QuantumCLEF 2026 - The Third Edition of the Quantum Computing Lab at CLEF

Andrea Pasin[1], Maurizio Ferrari Dacrema[2], Paolo Cremonesi[2], Washington Cunha[3], Marcos André Gonçalves[3], and Nicola Ferro[1]

[1] University of Padua, Padua, Italy
andrea.pasin.1@phd.unipd.it, nicola.ferro@unipd.it
[2] Politecnico di Milano, Milan, Italy
{maurizio.ferrari, paolo.cremonesi}@polimi.it
[3] Federal University of Minas Gerais, Belo Horizonte, Brazil
{washingtoncunha, mgoncalv}@dcc.ufmg.br

Abstract. *Quantum Computing (QC)* is a research field that has been in the limelight in recent years. In fact, this new paradigm has the potential to revolutionize the way we currently solve problems by leveraging quantum-mechanical phenomena, which allow quantum computers to solve specific problems more efficiently than traditional computers. As quantum computers are starting to become more available, our objective is to investigate the application of QC within the *Information Retrieval (IR)* and *Recommender Systems (RS)* fields. In fact, IR and RS systems perform computationally intensive operations on extensive datasets, and using QC in their pipeline could be useful to improve their efficiency and, in some cases, effectiveness.
Thus, in this work, we present the third edition of the QuantumCLEF lab, the first lab that allow participants to use **real quantum computers** for solving IR and RS tasks. The lab is composed of three main tasks that aim at discovering and evaluating *Quantum Annealing (QA)* approaches compared to their traditional counterpart while also establishing collaborations among researchers from different fields to harness their knowledge and skills to solve the considered challenges and promote the usage of QA. Moreover, based on the availability of quantum resources, we plan to introduce a small set of gate-based QC tasks for more-experienced researchers.

Keywords: Quantum Computing · Quantum Annealing · Feature Selection · Instance Selection · Clustering · Recommender Systems · Information Retrieval

1 Introduction

Quantum Computing (QC) is a paradigm that has the potential to revolutionize the way problems are solved across different fields. By leveraging quantum-mechanical phenomena, quantum computers are expected to solve certain problems more efficiently and, in some cases, more effectively than traditional ones.

Given the current scenario where *Information Retrieval (IR)* and *Recommender Systems (RS)* systems face ever-increasing amounts of data and rely on computationally demanding approaches, *Quantum Computing (QC)* could be integrated in the pipeline of such systems to improve their performance. However, while QC has already been applied in several domains, limited work has been done specifically for the IR and RS fields [4, 13, 20]. Indeed, the area of IR called Quantum IR [11, 23, 25] consists of exploiting the concepts of quantum mechanics to formulate IR models and problems but it does not deal with implementing IR and RS models and algorithms via QC technologies.
In this work, we focus on *Quantum Annealing (QA)*, which exploits special-purpose devices, called quantum annealers, which are able to rapidly find optimal (or close to optimal) solutions to optimization problems by leveraging quantum-mechanical effects. Our goal is therefore to understand QA can improve the efficiency and effectiveness of IR and RS systems by integrating it into their pipelines. Thus, we present here the third edition of the evaluation lab called *QuantumCLEF (qCLEF)*[4] [15], which aims at:

- evaluating the performance of QA with respect to traditional approaches;
- identifying new ways of formulating IR and RS algorithms and methods, so that they can be solved with QA;
- growing a research community around this new field in order to promote a wider adoption of QC technologies for IR and RS.

Working with QA does not require particular knowledge about how quantum physics works underneath it. There are, in fact, available tools and libraries that can be easily used to program and solve problems through this paradigm. Moreover, by participating in qCLEF, participants are provided access to the KIMERA infrastructure, allowing them to seamlessly use real state-of-the-art quantum computers [19] while monitoring resource usage for fairness, comparability, and reproducibility.

The paper is organized as follows: Section 2 introduces related work; Section 3 presents the tasks in the qCLEF lab; Section 4 considers some critical evaluation aspects; finally, Section 5 draws some conclusions and outlooks some future work.

2 Related Work

What is Quantum Annealing. QA is a QC paradigm that can be specifically used to solve optimization problems. A quantum annealer, which is a special-purpose quantum machine for QA, represents a problem as the energy of a physical system and then leverages quantum-mechanical phenomena to let the system find a state of minimal (or close to minimal) energy, corresponding to an optimal (or close to optimal) solution of the original problem.

[4] https://qclef.dei.unipd.it/

Since quantum annealers can be used to find the minimum energy state, problems must be formulated as minimization ones using the *Quadratic Unconstrained Binary Optimization (QUBO)* formulation, defined as follows:

$$\min \quad y = x^T Q x$$

where x is a vector of binary decision variables and Q is a matrix of constant values representing the problem to solve. Through QUBO formulations, it is possible to represent many problems [7]. Then, the *minor embedding* step maps the QUBO problem into the quantum annealer hardware, accounting for its topology. This can be done automatically, relying on some heuristics. A QUBO problem is usually solved by quantum annealers in few *milliseconds*.

Applications of Quantum Annealing. QA can be practically used to tackle problems in different fields due to its capabilities of solving *NP-Hard* integer optimization problems. QA has previously been employed to address IR and RS tasks, including Feature Selection [13], demonstrating both feasibility and competitiveness in terms of performance. Moreover, QA has also been explored in various *Machine Learning (ML)* tasks. For instance, Willsch et al. [26] introduce a kernel-based *Support Vector Machine (SVM)* formulation implemented on a D-Wave 2000Q quantum annealer, while Delilbasic et al. [5] propose a quantum multiclass SVM approach designed to reduce execution time for large training datasets. Additionally, QA has been applied to clustering problems; for example, Zaiou et al. [28] leverage it in a balanced K-means method, achieving improved performance as measured by the Davies-Bouldin Index.

Previous Editions. This is the third edition of the QuantumCLEF Lab. In the previous editions (2024 and 2025), there have been a total of 12 research teams that successfully participated and provided official submissions for the proposed tasks [16, 18]. The previous editions encompassed different tasks such as Feature Selection, Clustering, and Instance Selection using state-of-the-art quantum annealers. The results in the previous editions suggest that quantum annealers are overall able to maintain a comparable level of effectiveness with respect to more traditional approaches (e.g., Simulated Annealing) while being able to solve the problems efficiently, considering just the time required for the annealing phase [16–18]. We received very positive feedback from the participants, most of which never experienced using QC resources before.

3 Tasks

In this edition of the qCLEF lab, we plan to organize three different tasks, each with the following goals:

- find one or more possible QUBO formulations of the problem;
- evaluate the quantum annealer approach compared to a corresponding traditional approach to assess both its efficiency and its effectiveness.

In general, we expect QA to solve problems more quickly than traditional approaches, achieving results that are similar or better in terms of effectiveness.

Additionally, based on the availability of gate-based quantum computers, we plan to organize a further small set of tasks involving the use of these devices. These will likely be the target for more experienced participants having prior knowledge on QC.

3.1 Task 1 - Quantum Feature Selection

This task focuses on solving the *NP-Hard* Feature Selection problem with QA, similarly to other previous works [4, 13].

Feature Selection is a well-known problem for both IR and RS, which requires the identification of a subset of the available features to train a learning model more efficiently and effectively. This problem is very impacting since many IR and RS systems involve the optimization of learning models, and reducing the dimensionality and noise of the input data can improve their performance.

If the input data has n features, we can enumerate all the possible sets of input data having a fixed number k of features, thus obtaining $\binom{n}{k}$ possible subsets. Therefore, to find the best subset of k features the learning model should be trained on all the subsets of features, which is infeasible even for small datasets. So, in this task, we want to understand if QA can be used to solve this problem more efficiently and effectively.

We have identified some possible datasets, such as MQ2007 [22] or Istella S-LETOR [10]. These datasets contain pre-computed features, and the objective is to select a subset of these features to train a learning model, such as LambdaMART [1] or a content-based RS, and to achieve the best performance according to metrics such as nDCG@10.

3.2 Task 2 - Quantum Instance Selection

This task focuses on formulating and solving the Instance Selection problem with QA [14].

Currently, transformer-based architectures, including 1st and 2nd generation transformers (e.g., RoBERTa [9]) as well as current large language models (e.g., Llama3 [24]), are considered state-of-the-art in several fields. Given their high costs, one of the big challenges is to fine-tune these models efficiently. Instance Selection focuses on selecting a representative subset of instances from a dataset to make the training of these models faster while maintaining a high level of effectiveness of the trained model [2, 3].

In this task, we thus aim at using QA to find a good subset of instances in a dataset in an efficient way, for fine-tuning a **Llama3.1** model to perform a text classification task as effectively as it would on the entire original dataset.

We have identified some possible datasets, such as Vader NYT or Yelp Reviews, that will be provided in a five-fold cross-validation split. The extracted subsets will be then used to fine-tune the Llama3.1 model and the effectiveness will be measured with the Macro-F1 score [21].

3.3 Task 3 - Quantum Clustering

This task focuses on the formulation of the Clustering problem and solving it with a quantum annealer.

Clustering can be useful to organize large collections and help users to explore them. It can also be used to divide users according to their interests or build user models with the cluster centroids [27], boosting efficiency or effectiveness for users with limited data. In this task, each document or user can be represented as a vector in a similarity space, and it is possible to cluster documents based on their similarity between each other.

For the IR task, we have identified ANTIQUE [8] as a possible dataset. From the dataset, we will produce embeddings using models such as BERT [6]. The cluster quality will be measured with user queries that undergo the same embedding process. These queries will match only the most representative embeddings of the clusters, avoiding computing similarities on the whole collection, as in an approximate vector search scenario. For the recommendation task, the goal will be to partition the users into communities based on their past interactions, in such a way that users within a community share similar interests [12]. The quality of the communities will be assessed based on the effectiveness of a non-personalized RS algorithm trained on each community.

The cluster quality will be measured according to the Davies-Bouldin Index and nDCG@10.

3.4 Gate-Based Teaching Tasks

In addition to the QA tracks, the lab will include a small set of gate-model tasks conceived as teaching exercises for researchers who want to experiment with circuit-based paradigms in IR and RS. The tasks will be intentionally small to reflect limited resource availability and the cost of simulating gate-level systems.

The first task focuses on computing the similarity of documents with a swap test. Given a swap-test circuit, the participants will have to define an appropriate encoding for the documents with complex-valued embeddings (or quantum states) so that the test can measure their similarity. A second, more advanced task asks participants to design not only the encoding, but also the short quantum circuit able to build it. The goal is to identify a sequence of operations that yields good encodings for document similarity while keeping the number of operations small. Finally, we introduce a hybrid quantum-classical exercise based on the Variational Quantum Approximate Optimization Algorithm to solve small QUBO instances. Participants will implement a minimal QAOA loop, report solution quality and runtime.

4 Evaluation of Quantum Annealing

Using a quantum annealer requires several stages:

Formulation: compute the QUBO matrix Q;

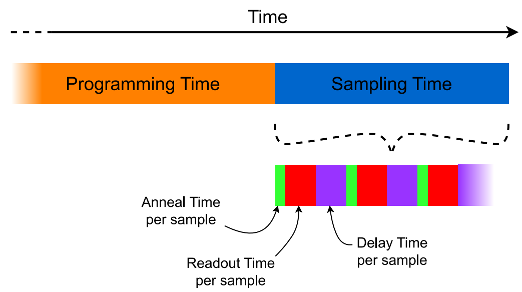

Fig. 1: The quantum annealer access time split in several steps.

Embedding: generate the *minor embedding* of the QUBO for the hardware;
Data Transfer: transfer the problem and the embedding to the data center that hosts the quantum annealer;
Annealing: run the quantum annealer itself.

In terms of effectiveness, at least two sources of stochasticity must be considered. First, during the embedding phase, heuristic methods transform the QUBO formulation into an equivalent version that fits the hardware. This process is non-deterministic: it can produce different embeddings for the same problem, which are theoretically equivalent but may lead to variations in practice. Second, the annealing phase samples a low-energy solution. However, obtaining a reliable solution may require many samples.

Regarding efficiency, problem transfer over the network can introduce significant delays, and generating the minor embedding may even take minutes for particularly large problems. Furthermore, the runtime can be divided into several stages, as illustrated in Fig. 1: first, the device must be programmed; next, the annealing process is executed; finally, the results are read. The annealing itself is extremely fast, requiring only *microseconds*, but it must be repeated multiple times due to the device's stochastic nature. Therefore, a complete efficiency assessment must account for the time required by all stages of the process.

5 Conclusions and Future Work

In this paper, we have outlined the third edition of the qCLEF lab, a lab composed of three practical tasks aiming at evaluating the performance of QC and especially QA applied to IR and RS. We have also discussed the potential benefits that QA can bring in terms of performance to the IR and RS fields, and we have also analyzed the different challenges that lie in the evaluation of these new quantum-based approaches.

qCLEF is the first initiative that provides access to real state-of-the-art quantum devices for researchers to tackle practical IR and RS tasks. It also represents a unique opportunity for researchers worldwide to start learning more about these new cutting-edge technologies that will likely have a big impact in the future. Finally, through our lab, it will also be possible to assess whether quantum-based methods can be employed to improve the current state-of-the-art approaches, hopefully delivering new performing solutions.

Acknowledgements

We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support. W. Cunha and M. A. Gonçalves have further been supported by the CNPq, CAPES, Instituto Nacional de Ciência e Tecnologia em Inteligência Artificial Responsável para Linguística Computacional, Tratamento e Disseminação de Informação (INCT-TILDIAR; grant # 408490/2024-1), FAPEMIG, AWS, Google, NVIDIA, CIIASaúde, and FAPESP.

Disclosure of Interests

The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Burges, C.J.C.: From RankNet to LambdaRank to LambdaMART: An Overview. Tech. rep., Microsoft Research, MSR-TR-2010-82 (June 2010)
2. Cunha, W., França, C., Fonseca, G., Rocha, L., Gonçalves, M.A.: An effective, efficient, and scalable confidence-based instance selection framework for transformer-based text classification. In: Chen, H., Duh, W.E., Huang, H., Kato, M.P., Mothe, J., Poblete, B. (eds.) Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023, pp. 665–674, ACM (2023), https://doi.org/10.1145/3539618.3591638, URL https://doi.org/10.1145/3539618.3591638
3. Cunha, W., Viegas, F., França, C., Rosa, T., Rocha, L., Gonçalves, M.A.: A comparative survey of instance selection methods applied to non-neural and transformer-based text classification. ACM Comput. Surv. **55**(13s), 265:1–265:52 (2023), https://doi.org/10.1145/3582000, URL https://doi.org/10.1145/3582000
4. Dacrema, M.F., Moroni, F., Nembrini, R., Ferro, N., Faggioli, G., Cremonesi, P.: Towards feature selection for ranking and classification exploiting quantum annealers. In: Amigó, E., Castells, P., Gonzalo, J., Carterette, B., Culpepper, J.S., Kazai, G. (eds.) SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022, pp. 2814–2824, ACM (2022), https://doi.org/10.1145/3477495.3531755, URL https://doi.org/10.1145/3477495.3531755

5. Delilbasic, A., Saux, B.L., Riedel, M., Michielsen, K., Cavallaro, G.: A single-step multiclass svm based on quantum annealing for remote sensing data classification. arXiv preprint arXiv:2303.11705 (2023)
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
7. Glover, F.W., Kochenberger, G.A., Du, Y.: Quantum bridge analytics I: a tutorial on formulating and using QUBO models. 4OR **17**(4), 335–371 (2019), https://doi.org/10.1007/s10288-019-00424-y
8. Hashemi, H., Aliannejadi, M., Zamani, H., Croft, W.B.: Antique: A non-factoid question answering benchmark. In: Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II 42, pp. 166–173, Springer (2020)
9. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pre-training approach. arXiv preprint 1907.11692 (2019)
10. Lucchese, C., Nardini, F.M., Orlando, S., Perego, R., Silvestri, F., Trani, S.: 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, pp. 949–952 (2016)
11. Melucci, M.: Introduction to Information Retrieval and Quantum Mechanics, The Information Retrieval Series, vol. 35. Springer (2015), ISBN 978-3-662-48312-1, https://doi.org/10.1007/978-3-662-48313-8, URL https://doi.org/10.1007/978-3-662-48313-8
12. Nembrini, R., Carugno, C., Ferrari Dacrema, M., Cremonesi, P.: Towards recommender systems with community detection and quantum computing. In: RecSys '22: Sixteenth ACM Conference on Recommender Systems, Seattle, WA, USA, September 18 - 23, 2022, pp. 579–585, ACM (2022), https://doi.org/10.1145/3523227.3551478
13. Nembrini, R., Dacrema, M.F., Cremonesi, P.: Feature selection for recommender systems with quantum computing. Entropy **23**(8), 970 (2021), https://doi.org/10.3390/E23080970, URL https://doi.org/10.3390/e23080970
14. Pasin, A., Cunha, W., Gonçalves, M.A., Ferro, N.: A quantum annealing instance selection approach for efficient and effective transformer fine-tuning. In: Oosterhuis, H., Bast, H., Xiong, C. (eds.) Proceedings of the 2024 ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR 2024, Washington, DC, USA, 13 July 2024, pp. 205–214, ACM (2024), https://doi.org/10.1145/3664190.3672515, URL https://doi.org/10.1145/3664190.3672515
15. Pasin, A., Dacrema, M.F., Cremonesi, P., Ferro, N.: qclef: A proposal to evaluate quantum annealing for information retrieval and recommender systems. In: Arampatzis, A., Kanoulas, E., Tsikrika, T., Vrochidis, S., Giachanou, A., Li, D., Aliannejadi, M., Vlachos, M., Faggioli, G., Ferro, N. (eds.) Experimental IR Meets Multilinguality, Multimodality, and Interaction - 14th In-

ternational Conference of the CLEF Association, CLEF 2023, Thessaloniki, Greece, September 18-21, 2023, Proceedings, Lecture Notes in Computer Science, vol. 14163, pp. 97–108, Springer (2023), https://doi.org/10.1007/978-3-031-42448-9_9, URL https://doi.org/10.1007/978-3-031-42448-9_9

16. Pasin, A., Dacrema, M.F., Cremonesi, P., Ferro, N.: Overview of quantumclef 2024: The quantum computing challenge for information retrieval and recommender systems at CLEF. In: Goeuriot, L., Mulhem, P., Quénot, G., Schwab, D., Nunzio, G.M.D., Soulier, L., Galuscáková, P., de Herrera, A.G.S., Faggioli, G., Ferro, N. (eds.) Experimental IR Meets Multilinguality, Multimodality, and Interaction - 15th International Conference of the CLEF Association, CLEF 2024, Grenoble, France, September 9-12, 2024, Proceedings, Part II, Lecture Notes in Computer Science, vol. 14959, pp. 260–282, Springer (2024), https://doi.org/10.1007/978-3-031-71908-0_12, URL https://doi.org/10.1007/978-3-031-71908-0_12
17. Pasin, A., Dacrema, M.F., Cremonesi, P., Ferro, N.: Quantumclef 2024: Overview of the quantum computing challenge for information retrieval and recommender systems at CLEF. In: Faggioli, G., Ferro, N., Galuscáková, P., de Herrera, A.G.S. (eds.) Working Notes of the Conference and Labs of the Evaluation Forum (CLEF 2024), Grenoble, France, 9-12 September, 2024, CEUR Workshop Proceedings, vol. 3740, pp. 3032–3053, CEUR-WS.org (2024), URL https://ceur-ws.org/Vol-3740/paper-297.pdf
18. Pasin, A., Dacrema, M.F., Cunha, W., Gonçalves, M.A., Cremonesi, P., Ferro, N.: Overview of quantumclef 2025: The second quantum computing challenge for information retrieval and recommender systems at CLEF. In: Carrillo-de-Albornoz, J., de Herrera, A.G.S., Gonzalo, J., Plaza, L., Mothe, J., Piroi, F., Rosso, P., Spina, D., Faggioli, G., Ferro, N. (eds.) Experimental IR Meets Multilinguality, Multimodality, and Interaction - 16th International Conference of the CLEF Association, CLEF 2025, Madrid, Spain, September 9-12, 2025, Proceedings, Lecture Notes in Computer Science, vol. 16089, pp. 412–435, Springer (2025), https://doi.org/10.1007/978-3-032-04354-2_22, URL https://doi.org/10.1007/978-3-032-04354-2_22
19. Pasin, A., Ferro, N.: KIMERA: from evaluation-as-a-service to evaluation-in-the-cloud. In: Ferro, N., Maistro, M., Pasi, G., Alonso, O., Trotman, A., Verberne, S. (eds.) Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2025, Padua, Italy, July 13-18, 2025, pp. 3584–3594, ACM (2025), https://doi.org/10.1145/3726302.3730298, URL https://doi.org/10.1145/3726302.3730298
20. Pilato, G., Vella, F.: A survey on quantum computing for recommendation systems. Inf. **14**(1), 20 (2023), https://doi.org/10.3390/info14010020
21. Powers, D.M.W.: Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation. CoRR **abs/2010.16061** (2020), URL https://arxiv.org/abs/2010.16061
22. Qin, T., Liu, T.Y.: Introducing LETOR 4.0 Datasets. arXiv.org, Information Retrieval (cs.IR) **arXiv:1306.2597** (June 2013)

23. van Rijsbergen, C.J.: The geometry of information retrieval. Cambridge University Press (2004), ISBN 978-0-521-83805-4
24. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., Lample, G.: Llama: Open and efficient foundation language models (2023)
25. Uprety, S., Gkoumas, D., Song, D.: A survey of quantum theory inspired approaches to information retrieval. ACM Comput. Surv. **53**(5), 98:1–98:39 (2021), https://doi.org/10.1145/3402179, URL https://doi.org/10.1145/3402179
26. Willsch, D., Willsch, M., De Raedt, H., Michielsen, K.: Support vector machines on the d-wave quantum annealer. Computer physics communications **248**, 107006 (2020)
27. Wu, Y., Cao, Q., Shen, H., Tao, S., Cheng, X.: INMO: A model-agnostic and scalable module for inductive collaborative filtering. In: SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, pp. 91–101, ACM (2022), https://doi.org/10.1145/3477495.3532000
28. Zaiou, A., Bennani, Y., Matei, B., Hibti, M.: Balanced k-means using quantum annealing. In: 2021 IEEE Symposium Series on Computational Intelligence (SSCI), pp. 1–7 (2021), https://doi.org/10.1109/SSCI50451.2021.9659997