Reproducibility Report for ACM SIGMOD 2023 Paper: “A Step Toward Deep Online Aggregation”

TIANXUN HU, Simon Fraser University, Canada
THOMAS HÜTTER, University of Salzburg, Austria
GUANGJING WANG, Michigan State University, USA

The source code package provided by the authors is well-organized. The core results of the paper can be successfully reproduced. The extensibility of the evaluation is facilitated by the well-designed scripts. We thank the authors for putting their efforts into DB Availability and Reproducibility.

1 INTRODUCTION

The original paper [1] deals with the challenge of *Deep Online Aggregation (Deep OLA)* to estimate the answer of nested analytic operations on data. An OLA system *WAKE* is proposed that is based on a novel data model, called *evolving data frames (edf)*. The proposed system is evaluated and compared to state-of-the-art OLA and non-OLA systems concerning latency, memory usage, approximation error, and recall.

2 SUBMISSION

The reproducibility submission mainly refers to the repository of the produced artifacts, which is available on GitHub: https://github.com/illinoisdata/DeepOLA/. The repository is self-contained and includes the source code, data generation, and documentation. A detailed list of the repository content is given:

- Documentation: A `README` and reproducibility instructions (cf. https://github.com/illinoisdata/DeepOLA/blob/main/sigmod2023-reproducibility.md) are included.
- Source code: The repository contains the source code of the developed approach, implemented in Rust. Next to the system, scripts are provided to facilitate the execution of the algorithms. Note that some competitors are not included due to license issues.
- Datasets: The data used in the experimental evaluation is based on the TPC-H dataset. To generate the data and the queries, the authors use the `tpch-dbgen` kit available at https://github.com/dragansah/tpch-dbgen. Scripts that reproduce their data are provided in the repository as well.

3 HARDWARE AND SOFTWARE ENVIRONMENT

Compared to the experiments in the original paper, we use different hardware for our reproducibility review. To facilitate the comparison of the original and the reproduced results, we summarize the differences in Table 1. Note that in the reproducibility review, we simulated the environment from the original paper by Docker containers that are set with runtime options to limit memory size to 64GB (`–memory=65536m`) and CPU cores to 16 (`–cpus=16`).

4 REPRODUCIBILITY EVALUATION

4.1 Process

To reproduce the experiments, we followed the approach described in the reproducibility instructions. In particular, we cloned the repository and used the `master_script.sh` with the suggested parameters, i.e., `$(pwd)/datasets 100 100 10`.

Make sure to equip yourself with a disk with more than 900GB of space left. For example, moving the Docker data directory to another partition with sufficient space is useful in some cases (e.g.,

Table 1. Hardware & Software environment

	Paper	Reproducibility Review
Machine	Azure Standard D16ads v5	Local Server
CPU		Intel Xeon Gold 6242R
cores	16 vCPU(s)	2 × 40 (× 2 threads) (limited to 16 in Docker)
GHz		3.1
RAM	64GB	375GB (limited to 64GB in Docker)
Storage		SSD

the Wanderjoin part). During the reproducibility process, we encountered "`The input device is not a TTY`" error when running Postgres to get latencies (cf. Figure 1). We fixed the error by removing the runtime option `-it` from `docker run` commands.

After the execution of the master script, the resulting figures are located in `results/viz` in the root directory of the repository. There were two plotting bugs detected during our evaluation. The first one was in Figure 7 where the WAKE-final variant was plotted twice instead of the WAKE-first variant. The second one was in Figure 9 where the MAPE value of Wanderjoin was a ratio instead of a percentage. Thus, we suggested to the authors to scale the values to percentages before plotting. We contacted the authors to fix these issues in the repository.

4.2 Results

Executing the master script aims to reproduce Figures {7-11} of the original paper. Note that Figure 12 is skipped on purpose. We next discuss the findings in our experimental results based on the reproduced plots (cf. Figures 1 - 5).

Summarized, our reproduced figures show similar trends to the figures in the original paper. However, we observed three exceptions: (1) In Figure 2, the plot for Q18 does not show the result for MAPE. (2) In Figure 3, the initial approximation errors of WAKE for modified Q3, Q7, and Q10 are not around 1%. (3) In Figure 4, the CI value deviates from the results in the original figure, e.g., the Max and P95 values are almost identical. Further, using the suggested parameters in the reproducibility instructions will miss the data point for 150 partitions as presented in the paper.

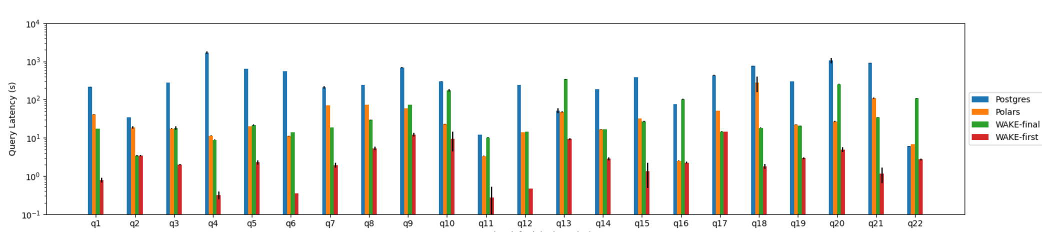

Fig. 1. Reproduced results matching Figure 7 of the original paper.

5 SUMMARY

In conclusion, the core results of the paper can be successfully reproduced. We want to thank the authors for their efforts to fix bugs and cooperation during this review process.

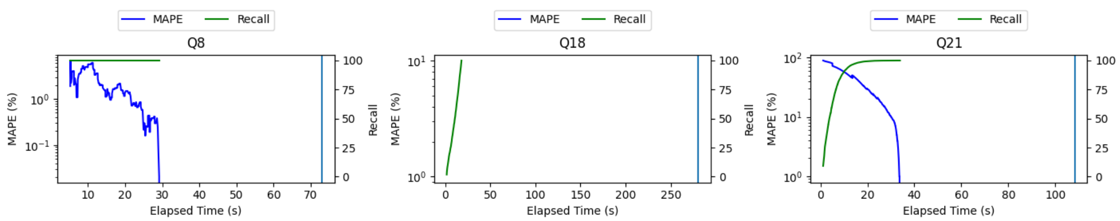

Fig. 2. Reproduced results matching Figure 8 of the original paper.

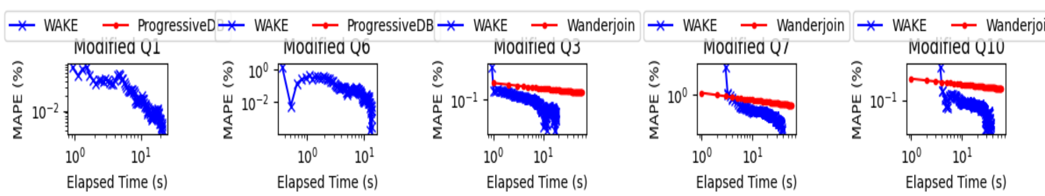

Fig. 3. Reproduced results matching Figure 9 of the original paper.

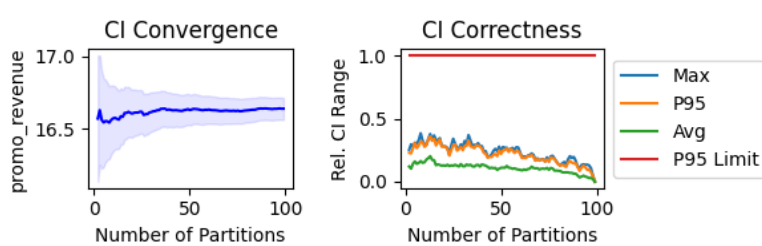

Fig. 4. Reproduced results matching Figure 10 of the original paper.

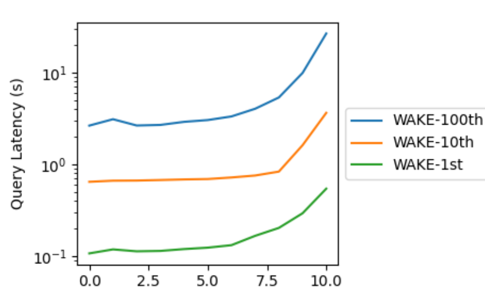

Fig. 5. Reproduced results matching Figure 11 of the original paper.

REFERENCES

[1] Nikhil Sheoran, Supawit Chockchowwat, Arav Chheda, Suwen Wang, Riya Verma, and Yongjoo Park. 2023. A Step Toward Deep Online Aggregation. *Proceedings of the ACM on Management of Data* 1, 2 (2023), 1–28.