

# ILMART: Interpretable Ranking with Constrained LambdaMART

Claudio Lucchese  
claudio.lucchese@unive.it  
Ca' Foscari University of Venice  
Italy

Franco Maria Nardini  
francomaria.nardini@isti.cnr.it  
ISTI-CNR  
Italy

Salvatore Orlando  
orlando@unive.it  
Ca' Foscari University of Venice  
Italy

Raffaele Perego  
raffaele.perego@isti.cnr.it  
ISTI-CNR  
Italy

Alberto Veneri  
alberto.veneri@unive.it  
Ca' Foscari University of Venice  
ISTI-CNR  
Italy

## ABSTRACT

Interpretable Learning to Rank (LtR) is an emerging field within the research area of explainable AI, aiming at developing intelligible and accurate predictive models. While most of the previous research efforts focus on creating post-hoc explanations, in this paper we investigate how to train effective and *intrinsically-interpretable* ranking models. Developing these models is particularly challenging and it also requires finding a trade-off between ranking quality and model complexity. State-of-the-art rankers, made of either large ensembles of trees or several neural layers, exploit in fact an unlimited number of feature interactions making them black boxes. Previous approaches on intrinsically-interpretable ranking models address this issue by avoiding interactions between features thus paying a significant performance drop with respect to full-complexity models. Conversely, ILMART, our novel and interpretable LtR solution based on LambdaMART, is able to train effective and intelligible models by exploiting a limited and controlled number of pairwise feature interactions. Exhaustive and reproducible experiments conducted on three publicly-available LtR datasets show that ILMART outperforms the current state-of-the-art solution for interpretable ranking of a large margin with a gain of nDCG of up to 8%.

## CCS CONCEPTS

• Information systems → Learning to rank; • Computing methodologies → Boosting.

## KEYWORDS

Interpretable Ranking, Interpretable Boosting, LambdaMART

### ACM Reference Format:

Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, and Alberto Veneri. 2022. ILMART: Interpretable Ranking with Constrained LambdaMART. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

SIGIR '22, July 11–15, 2022, Madrid, Spain

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8732-3/22/07...\$15.00

<https://doi.org/10.1145/3477495.3531840>

2022, Madrid, Spain. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3477495.3531840>

## 1 INTRODUCTION

Learning to Rank (LtR) is a vast research area focusing on learning effective models for solving ranking tasks. Although being effective, state-of-the-art LtR techniques do not deal with the *interpretability* of the learned model, i.e., the possibility given to humans to easily understand the reasoning behind predictions made by the model. The need for learning trustworthy models in LtR is becoming a topic of increasing interest, which tries to address important concerns about privacy, transparency, and fairness of ranking processes behind many everyday human activities.

For what regards the fairness of ranking algorithms, several solutions have been proposed to create fairer models, where the term *fair* takes different meanings according to the metric employed [2]. However, focusing only on devising a fairer model does not imply to be compliant with new regulations that are emerging. An example of that is *The Right of Explanation*, which allows each European citizen to claim a justification for every algorithm decision made with their data that affects their life [13]. Therefore, for highly-sensitive scenarios new LtR approaches have to be investigated to train ranking models that do not jeopardize interpretability in the name of effectiveness.

In literature, two main lines of work dealing with the interpretation of ranking models are investigated. The first concerns the creation of post-hoc explanations for black-box models, while the second focuses on the design of intrinsically-interpretable models. In the first line, the goal is the definition of techniques that allow, given a ranking model, to generate post-hoc explanations of the decisions made. Here, some of the explanation techniques used in regression or classification tasks have been adapted to the ranking scenario. For example, two recent contributions [17, 19] present two different approaches to adapt LIME [15] to the ranking task. The work by Singhs and Anand focuses only on pointwise LtR techniques and treats the ranking problem as a classification problem, where the probability of a document to be relevant is computed with respect to its score, and a simple linear SVM is used as a local explanation model [17]. The work by Verma and Ganguly evaluates the consistency and correctness of the explanations generated by LIME by varying the sampling strategies used to create the local neighborhood of a document [19]. In the same line, Fernando *et al.* propose an approach to reuse SHAP [12] for the ranking task [5].

Even though post-hoc explanations can be seen as a viable way to inspect and audit accurate ranking models, they can produce explanations that are not faithful to internal computations of the model, and thus they can be misleading [16]. For this reason, Zhuang *et al.* pioneer the problem of developing intrinsically-interpretable ranking models [20]. In detail, they propose *Neural RankGAM*, an adaptation of GAM [6] to the ranking task. Although Neural RankGAM nicely exploits the interpretability of GAMs, the authors show that the ranking performance achieved is significantly worse than those provided by state-of-the-art black-box models. However, it is worth noticing that in other contexts, such as in regression and classification tasks, other intrinsically interpretable models appear to be as accurate as other black-box models. A notable example of an interpretable and effective solution is Explainable Boosting Machine (EBM), a C++ implementation of the  $GA^2M$  algorithm [9], where each component of the Generalized Additive Models (GAM) are learned through a mixed bagging-boosting procedure [8].

In this paper, we contribute to the development of intrinsically-interpretable ranking models with ILMART, a novel LtR solution based on LambdaMART. ILMART creates effective and intelligible models by encompassing a limited and controlled number of pairwise feature interactions. We evaluate ILMART on three publicly-available LtR datasets. The results of our reproducible experiments show that ILMART outperforms the current state-of-the-art solution for interpretable ranking of a large margin, with a gain in terms of nDCG of up to 8% with respect to Neural RankGAM.

## 2 INTERPRETABLE LAMBDA MART

ILMART produces ensemble-based ranking models having the structure of a GAM, which is considered an interpretable model [8]. Before introducing ILMART, we review the main aspects of GAMs to highlight similarities and differences with respect to the newly proposed technique.

**Generalized Additive Models.** Let  $\mathbf{x} \in \mathbb{R}^d$  be an instance in a  $d$ -dimensional vector space, with  $x_j$  being the value of its  $j$ -th feature. An instance  $\mathbf{x}$  is associated with a target label  $y$  generated by some unknown target function  $g(\cdot)$ , i.e.,  $g(\mathbf{x}) = y$ . In its basic form, a GAM models the target variable  $y$  as a combination of  $p$  main effects and  $K$  non-zero interaction effects, as defined in [6]:

$$l(\mu(\mathbf{x})) = \alpha + \sum_{j=1}^p s_j(x_j) + \sum_{j=1, k=1}^{j=p, k=p} s_{jk}(x_j, x_k)$$

Where  $\mu(\mathbf{x}) = E[y|\mathbf{x}]$ , with  $y$  following a distribution of the exponential family, and  $l(\cdot)$  being the so-called *link function* that is able to describe the relationship between the  $p$  main effects and the  $K$  interactions effects with the expected value of  $y$ . In addition, the number of non-zero interaction effects is equal to  $K$ , i.e.  $|\{s_{jk}(\cdot)|s_{jk}(\cdot) \neq 0\}| = K$ . Usually, the  $p$  main effects and the  $K$  interactions are modeled through simple third order splines. From the explainability point of view, the analyst is promptly provided with a plot of each basis function  $s_j$  that clearly describes the relation between each feature and the target variable  $y$ , while  $s_{jk}$  values represent the contribution of the interactions between feature pairs. For classification and regression tasks, (shallow) decision trees instead of splines were investigated in [9].

**Interpretable LambdaMART.** We now introduce ILMART, our novel method based on LambdaMART to learn ranking models that are both accurate and interpretable. ILMART presents similarities with GAM, and builds the prediction  $\hat{y}$  of an unknown target variable  $y$  as follows:

$$\hat{y} = \underbrace{\sum_{j \in \mathcal{J}} \tau_j(x_j)}_{R \text{ trees}} + \underbrace{\sum_{(i,j) \in \mathcal{K}} \tau_{ij}(x_i, x_j)}_{T \text{ trees}} \quad (1)$$

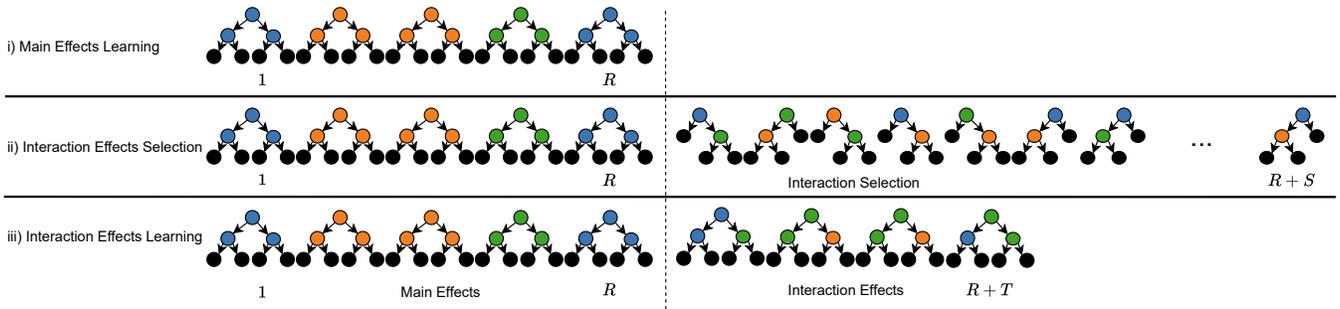
where  $\tau_j(x_j)$  and  $\tau_{ij}(x_i, x_j)$  are ensembles of trees. Each tree  $t \in \tau_j(x_j)$  models a *main effect* in which the only feature allowed for each split is  $x_j$ .

On the other hand, each tree  $t \in \tau_{ij}(x_i, x_j)$  models the *interaction effects*, where the only features allowed in a split of tree  $t$  are either  $x_i$  or  $x_j$ . Similarly to the notation used to describe a GAM,  $p$  denotes the number of distinct main effects  $x_i$ , whereas  $K$  denotes the number of distinct interaction effects  $(x_i, x_j)$ .

To learn the ensemble of trees modeled by Eq. (1), we adopt a modified version of the LambdaMART boosting algorithm. Unlike GAM, we do not make any assumption on the distribution of  $y$  to guide the learning process. The learning strategy to generate the main effects  $t \in \tau_j(x_j)$  and the interaction effects  $t \in \tau_{ij}(x_i, x_j)$  in Equation 1 is made of three steps depicted in Fig. 1: *i*) Main Effects Learning, which learns a set of trees, each one working on a main effect (single feature); *ii*) Interaction Effects Selection, which selects the top- $K$  most important interactions effects (feature pairs); *iii*) Interaction Effects Learning, which learns a set of trees, each one exploiting one of the interaction effects selected by the previous step.

*Main Effects Learning.* In the initial step, illustrated on top of Fig. 1, we learn an ensemble of  $R$  trees modeling the main effects. To this end, we constrain the LambdaMART boosting procedure to use a single feature per tree in the first  $R$  boosting rounds. In other words, at a given boosting round of LambdaMART, when feature  $x_j$  is chosen for the root's split of a tree  $t$ , we force the algorithm to use  $x_j$  for all the further splits in  $t$ , until a stopping criterion for tree growth is reached. LambdaMART stops its boosting rounds when no improvements are observed on the validation set. Eventually, we obtain an ensemble of  $R$  trees only using for their splits a set  $\mathcal{J}$  of  $p$  features,  $p \leq d$  and  $p \leq R$ , modeling the main effects of Equation 1. Thus,  $p$  is not an input parameter of ILMART, but it is the number of distinct features used by the ensemble of  $R$  trees modeling the main effects only. We name this initial model ILMART. It is worth noting that, as a side effect of this first step, we achieve the additional effect of potentially reducing the cardinality of the features, from  $d$  to  $p$ .

*Interaction Effects Selection.* In the second step, we still exploit the LambdaMART boosting procedure to *select* the top- $K$  interaction effects, as illustrated in the middle of Fig. 1. Specifically, we continue the learning of the ILMART model obtained so far, by enabling the interactions between all possible pairs of  $p$  features identified by the previous step, according to the interaction heredity principle [4]. This is accomplished by constraining LambdaMART to fit trees with only 3 leaves (and 2 splits), where the 2 splits use a pair of distinct



**Figure 1: ILMART learning process.** Each color corresponds to a different feature used in the splits, and nodes colored in black are leaves. The trees for the main and interaction effects are represented with the same depth for illustration purposes only.

features. The boosting round is stopped when the size of the set  $\mathcal{K}$ , composed of the distinct feature pairs used by the ensemble of  $S$  trees, is exactly  $K$ , where  $K \leq \binom{p}{2}$ . The  $S$  trees are generated solely for the purpose of identifying the interaction effects and they are discarded afterwards.

Note that for selecting the main  $K$  interaction effects, we rely on the capability of LambdaMART in identifying, at each boosting round, the features and the splits with the highest contributions in minimizing the loss function.

Thus, the interaction effects selected by the boosting procedure are likely to be generated in order of importance.

*Interaction effects learning.* Finally, the last step sketched on the bottom of Fig. 1 consists in adding to the initial ILMART model  $T$  new trees learned by constraining LambdaMART to use only the top- $K$  feature pairs identified in the previous step. The number  $T$  of new learned trees is not an input of algorithm, but it is chosen on the basis of the validation set. The final model, obtained by adding to ILMART the  $T$  trees working on the  $K$  interaction effects is named ILMART <sub>$i$</sub> .

### 3 EXPERIMENTAL EVALUATION

We experimentally assess the performance of ILMART and ILMART <sub>$i$</sub>  on three public datasets for LTR, namely ISTEELLA-S, WEB30K, and YAHOO. The ISTEELLA-S dataset [11] includes 33,018 queries with an average of 103 documents per query. Each document-query pair is represented by 220 features. The WEB30K dataset [14] is composed of more than 30,000 queries (in 5 folds), with an average of 120 documents per query and 136 features per query-document pair. The YAHOO dataset [3] is composed of two sets, including document-query pairs with 699 features. The feature vectors of the three datasets are labeled with relevance judgments ranging from 0 (irrelevant) to 4 (perfectly relevant). In our experiments, we use train/validation and test splits from fold 1 of the WEB30K dataset and data from “set 1” of the YAHOO dataset.

**Competitors.** We compare the performance of ILMART and ILMART <sub>$i$</sub>  with the following competitors:

- Neural RankGAM (NRGAM), a neural-based approach to learn GAMs for ranking. The approach exploits standalone neural networks to instantiate sub-models for each individual feature. It is the current state-of-the-art technique for learning interpretable models for ranking [20].

- EBM, an efficient implementation of  $GA^2M$  using a mixed bagging-boosting procedure [9]. In this case the learned model is used as a pointwise ranker, since the implementations available are only made to solve regression and classification tasks. We denote with EBM <sub>$i$</sub>  the EBM model using interaction effects.

**Metrics.** We measure the performance in terms of normalized Discounted Cumulative Gain (nDCG) at three different cutoffs, i.e.,  $\{1, 5, 10\}$ . We compute the nDCG metric by employing exponential weighing of the relevance [1]. By default, queries with missing relevant documents have been assigned a nDCG = 1.0. The statistical significance is computed with a two-sided Fisher’s randomization test [18] with significance level  $p < 0.05$ . The test is computed using the RankEval library [10].

**Implementation and training settings.** We implement ILMART as an extension of the LightGBM library<sup>1</sup> [7]. The extension works by adding the possibility to constraint the boosting procedure to use only a limited number of features per tree. The source code used in our experiments along with the models trained on the three public datasets is publicly available online<sup>2</sup>.

The training of ILMART and ILMART <sub>$i$</sub>  optimizes nDCG@10. We perform hyper-parameter tuning by varying the number of leaves in the trees in  $\{32, 64, 128\}$  and the learning rate in  $\{0.001, 0.01, 0.1\}$ . Early stopping, i.e., stopping boosting if no improvement on the validation set is observed for 100 consecutive rounds, is used for training in step one (main effects learning) and three (interaction effects learning). The optimal ILMART models learned with the settings above are made of 914, 895 and 1,969 trees for the WEB30K, YAHOO, and the ISTEELLA-S datasets, respectively. On the same datasets, the three optimal ILMART <sub>$i$</sub>  models are made of 1,365, 1,215 and 3,024 trees, respectively. Neural RankGAM is also trained by optimizing nDCG@10. We use the public implementation of NRGAM available in Keras<sup>3</sup> and made available by the authors of the original paper [20]. We train NRGAM by using the same settings reported by the authors. In addition, we employ a batch size of 128, and we train the networks with 3,000 epochs for WEB30K, 700 for YAHOO and 1,000 for ISTEELLA-S, with early stopping set to 100 epochs. We also learn EBM by optimizing nDCG@10. The performance of EBM is fine-tuned by varying the number of outer-bags in  $\{5, 10, 15\}$  for

<sup>1</sup><https://github.com/microsoft/LightGBM>

<sup>2</sup><https://github.com/veneres/ilmart>

<sup>3</sup><https://github.com/tensorflow/ranking/issues/202>

**Table 1: Performance comparison in terms of nDCG. Statistically significant improvements w.r.t. NRGAM are marked with an asterisk (\*). Best results are reported in bold.**

Dataset	Method	nDCG			$p$	$K$
		@1	@5	@10		
WEB30K	NRGAM	44.80	43.72	45.73	136	0
	EBM	41.70	43.14	45.66	136	0
	EBM <sub>i</sub>	46.51*	45.96*	48.01*	136	50
	ILMART	45.17	45.00*	47.05*	79	0
	ILMART <sub>i</sub>	<b>48.94*</b>	<b>47.76*</b>	<b>49.55*</b>	79	46
YAHOO	NRGAM	67.92	70.13	75.07	699	0
	EBM	70.33*	72.77*	77.29*	699	0
	EBM <sub>i</sub>	70.22*	72.85*	77.36*	699	50
	ILMART	70.57*	72.64*	77.21*	153	0
	ILMART <sub>i</sub>	<b>70.80*</b>	<b>73.20*</b>	<b>77.57*</b>	153	48
ISTELLA-S	NRGAM	63.37	63.04	69.18	220	0
	EBM	57.75	59.67	66.90	220	0
	EBM <sub>i</sub>	62.29	63.05	69.71	220	50
	ILMART	66.99*	66.29*	72.36*	106	0
	ILMART <sub>i</sub>	<b>69.06*</b>	<b>68.22*</b>	<b>74.04*</b>	106	50

the two versions with and without interactions. In our experiments, we employ a publicly-available implementation of EBM<sup>4</sup>. For what regards the two strategies encompassing interaction effects, i.e., ILMART<sub>i</sub> and EBM<sub>i</sub>, we limit the number of maximum interactions allowed to 50.

**Experimental results.** Table 1 reports the results of the evaluation of ILMART, ILMART<sub>i</sub> and their state-of-the-art competitors. For all the datasets considered, results show that ILMART achieves a statistically significant improvement in terms of nDCG for almost all the cutoffs except the case of nDCG@1 on WEB30K. More importantly, ILMART<sub>i</sub> always achieves the best performance in terms of nDCG that always result in a statistically significant improvement over NRGAM. The best improvements of ILMART<sub>i</sub> over *all* the competitors is obtained on ISTELLA-S, where nDCG@10 improves of up to 6% w.r.t. EBM<sub>i</sub>. On WEB30K, we observe that ILMART<sub>i</sub> improves nDCG@10 of up to 8% w.r.t NRGAM, although the best competitor is, in this case, EBM<sub>i</sub>, whose performance still lags behind ILMART<sub>i</sub>. On WEB30K, EBM<sub>i</sub> also achieves statistically significant improvements in terms of nDCG with respect to NRGAM. Even though EBM<sub>i</sub> is used as a pointwise ranker, these differences confirm the efficacy of tree-based approaches in Ltr and the need to include pairwise interactions to learn more effective ranking functions.

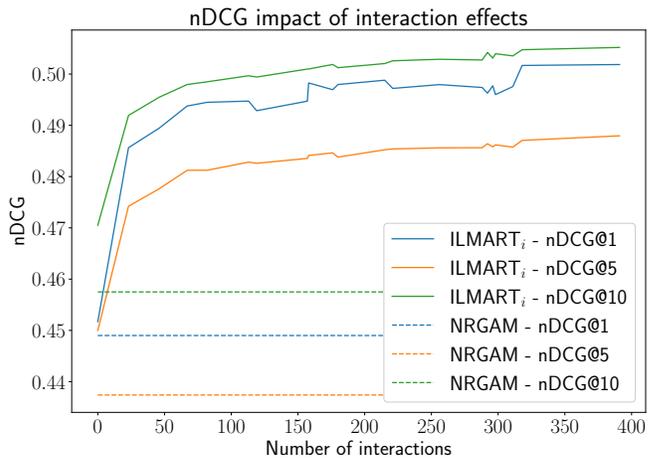
Furthermore, it is worth highlighting that our proposed techniques perform a feature selection during the *main effects learning* phase, while the competitors leave the feature selection to the final user. This is an important difference between ILMART and its competitors as a high number of features used do not allow to easily understand the global behavior of the model. For example, the YAHOO dataset comes with a large number of features, i.e., 699. ILMART and ILMART<sub>i</sub> effectively reduce the initial set of 699 features to 153 (column  $p$  in Table 1) with, at the same time, a significant

<sup>4</sup><https://interpret.ml/docs/ebm.html>

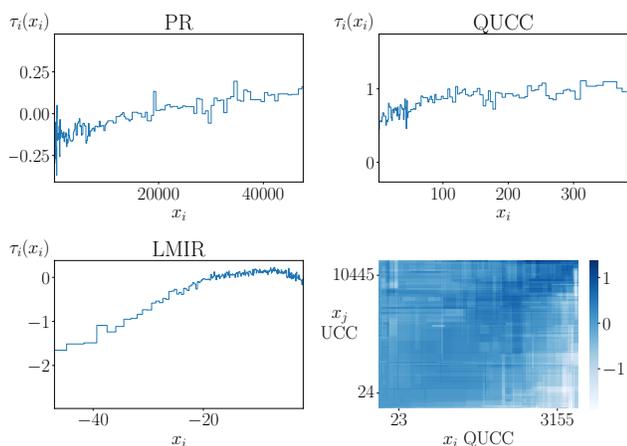
increase of the effectiveness of the learned model. In addition, we point out that eventually ILMART<sub>i</sub> may have a number of interaction effects  $K$  smaller than 50 due to the training of the new  $T$  trees which may not include all the  $K$  pairs selected.

We now investigate how the number of interactions added to ILMART<sub>i</sub> can influence its prediction accuracy. Fig. 2 compares the results of ILMART<sub>i</sub> with the ones of NRGAM by reporting their results in terms of nDCG. We observe that the ranking accuracy of ILMART<sub>i</sub> increases steeply when the first interaction effects are added; then, the gain obtained by adding new interaction trees saturates. This means that ILMART<sub>i</sub> is able to significantly improve the accuracy of an interpretable Ltr model by adding a small number of interaction effects, thus without hindering the overall interpretability of the model. Even though the addition of interactions effects improves a lot the accuracy of the model, we acknowledge that the difference from the full-complexity model LambdaMART is still marked. For example, on WEB30K an optimized black-box LambdaMART model achieves a nDCG@10 close to 52%.

Finally, to show how the main and interaction effects can be explored, Fig. 3 presents three main effects and one interaction effects learned with ILMART<sub>i</sub> trained over WEB30K. The three main effects selected are the ones associated with the features with highest feature importance, i.e., PageRank (PR), Query-url Click Count (QUCC), and Language Model for Information Retrieval (LMIR), while the interaction effect between URL Click Count (UCC) and QUCC is the one with the highest average contribution for each value of  $(x_i, x_j)$ . The 2D plots and the heatmap are built by aggregating all the trees using the same features, i.e., by representing the contribution of each  $\tau_j(x_j)$  and each  $\tau_{ij}(x_i, x_j)$ . From the plots, the analyst can easily understand the exact behavior of the model at varying values of each feature. For example, taking into account the interaction effect, the higher the value of UCC and QUCC the larger is the positive contribution of  $\tau_{ij}(x_i, x_j)$ .



**Figure 2: nDCG by varying the number of interaction effects in ILMART<sub>i</sub> on the WEB30K dataset.**



**Figure 3: Plots of the four most important main and interaction effects of ILMART<sub>i</sub> on the WEB30K dataset.**

## 4 CONCLUSIONS AND FUTURE WORK

We presented ILMART, a novel and interpretable Ltr solution based on LambdaMART, that allows to train effective and intelligible models. We experimentally assessed the performance of ILMART on three public datasets. Results show that it achieves statistically significant improvements in terms of nDCG of up to 8% w.r.t. the current state-of-the-art solution, i.e., NRGAM. We also showed that ILMART is able to effectively exploit feature interactions without hindering the overall interpretability of the model. With this work we laid the foundations for the development of new accurate Ltr models based on boosting methods that are also interpretable.

As future work, we plan to investigate novel techniques based on boosting for building effective and interpretable models for ranking and new approaches for effectively identifying main and interaction effects during the learning of the ranking models, with the main goal of filling the performance gap with the black-box models.

## REFERENCES

- [1] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. 2005. Learning to Rank Using Gradient Descent. In *Proceedings of the 22nd International Conference on Machine Learning* (Bonn, Germany) (ICML '05). Association for Computing Machinery, New York, NY, USA, 89–96. <https://doi.org/10.1145/1102351.1102363>
- [2] Carlos Castillo. 2019. Fairness and Transparency in Ranking. *ACM SIGIR Forum* 52, 2 (Jan. 2019), 64–71. <https://doi.org/10.1145/3308774.3308783>
- [3] Olivier Chapelle and Yi Chang. 2011. Yahoo! Learning to Rank Challenge Overview. In *Proceedings of the Learning to Rank Challenge (Proceedings of Machine Learning Research, Vol. 14)*, Olivier Chapelle, Yi Chang, and Tie-Yan Liu (Eds.). PMLR, Haifa, Israel, 1–24. <https://proceedings.mlr.press/v14/chapelle11a.html>
- [4] D. R. Cox. 1984. Interaction. *International Statistical Review / Revue Internationale de Statistique* 52, 1 (1984), 1–24. <https://doi.org/10.2307/1403235> Publisher: [Wiley, International Statistical Institute (ISI)].
- [5] Zeon Trevor Fernando, Jaspreet Singh, and Avishek Anand. 2019. A study on the Interpretability of Neural Retrieval Models using DeepSHAP. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'19)*. Association for Computing Machinery, New York, NY, USA, 1005–1008. <https://doi.org/10.1145/3331184.3331312>
- [6] Trevor Hastie and Robert Tibshirani. 1986. Generalized Additive Models. *Statist. Sci.* 1, 3 (1986), 297–310. <https://doi.org/10.1214/ss/1177013604>
- [7] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 3149–3157.
- [8] Yin Lou, Rich Caruana, and Johannes Gehrke. 2012. Intelligible models for classification and regression. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '12*. ACM Press, Beijing, China, 150. <https://doi.org/10.1145/2339530.2339556>
- [9] Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. 2013. Accurate intelligible models with pairwise interactions. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '13*. ACM Press, Chicago, Illinois, USA, 623. <https://doi.org/10.1145/2487575.2487579>
- [10] Claudio Lucchese, Cristina Ioana Muntean, Franco Maria Nardini, Raffaele Perego, and Salvatore Trani. 2017. RankEval: An Evaluation and Analysis Framework for Learning-to-Rank Solutions. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (Shinjuku, Tokyo, Japan) (SIGIR '17)*. Association for Computing Machinery, New York, NY, USA, 1281–1284. <https://doi.org/10.1145/3077136.3084140>
- [11] Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Fabrizio Silvestri, and Salvatore Trani. 2016. Post-Learning Optimization of Tree Ensembles for Efficient Ranking. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (Pisa, Italy) (SIGIR '16)*. Association for Computing Machinery, New York, NY, USA, 949–952. <https://doi.org/10.1145/2911451.2914763>
- [12] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).
- [13] Eugenia Politou, Efthimios Alepis, and Constantinos Patsakis. 2018. Forgetting personal data and revoking consent under the GDPR: Challenges and proposed solutions. *Journal of Cybersecurity* 4, 1 (03 2018). <https://doi.org/10.1093/cybsec/tyy001>
- [14] Tao Qin and Tie-Yan Liu. 2013. Introducing LETOR 4.0 Datasets. *CoRR* abs/1306.2597 (2013). <http://arxiv.org/abs/1306.2597>
- [15] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, San Francisco California USA, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- [16] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (May 2019), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- [17] Jaspreet Singh and Avishek Anand. 2019. EXS: Explainable Search Using Local Model Agnostic Interpretability. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. ACM, Melbourne VIC Australia, 770–773. <https://doi.org/10.1145/3289600.3290620>
- [18] Mark D. Smucker, James Allan, and Ben Carterette. 2007. A comparison of statistical significance tests for information retrieval evaluation. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management - CIKM '07*. ACM Press, Lisbon, Portugal, 623. <https://doi.org/10.1145/1321440.1321528>
- [19] Manisha Verma and Debasis Ganguly. 2019. LIRME: Locally Interpretable Ranking Model Explanation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'19)*. Association for Computing Machinery, New York, NY, USA, 1281–1284. <https://doi.org/10.1145/3331184.3331377>
- [20] Honglei Zhuang, Xuanhui Wang, Michael Bendersky, Alexander Grushetsky, Yonghui Wu, Petr Mitrichev, Ethan Sterling, Nathan Bell, Walker Ravina, and Hai Qian. 2021. Interpretable Ranking with Generalized Additive Models. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM '21)*. Association for Computing Machinery, New York, NY, USA, 499–507. <https://doi.org/10.1145/3437963.3441796>