RECENT DEVELOPMENTS IN SELF-DRIVING DATA MANAGEMENT WITH DEEP REINFORCEMENT LEARNING

TUTORIAL PROPOSAL

Gabriel Campero Durand
University of Magdeburg
Universitaetsplatz 2
39106, Magdeburg
campero@ovgu.de

1 Description

The efficiency of data management tools is predicated on how well their configuration (e.g. physical design, query optimization mechanisms) matches the workload that they process. Database administrators are commonly responsible for manually defining such configuration, but even for the most experienced practitioners finding the optimal remains challenging when considering: a) the high number of configurable knobs and possible configurations, b) rapidly changing workloads, and c) the uncertainty in predicting the impact of choices when based on cost models or assumptions (e.g. knob independence) that might not fully match real-world systems[VAPGZ17].

To alleviate these challenges either fully or partially automated tools are used, supporting the selection of a configuration given a workload (e.g. physical design advisory tools[CN07]). Specially relevant for building such kind of tools is the incorporation of machine learning models, since these models can help the tools to learn from experience, reducing the reliance on assumptions.

In recent years, building on the success of reinforcement learning methods (RL) at outperforming humans in highly complex game scenarios, both academia[SSD18, MP18, KYG+18, DPP+18, OBGK18, TMM+18] and industry[1] have proposed several self-driving data management solutions that learn from real-world signals using RL or deep reinforcement learning (DRL, i.e., the combination of reinforcement learning methods, with neural networks for function approximation[ADBB17, FLHI+18]). In this context, DRL is specially valuable because it enables models to have a limited memory footprint, a competitive inference process and, when properly developed, models can generalize well from past experiences to unknown states. Hence, this approach can have a large impact on how autonomous systems are built.

Nowadays, with the increasing availability of open-source frameworks to support DRL, it is possible for data engineers and developers to rapidly prototype self-driving DRL solutions, and embed them into the lifecycle of their applications.

In this tutorial we introduce the audience to the field by bringing together three perspectives. First, a theoretical perspective, covering in detail the most commonly used DRL models. Second, a practical perspective, introducing off-the-shelf DRL frameworks (such as Google Dopamine[CMG+18][2], Facebook Horizon[GCL+18][3] and Amazon SageMaker[4]) and providing examples for tackling a simple data management use case (i.e., building an index advisor) with these frameworks. Finally, we take an application perspective, reviewing recent developments in state-of-the-art self-driving data management with DRL. To conclude we consider issues relevant to adopting self-driving DRL solutions in production systems, and we highlight open research directions.

[1] https://blogs.oracle.com/oracle-database/oracle-database-19c-now-available-on-oracle-exadata
[2] https://github.com/google/dopamine/
[3] https://github.com/facebookresearch/Horizon
[4] https://aws.amazon.com/sagemaker

2 Structure of the Tutorial

1. Motivation and Background on RL (10 minutes)
2. Introduction to DRL (10 minutes)
3. Main DRL methods (30 minutes)
   (a) Value-based methods
   (b) Policy gradient methods
   (c) Model-based methods
4. Building an index advisor using off-the-shelf DRL frameworks (20 minutes)
5. Recent developments in self-driving data management with DRL (30 minutes)
   (a) A general product development process
   (b) Applications in storage engine management
   (c) Applications in query processing
   (d) Other applications
6. Assessing the production-readiness of self-driving data management DRL solutions (10 minutes)
7. Open directions (10 minutes)

3 Short Bio

Gabriel Campero Durand[5] is a PhD candidate at the University of Magdeburg, Germany, supervised by Prof. Dr. Gunter Saake. Gabriel holds a Master on Data and Knowledge Engineering from the University of Magdeburg, and a Systems Engineering Degree from the Universidad de Los Andes, Venezuela. He's worked with IBM's R&D Germany (2015-2017), developing graph technologies supporting Bluemix use cases. Currently he researches on production-readiness for self-driving data management with deep reinforcement learning.

References

[ADBB17] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017.

[CMG+18] Pablo Samuel Castro, Subhodeep Moitra, Carles Gelada, Saurabh Kumar, and Marc G Bellemare. Dopamine: A research framework for deep reinforcement learning. *arXiv preprint arXiv:1812.06110*, 2018.

[CN07] Surajit Chaudhuri and Vivek Narasayya. Self-tuning database systems: a decade of progress. In *Proceedings of the 33rd international conference on Very large data bases*, pages 3–14. VLDB Endowment, 2007.

[DPP+18] Gabriel Campero Durand, Marcus Pinnecke, Rufat Piriyev, Mahmoud Mohsen, David Broneske, Gunter Saake, Maya S Sekeran, Fabián Rodriguez, and Laxmi Balami. Gridformation: Towards self-driven online data partitioning using reinforcement learning. In *Proceedings of the First International Workshop on Exploiting Artificial Intelligence Techniques for Data Management*, page 1. ACM, 2018.

[FLHI+18] Vincent François-Lavet, Peter Henderson, Riashat Islam, Marc G Bellemare, Joelle Pineau, et al. An introduction to deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 11(3-4):219–354, 2018.

[GCL+18] Jason Gauci, Edoardo Conti, Yitao Liang, Kittipat Virochsiri, Yuchen He, Zachary Kaden, Vivek Narayanan, and Xiaohui Ye. Horizon: Facebook's open source applied reinforcement learning platform. *arXiv preprint arXiv:1811.00260*, 2018.

[KYG+18] Sanjay Krishnan, Zongheng Yang, Ken Goldberg, Joseph Hellerstein, and Ion Stoica. Learning to optimize join queries with deep reinforcement learning. *arXiv preprint arXiv:1808.03196*, 2018.

[MP18] Ryan Marcus and Olga Papaemmanouil. Deep reinforcement learning for join order enumeration. In *Proceedings of the First International Workshop on Exploiting Artificial Intelligence Techniques for Data Management*, page 3. ACM, 2018.

[5] `http://www.dbse.ovgu.de/Mitarbeiter/Gabriel+Campero+Durand.html`

[OBGK18] Jennifer Ortiz, Magdalena Balazinska, Johannes Gehrke, and S Sathiya Keerthi. Learning state representations for query optimization with deep reinforcement learning. *arXiv preprint arXiv:1803.08604*, 2018.

[SSD18] Ankur Sharma, Felix Martin Schuhknecht, and Jens Dittrich. The case for automatic database administration using deep reinforcement learning. *arXiv preprint arXiv:1801.05643*, 2018.

[TMM+18] Immanuel Trummer, Samuel Moseley, Deepak Maram, Saehan Jo, and Joseph Antonakakis. Skinnerdb: regret-bounded query evaluation via reinforcement learning. *Proceedings of the VLDB Endowment*, 11(12):2074–2077, 2018.

[VAPGZ17] Dana Van Aken, Andrew Pavlo, Geoffrey J Gordon, and Bohan Zhang. Automatic database management system tuning through large-scale machine learning. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1009–1024. ACM, 2017.