 1232–1240
28. Lange K (2010) Singular value decomposition. In: *Numerical Analysis for Statisticians*, Springer, pp 129–142
29. Laporte L, Flamary R, Canu S, Déjean S, Mothe J (2014) Nonconvex regularizations for feature selection in ranking with sparse SVM. *IEEE Transactions on Neural Networks and Learning Systems* 25(6):1118–1130
30. Lawson C, Hanson R (1995) *Solving Least Square Problems*, Classics in Applied Mathematics, vol 15. SIAM, Philadelphia
31. Li P, Wu Q, Burges CJ (2007) McRank: Learning to rank using multiple classification and gradient boosting. In: *NIPS*, pp 897–904
32. Liu TY (2009) Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval* 3(3):225–331
33. Liu TY (2011) *Learning to Rank for Information Retrieval*. Springer Science & Business Media
34. Metzler DA (2007) Automatic feature selection in the markov random field model for information retrieval. In: *CIKM*, ACM, pp 253–262
35. Motoda H, Liu H (2002) Feature selection, extraction and construction. *Communication of IICM (Institute of Information and Computing Machinery, Taiwan)* 5:67–72
36. Mukuta Y, Harada T (2014) Probabilistic partial canonical correlation analysis. In: *ICML*, pp 1449–1457
37. Naini KD, Altingövde IS (2014) Exploiting result diversification methods for feature selection in learning to rank. In: *ECIR*, Springer, pp 455–461
38. Ng AY (2004) Feature selection, L1 vs. L2 regularization, and rotational invariance. In: *ICML*, pp 78–82
39. Niu S, Guo J, Lan Y, Cheng X (2012) Top-K learning to rank: Labeling, ranking and evaluation. In: *SIGIR*, pp 751–760
40. Pan F, Converse T, Ahn D, Salvetti F, Donato G (2009) Feature selection for ranking using boosted trees. In: *CIKM*, ACM, pp 2025–2028
41. Platt JC, Barr AH (1988) *Constrained differential optimization for neural networks*. Tech. Rep. TR-88-17, Department of Computer Science, California Institute of Technology
42. Qin T, Liu TY (2013) Introducing LETOR 4.0 datasets. arXiv 1306.2597
43. Qin T, Liu TY, Xu J, Li H (2010) LETOR: A benchmark collection for research on learning to rank for information retrieval. *Information Retrieval* 13(4):346–374

44. Roweis ST, Saul LK (2000) Nonlinear dimensionality reduction by locally linear embedding. *Science* 290:2323–2326
45. Schölkopf B, Smola A, Müller KR (1998) Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation* 10(5):1299–1319
46. Schuth A, Oosterhuis H, Whiteson S, de Rijke M (2016) Multileave gradient descent for fast online learning to rank. In: *WSDM 2016: The 9th International Conference on Web Search and Data Mining*, ACM, pp 457–466
47. Severyn A, Moschitti A (2015) Learning to rank short text pairs with convolutional deep neural networks. In: *SIGIR*, ACM, pp 373–382
48. Shalit U, Chechik G (2014) Coordinate-descent for learning orthogonal matrices through Givens rotations. In: *ICML*, pp 548–556
49. Shivanna R, Bhattacharyya C (2014) Learning on graphs using orthonormal representation is statistically consistent. In: *NIPS*, pp 3635–3643
50. Tenenbaum JB, de Silva V, Langford JC (2000) A global geometric framework for nonlinear dimensionality reduction. *Science* 290(5500):2319–2323
51. Tipping ME, Bishop CM (1999) Probabilistic principal component analysis. *Journal of the Royal Statistical Society, Series B* 61(3):611–622
52. Tsai MF, Liu TY, Qin T, Chen HH, Ma WY (2007) FRank: A ranking method with fidelity loss. In: *SIGIR*, ACM, pp 383–390
53. Valizadegan H, Jin R, Zhang R, Mao J (2009) Learning to rank by optimizing NDCG measure. In: *NIPS*, pp 1883–1891
54. Volkovs M, Zemel RS (2009) Boltzrank: Learning to maximize expected ranking gain. In: *ICML*, pp 1089–1096
55. Wang S, Wu Y, Gao BJ, Wang K, Lauw HW, Ma J (2015) A cooperative coevolution framework for parallel learning to rank. *IEEE Transactions on Knowledge and Data Engineering* 27(12):3152–3165
56. Weston J, Mukherjee S, Chapelle O, Pontil M, Poggio T, Vapnik V (2000) Feature selection for SVMs. In: *NIPS*, pp 668–674
57. Wolf L, Bileschi S (2005) Combining variable selection with dimensionality reduction. In: *CVPR*, pp 801–806
58. Wyse N, Dubes R, Jain A (1980) A critical evaluation of intrinsic dimensionality algorithms. In: Gelsema E, Kanal L (eds) *Pattern Recognition in Practice*. Proc. workshop Amsterdam, May 1980, North-Holland, pp 415–425
59. Xu J, Li H (2007) AdaRank: A boosting algorithm for information retrieval. In: *SIGIR*, ACM, pp 391–398
60. Yu H, Oh J, Han W (2009) Efficient feature weighting methods for ranking. In: *CIKM*, ACM, pp 1157–1166
61. Yue Y, Finley T, Radlinski F, Joachims T (2007) A support vector method for optimizing average precision. In: *SIGIR*, ACM, pp 271–278