

Quantum Feature Selection from Interpretable Models using a QUBO Formulation

Notebook for the QuantumCLEF Lab at CLEF 2025

Flavio Giobergia^{1,*}, Claudio Savelli¹, Alkis Koudounas¹ and Elena Baralis¹

¹Politecnico di Torino, Corso Duca degli Abruzzi, 24, 10129 Torino (TO), Italy

Abstract

In this work, we tackle the feature selection problem for content-based recommender systems using Quadratic Unconstrained Binary Optimization (QUBO). Our approach, submitted as Team MALTO to the QuantumCLEF 2025 challenge, aims to improve the performance of an Item-Based K-Nearest Neighbors (Item-KNN) model by selecting a compact and informative subset of item features. We formulate a QUBO objective that combines *feature relevance* – estimated via Random Forest (RF) importance scores – and *feature redundancy* – captured through pairwise Pearson correlations. We compare our method against a collaborative-driven QUBO baseline and a random selection strategy. Experiments on the official QuantumCLEF dataset demonstrate that our relevance-aware strategy outperforms the other methods regarding recommendation quality, especially in low-dimensional feature regimes. Our results highlight the potential of combining machine learning and quantum optimization for effective feature selection in recommender systems.

Keywords

Feature selection, Simulated annealing, Quantum annealing, QUBO, Recommender systems, Item-KNN

1. Introduction

Recommender systems help users discover relevant information in many domains. These include e-commerce, video and music streaming, and digital libraries [1, 2, 3, 4]. A widely adopted class of recommendation models is called *neighborhood-based collaborative filtering* [5, 6]. These models generate personalized rankings by analyzing similarities between items. One popular method in this category is the *Item-Based k-Nearest Neighbors (Item-KNN)* algorithm [7]. Item-KNN is valued for being easy to understand, scalable to large datasets, and effective even when user interactions are sparse.

A key factor that affects the performance of these models is the choice of item features, which are often stored in what is known as an *Item Content Matrix (ICM)*. In real-world applications, ICMs may contain hundreds of different content descriptors [8]. However, not all features in the ICM contribute equally to meaningful item similarity. Some features may be irrelevant, redundant, or noisy.

Using such features can reduce recommendation quality. They can also make similarity computations more expensive. Therefore, selecting the right subset of features becomes a critical task. This process is known as *feature selection* [9, 10, 11]. In this work, we present a feature selection approach based on *Quadratic Unconstrained Binary Optimization (QUBO)* [12]. We aim to choose a subset of features that improves the performance of the Item-KNN model. QUBO is a powerful mathematical framework for solving combinatorial optimization problems. It allows us to model both the importance of individual features and the redundancy between pairs of features.

This paper is the official submission of Team MALTO (Machine Learning @ PoliTO) to the QuantumCLEF 2025 challenge [13, 14], which focuses on *Feature Selection for Recommendation Systems*. We study feature selection in the context of content-based item-KNN. We are given a User Rating Matrix (URM) and two versions of ICMs, one with 100 features and another with 400. Our goal is to select a subset of features that leads to better recommendation quality.

CLEF 2025 Working Notes, 9 – 12 September 2025, Madrid, Spain

*Corresponding author.

✉ flavio.giobergia@polito.it (F. Giobergia); claudio.savelli@polito.it (C. Savelli); alkis.koudounas@polito.it (A. Koudounas); elena.baralis@polito.it (E. Baralis)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

To achieve this, we use a supervised learning model to estimate the relevance of each feature. We then encode these relevance scores into a QUBO matrix. We also include feature redundancy terms in the QUBO, based on pairwise feature correlations. The result is a QUBO problem with a cardinality constraint that limits the number of selected features.

This QUBO-based formulation can be tackled using different optimization strategies. *Simulated Annealing* [15] is a classical metaheuristic commonly employed for solving combinatorial problems. In contrast, *Quantum Annealing* [16] is a quantum-inspired method that exploits quantum tunneling to escape local minima and potentially navigate the solution space more efficiently. These solvers can be compared in terms of the selected features’ quality and the search’s computational efficiency, offering insights into the relative advantages of classical and quantum approaches to QUBO optimization.

The rest of this paper is organized as follows. Section 2 reviews related work and provides context for the problem. Section 3 presents our proposed methodology in detail. Section 4 describes the experimental setup used in our study, while Section 5 reports and discusses the results. Finally, Section 6 concludes the paper by summarizing the proposed approach and its performance.

2. Related Works

QUBO Formulation. The Quadratic Unconstrained Binary Optimization [12] is a general mathematical formulation for combinatorial optimization problems where the goal is to minimize a quadratic polynomial over binary variables:

$$\min_{x \in \{0,1\}^n} x^T Q x, \quad (1)$$

Where $Q \in \mathbb{R}^{n \times n}$ is a symmetric matrix, and x is a binary vector representing the inclusion or exclusion of elements – in our case, features. QUBO formulations are particularly well-suited for feature selection, as they allow the expression of both feature relevance (via linear terms Q_{ii}) and feature redundancy (via quadratic terms Q_{ij} for $i \neq j$). Relevance can be derived from importance weights computed through a predictive model, while redundancy is often modeled using pairwise correlations or mutual information between features. To guide the selection of a fixed number k of features, it is possible to incorporate a *cardinality constraint* [17]. Since QUBO does not support hard constraints directly, we encode this as a penalty term added to the objective:

$$\min_{\mathbf{x} \in \{0,1\}^n} \mathbf{x}^T Q \mathbf{x} + \lambda \left(\sum_{i=1}^n x_i - k \right)^2, \quad (2)$$

Where $\lambda > 0$ is a hyperparameter controlling the trade-off between optimizing the QUBO objective and enforcing the selection of exactly k features. This penalty is minimized when the sum of selected features is exactly k , effectively turning the constraint into a soft requirement.

Due to QUBO’s expressive power, a growing body of work has explored solving such formulations using classical heuristics (e.g., simulated annealing, tabu search) and quantum methods (e.g., quantum annealing). In particular, quantum annealers such as those developed by D-Wave Systems have shown promise in exploring complex energy landscapes efficiently, offering a potentially more effective alternative to classical solvers for specific NP-hard problems.

Feature Selection in Recommender Systems. Feature selection has long been studied to improve the accuracy and efficiency of machine learning models. In recommender systems, especially those incorporating item content, selecting a compact and informative subset of features can reduce overfitting, improve model interpretability, and decrease computational costs. Payares et al. [18] proposed a quantum annealing-based approach to feature selection for Information Retrieval, applying it to the QuantumCLEF 2024 challenge [19]. They explored multiple QUBO formulations, including mutual information, conditional mutual information, and correlation coefficients, and compared quantum, hybrid, and simulated annealing solvers. Almeida and Matos [20] proposed a hyperparameter-free

QUBO formulation for feature selection in learning-to-rank models. Their method balances relevance and redundancy without requiring manual tuning and demonstrates competitive results on the MQ2007 [21] and ISTECLA [22] datasets. Niu et al. [23] introduced a QUBO-based feature selection framework that integrates *Counterfactual Analysis* to enhance the effectiveness of item-based recommendation models. Unlike traditional Mutual Information-based approaches, their method explicitly incorporates the impact of each feature on the model’s performance, leading to a more goal-aligned optimization. Nembrini et al. [24] proposed a collaborative-driven feature selection method that aligns content-based similarity with a pre-trained collaborative model. Their approach selects item features by comparing the similarity structure derived from user interactions with that derived from content metadata. Features that produce misleading or unsupported similarities are penalized, and the final feature selection is formulated as a QUBO problem and solved using quantum annealing. Our work builds on this line of research by formulating feature selection for an item-kNN model as a QUBO problem and comparing the results obtained via simulated annealing and quantum annealing.

3. Methodology

We approach the feature selection problem by formulating it as a QUBO task. We assume we are given a tabular dataset that includes d features; each feature related to the user for which items should be recommended. Our goal is to select a subset of $d' \leq d$ features that are most useful for the task of building a recommender system. We express this selection problem as a minimization problem, defined in Equation 2. Each feature is associated with a binary variable x_i . If $x_i = 1$, the i -th feature is selected. If $x_i = 0$, the feature is not selected.

The QUBO objective function is defined by a matrix Q , which contains both diagonal and off-diagonal terms. The diagonal entry Q_{ii} represents the cost or penalty of selecting feature i . The off-diagonal entry Q_{ij} represents the penalty for selecting both features i and j together.

In this work, we adopt the same general structure as the one proposed by Mücke et al. [17]. We treat feature selection as a trade-off between relevance and redundancy. Specifically, the diagonal values in Q are set to the negative importance scores of the features. This means that selecting more important features will lower the overall objective function. The off-diagonal values in Q are used to encode redundancy between pairs of features. If two features are highly redundant, selecting both will increase the objective value. In [17], feature importance is measured using mutual information *between each feature and the target*. Redundancy is measured as the mutual information *between pairs of features*.

While the target is clear for a classification or regression problem, the goal is not as clearly defined for recommender systems. As such, our approach uses a different way to quantify both importance and redundancy. We estimate the relevance of features using the importance scoring mechanism of a Random Forest [25] model. This allows us to capture complex, non-linear dependencies between features and the target. For redundancy, we compute the Pearson correlation between each pair of features. This gives a simple but effective way to penalize the selection of highly correlated features. By combining these components into a QUBO matrix, we aim to select a compact and informative set of features. This set should balance high predictive value with low redundancy.

The rest of the section details (i) how the features are defined for each user, (ii) how the feature importance is computed, (iii) how the redundancy of features is estimated, and (iv) how these quantities are framed, along with constraints, as a QUBO problem.

User feature extraction In this recommendation problem, we consider a setting with n users and m items, where each item is described by d features. Two main sources of information are available. The first one is the binary User-Rating Matrix $U \in \{0, 1\}^{n \times m}$, where $U_{ij} = 1$ if user i rated item j , 0 otherwise. The second one is the Item Content Metrics matrix $C \in \mathbb{R}^{m \times d}$, which represents each item through a set of d descriptive features..

We aim to represent each user with a set of features. Given the available information, we produce a representation of each user as the sum of the representations of the items that the user rated. We can

efficiently compute this new representation of users as $U_c = UC^\top$.

Feature importance Random Forests [25] are ensemble learning methods widely employed for supervised tasks such as classification and regression. An RF aggregates the predictions of multiple Decision Trees, each trained on a bootstrap dataset sample and using a random subset of features at each split. This randomness introduces model diversity and improves generalization, making RFs more robust to overfitting than individual trees. Beyond their predictive performance, Decision Trees offer interpretability, as each path from root to leaf corresponds to a sequence of feature-based decisions. A key advantage of both DTs and RFs is the ability to quantify *feature importance* – a measure of each feature’s contribution to the model’s predictive performance.

In RFs, feature importance is typically computed by summing the reductions in impurity that a feature contributes across all the splits in all trees of the forest. When a feature is used to split a node, the reduction in impurity (e.g., measured via Gini impurity or entropy) is evaluated as the difference between the parent node’s impurity and the weighted sum of the child nodes’ impurities. These reductions are accumulated across all trees, yielding a score for each feature. Higher scores indicate greater relevance.

In this work, we use the RF-derived feature importance scores to define the diagonal terms of the QUBO matrix. Specifically, if the importance of feature i is denoted by f_i , we set $Q_{ii} = -\alpha \cdot f_i$, where α is a scaling factor that balances the trade-off between relevance and redundancy. The negative sign reflects the minimization objective of the QUBO formulation. As RF importance is normalized to sum to 1, we note that $0 \leq f_i \leq 1$ for all i .

Unlike typical supervised learning tasks, our recommendation setting does not have a single target to predict. Instead, we are interested in estimating the likelihood of each user interacting with different items. We frame the task as a multi-output classification problem to compute feature importance in this context. For each user, the model predicts whether they would interact with a selection of items. RFs naturally handle this setting by building separate trees for each item. We restrict the predictions to a random subset of 100 items to keep computation efficient. Empirically, we observed that this approximation does not significantly affect the quality of the resulting feature importance scores.

Redundancy quantification To estimate redundancy between pairs of features, we compute the Pearson correlation coefficient between their corresponding user-level representations. This metric captures the degree of linear association between two features, whether positive or negative. In our formulation, we consider any strong correlation (including negative ones) as a form of redundancy. For this reason, since we are interested in the magnitude of the correlation, we use the absolute value of the correlation as the redundancy term, setting $Q_{ij} = |r_{ij}|$. This ensures that positive and negative correlations contribute equally to the redundancy penalty.

QUBO problem framing We define the QUBO matrix Q by combining the relevance and redundancy components described in the previous sections. Each diagonal element Q_{ii} captures the relevance of feature i and is set to $Q_{ii} = -\alpha f_i$, where f_i is the feature importance score computed via RF and α is a scaling parameter that controls the trade-off between relevance and redundancy. The off-diagonal elements Q_{ij} for $i \neq j$ encode redundancy between features and are defined as $Q_{ij} = |r_{ij}|$, where r_{ij} is the Pearson correlation between the user-level representations of features i and j . The relevance/redundancy trade-off is controlled by the parameter α , which is tuned through validation.

4. Experimental Setup

This section describes the experimental setup used to evaluate our QUBO-based feature selection strategy. We detail the dataset used, the methods compared, and the evaluation procedure.

4.1. Dataset

We conduct our experiments on the benchmark dataset provided by the QuantumCLEF 2025 challenge, designed for evaluating feature selection strategies in recommender systems. The dataset includes:

- A **User-Rating Matrix** (U), representing implicit feedback where $U_{ij} = 1$ indicates that user i interacted with item j , and 0 otherwise.
- Two **Item Content Matrices** (C): one with 100 features and one with 400 features, representing item-level descriptors such as genres, tags, or metadata.

We generate a feature representation for each user by aggregating the features of the items they interacted with, effectively computing $U_c = U C^\top$.

4.2. Methods

We compare three feature selection strategies, each implemented as a QUBO optimization problem:

- **Random**: A naive baseline where k features are randomly selected without any optimization.
- **Baseline**: As a baseline, we implement the method proposed by Nembrini et al. [24], which formulates feature selection as a QUBO problem to align content-based similarities with a collaborative model.
- **RF-QUBO (ours)**: Our method encodes feature relevance using importance scores extracted from the RF classifier and feature redundancy using Pearson correlation. The QUBO matrix Q combines both components with a soft penalty enforcing the selection of exactly k features.

Each QUBO problem can be solved using two solvers:

- **Simulated Annealing (SA)**: A probabilistic metaheuristic inspired by the physical annealing process, applied to solve the QUBO minimization problem.
- **Quantum Annealing (QA)**: A quantum optimization technique that solves QUBO problems by evolving a quantum system toward its ground state. Unlike classical methods, QA leverages quantum tunneling to escape local minima, potentially offering advantages in exploring complex energy landscapes.

Hyperparameters α (relevance-redundancy trade-off) and λ (cardinality penalty) are tuned via grid search on the validation set. We set $\alpha = 3$ and $\lambda = 0.01$ in the final configuration. The number of selected features k ranges from 1 to 90 (in the 100-feature scenario) and up to 390 (in the 400-feature scenario).

4.3. Evaluation

The selected features are used to build an **Item-KNN recommender**, where item-item similarity is computed using cosine similarity on the reduced ICM. The number of nearest neighbors is fixed to 100, and a shrinkage factor of 5 is applied. We use **nDCG@10** (normalized Discounted Cumulative Gain at rank 10) as the primary performance metric, which measures recommended items' relevance and ranking position in the top-10 list. This is the same evaluation metric adopted by the QuantumCLEF 2025 challenge, ensuring consistency with the official ranking criteria. All feature selection methods are evaluated under identical experimental conditions to ensure fairness. Specifically, each method is applied to the same dataset split, and the corresponding reduced ICM is used to generate recommendations. To assess the variability of results, each experiment is repeated 3 times, and the reported values reflect the average performance and variability (standard deviation) across runs. In addition to performance metrics, we track the number of effective features selected after QUBO optimization. This is compared against the feature budget constraint imposed during optimization to assess the degree of compliance or deviation.




Figure 1: nDCG@10 performance of the different feature selection strategies as a function of the number of features enforced by the QUBO constraint on the 100 features scenario.




Figure 2: nDCG@10 performance of the different feature selection strategies as a function of the number of features enforced by the QUBO constraint on the 400 features scenario.

5. Results

In this section, we compare the performance of the proposed **RF-QUBO** method against the **Random** and **Baseline** [24] strategies, using the *Simulated Annealing* solver. Although *Quantum Annealing* was part of our intended evaluation, it could not be executed due to technical issues encountered during the QuantumCLEF 2025 challenge. Nonetheless, prior studies suggest that QA typically yields results comparable to SA, often with significantly reduced execution time [23, 20, 18]. We report results for both the 100-feature and 400-feature scenarios from the QuantumCLEF 2025 dataset.

Figure 1 shows the nDCG@10 achieved by the three methods in the 100-feature setting as the number of selected features increases, as expected. However, the RF-QUBO method consistently outperforms both the random and the Baseline, particularly in the low-dimensional regime (up to 1-40 features). These results show that including feature relevance information in the QUBO formulation helps the model make better recommendations, especially when only a small number of features can be selected. In other words, choosing truly informative features – rather than just any set of features – makes a clear difference in performance. As expected, the advantage becomes smaller as the number of selected features increases. This is because when most features are included, even random or less-informed

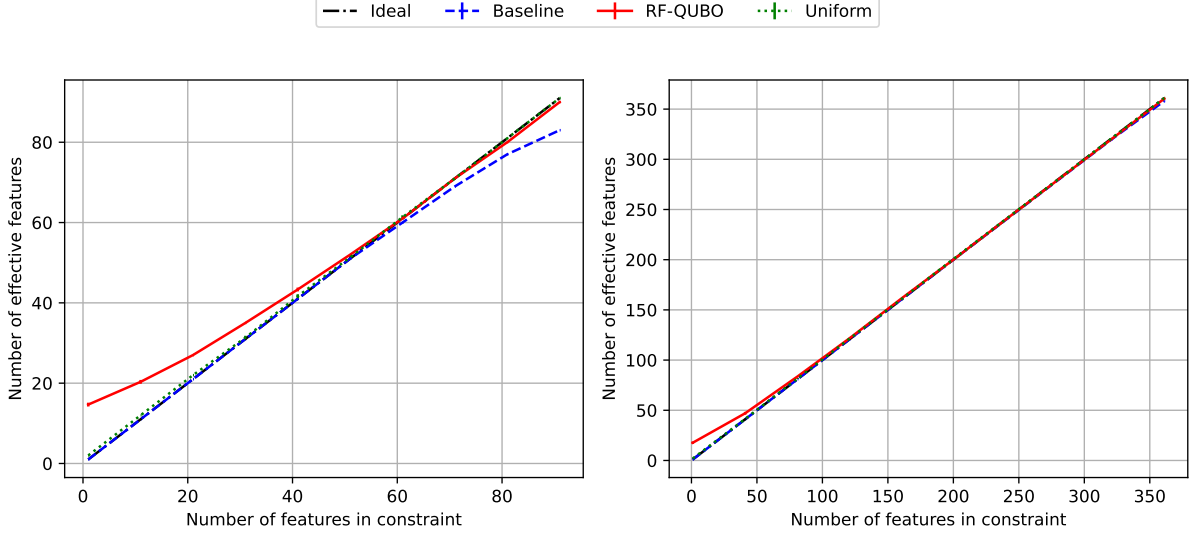


Figure 3: Comparison between the number of features specified as a constraint and the number of features selected after QUBO optimization on the 100 (left) and 400 (right) features scenario. The dotted “Ideal” line indicates perfect compliance with the constraint.

selections start to resemble the complete feature set, and all methods tend to achieve similar results.

Figure 2 presents the results for the 400-feature version of the dataset. While the overall trends are consistent with those observed in the 100-feature setting, the performance gains achieved by the RF-QUBO method are less pronounced in the lower-dimensional case. In particular, the margin of improvement over the Baseline and random strategies is narrower when the number of available features is limited. Nonetheless, the RF-QUBO method remains the best-performing approach across most constraint levels, confirming its robustness and effectiveness in identifying relevant features even under different dimensional settings.

Figure 3 shows the effective number of features selected by each method compared to the number specified in the constraint. As expected, slight deviations are observed because the QUBO objective includes a soft penalty term rather than a hard constraint. In both cases, we observe that RF-QUBO tends to select a larger number of features when the number of features in the constraint is low. This occurs because the penalty for the soft constraint is low enough that it allows for small deviations from the requested value. If the requested number of features is a strict constraint, the penalty term λ can be increased: indeed, we observe perfect agreement when this value is large enough (at the cost of a lower quality of the resulting nDCG). Potentially, an additional post-processing step could be adopted to strictly enforce the constraint. For instance, a subset of features could be discarded when the number of selected features exceeds the number of allowed ones.

Overall, these results demonstrate that our RF-QUBO approach offers a reliable and effective strategy for feature selection in content-based recommender systems. We additionally report the official results obtained during the challenge, in Table 1. First, we note, in general, a drop in performance w.r.t. the results we obtained locally. We expect this to be the case due to a different (potentially more complex) test set being used – though this step was opaque to the participants.

The other aspect standing out is the fact that the solution using all features outperforms our approach. While not explicitly discussed in previous results, this was to be expected, given the (approximately) monotone trends in performance, as a function of the number of selected features (not only for RF-QUBO, but also for the other baselines). This implies that a richer representation appears to produce better recommendations. However, we argue that the usefulness of feature selection still holds: many scenarios (e.g., low-resource settings) can only process a limited number of features; in other cases, the curse of dimensionality [26] prevents large numbers of features from being used.

N. features	Method	N. selected	nDCG@10
100	RF-QUBO	51	0.0207
	All features	100	0.0226
400	RF-QUBO	53	0.0182
	RF-QUBO	200	0.0294
	All features	400	0.0328

Table 1

Comparison in terms of nDCG@10 for RF-QUBO and the solution that applies no feature selection.

6. Conclusion

We presented a feature selection method for content-based recommender systems, formulated as a Quadratic Unconstrained Binary Optimization problem. The approach incorporates feature relevance, which is derived from RF importance scores, and feature redundancy, which is measured through Pearson correlation, within a unique optimization framework.

The method was evaluated on the official QuantumCLEF 2025 dataset across two feature dimensionalities (100 and 400 features). Results show that the proposed RF-QUBO strategy consistently outperforms both random and Baseline approaches, especially under tight feature selection constraints. Despite not outperforming the baseline that uses all features, we argue that many scenarios can benefit from reducing the number of features (e.g., to reduce the computational cost).

Acknowledgements

This study was carried out within the FAIR - Future Artificial Intelligence Research and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1555 11/10/2022, PE00000013). This manuscript reflects only the authors’ views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to paraphrase and reword the text. After using this tool, authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

- [1] R. Burke, A. Felfernig, M. H. Göker, Recommender systems: An overview, *Ai Magazine* 32 (2011) 13–18.
- [2] Z. Wang, L. Sun, W. Zhu, S. Yang, H. Li, D. Wu, Joint social and content recommendation for user-generated videos in online social network, *IEEE Transactions on Multimedia* 15 (2012) 698–709.
- [3] C. A. Gomez-Urbe, N. Hunt, The netflix recommender system: Algorithms, business value, and innovation, *ACM Transactions on Management Information Systems (TMIS)* 6 (2015) 1–19.
- [4] B. Smith, G. Linden, Two decades of recommender systems at amazon. com, *Ieee internet computing* 21 (2017) 12–18.
- [5] J. A. Konstan, B. N. Miller, D. Maltz, J. L. Herlocker, L. R. Gordon, J. Riedl, Grouplens: Applying collaborative filtering to usenet news, *Communications of the ACM* 40 (1997) 77–87.

- [6] A. N. Nikolakopoulos, X. Ning, C. Desrosiers, G. Karypis, Trust your neighbors: A comprehensive survey of neighborhood-based methods for recommender systems, *Recommender systems handbook* (2021) 39–89.
- [7] M. Deshpande, G. Karypis, Item-based top-n recommendation algorithms, *ACM Transactions on Information Systems (TOIS)* 22 (2004) 143–177.
- [8] P. Castells, D. Jannach, Recommender systems: A primer, *arXiv preprint arXiv:2302.02579* (2023).
- [9] H. Liu, H. Motoda, Computational methods of feature selection, CRC press, 2007.
- [10] F. Giobergia, E. Baralis, M. Camuglia, T. Cerquitelli, M. Mellia, A. Neri, D. Tricarico, A. Tuninetti, Mining sensor data for predictive maintenance in the automotive industry, in: *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, IEEE, 2018, pp. 351–360.
- [11] M. Rahat, P. S. Mashhadi, S. Nowaczyk, S. Choudhury, L. Petrin, T. Rognvaldsson, A. Voskou, C. Metta, C. Savelli, et al., Volvo discovery challenge at ecml-pkdd 2024, *arXiv preprint arXiv:2409.11446* (2024).
- [12] F. Glover, G. Kochenberger, R. Hennig, Y. Du, Quantum bridge analytics i: a tutorial on formulating and using qubo models, *Annals of Operations Research* 314 (2022) 141–183.
- [13] A. Pasin, M. F. Dacrema, W. Cuhna, M. A. Gonçalves, P. Cremonesi, N. Ferro, Quantumclef 2025: Overview of the second quantum computing challenge for information retrieval and recommender systems at CLEF, in: G. Faggioli, N. Ferro, P. Rosso, D. Spina (Eds.), *Working Notes of CLEF 2025 - Conference and Labs of the Evaluation Forum, CEUR Workshop Proceedings*, 2025.
- [14] A. Pasin, M. F. Dacrema, W. Cuhna, M. A. Gonçalves, P. Cremonesi, N. Ferro, Overview of quantumclef 2025: The second quantum computing challenge for information retrieval and recommender systems at CLEF, in: J. Carrillo-de-Albornoz, J. Gonzalo, L. Plaza, A. G. S. de Herrera, J. Mothe, F. Piroi, P. Rosso, D. Spina, G. Faggioli, N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the Sixteenth International Conference of the CLEF Association (CLEF 2025)*, *Lecture Notes in Computer Science*, 2025.
- [15] D. Bertsimas, J. Tsitsiklis, Simulated annealing, *Statistical science* 8 (1993) 10–15.
- [16] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, J. D. Doll, Quantum annealing: A new method for minimizing multidimensional functions, *Chemical physics letters* 219 (1994) 343–348.
- [17] S. Mücke, R. Heese, S. Müller, M. Wolter, N. Piatkowski, Feature selection on quantum computers, *Quantum Machine Intelligence* 5 (2023) 11.
- [18] E. Payares, E. Puertas, J. C. Martinez Santos, Team qtb on feature selection via quantum annealing and hybrid models, *Working Notes of CLEF* (2024).
- [19] A. Pasin, M. Ferrari Dacrema, P. Cremonesi, N. Ferro, Overview of quantumclef 2024: The quantum computing challenge for information retrieval and recommender systems at clef, in: *International Conference of the Cross-Language Evaluation Forum for European Languages*, Springer, 2024, pp. 260–282.
- [20] T. Almeida, S. Matos, Towards a hyperparameter-free qubo formulation for feature selection in ir, *Working Notes of CLEF* (2024).
- [21] T. Qin, T.-Y. Liu, Introducing letor 4.0 datasets, *arXiv preprint arXiv:1306.2597* (2013).
- [22] D. Dato, S. MacAvaney, F. M. Nardini, R. Perego, N. Tonellotto, The istella22 dataset: Bridging traditional and neural learning to rank evaluation, in: *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022, pp. 3099–3107.
- [23] J. Niu, J. Li, K. Deng, Y. Ren, Cruise on quantum computing for feature selection in recommender systems, *arXiv preprint arXiv:2407.02839* (2024).
- [24] R. Nembrini, M. Ferrari Dacrema, P. Cremonesi, Feature selection for recommender systems with quantum computing, *Entropy* 23 (2021) 970.
- [25] L. Breiman, Random forests, *Machine learning* 45 (2001) 5–32.
- [26] R. Bellman, Dynamic programming, *science* 153 (1966) 34–37.