MO-LightGBM: A Library for Multi-objective Learning to Rank with LightGBM

Chaosheng Dong*
chaosd@amazon.com
Amazon
Seattle, Unite States

Michinari Momma
michi@amazon.com
Amazon
Seattle, Unite States

Abstract

This paper introduces MO-LightGBM, an open-source library built upon LightGBM, specifically designed to offer an integrated, versatile, and easily adaptable framework for Multi-objective Learning to Rank (MOLTR). MO-LightGBM supports diverse Multi-objective optimization (MOO) settings and incorporates 12 state-of-the-art optimization strategies. Its modular architecture enhances usability and flexibility, allowing researchers and practitioners to easily develop new MOO methodologies, perform rigorous comparisons with existing techniques, and effectively deploy MOO algorithms in practical ranking applications. We illustrate the utility of MO-LightGBM through a Bi-objective Learning to Rank example and present visualizations of the results. MO-LightGBM is available at https://github.com/amazon-science/MO-LightGBM.

CCS Concepts

• **Information systems → Information retrieval**.

Keywords

Multi-objective optimization, Learning to Rank, LightGBM

ACM Reference Format:
Chaosheng Dong and Michinari Momma. 2025. MO-LightGBM: A Library for Multi-objective Learning to Rank with LightGBM. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25), July 13–18, 2025, Padua, Italy*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3726302.3731937

1 Introduction

Multi-objective Learning to Rank (MOLTR) has gained increasing attention in industry as it acknowledges the multifaceted nature of relevance in ranking tasks [4, 10, 14, 15, 18, 22, 26]. Traditional LTR methods typically focus on a single relevance criterion, limiting their ability to capture the complex trade-offs inherent in real-world applications. By incorporating multiple relevance criteria, MOLTR offers a more comprehensive and accurate representation of user preferences, leading to improved ranking quality across diverse objectives.

In the field of MOLTR, numerous methods have been developed to address the complexity of Multi-objective data. Early approaches primarily focused on differentiating between relevant and irrelevant labels, as seen in works such as [5, 6], and analyzing pairwise label interactions [8]. Further advancements integrated tree-based methods with MOO algorithms [3, 14, 20], enhancing their ability to model trade-offs effectively. Specifically, gradient boosted decision trees (GBDT) [7] demonstrated their strength in capturing complex non-linear relationships while maintaining efficiency, accuracy, and interpretability. When combined with MOO techniques, these methods provided researchers with powerful tools to navigate competing objectives in ranking tasks. This integration marked a significant step forward in developing more sophisticated solutions for MOLTR challenges.

Related work As Multi-objective machine learning problems become increasingly prevalent, tree-based methods have been developed to handle multiple outputs. [1, 21] extended impurity measures for binary classification and ranking to a Multi-objective setting for node splitting, but their approaches rely on random forests rather than boosting. In contrast, [27] introduced GBDT-MO, a general GBDT approach for multi-objective learning, where each leaf predicts all or a subset of variables. However, this method aggregates objective gains across outputs using linear scalarization [17] with equal weights for different objectives, limiting its ability to capture the full Pareto-optimal front. To date, no open-source MOLTR frameworks leverage GBDT with general MOO algorithms. To bridge this gap, we propose a new framework that integrates MOO algorithms within GBDT-based ranking models. LightGBM [11] has emerged as a leading GBDT variant due to its efficient histogram-based training, making it an ideal foundation for MOLTR. Consequently, we develop the MO-LightGBM framework to support multi-objective LTR in a scalable and extensible manner.

This paper presents two main contributions to the field:

- **First open-source MOLTR framework based on GBDT:** Unlike existing ranking methods that optimize a single objective, MO-LightGBM integrates state-of-the-art MOO strategies, enabling a flexible and efficient approach to handling competing ranking criteria. The framework facilitates fair comparisons among MOO methods and provides a modular design that simplifies the development and application of new multi-objective LTR algorithms.
- **Demonstration of MOO in Learning to Rank:** MO-LightGBM is empirically validated through a Bi-objective LTR task, showcasing the effectiveness of different preference- and constraint-based MOO methods. Using the Istella LETOR dataset, the paper demonstrates how MO-LightGBM supports various trade-off specifications, visualizes Pareto-optimal

*Corresponding author.

SIGIR '25, Padua, Italy

ACM ISBN 979-8-4007-1592-1/2025/07
https://doi.org/10.1145/3726302.3731937

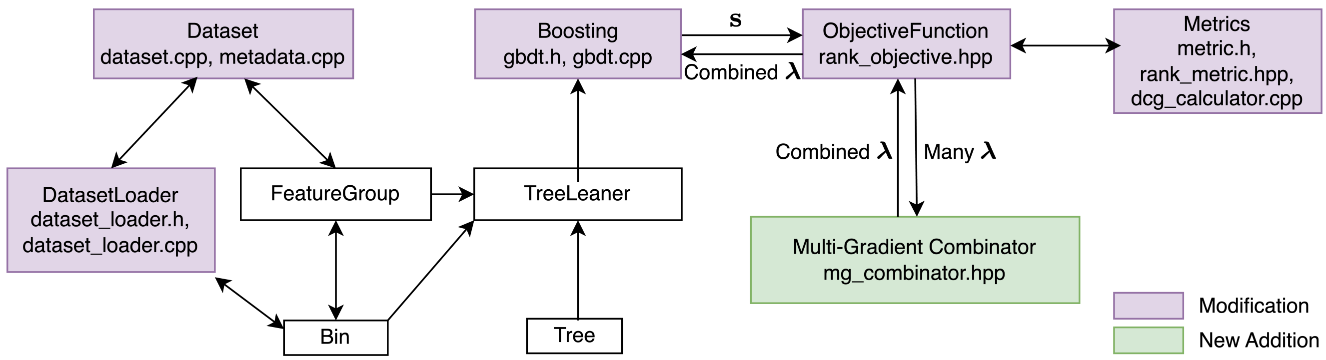

Figure 1: Architecture of MO-LightGBM builds upon LightGBM with key modifications to support multi-objective LTR. Modules modified from LightGBM are highlighted in purple, while newly added modules are shown in green.

solutions. These results establish MO-LightGBM as a practical and extensible tool for real-world multi-objective ranking applications in information retrieval.

2 Overview and Design of MO-LightGBM

2.1 Multi-objective Learning to Rank

In MOLTR, we consider a setting where multiple labels are provided for each query-item pair. The input for (q, d_i) is represented by a feature vector $\mathbf{x}_i^q \in \mathbb{R}^p$ for $i \in I_q$, where I_q denotes the index set of the match set for the query. The labels associated with (q, d_i) are y_{ik}^q for $k \in [K]$, where K is the number of labels.

MOLTR aims to learn a scoring function that assigns a score s_i^q to each (q, d_i) pair based on its input $\mathbf{x}_i^q$. Denote $\mathbf{s}^q$ as the vector whose i-th element is s_i^q and y_k^q as the vector whose i-th element is y_{ik}^q. In LightGBM, the scoring function is modeled using a GBM [7], where decision trees are built sequentially. The training cost of GBM is defined as a function of the scores:

$$c_k(\mathbf{s}) = \frac{1}{|\mathbb{Q}|} \sum_{q \in \mathbb{Q}} \ell(\mathbf{s}^q, \mathrm{y}_k^q), \quad \forall k \in [K].$$

Thus, the training cost in MOLTR is a vector-valued function: $\mathbf{c}(\mathbf{s}) = [c_1(\mathbf{s}), \cdots, c_K(\mathbf{s})]^T$, naturally making it an MOO problem.

2.2 Supported MOO Algorithms

The cost function in MOLTR gives rise to K score-gradients, $\nabla_{\mathbf{s}} c_k$ for $k \in [K]$. However, for training the GBM-based scoring function, each decision tree requires exactly one score-gradient as labels, not K score-gradients. We combine the K score-gradients as

$$\boldsymbol{\lambda} = \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{s}} c_k, \quad \text{s.t.} \sum_{k=1}^{K} \alpha_k = 1, \quad \boldsymbol{\alpha} \in \mathbb{R}_+^K, \tag{1}$$

where $\boldsymbol{\lambda}$ serves as the labels for training the trees in GBM and $\boldsymbol{\alpha}$ are combination coefficients.

In practice, we observe that MOO methods may exhibit oscillatory behavior in their cost functions during training. To address this issue and achieve a smoother trajectory, we apply a moving average

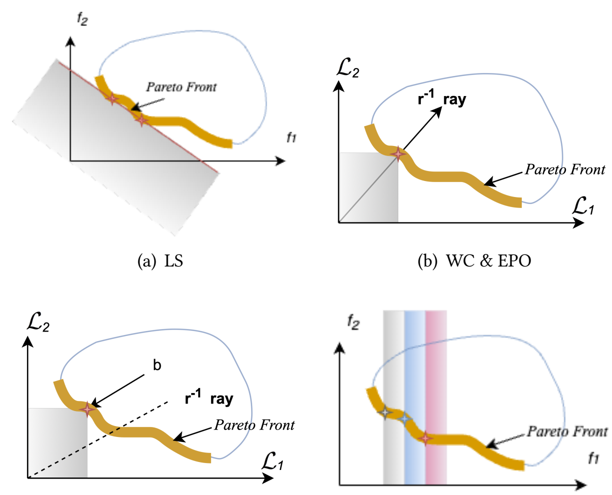

(a) LS (b) WC & EPO (c) WCMGDA (d) EC

Figure 2: Illustration of trade-off specifications.

filter to $\boldsymbol{\alpha}$ between consecutive iterations: $\boldsymbol{\alpha}^t \leftarrow \nu \boldsymbol{\alpha}^t + (1-\nu)\boldsymbol{\alpha}^{t-1}$, where $\nu \in (0, 1)$ is a smoothing factor.

Currently, MO-LightGBM supports 12 MOO strategies, as listed in Table 1. We also visualize different preference and constraints-based methods in Figure 2.

2.3 The Modular Design of MO-LightGBM

Figure 1 illustrates the overall framework of MO-LightGBM, which is designed with modularity to enable users to flexibly add custom designs. Modules modified from LightGBM are highlighted in purple, while newly added modules are shown in green. Additionally, the communication of scores (s) and gradients ($\boldsymbol{\lambda}$) between different modules is emphasized to illustrate the information flow.

Table 1: List of integrated MOO algorithms in MO-LightGBM.

MOO	Type	Smoothing	Reference
linear_scalarization	Preference	N	[17]
stochastic_label_aggregation	Preference	N	[3]
chebychev_scalarization	Preference	N	[17]
chebychev_scalarization_decay	Preference	Y	[19]
epo_search	Preference	N	[16]
epo_search_decay	Preference	Y	[16]
w_mgda	Preference	N	[19]
wc_mgda	Preference	N	[19]
wc_mgda_decay	Preference	Y	[19]
e_constraint	Constraint	N	[20]
ec_mgda	Constraint	N	[9]
ec_mgda_decay	Constraint	Y	[19]

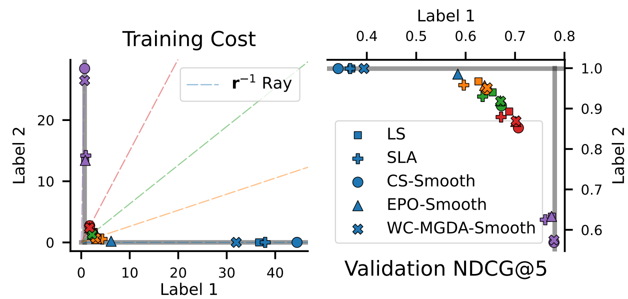

(a)

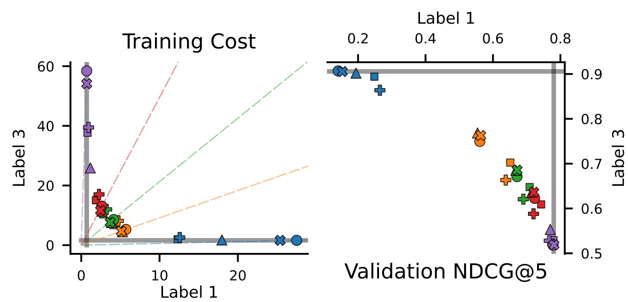

(b)

Figure 3: Illustration of results for bi-objective LTR experiments on Istella dataset with preference-based MOO. Colored lines and points represent different trade-off specifications and the corresponding solutions, respectively.

3 Demo of MO-LightGBM

Bi-objective Learning to Rank: This example shows how to train and evaluate a Bi-objective LTR model using various preference and constraint-based MOO methods with minimum code. With these few lines, the model is trained for each of the 5 preferences or upper bounds specified by the user with all the MOO methods provided in the mg_combinators. Evaluation results on selected metrics (Losses@k, NDCG@k) can be easily obtained and visualized for both training and test datasets. Refer to the GitHub page for further examples on training Multi-objective Learning to Rank models and customization options.

```
# Example: Config file for bi-objective experiment
dataset:
  name: istella
  train_file: "dataset/train.tsv"
  valid_file: "dataset/test.tsv"
  query_column: "1" # column of the query group id
  main_label: "0" # only for constraint-based moo
  all_labels: ['0', '11', '194'] # list of columns
       of labels
  bilabels_idx: [[0,1], [0,2]] # list of pair of
       columns for bi-objective tasks
  ignore_columns: ['0', '11', '194']  # all_lables +
        extra columns that need to be ignored

lightgbm_parameters:
  objective_type: lambdarank # lambdarank or ranknet
  num_iterations: '1000'
  num_thread: '0'
  ndcg_eval_at: '5,30'

lightgbm_path: "../LightGBM/lightgbm"
sample_lightgbm_config: "sample_config.conf"
num_tradeoffs: 5

mg_combinators:
  preference_based:
    - linear_scalarization
    - stochastic_label_aggregation
    - chebychev_scalarization_decay
    - chebychev_scalarization
    - epo_search_decay
    - epo_search
    - wc_mgda_decay
    - wc_mgda
    - w_mgda
  constraint_based:
    - e_constraint
    - ec_mgda_decay
    - ec_mgda
```

4 Experiment

We ran experiments on the Istella-S LETOR dataset [13]. Each query-url pair is represented by 200 features. Although these features are engineered, their descriptions are not publicly released. To define additional objectives, we selected features with more than five distinct values and identified those with minimal correlation among them. The final set of selected labels includes [0', 11', '194']. In total, we incorporated three objectives, including the original relevance label, and constructed two bi-objective cases. For model hyperparameter tuning, we used 1000 trees and a learning rate of 0.25 as an example.

We present the results of a single run for the Istella LETOR dataset. To maintain clarity, we separately plot results from preference-based MOO methods in Figure 3 and constraint-based MOO methods in Figure 4, preventing overcrowding. For preference-based MOO methods, all solutions exhibit compliance with the specified preferences, ensuring alignment with the intended ranking objectives. In contrast, for constraint-based methods, an equidistributed upper bounds in the cost space does not necessarily translate to an

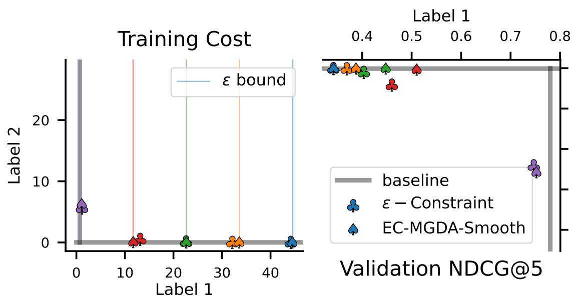

(a)

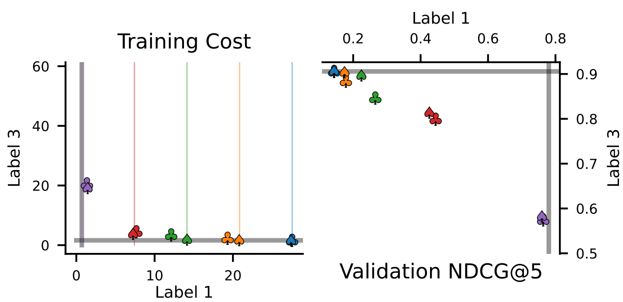

(b)

Figure 4: Illustration of results for bi-objective LTR experiments on Istella dataset with constraint-based MOO. Colored lines and points represent different trade-off specifications and the corresponding solutions, respectively.

equidistributed solution frontier in the NDCG space. This discrepancy is problem-dependent and arises from the loose approximation of LambdaRank [2] cost to the NDCG metric, rather than an inherent limitation of the constraint-based approach. Nonetheless, the MOO method effectively adheres to the prescribed trade-off specifications in the cost space.

5 Conclusion and Future Directions

We introduced MO-LightGBM, a comprehensive library for Multi-objective LTR within the LightGBM framework. By integrating 12 popular MOO methods, MO-LightGBM simplifies the training and evaluation of multi-objective ranking models, making advanced optimization techniques more accessible. These enhancements empower both researchers and practitioners by facilitating seamless model integration, enhancing usability, and lowering the barriers to deploying effective MOLTR solutions in diverse applications.

Several opportunities remain for further development. One promising direction is to explore the integration of various gradient manipulation methods designed to find balanced and fair solutions by optimizing conflict-aware update directions [23, 25]. Another avenue involves extending the framework to support automated hyper-parameter tuning and adaptive pipelines that dynamically adjust to shifting data distributions using Bayesian Optimization. Additionally, improving computational scalability and incorporating distributed training strategies [12, 24] could further enhance MO-LightGBM's ability to handle complex datasets more efficiently.

Short Bio of the Main Presenter

Dr. Chaosheng Dong is a machine learning scientist specializing in search relevance, generative information retrieval, and multi-objective optimization. Currently a Tech Lead at Amazon, Dr. Dong has played a pivotal role in transitioning Amazon Search from a single-task to a multi-task system, significantly improving search performance metrics. With a Ph.D. in Operations Research from the University of Pittsburgh, Dr. Dong has authored numerous papers in top-tier conferences such as ICML, KDD, and NeurIPS. His expertise spans multi-task learning, ranking models, and personalized recommendation systems, positioning him as an active contributor in search and recommendation advancements.

References

[1] Rahul Agrawal, Archit Gupta, Yashoteja Prabhu, and Manik Varma. 2013. Multi-label learning with millions of labels: Recommending advertiser bid phrases for web pages. In *Proceedings of the 22nd international conference on World Wide Web*. 13–24.

[2] C. Burges. 2010. From RankNet to LambdaRank to LambdaMART: An Overview.

[3] David Carmel, Elad Haramaty, Arnon Lazerson, and Liane Lewin-Eytan. 2020. Multi-Objective Ranking Optimization for Product Search Using Stochastic Label Aggregation. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) *(WWW '20)*. Association for Computing Machinery, New York, NY, USA, 373–383. https://doi.org/10.1145/3366423.3380122

[4] Chaosheng Dong, Peiyao Xiao, and Kaiyi Ji. 2025. Deep Multi-task Learning to Rank. (2025).

[5] André Elisseeff and Jason Weston. 2001. Kernel methods for Multi-labelled classification and Categ orical regression problems. In *NIPS 2001*. https://api.semanticscholar.org/CorpusID:14034941

[6] Mauri Ferrandin and Ricardo Cerri. 2023. Multi-label classification via closed frequent labelsets and label taxonomies. *Soft Computing* (2023), 1–34.

[7] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.

[8] Nadia Ghamrawi and Andrew McCallum. 2005. Collective Multi-Label Classification *(CIKM '05)*. Association for Computing Machinery, New York, NY, USA, 195–200. https://doi.org/10.1145/1099554.1099591

[9] Chengyue Gong, Xingchao Liu, and qiang liu. 2021. Automatic and Harmless Regularization with Constrained and Lexicographic Optimization: A Dynamic Barrier Approach. In *Advances in Neural Information Processing Systems*.

[10] Sariel Har-Peled, Dan Roth, and Dav Zimak. 2002. Constraint Classification for Multiclass Classification and Ranking. In *Proceedings of the 15th International Conference on Neural Information Processing Systems (NIPS'02)*. MIT Press, Cambridge, MA, USA, 809–816.

[11] 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. *Advances in neural information processing systems* 30 (2017).

[12] Zhuqing Liu, Chaosheng Dong, Michinari Momma, Simone Shao, Shaoyuan Xu, Haibo Yang, Yan Gao, and Jia Liu. 2025. STIMULUS: Achieving Fast Convergence and Low Sample Complexity in Stochastic Multi-Objective Learning. (2025).

[13] 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*.

[14] Debabrata Mahapatra, Chaosheng Dong, Yetian Chen, Deqiang Meng, and Michinari Momma. 2022. Multi-label learning to rank through multi-objective optimization. *arXiv preprint arXiv:2207.03060* (2022).

[15] Debabrata Mahapatra, Chaosheng Dong, and Michinari Momma. 2023. Querywise Fair Learning to Rank through Multi-Objective Optimization. In *Proceedings of the 29th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.

[16] Debabrata Mahapatra and Vaibhav Rajan. 2020. Multi-Task Learning with User Preferences: Gradient Descent with Controlled Ascent in Pareto Optimization. In *Proceedings of the 37th International Conference on Machine Learning*.

[17] Kaisa Miettinen. 1999. *Nonlinear multiobjective optimization*. Vol. 12. Springer Science & Business Media.

[18] Michinari Momma, Chaosheng Dong, and Yetian Chen. 2022. Multi-objective ranking with directions of preferences. *SIGIR Workshop On eCommerce* (2022).

[19] Michinari Momma, Chaosheng Dong, and Jia Liu. 2022. A multi-objective / multi-task learning framework induced by Pareto stationarity. In *Proceedings of the 39th International Conference on Machine Learning*.

[20] Michinari Momma, Alireza Bagheri Garakani, and Yi Sun. 2019. Multi-objective Relevance Ranking.. In *eCOM@ SIGIR*.

[21] Yashoteja Prabhu and Manik Varma. 2014. Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 263–272.

[22] Krysta M. Svore, Maksims N. Volkovs, and Chris J.C. Burges. 2011. Learning to Rank with Multiple Objective Functions. In *Proceedings of WWW 2011* (proceedings of www 2011 ed.).

[23] Peiyao Xiao, Chaosheng Dong, Shaofeng Zou, and Kaiyi Ji. 2025. Scalable Bilevel Loss Balancing for Multi-Task Learning. *arXiv preprint arXiv:2502.08585* (2025).

[24] Haibo Yang, Zhuqing Liu, Jia Liu, Chaosheng Dong, and Michinari Momma. 2023. Federated multi-objective learning. *Advances in neural information processing systems* (2023).

[25] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. Gradient surgery for multi-task learning. *Advances in neural information processing systems* 33 (2020), 5824–5836.

[26] Min-Ling Zhang and Zhi-Hua Zhou. 2014. A Review on Multi-Label Learning Algorithms. *IEEE Transactions on Knowledge and Data Engineering* 26, 8 (2014), 1819–1837. https://doi.org/10.1109/TKDE.2013.39

[27] Zhendong Zhang and Cheolkon Jung. 2020. GBDT-MO: Gradient-boosted decision trees for multiple outputs. *IEEE transactions on neural networks and learning systems* 32, 7 (2020), 3156–3167.