Reproducibility Report for ACM SIGMOD 2023 Paper: "A Universal Question-Answering Platform for Knowledge Graphs"

KEMING LI, University of California, Irvine, USA
NIMA SHAHBAZI, University of Illinois Chicago, USA
IMMANUEL HAFFNER, Firebolt Analytics Inc., Germany
REHAM OMAR, Concordia University, Canada

We have been able to reproduce the core results of the paper. We used machines identically configured to the author's machine to reproduce the response time experiment.

1 INTRODUCTION

The reproduced paper [1] is written by Omar, Dhall, and Mansour from Concordia University, Canada and Kalnis from KAUST, Saudi Arabia. Source codes are provided in the github repository, scripts are provided to run the experiments of the paper, and experiment results are matched using similar machines.

2 SUBMISSION

To reproduce the experimental results we used the following artifacts:

- GitHub repository with code and scripts at: https://github.com/CoDS-GCS/KGQAn
- A detailed readme file at: https://github.com/CoDS-GCS/KGQAn/blob/dev/reproducability/Instructions.md
- A master script for experiments at: https://github.com/CoDS-GCS/KGQAn/blob/dev/docker_run.sh
- A master figure generation script at: https://github.com/CoDS-GCS/KGQAn/blob/dev/figures_generation.sh

3 HARDWARE AND SOFTWARE ENVIRONMENT

The paper itself uses two machines for experiments, as shown in Table 1. One is named "Paper-main" that means most of the experiments in this paper use this environment, and the other one is "Paper-index" corresponding to the indexing experiment in Table 2 of the paper [1], because it requires a very large amount of memory.

Table 1. Hardware & Software environment

	Paper-main	Paper-index	Repro Review
CPU	16	16	Intel Gold 5122
cores	16	32	4
OS	Ubuntu 20.04	Rocky 9	Ubuntu 18.04
RAM	500GB	3TB	512GB
GPU	NVIDIA Tesla P100		Titan XP * 2

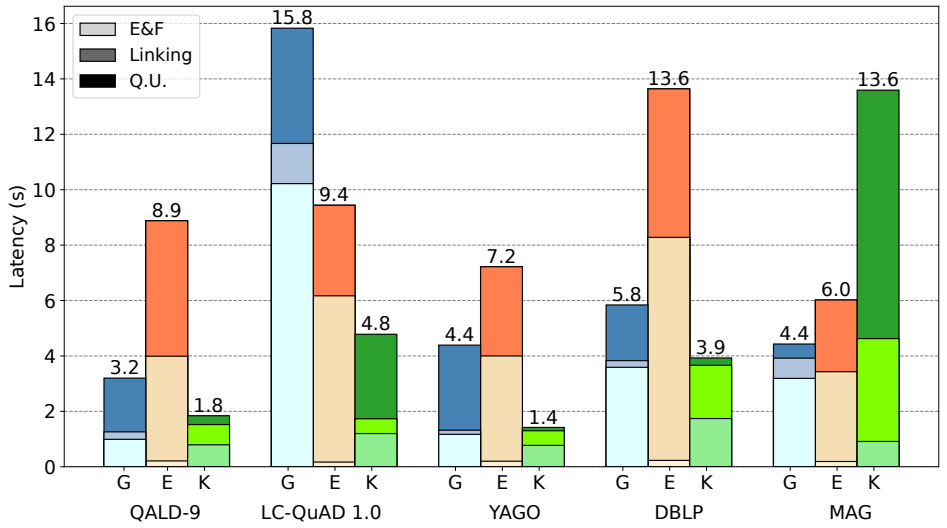

Fig. 1. Reproduce of paper's Figure 7

4 REPRODUCIBILITY EVALUATION

4.1 Process

The most challenging part of the reproducibility effort was to find a powerful machine with large-enough memory. Further, the experiments require a machine with high GPU computation power. Having said that, a GPU is not necessary but to run the experiments, but it has important performance implications. Once the experimental setup is prepared, the experiments can be reproducing from the "reproducibility" folder of the authors' GitHub repository, following the instructions in "instruction.md".

4.2 Results

In order to match the same response time as mentioned in the paper, using a similar high-end setup as the one available to the authors is required. The other experimental results were successfully collected following the instructors in "instruction.md". What we reproduced is as shown in Figures 1-4. These results represent the core experiments needed to support the main thesis of the paper.

4.3 Brief Analysis of Reproduced Experiments

- When running the experiments using a less powerful machine than the one used by the authors, we reproduced Table 3, and Figures 9 and 10. However, some experiments in Figure 7 related to MAG were affected by the computing power of the machine.
- Specifically, it is expected that MAG would take more time than the other graphs because the size of MAG is 60 times larger than the second-largest graph and the computation power of the machine skews the results compared to the original paper.
- The response time experiment is significantly affected when running on a machine without GPU, however, even a CPU-based execution produces manageable times since the experiment only performs inference.

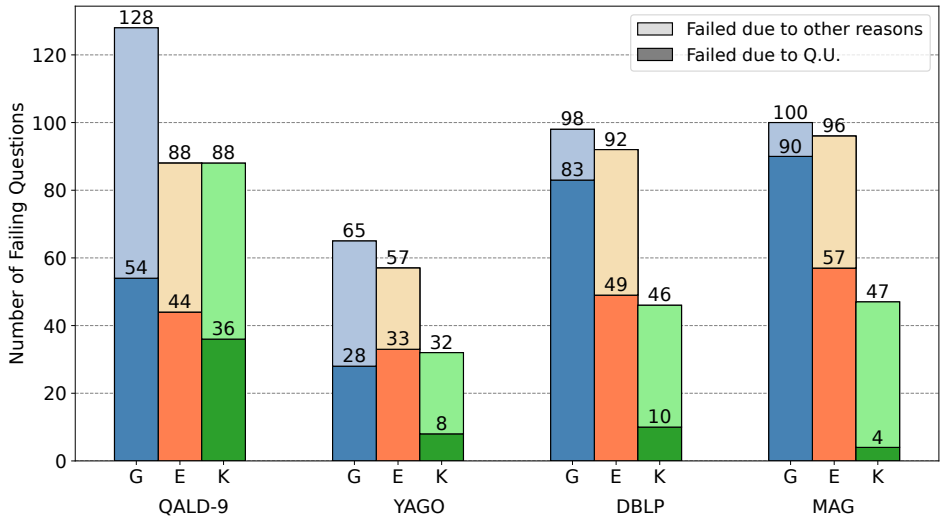

Fig. 2. Reproduce of paper's Figure 8

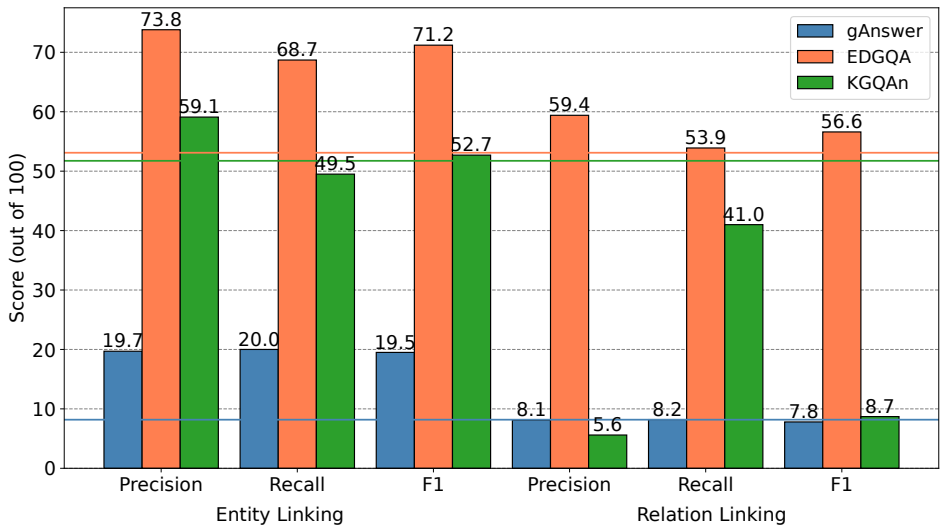

Fig. 3. Reproduce of paper's Figure 9

- The authors worked with the reviewers to update the quality of the GitHub repository and the final artifact, the changes included updating the code to handle the most recent docker version, fixing path issues, typos, and some organizational changes in the repository.

5 SUMMARY

The core thesis of the paper is reproduced. There was a discrepancy in one of the bars (MAG) as shown in Figure 1 (Figure 7 of the paper), which is discussed in Section 4.3.

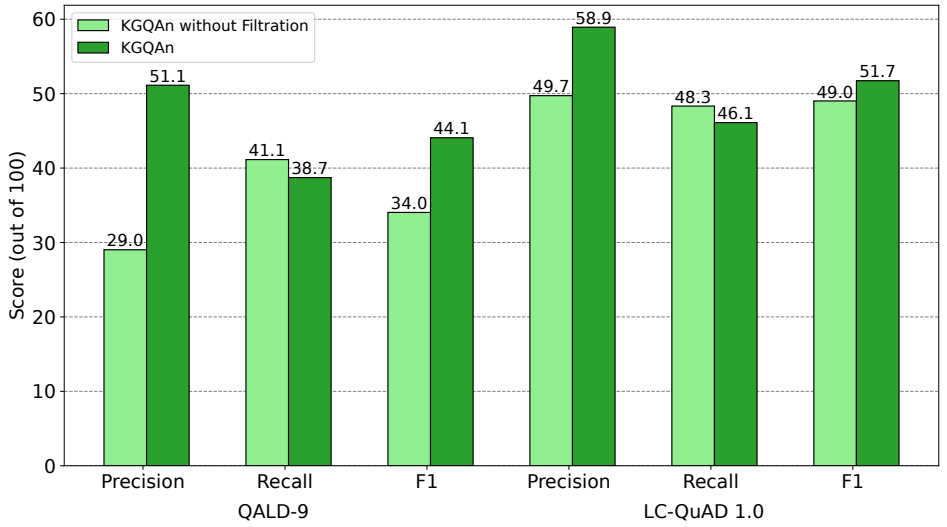

Fig. 4. Reproduce of paper's Figure 10

REFERENCES

[1] Reham Omar, Ishika Dhall, Panos Kalnis, and Essam Mansour. 2023. A universal question-answering platform for knowledge graphs. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–25.