An Inquiry into Machine Learning-based Automatic Configuration Tuning Services on Real-World Database Management Systems

Dana Van Aken, Dongsheng Yang, Sebastien Brillard, Ari Fiorino,
Bohan Zhang, Christian Bilien, Andrew Pavlo

Motivation

Problem

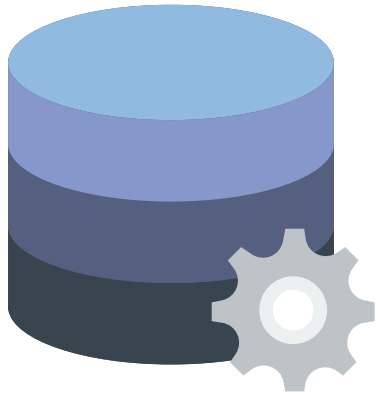

Modern DBMS have hundreds of tunable configuration knobs

Source: Aken et al., 2021

Existing Approaches

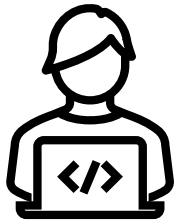

DBA

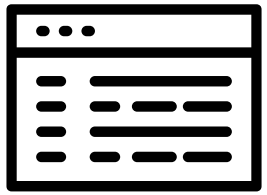

Static rules

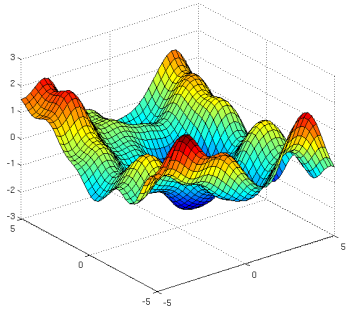

Machine learning

Limitations

Open-source DBMSs

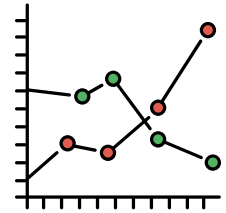

Synthetic benchmarks

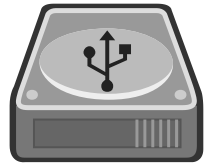

Dedicated local storage

Machine Learning Based Knob Tuning

Algorithms

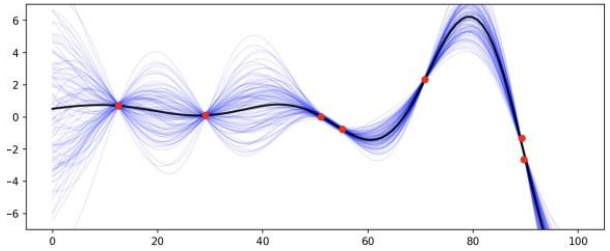

Gaussian Process Regression (GPR)

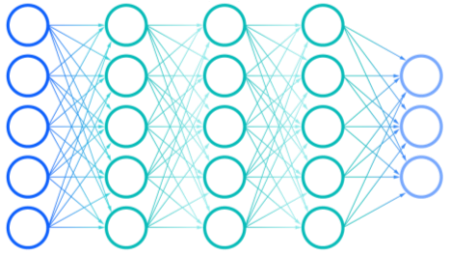

Deep Neural Network (DNN)

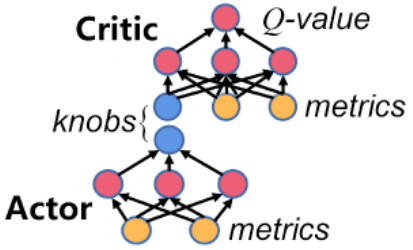

Deep Deterministic Policy Gradient (DDPG++)

Real-world case study

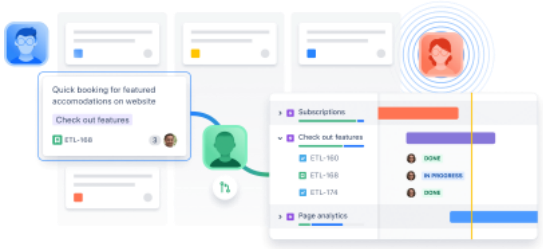

Real-world ticketing application

Oracle database

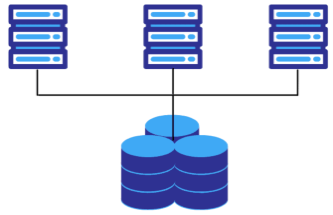

Shared storage

Source: Aken et al., 2021

OtterTune Architecture (Aken et al., 2017)

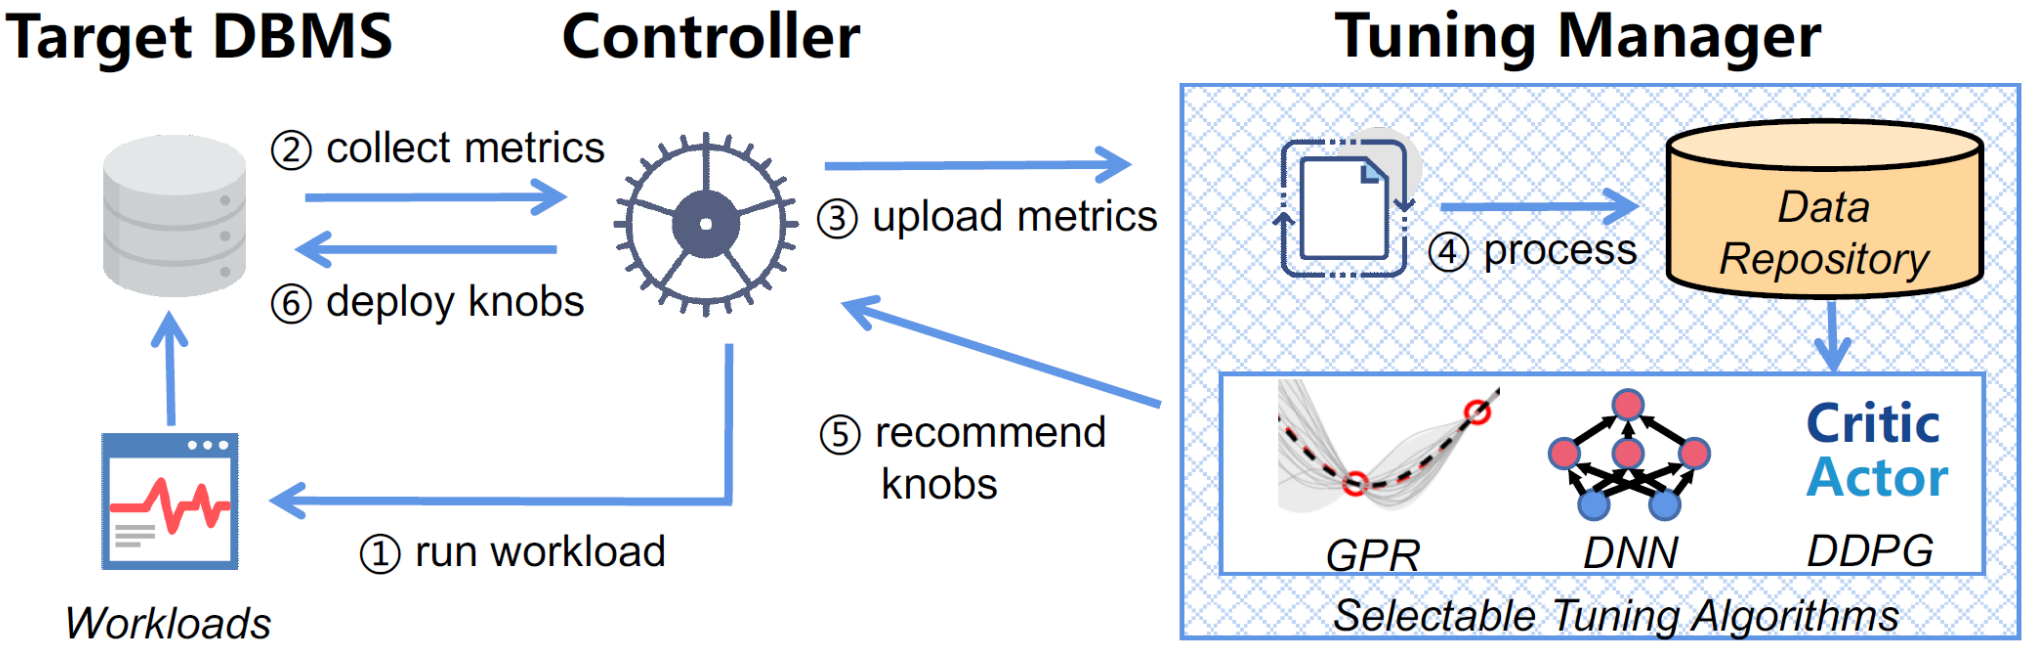

Source: Aken et al., 2021

GPR/DNN Tuning Pipeline (OtterTune)

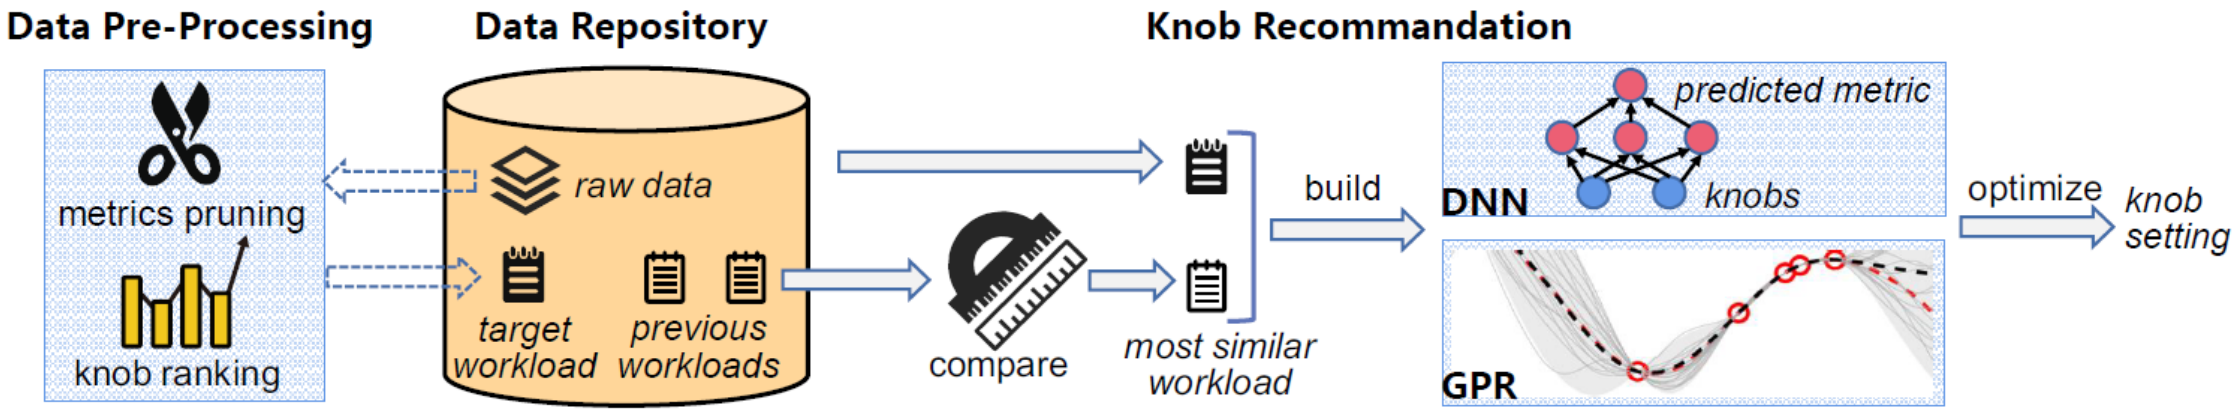

Source: Aken et al., 2021

DDPG++ Tuning Pipeline (CDBTune)

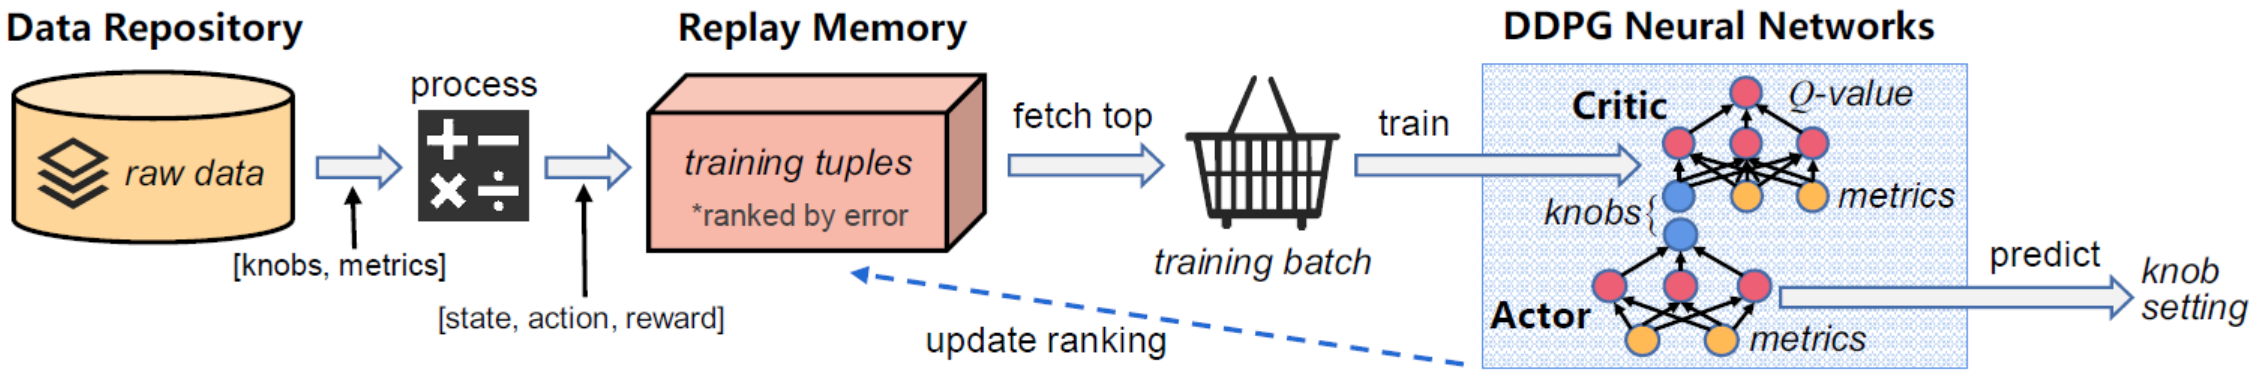

Source: Aken et al., 2021

Machine Learning Based Knob Tuning

Algorithms

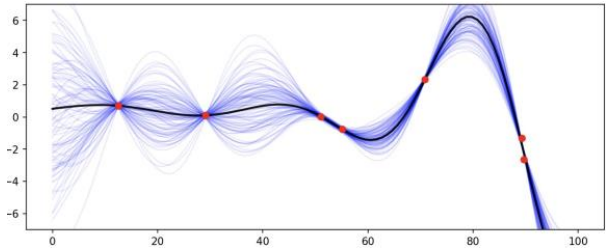

Gaussian Process Regression (GPR)

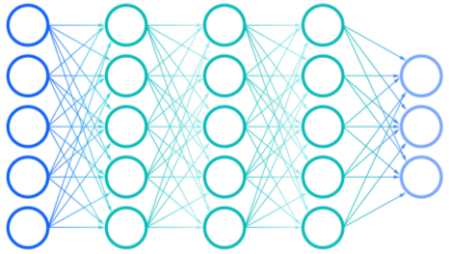

Deep Neural Network (DNN)

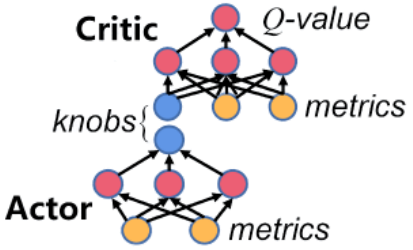

Deep Deterministic Policy Gradient (DDPG++)

Real-world case study

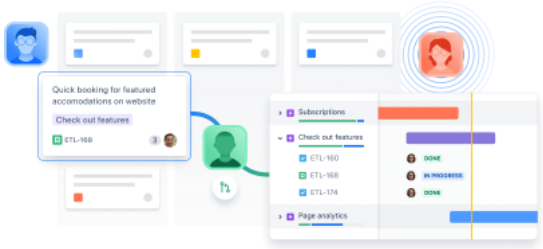

Real-world ticketing application

Oracle database

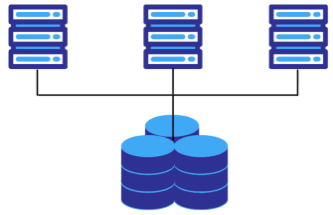

Shared storage

Source: Aken et al., 2021

Tuning Knobs Selected by DBA

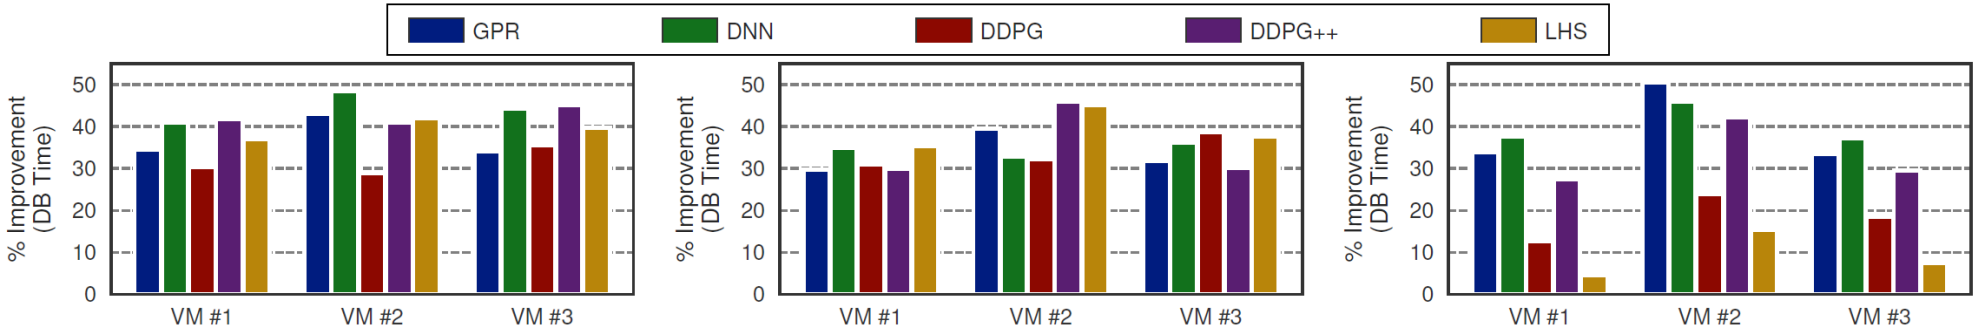

(a) 10 Knobs **(b)** 20 Knobs **(c)** 40 Knobs

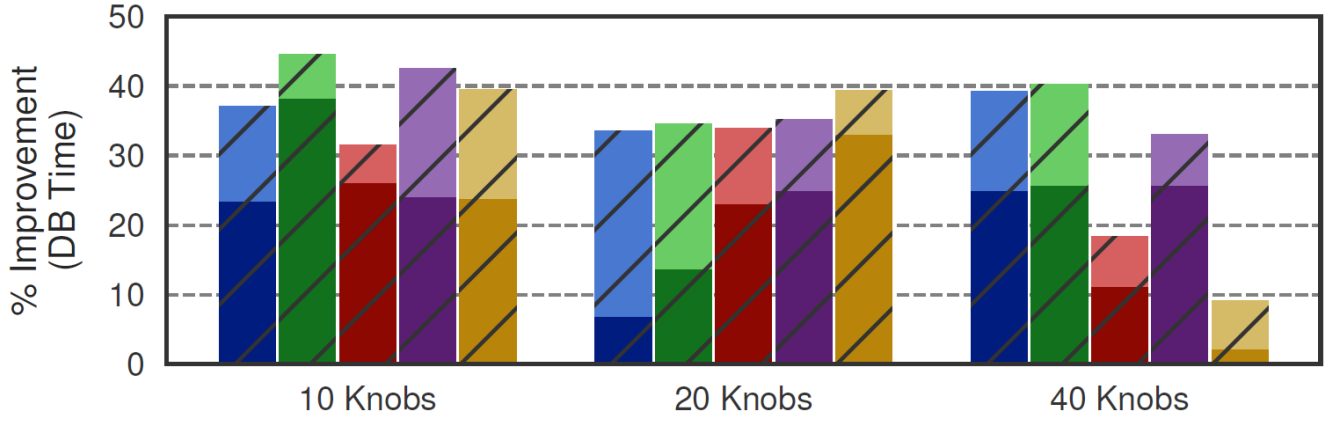

Source: Aken et al., 2021

Tuning Knobs Ranked by OtterTune

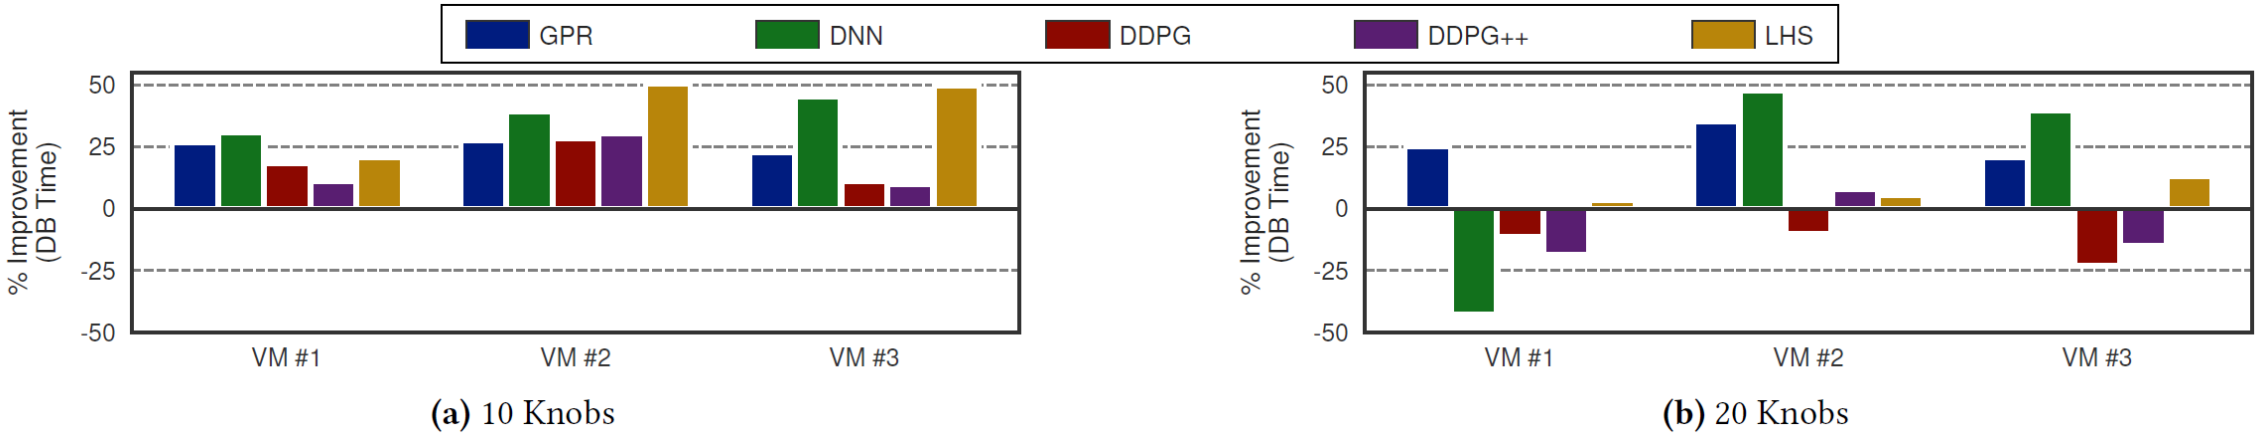

(a) 10 Knobs

(b) 20 Knobs

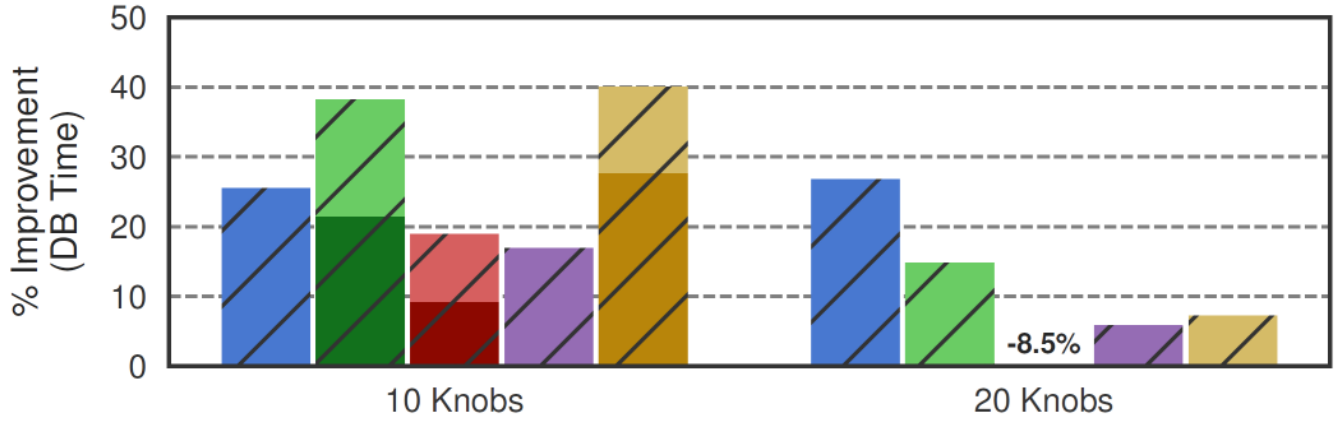

Source: Aken et al., 2021

Adaptability to Different Workloads

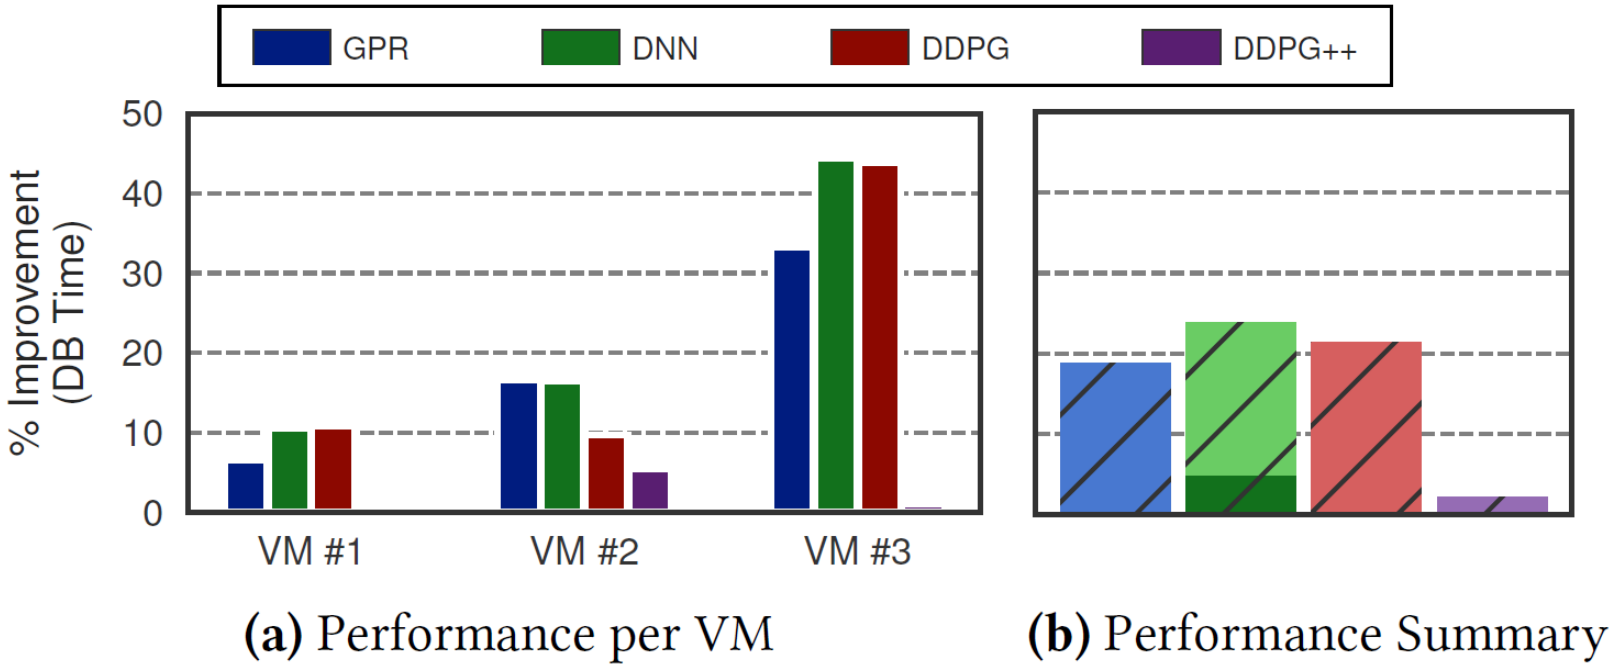

(a) Performance per VM

(b) Performance Summary

Source: Aken et al., 2021

Summary

- OtterTune extension: DNN and DDPG++
- Real-world case study: ticketing application running on Oracle database with shared storage
- All approaches significantly outperform the baseline, however, there is no clear ranking between the different approaches
- Hybrid approach combining DBA-selected and ML-tuned knobs outperforms fully automated approach
- Latin hypercube sampling performs (surprisingly) well for small knob sets
- The underlying hardware can have a significant impact on the performance

Source: Aken et al., 2021

Some Thoughts on the Paper

- Original OtterTune paper 👍, follow-up paper 👎
- Paper contains errors

"As an algorithm learns more, it is less likely to select poor configurations. Thus, the number of long-running replays decreases as the algorithm nears convergence. Table 4 shows that GPR has the fewest canceled replays. ***DDPG++ has fewer canceled replays than DDPG due to its improved convergence rate*** *(see Section 4.3). LHS has the highest workload execution time and percentage of canceled replays because it is a sampling technique and never converges."* [p. 11]

	GPR	DNN	DDPG	DDPG++	LHS
Execute (sec)	762	1006	1021	1274	1311
% Canceled	1.8%	8.7%	12.9%	26.8%	32.4%

Table 4: Workload Replay Time per Algorithm – The median workload execution time and the percentage of replays canceled for the algorithms.

Source: Aken et al., 2021

Some Thoughts on the Paper

- Industry case-study, few novel contributions to the field
- Issues raised in the introduction are only partially addressed
  - Open-source vs. enterprise DBMS
  - Synthetic vs. real-world workloads
  - Dedicated vs. shared storage
- Results provide limited insights for readers and practioners
- Extend evaluation to other real-world applications (e.g., OLAP workloads)
- Extend comparison of DBA-tuned knobs to DNN and DDPG++ approaches

OtterTune in 2022

Commercial Service

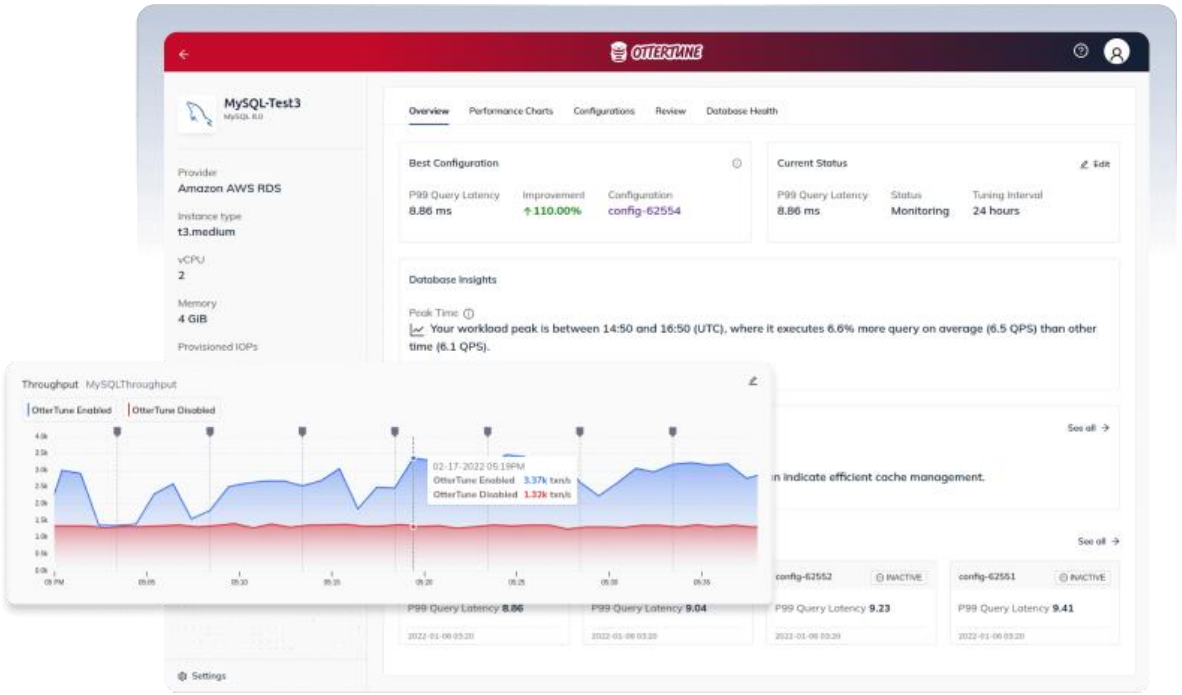

Record Label ☺

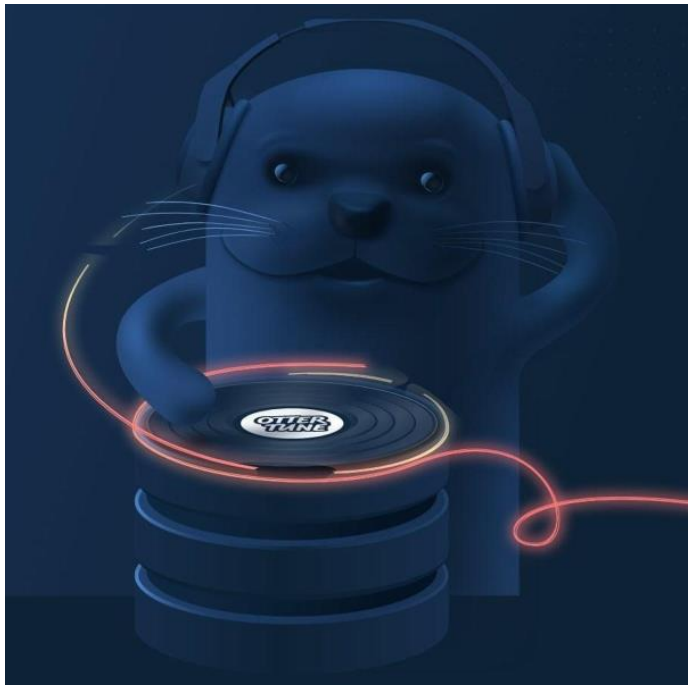

Questions / Discussion

References

- Dana Van Aken, Dongsheng Yang, Sebastien Brillard, Ari Fiorino, Bohan Zhang, Christian Bilien, and Andrew Pavlo. 2021. *An inquiry into machine learning-based automatic configuration tuning services on real-world database management systems*. Proc. VLDB Endow. 14, 7 (March 2021), 1241–1253. https://doi.org/10.14778/3450980.3450992
- Dana Van Aken, Andrew Pavlo, Geoffrey J. Gordon, and Bohan Zhang. 2017. *Automatic Database Management System Tuning Through Large-scale Machine Learning*. In Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD '17). Association for Computing Machinery, New York, NY, USA, 1009–1024. https://doi.org/10.1145/3035918.3064029

Image Sources

- Dana Van Aken, Dongsheng Yang, Sebastien Brillard, Ari Fiorino, Bohan Zhang, Christian Bilien, and Andrew Pavlo. 2021. *An inquiry into machine learning-based automatic configuration tuning services on real-world database management systems*. Proc. VLDB Endow. 14, 7 (March 2021), 1241–1253. https://doi.org/10.14778/3450980.3450992
- Dana Van Aken, Andrew Pavlo, Geoffrey J. Gordon, and Bohan Zhang. 2017. *Automatic Database Management System Tuning Through Large-scale Machine Learning*. In Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD '17). Association for Computing Machinery, New York, NY, USA, 1009–1024. https://doi.org/10.1145/3035918.3064029
- https://www.svgrepo.com/svg/95966/database
- https://www.svgrepo.com/svg/61151/web-code
- https://thegradient.pub/content/images/size/w1600/2019/11/kernel_cookbook-2.png
- https://d1.awsstatic.com/asset-repository/products/amazon-rds/1024px-MySQL.ff87215b43fd7292af172e2a5d9b844217262571.png
- https://images.g2crowd.com/uploads/product/image/large_detail/large_detail_251be2af3ae607c45c14e816eaa1cf41/postgresql.png
- https://www.svgrepo.com/svg/254993/stats
- https://www.svgrepo.com/svg/366187/harddisk-usb
- https://i.ytimg.com/vi/5ZOPtHKzHrM/maxresdefault.jpg
- https://1.cms.s81c.com/sites/default/files/2021-01-06/ICLH_Diagram_Batch_01_03-DeepNeuralNetwork-WHITEBG.png
- https://wac-cdn.atlassian.com/dam/jcr:4ddaea11-b00e-4546-8241-25f45689fed6/hero-illustration.png?cdnVersion=638
- https://www.instana.com/media/01_INSTANA_IconSet_oracle.svg
- https://thinketl.com/wp-content/uploads/2022/02/67-Shared-Disk.png
- https://www.svgrepo.com/svg/344899/hand-thumbs-up
- https://www.svgrepo.com/svg/344898/hand-thumbs-down
- https://ottertune.com/wp-content/uploads/2022/05/mockup-2.png
- https://www.zdnet.com/a/img/resize/3e223c1a202cf5eea3c9b7d906f1473582f743e0/2021/05/12/885e5754-1e78-495e-a6f8-d7ffd3776245/ottertune.jpg?auto=webp&fit=crop&height=1200&width=1200
- https://images.squarespace-cdn.com/content/v1/596f25c2725e25fb89b3a6f4/1544176431098-WT0ELR52Q0X1IQC89YRH/discussion+cc3.0.png?format=1500w