



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Machine Learning over Static and Dynamic Relational Data

Citation for published version:

Kara, A, Nikolic, M, Olteanu, D & Zhang, H 2021, Machine Learning over Static and Dynamic Relational Data. in *Proceedings of the 15th ACM International Conference on Distributed and Event-based Systems*. DEBS '21, ACM, Virtual Event, Italy, pp. 160--163, The 15th ACM International Conference on Distributed and Event-based Systems, Milan, Italy, 28/06/21. <https://doi.org/10.1145/3465480.3467843>

Digital Object Identifier (DOI):

[10.1145/3465480.3467843](https://doi.org/10.1145/3465480.3467843)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 15th ACM International Conference on Distributed and Event-based Systems

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Machine Learning over Static and Dynamic Relational Data

Ahmet Kara
University of Zurich
Switzerland

Dan Olteanu
University of Zurich
Switzerland

Milos Nikolic
University of Edinburgh
United Kingdom

Haozhe Zhang
University of Zurich
Switzerland

ABSTRACT

This tutorial overviews principles behind recent works on training and maintaining machine learning models over relational data, with an emphasis on the exploitation of the relational data structure to improve the runtime performance of the learning task.

The tutorial has the following parts:

- (1) Database research for data science
- (2) Three main ideas to achieve performance improvements
 - (2.1) Turn the ML problem into a DB problem
 - (2.2) Exploit structure of the data and problem
 - (2.3) Exploit engineering tools of a DB researcher
- (3) Avenues for future research

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; *Learning linear models*; • **Information systems** → **Data management systems**; **Database query processing**; *Query optimization*; *Database views*; *Stream management*.

KEYWORDS

machine learning models, incremental maintenance, in-database machine learning

ACM Reference Format:

Ahmet Kara, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. 2021. Machine Learning over Static and Dynamic Relational Data. In *The 15th ACM International Conference on Distributed and Event-based Systems (DEBS '21)*, June 28–July 2, 2021, Virtual Event, Italy. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3465480.3467843>

The starting points of this tutorial are a VLDB 2020 keynote [34] and a short tutorial at SUM 2019 [44]. This tutorial goes beyond these prior tutorials, in particular on learning over dynamic relational data. The interested reader may also want to explore further work on machine learning over databases overviewed in two prior tutorials [23, 39].

In the following, we briefly overview each part of the tutorial.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DEBS '21, June 28–July 2, 2021, Virtual Event, Italy

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8555-8/21/06...\$15.00
<https://doi.org/10.1145/3465480.3467843>

1 DATABASE RESEARCH FOR DATA SCIENCE

Database research is thriving in the data science era. Relational data remains ubiquitous. According to a recent Kaggle survey [16], most data scientists use relational data. The widespread use of relational data maintains the relevance of existing relational processing techniques. Furthermore, the new requirements brought by the machine learning workloads have led to new relational processing techniques. In this tutorial, we overview some of these existing and new techniques. They rely to a varying degree on an integration of relational database engines and machine learning libraries.

This first part of the tutorial overviews two main approaches at the interface of data processing engines and machine learning libraries. The main message is that this interface provides a fruitful and exciting opportunity for database research to shine. A tighter integration of the database and machine learning computation uncovers new research challenges and can lead to significant performance improvements.

A typical approach to learning over relational data involves the construction of the training dataset using a feature extraction query that joins the input relations and constructs new features using aggregates over the data columns. This query can be expressed in SQL and executed using a database system, e.g., PostgreSQL or SparkSQL [50], or Python Pandas [26] for Jupyter notebooks. The desired model is then learned using a machine learning library, e.g., scikit-learn [38], R [40], TensorFlow [1], or MLlib [27]. This approach ignores the structure of the underlying relational data at the expense of runtime performance. It is dubbed *structure-agnostic* in this tutorial. It puts together two black-box specialised systems for data processing and machine learning and may work for virtually any dataset and model. There are several significant downsides of this approach, including: the materialisation of the result of the feature extraction query; the export of this result from the data processing system to the machine learning library; high maintenance cost in case of changes to the underlying data; the limitations of each of the two systems become a limitation of their combination. These downsides hinder the runtime performance of a data science solution using this approach.

The tutorial will also present an alternative approach, dubbed *structure-aware* learning. It exploits the data sparsity and the structure of the relational data, in particular the various dependencies in the data and the result of the feature extraction query (multi-valued, functional). In this approach, the learning algorithm is opened up and rewritten such that its data-intensive components are moved closer to the data inside the relational query processor. Such components can be expressed as (group-by) aggregates over the feature extraction query. Part 2.1 of the tutorial is dedicated to them. Their

output size is much smaller than that of the feature extraction query and can be computed asymptotically faster than the feature extraction query itself. Since this approach avoids the materialisation of the training dataset and its move between the two systems, it enjoys excellent runtime performance. Opening up the learning black box also allows to consider known mechanisms to maintain the data-intensive components under data updates [33].

The tutorial will contrast the aforementioned two approaches and present experimental evidence for the superior runtime performance of the structure-aware over structure-agnostic [42]. This runtime performance can in fact be translated into accuracy performance: Within the time budget of training one model with the structure-agnostic approach, the structure-aware approach may train many possible models and eventually choose one with the best accuracy. This part will conclude with a brief account to several instantiations of these two approaches in the literature.

2 STRUCTURE-AWARE LEARNING

Part 1 of the tutorial sets up the scene for Part 2, which overviews principles behind the structure-aware approach for learning over relational data. We divide this second part of the tutorial into three blocks. The first block looks at how to turn the learning problem into a database problem.

The second block overviews techniques for computing and maintaining batches of aggregates that arise from mapping the learning problem to a database problem. It considers a variety of types of structure that can be exploited for improving the runtime performance of such techniques, including algebraic, combinatorial, statistical, and geometric structure. Techniques presented in this block aim at lowering the asymptotic computational complexity.

Finally, the third block overviews engineering tools that proved effective for improving runtime performance, such as parallelisation, specialisation to workload and dataset, and low-level sharing of data access between different components of the code. Such tools can effectively lower the constant factors of the computation time.

2.1 Turn the ML Problem into a DB Problem

The tutorial will explain how to express the data-intensive computation of the learning task using various forms of aggregation over the data matrix that are readily supported by database query languages. A particular focus will be on the learning task for ridge linear regression, support vector machines, and tree-structured Bayesian networks (Chow-Liu trees) – as representatives for wider classes of models and objective functions for the learning task. An in-depth treatment for specific models and objective functions is provided in the literature, e.g., [2–5, 20, 24, 42, 43].

For learning using the least-squares loss function, the gradient vector of this function is built up using sum-product aggregates over the model features and parameters: For each pair of features (database attributes) x_i and x_j , there is one such aggregate $\text{sum}(x_i * x_j)$ that sums over all tuples in the training dataset the product of the values of the two attributes. In case an attribute corresponds to a categorical feature, then it is promoted from the sum to the group-by clause. If both attributes are in the group-by clause, then the sum becomes the count of the number of occurrences of each pair of categories of the two attributes: $\text{sum}(1 \text{ group by } x_i, x_j)$.

This approach applies to ridge linear regression [4, 43] and polynomial regression models in general [35], factorisation machines [3–5], sum-product networks [10], principal component analysis [4], quadratically regularised low-rank models [15], and QR and SVD decompositions [49].

For decision trees, the computation of the cost functions for each attribute and condition at a decision tree node can be expressed by a sum-product aggregate with a filter condition. The cost functions used by algorithms such as CART [9] for constructing regression trees rely on aggregates that compute the variance of the response y conditioned on a filter restricting the value of an attribute x_i to be equal, less than, or greater than a given constant (threshold c_j): $\text{VARIANCE}(y) \text{ WHERE } x_i \text{ op } c_j$. For a categorical attribute, the filter condition expresses its membership in a set of possible categories. The thresholds and categories are decided in advance based on the distribution of values for x_i . The variance aggregate is expressed using the sum of squares, the square of sum, and the count. For classification trees, the aggregates encode the entropy or the Gini index using group-by counts to compute value frequencies in the data matrix.

A large class of models, including support vector machines, are trained using sub-gradient descent. They use non-polynomial loss functions, such as (ordinal) hinge, Huber, scalene, and epsilon insensitive, that are defined by multiple cases conditioned on additive inequalities of the form $\sum_i x_i \cdot w_i > c$, where w_i and c are constants and x_i are the features. The efficient computation of aggregates conditioned on additive inequalities calls for new algorithms beyond the classical ones for theta joins [2, 20]. Similar aggregates are derived for k -means clustering [2].

2.2 Exploit Structure of Data and Problem

This part of the tutorial overviews principles behind new efficient algorithms for batches of group-by aggregates, worst-case optimal equality joins, and additive inequality joins. Such algorithms power the structure-aware learning paradigm by systematically exploiting the structure of the relational data to lower the computational complexity and improve the runtime performance. The tutorial discusses the algebraic, combinatorial, statistical, and geometric structure of relational data, with a focus on the algebraic structure.

A. Algebraic structure. A relation is a sum-product expression, where the sum is the set union and the product is the Cartesian product. The computation expressed using relational algebra can be captured using (semi)rings. There is extensive work in the literature on k -relations over provenance semirings [13], generalised multiset relations [21], and factors over multiple semirings [6]. The tutorial will overview particular properties that make the rings effective for computing and maintaining aggregates, as required by structure-aware learning.

A.1. Distributivity of product over sum. This law allows to factor out data blocks common to several tuples in a relation, represent them once, and compute over them once. It is the main conceptual ingredient of factorised databases [35, 36]. By systematically applying the distributivity law, relations can be represented more succinctly yet losslessly as directed acyclic graphs with fewer data value repetitions. The factorisations of relations representing the answers to relational queries can be asymptotically smaller than the

standard representation as list of tuples. This applies to a rich class of queries made up of joins [8, 36], selections, projections, unions, group-by aggregates [7], and order-by clauses [7]. More relevant to our tutorial, factorisation can also lower the computational complexity of machine learning over feature extraction queries [2, 4, 35, 43]. The tutorial will exemplify factorised computation for joins and aggregates such as those discussed in Part 2.1. The framework of Functional Aggregate Queries [6] generalises factorised databases to semirings beyond sum-product and shows that many problems across Computer Science can benefit from factorised computation. LMFAO [41, 42], F-IVM [32, 33], and IVM^c [18, 19] employ factorised query computation and maintenance. These algorithms factorise the query into a hierarchy of increasingly simpler views, which are maintained bottom-up under data updates.

A.2. Sum-product abstraction. By conveniently overloading the sum and product operations in a (semi)ring, we can capture the computation for many different tasks over relational data, including the type of aggregates needed for training models. This is exemplified extensively in the FAQ framework [6], where the same structure of computation but with possibly different definitions of sum and product can be used to compute, among others, database queries, matrix chains, and inference queries (marginals and MAP) in probabilistic graphical models. The tutorial will exemplify this for covariance matrices [32] and mutual information [33] over feature extraction queries, used in the context of training models over relational data.

A.3. Additive inverse. To accommodate efficient mechanisms for incremental maintenance of learned models over feature extraction queries, we extend tuples to carry along payloads, which are elements from a ring. These payloads can for instance capture the multiplicities of the corresponding tuples or, in case of tuples in the query result, the number of derivations of that tuple from the input tuples via the query. Similarly, the aggregates needed for training a specific model may represent payloads. Tuple inserts and deletes to the underlying data can be modelled uniformly as inserts with appropriate payloads a and respectively $-a$ that follow the additive inverse law of the ring: $a + (-a) = 0$, where 0 is the neutral element for summation in the ring. Whenever a tuple's payload becomes 0 , then it is not anymore part of the data. In the simplest case, a is an integer representing the multiplicity [17, 21, 22]. Further examples discussed in the literature consider rings for factorised data representation and covariance matrices [32, 33]. The efficient maintenance of covariance matrices makes it possible to keep ridge linear regression models fresh under high-throughput data changes [32]. A recent tutorial overviews advances in incremental view maintenance [12]. Our tutorial will go through several examples showing how rings can be used to maintain models under data updates.

B. Combinatorial structure. The combinatorial structure prevalent in relational data is captured by notions such as the width measure of the query and the degree of a data value. For reasons of time limitation, the tutorial will only mention this type of structure. A brief overview is given in a recent keynote [34]. If a feature extraction query has width w , then its data complexity is $\tilde{O}(N^w)$ for a database of size N , where \tilde{O} hides logarithmic factors in N . Similar complexities have been shown for learning a variety of models over feature extraction queries [2, 35, 43] using prototypes such as F [43] and LMFAO [42]. There are several width measures

proposed in the literature yet they are beyond the scope of our tutorial. The degree information captures the number of occurrences of a data value in the input database [31]. Several existing query processing and maintenance algorithms, e.g., worst-case optimal join algorithms [30] and worst-case optimal incremental maintenance of triangle [17, 18] and hierarchical [19] queries, adapt their execution strategy depending on the degree of data values, with different strategies for high-degree and low-degree values. A special form of bounded degree is given by functional dependencies. They can be used to lower the learning runtime for ridge polynomial regression models and factorisation machines. Instead of learning a given model, we can instead learn a reparameterised model with fewer parameters and then map it back to the original model [4, 5].

C. Statistical structure. In case of very large datasets, a feasible approach is to learn approximately over data samples [28]. When learning over feature extraction queries, we would like to sample from the input data through the queries. Prior work considered the problem of sampling through selection conditions and joins, e.g., the ripple joins [14] and the wander joins [25], and for specific classes of machine learning models [37].

D. Geometric structure. This type of structure becomes relevant whenever we use distance measures. Clustering algorithms can exploit such measures, e.g., the optimal transport distance between two probability measures, and the distance-based triangle inequality. The relational k -means (Rk-means) [11] is a prime example of structure-aware learning approaches that exploits the geometric structure of the underlying data. It achieves a constant-factor approximation of the k -means objective by clustering over a small coresot instead of the full result of the feature extraction query.

2.3 Engineering Tools of DB Researcher

Towards taming the computational challenge raised by structure-aware learning, a large effort focuses on a toolbox of systems techniques such as: specialisation for workload, data, and hardware; observing the memory hierarchy and blocking operations; distribution and parallelisation. The tutorial will highlight several recent efforts in this space, in particular on compiling the task of learning specific models over feature extraction queries into efficiently executable low-level code. Such techniques can lead to significant runtime performance improvements, as reported for the AC/DC [47], F-IVM [32, 33], LMFAO [41, 42], and IFAQ [47, 48] prototypes. The tutorial will make the case for such a compilation approach. These systems are based on prior work on specialisation for queries and database schema [29, 45, 46]. The LMFAO system also systematically shares computation across the batch of aggregates for structure-aware learning and takes advantage of multi-core CPU architectures for domain and task parallelism.

3 FUTURE RESEARCH

The tutorial will conclude with reflections on the state of machine learning over relational data and will pinpoint several directions of future research in systems and theory for structure-aware learning. In particular, it will consider questions on the limits of structure-aware learning and how to make it readily useful for practical data science projects. It will also give a glimpse of ongoing work of the authors on maintaining machine learning models under updates.

Acknowledgements. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 682588.

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, and et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*. 265–283.
- [2] Mahmoud Abo Khamis, Ryan R. Curtin, Benjamin Moseley, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2019. On Functional Aggregate Queries with Additive Inequalities. In *PODS*. 414–431.
- [3] Mahmoud Abo Khamis, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2018. AC/DC: In-Database Learning Thunderstruck. In *DEEM*. 8:1–8:10.
- [4] Mahmoud Abo Khamis, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2018. In-Database Learning with Sparse Tensors. In *PODS*. 325–340.
- [5] Mahmoud Abo Khamis, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2020. Learning Models over Relational Data Using Sparse Tensors and Functional Dependencies. *ACM Trans. Database Syst.* 45, 2, Article 7 (2020), 66 pages.
- [6] Mahmoud Abo Khamis, Hung Q. Ngo, and Atri Rudra. 2016. FAQ: Questions Asked Frequently. In *PODS*. 13–28.
- [7] Nurzhan Bakibayev, Tomáš Kociský, Dan Olteanu, and Jakub Závodný. 2013. Aggregation and Ordering in Factorised Databases. *PVLDB* 6, 14 (2013), 1990–2001.
- [8] Nurzhan Bakibayev, Dan Olteanu, and Jakub Závodný. 2012. FDB: A Query Engine for Factorised Relational Databases. *PVLDB* 5, 11 (2012), 1232–1243.
- [9] L. Breiman, J. Friedman, R. Olshen, and C. Stone. 1984. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA.
- [10] George Chichirim. 2020. Structured Sum-Product Networks. BSc thesis, University of Oxford.
- [11] Ryan R. Curtin, Benjamin Moseley, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2020. Rk-means: Fast Clustering for Relational Data. In *AISTATS (Proc. ML Res., Vol. 108)*. 2742–2752.
- [12] Iman Elghandour, Ahmet Kara, Dan Olteanu, and Stijn Vansummeren. 2018. Incremental Techniques for Large-Scale Dynamic Query Processing. In *CIKM*. 2297–2298. Tutorial.
- [13] Todd J. Green and Val Tannen. 2017. The Semiring Framework for Database Provenance. In *PODS*. 93–99.
- [14] Peter J. Haas and Joseph M. Hellerstein. 1999. Ripple Joins for Online Aggregation. In *SIGMOD*. 287–298.
- [15] Gabriel-Robert Inelus. 2019. Quadratically Regularised Principal Component Analysis over Multi-Relational Databases. MSc thesis, University of Oxford.
- [16] Kaggle. 2017. The State of Data Science and Machine Learning. <https://www.kaggle.com/surveys/2017> <https://www.kaggle.com/surveys/2017>.
- [17] Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. 2019. Counting Triangles under Updates in Worst-Case Optimal Time. In *ICDT*. 4:1–4:18.
- [18] Ahmet Kara, Hung Q. Ngo, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. 2020. Maintaining Triangle Queries under Updates. *ACM Trans. Database Syst.* (2020).
- [19] Ahmet Kara, Milos Nikolic, Dan Olteanu, and Haozhe Zhang. 2020. Trade-offs in Static and Dynamic Evaluation of Hierarchical Queries. In *PODS*. 375–392.
- [20] Mahmoud Abo Khamis, Ryan R. Curtin, Benjamin Moseley, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2020. Functional Aggregate Queries with Additive Inequalities. *ACM Trans. Database Syst.* 45, 4 (2020), 17:1–17:41.
- [21] Christoph Koch. 2010. Incremental query evaluation in a ring of databases. In *PODS*. 87–98.
- [22] Christoph Koch, Yanif Ahmad, Oliver Kennedy, Milos Nikolic, Andres Nötzli, Daniel Lupei, and Amir Shaikhha. 2014. DBToaster: higher-order delta processing for dynamic, frequently fresh views. *VLDB J.* 23, 2 (2014), 253–278.
- [23] Arun Kumar, Matthias Boehm, and Jun Yang. 2017. Data Management in Machine Learning: Challenges, Techniques, and Systems. In *SIGMOD*. 1717–1722.
- [24] Arun Kumar, Jeffrey F. Naughton, and Jignesh M. Patel. 2015. Learning Generalized Linear Models Over Normalized Data. In *SIGMOD*. 1969–1984.
- [25] Feifei Li, Bin Wu, Ke Yi, and Zhuoyue Zhao. 2019. Wander Join and XDB: Online Aggregation via Random Walks. *ACM Trans. Database Syst.* 44, 1 (2019), 2:1–2:41.
- [26] Wes McKinney. 2011. Pandas: a foundational Python library for data analysis and statistics. *Python for High Performance and Scientific Computing* 14 (2011).
- [27] Xiangrui Meng, Joseph Bradley, et al. 2016. MLlib: Machine Learning in Apache Spark. *J. Mach. Learn. Res.* 17, 1 (2016), 1235–1241.
- [28] Kevin P. Murphy. 2013. *Machine learning: a probabilistic perspective*. MIT Press, Cambridge, Mass.
- [29] Thomas Neumann. 2011. Efficiently Compiling Efficient Query Plans for Modern Hardware. *PVLDB* 4, 9 (2011), 539–550.
- [30] Hung Q. Ngo, Ely Porat, Christopher Ré, and Atri Rudra. 2012. Worst-case Optimal Join Algorithms. In *PODS*. 37–48.
- [31] Hung Q. Ngo, Christopher Ré, and Atri Rudra. 2013. Skew Strikes Back: New Developments in the Theory of Join Algorithms. In *SIGMOD Rec.* 5–16.
- [32] Milos Nikolic and Dan Olteanu. 2018. Incremental View Maintenance with Triple Lock Factorization Benefits. In *SIGMOD*. 365–380.
- [33] Milos Nikolic, Haozhe Zhang, Ahmet Kara, and Dan Olteanu. 2020. F-IVM: Learning over Fast-Evolving Relational Data. In *SIGMOD*. 2773–2776.
- [34] Dan Olteanu. 2020. The Relational Data Borg is Learning. *Proc. VLDB Endow.* 13, 12 (2020), 3502–3515.
- [35] Dan Olteanu and Maximilian Schleich. 2016. Factorized Databases. *SIGMOD Rec.* 45, 2 (2016), 5–16.
- [36] Dan Olteanu and Jakub Závodný. 2015. Size Bounds for Factorised Representations of Query Results. *ACM Trans. Database Syst.* 40, 1 (2015), 2.
- [37] Yongjoo Park, Jingyi Qing, Xiaoyang Shen, and Barzan Mozafari. 2019. BlinkML: Efficient Maximum Likelihood Estimation with Probabilistic Guarantees. In *SIGMOD*. 1135–1152.
- [38] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, and et al. 2011. Scikit-learn: Machine Learning in Python. *J. Machine Learning Research* 12 (2011), 2825–2830.
- [39] Neoklis Polyzotis, Sudip Roy, Steven Euijong Whang, and Martin Zinkevich. 2017. Data Management Challenges in Production Machine Learning. In *SIGMOD*. 1723–1726.
- [40] R Core Team. 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Stat. Comp., www.r-project.org.
- [41] Maximilian Schleich and Dan Olteanu. 2020. LMFAO: An engine for batches of group-by aggregates. *PVLDB* 13, 12 (2020), 2945–2948.
- [42] Maximilian Schleich, Dan Olteanu, Mahmoud Abo Khamis, Hung Q. Ngo, and XuanLong Nguyen. 2019. A Layered Aggregate Engine for Analytics Workloads. In *SIGMOD*. 1642–1659.
- [43] Maximilian Schleich, Dan Olteanu, and Radu Ciucanu. 2016. Learning Linear Regression Models over Factorized Joins. In *SIGMOD*. 3–18.
- [44] Maximilian Schleich, Dan Olteanu, Mahmoud Abo Khamis, Hung Q. Ngo, and XuanLong Nguyen. 2019. Learning Models over Relational Data: A Brief Tutorial. In *SUM*. 423–432.
- [45] Amir Shaikhha, Yannis Klonatos, and Christoph Koch. 2018. Building Efficient Query Engines in a High-Level Language. *ACM Trans. Database Syst.* 43, 1, Article 4 (2018), 4:1–4:45 pages.
- [46] Amir Shaikhha, Yannis Klonatos, Lionel Parreaux, Lewis Brown, Mohammad Dashti, and Christoph Koch. 2016. How to Architect a Query Compiler. In *SIGMOD*. 1907–1922.
- [47] Amir Shaikhha, Maximilian Schleich, Alexandru Ghita, and Dan Olteanu. 2020. Multi-layer optimizations for end-to-end data analytics. In *CGO*. 145–157.
- [48] Amir Shaikhha, Maximilian Schleich, and Dan Olteanu. 2021. An Intermediate Representation for Hybrid Database and Machine Learning Workloads. *PVLDB* to appear (2021).
- [49] Bas van Geffen. 2018. QR Decomposition of Normalised Relational Data. MSc thesis, University of Oxford.
- [50] Matei Zaharia, Mosharaf Chowdhury, et al. 2012. Resilient Distributed Datasets: A Fault-tolerant Abstraction for In-memory Cluster Computing. In *NSDI*. 2–2.