

# The Relational Data Borg is Learning

Dan Olteanu

Department of Informatics, University of Zurich

olteanu@ifi.uzh.ch

<https://fdbresearch.github.io> <https://www.relational.ai>

<https://www.ifi.uzh.ch/en/dast.html>

## ABSTRACT

This paper overviews an approach that addresses machine learning over relational data as a database problem. This is justified by two observations. First, the input to the learning task is commonly the result of a feature extraction query over the relational data. Second, the learning task requires the computation of group-by aggregates. This approach has been already investigated for a number of supervised and unsupervised learning tasks, including: ridge linear regression, factorisation machines, support vector machines, decision trees, principal component analysis, and k-means; and also for linear algebra over data matrices.

The main message of this work is that the runtime performance of machine learning can be dramatically boosted by a toolbox of techniques that exploit the knowledge of the underlying data. This includes theoretical development on the algebraic, combinatorial, and statistical structure of relational data processing and systems development on code specialisation, low-level computation sharing, and parallelisation. These techniques aim at lowering both the complexity and the constant factors of the learning time.

This work is the outcome of extensive collaboration of the author with colleagues from RelationalAI, in particular Mahmoud Abo Khamis, Molham Aref, Hung Ngo, and XuanLong Nguyen, and from the FDB research project, in particular Ahmet Kara, Milos Nikolic, Maximilian Schleich, Amir Shaikhha, Jakub Závodný, and Haozhe Zhang. The author would also like to thank the members of the FDB project for the figures and examples used in this paper.

The author is grateful for support from industry: Amazon Web Services, Google, Infor, LogicBlox, Microsoft Azure, RelationalAI; and from the funding agencies EPSRC and ERC. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 682588.

### PVLDB Reference Format:

Dan Olteanu. The Relational Data Borg is Learning. *PVLDB*, 13(12): 3502-3515, 2020.

DOI: <https://doi.org/10.14778/3415478.3415572>

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

*Proceedings of the VLDB Endowment*, Vol. 13, No. 12

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3415478.3415572>

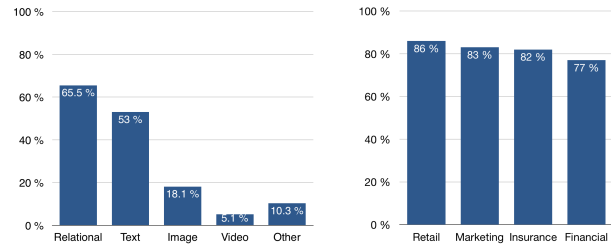


Figure 1: Kaggle survey of 16,000 practitioners on the state of data science and machine learning: Use of relational data: overall (left), by industry (right).

## 1. WHY BOTHER?

As we witness the data science revolution, each research community legitimately reflects on its relevance and place in this new landscape. The database research community has at least three reasons to feel empowered by this revolution. This has to do with the pervasiveness of relational data in data science, the widespread need for efficient data processing, and the new processing challenges posed by data science workloads beyond the classical database workloads. The first two aforementioned reasons are widely acknowledged as core to the community's *raison d'être*. The third reason explains the longevity of relational database management systems success: Whenever a new promising data-centric technology surfaces, research is under way to show that it can be captured naturally by variations or extensions of the existing relational techniques. Prime examples are the management of XML data [30], probabilistic data [71], graph data [11], code [13], and ontologies [75], as well as prior work on migrating data mining inside the database [18].

*Like the Star Trek's Borg Collective co-opting technology and knowledge of alien species, the Relational Data Borg a.k.a. RDBMS assimilates ideas and applications from connex fields to adapt to new requirements and become ever more powerful and versatile. Unlike the former, the latter moves fast, has great skin complexion, and is reasonably happy. Resistance is futile in either case.*

This paper overviews work by the author on learning over relational data. Along with a growing number of recent contributions, e.g., a database theory manifest [3] and SIGMOD tutorials overviewing work at the interface of machine learning and databases [61, 40], this effort is another episode of the relentless learning process for the Relational Data Borg.

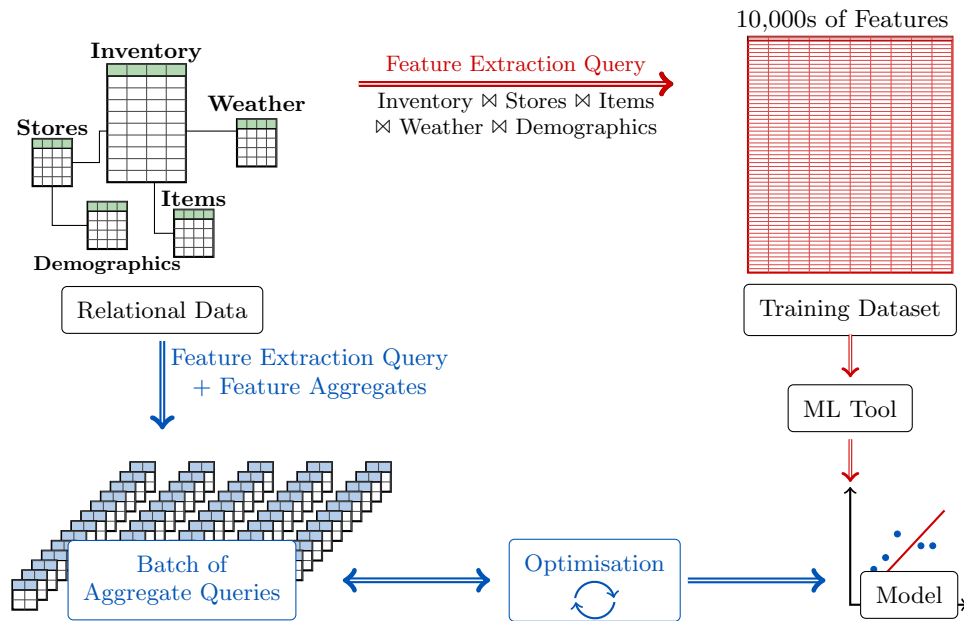


Figure 2: Learning over relational data. In **structure-agnostic learning** (the top flow in red), a feature extraction query constructs the data matrix on which the model is trained using a machine learning library. In **structure-aware learning** (the bottom flow in blue), sufficient statistics is computed over the input database using a batch of aggregate queries and then optimisation is performed on this statistics to obtain the model.

## 1.1 Relational data is ubiquitous

*According to a recent Kaggle survey [34], most data scientists use relational data (Figure 1).*

This is a testament to the usability of the relational data model. Many human hours are invested in building relational databases that are curated and enriched with knowledge of the underlying domains and integrity constraints.

Figure 2 depicts a dataset used for data science in the retailer domain. Features are gathered from a number of relations detailing information on items in stores, store information, demographics for areas around the stores, inventory units for items in stores on particular dates, and weather. The data matrix, which is the input to the learning task, is the result of a *feature extraction query* that joins these relations on keys for dates, locations, zipcode, and items. The query may construct further features using aggregates, e.g., running aggregates over days, weeks, months; min, max, average, median aggregates, or aggregates over many-to-many relationships and categorical attributes with high cardinality (e.g., identifiers) [60]. A typical model of interest would predict sales or inventory units for next month.

The author’s observation based on interactions with data scientists at LogicBlox and RelationalAI is that similar datasets of up to a few hundreds of millions of rows across a dozen relations are common in data science projects.

## 1.2 Ignorance is not bliss

*The mainstream data science solutions ignore the structure of the underlying relational data at the expense of runtime performance.*

The structure of relational data is rich. A large body of work in database research is on discovering, enforcing,

and exploiting structure such as dependencies and integrity constraints. It is therefore rather perplexing that the current state of affairs in learning over relational data is to ignore its structure! The flow of such **structure-agnostic** mainstream solutions is depicted in Figure 2: They first construct the data matrix using a feature extraction query expressed in SQL for a database system, e.g., PostgreSQL or SparkSQL [76], or in Python Pandas [45] for jupyter notebooks. The data matrix is then passed on to a machine learning library, e.g., scikit-learn [59], R [62], TensorFlow [1], or MLlib [47], which then learns the model.

The structure-agnostic solutions that put together black-box specialised systems for data processing and machine learning may work for virtually any dataset and model. The uncomfortable drawback of this marriage of convenience is that the two systems were not originally designed to work together. They may therefore suffer from some of the following shortcomings: (1) the feature extraction query is fully materialised; (2) its result is exported from the query engine and imported into the statistical library; (3) the categorical variables are one-hot encoded; (4) the pipeline may suffer from impedance mismatch and high maintenance cost in the face of changes to the underlying dataset; (5) and the pipeline inherits the limitations of both systems.

In practice, all of these shortcomings may significantly hinder the runtime performance of the data science solution.

**1. Data matrix materialisation.** The materialisation of the feature extraction query may take non-trivial time as it typically produces data matrices whose sizes are orders of magnitude larger than the input data. This departs from the classical database setting, where queries are over-constrained by highly selective joins and filter conditions. Here, the goal is to bring together the available data and use it to train accurate models. Figure 3 shows that for our

| Relation  | Cardinality/Attrs | CSV Size | PostgreSQL+TensorFlow |                | LMFAO [65] |           |        |
|-----------|-------------------|----------|-----------------------|----------------|------------|-----------|--------|
|           |                   |          | Time                  | CSV Size       | Time       | CSV Size  |        |
| Inventory | 84,055,817 / 4    | 2 GB     | Database              | –              | 2.1 GB     | –         | 2.1 GB |
| Items     | 5,618 / 5         | 129 KB   | Join                  | 152.06 secs    | 23 GB      | –         | –      |
| Stores    | 1,317 / 15        | 139 KB   | Export                | 351.76 secs    | 23 GB      | –         | –      |
| Demogr.   | 1,302 / 16        | 161 KB   | Shuffling             | 5,488.73 secs  | 23 GB      | –         | –      |
| Weather   | 1,159,457 / 8     | 33 MB    | Query batch           | –              | –          | 6.08 secs | 37 KB  |
| Join      | 84,055,817 / 44   | 23GB     | Grad Descent          | 7,249.58 secs  | –          | 0.05 secs | –      |
|           |                   |          | Total time            | 13,242.13 secs |            | 6.13 secs |        |

Figure 3: Left: characteristics of our retailer dataset. Right: Runtime performance experiment (Intel i7-4770, 3.4GHz, 32GB, 8 cores). The **structure-aware (LMFAO)** approach is 2,160x faster than the **structure-agnostic (PostgreSQL+TensorFlow)** approach, while producing a more accurate solution (RMSE on 2% test data). TensorFlow runs one epoch with 100K tuple batch. Both approaches train a linear regression model to predict the inventory given all the other features from the data matrix defined by the join of the relations.

retailer dataset, the result of the key-fkey join is one order of magnitude larger than the input data<sup>1</sup>.

**2. Data move.** Moving the data matrix from the query engine to the learning library is without doubt a time-consuming step. Furthermore, the data formats of the two systems may be different and require a conversion. In our experiment reported in Figure 3, moving the data matrix takes more than twice the time to compute it.

**3. One-hot encoding.** The data is standardised and one-hot encoded before the learning step. Done naïvely, the one-hot encoding may increase the size of the data matrix, turning it from lean (tall-and-thin) into chubby and thereby blurring the typical database distinction between the large number of rows and the small number of columns.

**4. Maintenance cost.** The structure-agnostic data science pipeline is not designed to cope effectively with on-the-fly changes to the underlying data. Data updates in the form of tuple insertions and deletions, as well as schema updates in the form of addition and removal of relations, columns, derived aggregate features, would require recomputation of the data matrix as well as the model learned over it.

**5. Inherited limitations.** The two systems have been designed for different use cases that are not necessarily aligned. For instance, the maximum data frame size in R and the maximum number of columns in PostgreSQL are much less than typical database sizes and respectively number of model features. The mainstream solutions thus inherit the limitations of both constituent systems.

All aforementioned shortcomings pose significant runtime performance challenges. Whereas query processing engines are designed to cope with large datasets, this is not the case for the bulk of popular learning libraries. By having the learning libraries work on orders of magnitude larger data than the query engine, the mainstream solutions exacerbate the runtime performance challenge. Even for medium-sized input datasets such as our retailer dataset, working solutions are forced to employ distributed platforms to train models. This accidental complexity comes with a high, rather unnecessary compute cost, as observed repeatedly in recent studies on the cost of scalability [46], and also non-trivial energy cost [21]. Distributed platforms should be reserved for truly very large problem instances to justify the cost.

<sup>1</sup>A notable exception is a factorised query engine [15, 14, 65]. For our retailer dataset, the factorised/non-factorised join is 26x smaller/10x larger than the input data size.

### 1.3 Structure awareness to the rescue!

*Learning over the data matrix defined by a feature extraction query may take much less time than materialising the data matrix and is feasible on a commodity machine.*

A natural question is whether the shortcomings highlighted in Section 1.2 can be avoided. There is increasing effort to address these shortcomings by migrating from the mainstream solution to an ever tighter integration of the querying and the learning tasks into a single execution plan that avoids moving the data matrix. Opening up the learning task and expressing its data-intensive computation steps as a *batch of aggregate queries* is another source of radical performance improvements. Computing the batch of queries over the feature extraction queries is now a purely database workload. We can exploit the join structure in the feature extraction query to optimise this database workload. In fact, an entire toolbox of database theory and systems techniques is at our disposal to tackle the performance challenges of this workload, as discussed in the next sections. This forms the basis of the **structure-aware learning** paradigm, as opposed to the **structure-agnostic** paradigm discussed before.

Figure 2 depicts the flow of a structure-aware learning solution: By analysing the learning task, i.e., the features of the model and the objective function used for learning the model, a batch of aggregates is synthesised manually [66] or automatically [70]. This batch is then composed with the feature extraction query, optimised, and evaluated. Its result is typically much smaller than that of the feature extraction query and its evaluation may be much faster than that of the feature extraction query alone! This result is sufficient to compute the model parameters using an optimisation algorithm tailored at the desired model such as gradient descent or CART. Section 2 exemplifies aggregates for several models and objective functions.

Our structure-aware learning system LMFAO (Layered Multiple Functional Aggregates Optimisation) [65] avoids the materialisation of the join and computes the batch of aggregates directly over the input database. As shown in Figure 3, LMFAO only takes 6 seconds to compute the batch. It then takes 50 milliseconds to compute the model parameters of a ridge linear regression model using gradient descent, where the gradient vector is built up using the computed aggregates and the current parameters.

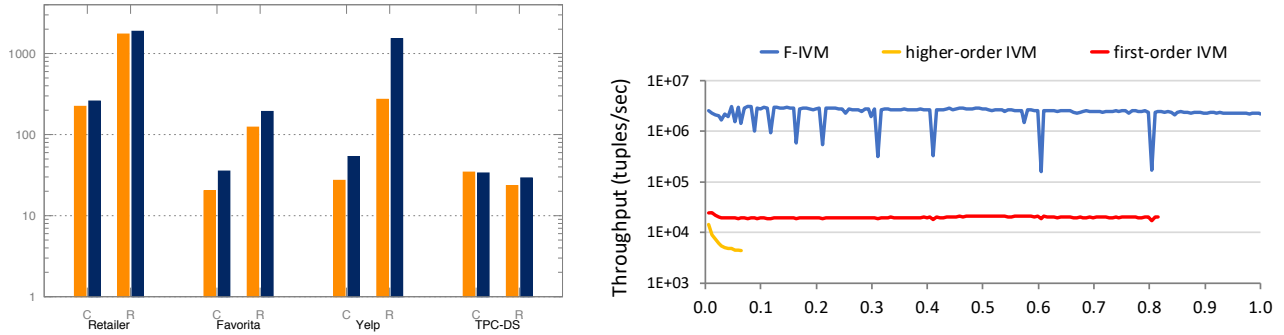


Figure 4: Left: Speedup of LMFAO over existing DBMSs (commercial **DBX** and **MonetDB**). The task is to compute query batches **C** (covariance matrix) and **R** (regression tree node) over four datasets (AWS d2.xlarge, 4 vCPUs, 32GB). Right: Throughput of F-IVM over first-order IVM (delta processing) and higher-order IVM (delta processing with intermediate views). The task is to maintain the covariance matrix in the retailer dataset (Azure DS14, Intel Xeon, 2.40 GHz, 112GB, 1 thread, one hour timeout).

In contrast, the structure-agnostic solution takes significantly more time. It needs 152 seconds only to compute the join using PostgreSQL. While commercial database systems may be faster than PostgreSQL, they would still need non-trivial time to create the 23GB join result, whereas the sufficient statistics computed by LMFAO is only 37KB. The structure-agnostic solution exports the data matrix from PostgreSQL, imports it into TensorFlow, and shuffles it. Finally, the model is learned using a variant of stochastic gradient descent in one epoch (i.e., one pass over the data matrix). Both times to shuffle and learn using TensorFlow dwarf the join time by a factor of 50x. Overall, the structure-aware solution is 2,160x faster than the structure-agnostic solution in our experiment. Both systems were given the same set of input features and the resulting model was validated against a training dataset. LMFAO computed a model very close to that given by the closed-form ordinary least squares solution, whereas TensorFlow returned a slightly less accurate model trained in one epoch (pass) over the data matrix.

This lack of efficiency can be observed beyond the choice of model and dataset in our experiment. Prior work [66, 5, 65] reports on similar stories with a range of public datasets (Favorita, TPC-DS, Yelp), models (decision trees, factorisation machines, k-means), and learning libraries (R [62], scikit-learn [59], Python StatsModels [72], mlpack [22], and XGBoost [19]). They often fail to process the data matrix on commodity machines due to out-of-memory errors.

Over the past years, we have investigated the gap between the two types of solutions from different angles and for different models. We have built a number of prototypes for learning a variety of models over relational data: F [66, 55], F-IVM<sup>2</sup> [53, 54], AC/DC [5], LMFAO [65, 64], IFAQ [70], and Rk-means [23]. We have also investigated theoretical aspects including: the use of sparse tensors and functional dependencies to improve the time complexity of learning polynomial regression models, factorisation machines, and PCA over relational data [6, 7]; algorithms and complexity analysis for queries with group-by aggregates and joins with equality and additive inequality conditions, and their application to k-means clustering and learning linear support vector machines and models with non-polynomial loss [4];

and the incremental maintenance of linear regression models [53] and of relational queries [35, 37] under data updates. Sections 3, 4, and 5 highlight some of this development.

Our effort is not alone. There is an increasingly larger body of work at the interface of databases and machine learning towards addressing the shortcomings listed in Section 1.2. For instance, increasingly more solutions, starting with MADlib [32] and Bismarck [26], integrate the two systems into one system to avoid Shortcomings (2), (4), and (5). This mirrors two-decades old efforts on integrating data mining and databases [18]. Shortcoming (3) is commonly addressed by compressed feature vectors that avoid explicit representation of zero values as done by, e.g., TensorFlow [1] and libFM [63]. Our previous work [6, 7] proposed a sparse tensor representation for one-hot encoded features that can be expressed using group-by queries; this is used by LMFAO [65]. A few solutions aim at addressing Shortcoming (1) by pushing the learning task past the joins. Whereas most solutions only handle restricted joins, e.g., star key-key joins [41], others allow for arbitrary join schemas [66, 65]. A brief and incomplete overview of such approaches is given in a recent tutorial [67]. A promising area of research is understanding the connection between learning and query languages, e.g., the relative expressiveness of graph neural networks and fragments of first-order logic and SQL [16].

## 1.4 Wait! This problem was already solved

*Death dealers, who categorise research problems as either "done/dead" or "not done", might say that the problem of learning over relational data is done. Yet existing database technology is not ready for workloads specific to training models.*

Given that the structure-aware solution essentially relies on casting data-dependent computation as queries, why not use existing database systems to compute them?

Figure 4 depicts two of our experimental findings.

The first experiment considers the computation of the aggregates that define the covariance matrix used to learn a linear regression model and the costs used by CART to decide on a node in a regression tree [65]. The number of distinct aggregates varies from hundred to thousands for the four

<sup>2</sup><https://github.com/fdbresearch/FIVM>

considered datasets, yet they are very similar and amenable to significant shared computation. The speedup of our prototype LMFAO over two competitive database systems is on par with the number of aggregates, suggesting that these systems might only provide limited support for sharing computation across the aggregates in a batch. These systems performed worse on SQL encodings of LMFAO’s shared execution plans, which compute many intermediate aggregates. This hints at LMFAO’s distinct design that departs from mainstream query processing.

The second experiment considers the maintenance of the covariance matrix under tuple insertions into an initially empty retailer dataset [53]. The throughput of our system F-IVM (Factorised Incremental View Maintenance) is orders of magnitude higher than of first-order (classical delta processing) and higher-order (delta processing with intermediate views) incremental view maintenance approaches. The three approaches are implemented on top of DBToaster’s backend [39] and take advantage of its intermediate representation language and code compilation. The main reason for the performance difference lies again in sharing the maintenance across the aggregates within a batch.

These experiments suggest that the task of learning over relational data comes with new database workloads that require novel optimisation and evaluation techniques. LMFAO and F-IVM are just two points of reference in a large design space for systems to cope with the new workloads.

## 1.5 Wait, wait! This effort is irrelevant

*This effort misses the point: Accuracy is the holy grail of machine learning, not runtime efficiency. Yet faster training can mean better accuracy.*

The term “performance” is mostly used to mean accuracy in the machine learning literature and runtime efficiency in the database community. How can a database mindset ever make a difference to the machine learning battle field? It is obviously pointless to quickly train models with poor accuracy and therefore not useful. While this is a valid point, being able to train much faster can in fact help improve the accuracy in several scenarios.

First, it raises the possibility to train several models within the same time that it takes a slower system to train one of these models. For instance, in our experiment in Figure 3, once we computed the covariance matrix over the given features, we can train a new model over a subset of these features in 50 milliseconds. In contrast, TensorFlow would use separate scans of the data matrix to train each model, with each scan taking over 7,000 seconds. This would effectively allow the faster approach to perform model selection by looking over a large set of possible models defined by subsets of the set of all given features within a fraction of the time it takes TensorFlow to complete one model.

Second, it allows to train more complex and expensive models in a reasonable time.

Third, it opens the door to keeping models fresh so that they accurately reflect the new evidence: Instead of updating the models say every day solely due to the abysmal runtime performance of the learning library, one can update it every hour or even every minute. In our streaming experiment in Figure 4 (right), F-IVM can maintain the covariance

|               | Retailer | Favorita | Yelp  | TPC-DS |
|---------------|----------|----------|-------|--------|
| Covar. matrix | 937      | 157      | 730   | 3,299  |
| Decision node | 3,150    | 273      | 1,392 | 4,299  |
| Mutual inf.   | 56       | 106      | 172   | 254    |
| $k$ -means    | 44       | 19       | 38    | 92     |

**Figure 5: Number of aggregates for various datasets and workloads (covariance matrix for linear regression; decision tree node; mutual information for model selection and Chow-Liu trees;  $k$ -means).**

matrix with a throughput of over one million tuples per second. Refreshing the linear regression model after each update (we considered a bulk of 1000 inserts at a time) would take tens of milliseconds only; it takes less than the 50 milliseconds for computing the parameters from scratch given the covariance matrix, since we resume the convergence procedure of the gradient descent already with parameter values that are close to the final ones.

Besides the runtime aspect, the structure of the relational data and the human-added knowledge in the form of constraints and ontologies may help synthesise more accurate models. This is a relevant subject that is out of the scope of this paper and deserves proper treatment on its own.

## 2. FROM LEARNING TO AGGREGATES

*Aggregation is the aspirin to all problems.*

As argued in Section 1, turning the machine learning problem into a database problem follows a database community’s recipe for success. The learning task relies on various forms of aggregation over the data matrix that are readily supported by database query languages. The types of aggregates depend on the model and objective function. In this section, we give such aggregates for several models; it is beyond the scope of this overview paper to show how to derive them. For an in-depth treatment, please refer to our work [6, 65, 7, 4]. Figure 5 shows that their number is much more than in a typical database query. The efficient computation of batches of such aggregates requires new techniques that exploit their structure and similarities.

We next consider that the data matrix  $D$ , which is defined by a feature extraction query, consists of tuples  $(\mathbf{x}, y)$  of a feature vector  $\mathbf{x}$  and a response  $y$ .

### 2.1 Least-squares loss function

*The gradient vector of the least-squares loss function is built up using sum-product aggregates over the model features and parameters.*

For linear regression models trained by minimising the least-squares loss function using gradient descent optimisation, the data-dependent aggregates that come up in the gradient vector have one of the following forms:

$$\begin{array}{ll} \text{SUM}(X_i * X_j) & \\ \text{SUM}(X_i) & \text{GROUP BY } X_j \\ \text{SUM}(1) & \text{GROUP BY } X_i, X_j \end{array}$$

where  $X_i$  and  $X_j$  are database attributes whose values make up the features  $x_i$  and respectively  $x_j$ . In case both  $x_i$  and



$x_j$  are continuous, we sum their products over all tuples in the data matrix  $D$ . In case  $x_i$  is continuous and  $x_j$  is categorical, we report the sum of  $x_i$  for each category of  $x_j$ . Finally, in case both  $x_i$  and  $x_j$  are categorical, we count the number of occurrences for each pair of their categories that appears in  $D$ . Each pair of features, or a pair of a feature and the response, defines such an aggregate; for  $n$  features, we thus have  $(n + 1)^2$  such aggregates (although the aggregate  $\text{SUM}(y^2)$  over the response  $y$  is not needed in the gradient vector). We arrange these aggregates in an  $(n + 1) \times (n + 1)$  symmetric matrix, where the  $(i, j)$  entry is the aggregate for  $X_i$  and  $X_j$ . If the features are standardised, this matrix is a compact encoding of the covariance matrix of the features.

Two remarks are in order. First, the data matrix  $D$  need not be materialised, so the above aggregates are expressed directly over the feature extraction query  $Q$ :

```
SELECT X, agg FROM Q GROUP BY X;
```

where  $X$  represents the categorical feature(s) and **agg** is one of the sum aggregates discussed before. Parts of the aggregates can be pushed past the joins and shared across the aggregates. This reduces the sizes of the intermediate results. For the evaluation of such queries, we resort to factorised query processing [57, 56] that combines aggregate pushdown with worst-case optimal join computation [74] to achieve a runtime asymptotically lower than classical query plans with binary joins [51]. The feature extraction queries are typically acyclic, in which case our execution strategy takes time linear in the *input* data instead of output data.

Second, these aggregates represent a sparse tensor encoding of the interactions of any two categorical features: Instead of one-hot encoding them, we only represent the pairs of categories of the features that appear in the data matrix  $D$  [7]. This is naturally expressed using the **group-by** clause.

Similar aggregates can be derived for polynomial regression models [56], factorisation machines [6], sum-product networks [20], quadratically regularised principal component analysis [33], and QR and SVD decompositions [73].

## 2.2 Cost functions 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 [17] for constructing regression trees rely on aggregates that compute the variance of the response conditioned on a filter:

```
VARIANCE(Y) WHERE Xi op cj
```

For a continuous feature  $x_i$ , the filter asks whether its value is above a threshold  $c_j$ :  $X_i \geq c_j$ . For a categorical feature  $x_i$ , the filter asks whether its value is in a set of possible categories:  $X_i$  **in**  $(v_1, \dots, v_k)$  or  $X_i = v$ . 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.

The aggregates for decision trees are thus very similar to those in Section 2.1, with the addition of filters.

## 2.3 Non-polynomial loss functions

*The sub-gradients of non-polynomial loss functions require a new type of theta joins with additive inequality conditions. Such joins call for new algorithms beyond the classical ones.*

A large class of models including support vector machines and robust regression are trained using sub-gradient descent. They use non-polynomial loss functions, such as (ordinal) hinge, 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. Huber loss admits a gradient with additive inequalities. The (sub)gradients of these loss functions can be formulated using aggregates of the form [4]:

```
SUM(X) WHERE X1 * w1 + ... + Xn * wn > c GROUP BY Z
```

where  $X$  may be the value one, a feature, or the product of two attributes for continuous features, while  $Z$  is none, one, or two attributes for categorical features.

This aggregate is structurally different from the previous ones as it involves an *additive inequality* join condition. This is a new type of theta join. Existing database systems evaluate it by iterating over the data matrix and checking the additive inequality condition for each tuple. As we show in recent work [4], there are better algorithms that do not need to iterate over the entire data matrix and that may need polynomially less time to compute such aggregates with additive inequalities. This is a prime example of a new database workload motivated by a machine learning application.

Similar aggregates are derived for  $k$ -means clustering [4].

## 3. LOWERING THE ASYMPTOTICS

*Over decades, the database community has developed a toolbox of techniques to understand and use the algebraic, combinatorial, statistical, and geometric structure of relational data to lower the computational complexity of database problems.*

In this section, we overview principles behind structure-aware learning that led to new algorithms for queries with joins, additive inequality conditions, and group-by aggregates, and the analysis of their computational complexity.

### 3.1 Algebraic structure

Relational data exhibits an algebraic structure: A relation is a sum-product expression, where the sum is the set union and the product is the Cartesian product. Relational algebra computation can be captured using (semi)rings<sup>3</sup>, such

<sup>3</sup>A ring  $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$  is a set  $\mathbf{D}$  with closed binary operations  $+$  and  $*$ , the additive identity  $\mathbf{0}$ , and the multiplicative identity  $\mathbf{1}$  satisfying the axioms  $(\forall a, b, c \in \mathbf{D})$ :

1.  $a + b = b + a$ .
2.  $(a + b) + c = a + (b + c)$ .
3.  $\mathbf{0} + a = a + \mathbf{0} = a$ .
4.  $\exists -a \in \mathbf{D} : a + (-a) = (-a) + a = \mathbf{0}$ .
5.  $(a * b) * c = a * (b * c)$ .
6.  $a * \mathbf{1} = \mathbf{1} * a = a$ .
7.  $a * (b + c) = a * b + a * c$  and  $(a + b) * c = a * c + b * c$ .

as work on  $k$ -relations over provenance semirings [27], generalised multiset relations [38], and factors over semirings [9].

The following properties make the rings particularly effective for computing and maintaining aggregates.

**Distributivity law.** This powerful law underlies our work on factorised databases [57, 56]. In particular, it allows to factor out data blocks common to several tuples, represent them once and compute over them once. Using factorisation, relations can be represented more succinctly as directed acyclic graphs. For instance, we can trade the explicit representation of the Cartesian product of two sets  $R = \cup_{i \in [n]} r_i$  and  $S = \cup_{j \in [m]} s_j$  for a symbolic representation:

$$\cup_{i \in [n], j \in [m]} (r_i \times s_j) = \left( \cup_{i \in [n]} r_i \right) \times \left( \cup_{j \in [m]} s_j \right)$$

Whereas the left expression has  $2nm$  values, the right expression only has  $n + m$  values. This can be generalised to joins, which are by definition unions of Cartesian products, and also to queries with select, project, join, group-by aggregates, and order-by clauses. A join query may admit several factorisations and a question is which ones are optimal, i.e., there is no factorisation whose size is smaller in worst-case. Can such factorisations be computed optimally, i.e., in time proportional to their sizes? These questions have been answered in the affirmative [57]. Factorisation can lower the computational complexity of joins [15], aggregates [14], and machine learning [66, 56, 6, 4]. Section 5.1 exemplifies factorised computation for joins and aggregates.

As a representation formalism, factorised databases are a special class of multi-valued ordered decision diagrams that have been studied in knowledge compilation, a field in AI literature that aims at exploiting structure in (worst-case hard) problems to solve them efficiently in practice [24]. The framework of Functional Aggregate Queries [9] generalises factorised databases to semirings beyond sum-product and shows that many problems across Computer Science can benefit from factorised computation. LMFAO [65, 64], F-IVM [53, 54], and IVM<sup>ε</sup> [37, 36] employ factorised query computation and maintenance. These systems factorise the query into a hierarchy of increasingly simpler views, which are maintained bottom-up under data updates.

**Additive inverse.** The additive inverse of rings allows to treat uniformly data updates (inserts and deletes). This would require relations to map tuples to their multiplicities, i.e., using the ring of integers. Then, an insert/delete of a tuple  $r$  into/from a relation  $R$  is modelled as the addition of the mapping  $\{r \mapsto 1\}$  and respectively  $\{r \mapsto -1\}$  to the map  $R$ . If the key  $r$  is already present in  $R$ , then we sum up their multiplicities so that each key appears once in the map  $R$ . Tuples with multiplicity zero are not represented.

This uniform treatment of updates simplifies the maintenance procedure for queries [39, 35] and models [53, 54].

**Sum-product abstraction.** The sum-product abstraction in rings allows to use the same processing (computing and maintaining) mechanism for seemingly disparate tasks, such as database queries, covariance matrices, inference in probabilistic graphical models, and matrix chain multiplication [9, 53]. Section 5.2 exemplifies a ring designed for covariance matrices. The efficient maintenance of covariance

A semiring  $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$  satisfies all of the above properties except the additive inverse property (4) and adds the axiom  $\mathbf{0} * a = a * \mathbf{0} = \mathbf{0}$ . A (semi)ring for which  $a * b = b * a$  is commutative. Axiom (7) is the distributivity law.

matrices makes it possible to keep models fresh under high-throughput data changes [53]. A recent tutorial overviews advances in incremental view maintenance [25].

## 3.2 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. 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  $N$ .

**Width measures.** Various increasingly more refined and smaller width measures have been proposed recently:

- The fractional edge cover number [29] captures the asymptotic size of the results for join queries [29, 12] and the time to compute them [50, 51, 74].
- The fractional hypertree width [43] is defined using the fractional edge cover number of the bags in the hypertree decompositions of the query.
- The factorisation width [57, 56] generalises the fractional hypertree width to non-Boolean queries. This width defines the size of factorised query results.
- The FAQ-width [9] generalises the factorisation width to functional aggregate queries over several semirings.
- The submodular width [44, 10] is the smallest known width for Boolean conjunctive queries.
- The relaxed versions [4] of FAQ-width and submodular width capture the complexity of queries with additive inequalities discussed in Section 2.3.

This theoretical development led to a number of algorithms that attain the above width measures: LFTJ [74], F [66, 56], InsideOut [9], EmptyHeaded [2], PANDA [10], and #PANDA [4]. Our prototypes F [66] and LMFAO [65] for learning over relational data use an underlying query engine that achieves the factorisation width.

**Data degree.** The degree information captures the number of occurrences of a data value in the input database [51]. Several existing query processing algorithms adapt their execution strategy depending on the degree of data values, with different strategies for heavy and light values, where a value is heavy/light if its degree is above/below a certain chosen threshold. This adaptive processing has been used by worst-case optimal join algorithms [50], worst-case optimal incremental maintenance of triangle queries [35, 36], and to achieve the submodular width for Boolean conjunctive queries [10]. The complexity of query evaluation can reduce dramatically in the presence of bounded degrees or even more refined notions of data sparsity, such as nowhere denseness [28]. This is the case, for instance, for computing queries with negated relations of bounded degree [8].

**Functional dependencies.** 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 [6, 7]. Take for instance a model whose parameters include a parameter  $\theta_{\text{city}}$  for each city and  $\theta_{\text{country}}$

for each country in the input data, and assume the functional dependency  $\text{city} \rightarrow \text{country}$  holds in the database. We can replace each parameter pair  $(\theta_{\text{city}}, \theta_{\text{country}})$  with a new parameter  $\theta_{(\text{city}, \text{country})}$ . The relationship between the original and new parameters can be expressed in closed form for ridge polynomial regression and factorisation machines, hence  $(\theta_{\text{city}}, \theta_{\text{country}})$  can be recovered from  $\theta_{(\text{city}, \text{country})}$ .

### 3.3 Statistical and geometric structure

An alternative approach to computing over the entire input data is to compute over data samples. There is a solid body of work in sampling for machine learning [48]. In our database setting, the major challenge is that we would need to sample through the feature extraction query. The most notable advances in the database literature is on sampling through selection conditions and joins, e.g., the ripple joins [31] and the wander joins [42]. However, the feature extraction query may also compute derived aggregate features, which means that sampling through aggregates would be also needed. More recent work considers sampling for specific classes of machine learning models [58]. This is a challenging and interesting subject to future work.

Sampling is employed whenever the input database is too large to be processed within a given time budget. It may nevertheless lead to approximation of both steps in the end-to-end learning task, from the computation of the feature extraction query to the subsequent optimization task that yields the desired model. Work in this space quantifies the loss in accuracy of the obtained model due to sampling [48].

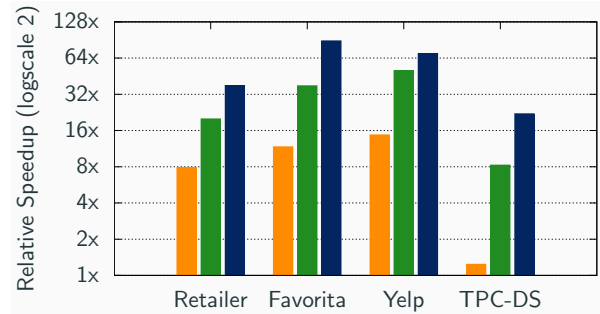
Geometric structure is relevant whenever we use distance measures. Clustering algorithms can exploit such measures, e.g., the optimal transport distance between two probability measures, and inequalities between them, e.g., the triangle inequality. One example is Relational  $k$ -means [23], which achieves constant-factor approximations of the  $k$ -means objective by clustering over a small coreset instead of the full result of the feature extraction query.

## 4. LOWERING THE CONSTANT FACTORS

*The database researcher continuously refines a toolbox of clever system tricks, including: specialisation for workload, data, and hardware; observing the memory hierarchy and blocking operations; distribution and parallelisation. These tools can bring otherwise computationally hard problems into the realm of the possible.*

While asymptotic runtime improvements are desirable, their successful implementation in practical cases is not a given. They also aim at upper bounding the worse-case scenario, which may not happen in practice. While worst-case optimal join algorithms [74, 11, 52] and factorised computation and maintenance of models [65, 53] have been validated in practice, the jury is still out on adaptive query strategies [35, 37] and the submodular width [10]. At least of equal relevance are systems aspects of theoretically-efficient algorithms. Our prototyping effort considered code specialisation, shared scans, and parallelisation, which contribute to LMFAO’s runtime performance as shown in Figure 6. The baseline for this comparison (1x speedup) is AC/DC [5], the precursor of LMFAO that has none of these optimisations.

We next mention aspects of lowering the constants in our LMFAO and IFAQ systems.



**Figure 6: Speedup of code optimisations in LMFAO for computing the covariance matrix relative to baseline implementation. The optimisations are added in this order: specialisation, sharing, parallelization (AWS d2.xlarge, 4 vCPUs, 32GB).**

**Workload compilation.** LMFAO generates code specific to the query and database schema [49, 68, 69], and also specific to the model.

For a more systematic study of the code optimisation and generation effort in LMFAO, we recently initiated a new project called IFAQ (Iterative Functional Aggregate Queries) [70]. IFAQ is a framework that allows to specify in a unified domain-specific language DB+ML workloads and provides a unified optimisation and compilation approach to programs in this language. Example of such programs include: gradient descent optimisation for learning a variety of models over relational data, the CART algorithm for decision trees, model selection, and cross validation. Section 5.3 exemplifies IFAQ transformations for learning a linear regression model over a join using gradient descent.

**Sharing computation.** Both LMFAO and IFAQ exhibit sharing across the batch of aggregates at different stages. We next sketch two common sharing scenarios.

Consider a join tree of the feature extraction query (if the query is cyclic, it is first partially evaluated to an acyclic query). We decompose each aggregate in the batch following a top-down traversal of this tree as follows. We assign at a node in the tree a restriction of the aggregate that only refers to the attributes within the subtree rooted at that node. This means that we assign the entire aggregate at the root. It also means that if the subtree rooted at a node has no attributes in the aggregate, then we assign a count aggregate at the node. The reason is that this restriction of the aggregate can be answered at each node in the tree using the partial aggregates computed at its children, if any. At each node, we then consolidate the synthesised aggregates into views. For the batches discussed in Section 2, it is often the case that several aggregates have the same partial aggregates at a node in the tree, so we only compute these partial common aggregates once and use them several times.

Furthermore, to compute the views at the same node, we can share the scan of the relation at the node. To increase sharing and minimise the sizes of the intermediate views, LMFAO may decompose different aggregates starting at different root nodes in the join tree.

**Parallelisation.** LMFAO supports parallelism by computing in parallel groups of views that do not depend on each other. It supports domain parallelism by deploying threads to compute view groups on disjoint parts of a relation.



| Orders   |        |        | Dish   |         | Items   |       | Natural join of Orders, Dish, and Items |          |          |         |         |
|----------|--------|--------|--------|---------|---------|-------|---|----------|----------|---------|---------|
| customer | day    | dish   | dish   | item    | item    | price | customer ×                              | day ×    | dish ×   | item ×  | price ∪ |
| Elise    | Monday | burger | burger | patty   | patty   | 6     | Elise ×                                 | Monday × | burger × | patty × | 6 ∪     |
| Elise    | Friday | burger | burger | onion   | onion   | 2     | Elise ×                                 | Monday × | burger × | onion × | 2 ∪     |
| Steve    | Friday | hotdog | burger | bun     | bun     | 2     | Elise ×                                 | Monday × | burger × | bun ×   | 2 ∪     |
| Joe      | Friday | hotdog | hotdog | bun     | sausage | 4     | Elise ×                                 | Friday × | burger × | patty × | 6 ∪     |
|          |        |        | hotdog | onion   |         |       | Elise ×                                 | Friday × | burger × | onion × | 2 ∪     |
|          |        |        | hotdog | sausage |         |       | Elise ×                                 | Friday × | burger × | bun ×   | 2 ∪     |

Figure 7: Example database used in Section 5.1 with relations Orders, Dish, and Items, and a fragment of their natural join expressed as a relational algebra expression with union and Cartesian product.

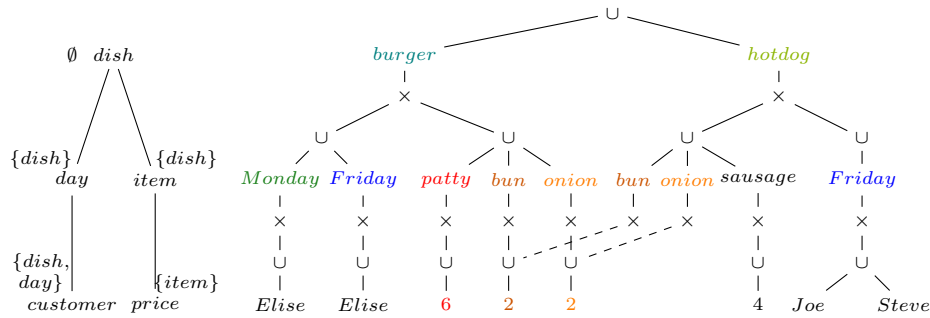


Figure 8: Factorised representation of the join result in Figure 7. Left: Order on the query variables, where each variable is adorned with its ancestors on which it depends. Right: Factorisation of the natural join from Figure 7 modelled on the variable order.

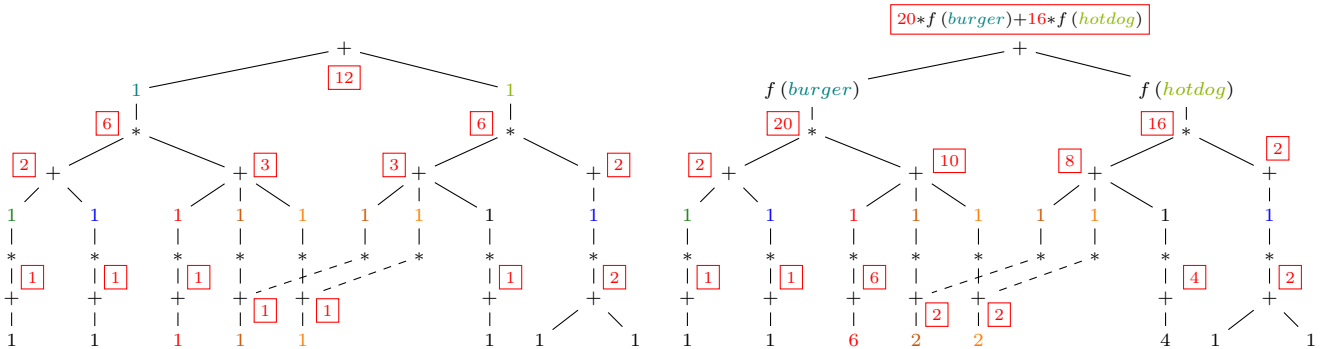


Figure 9: Aggregate computation computed in one pass over the factorised join in Figure 8. Left: SUM(1). We map all values to 1, ∪ to + and × to \*. Right: SUM(dish \* price). We assume there is a function  $f$  that turns dish into numbers (e.g., one-hot encoding) or we group by dish: SUM(price) GROUP BY dish. We map all values except for dish and price to 1, ∪ to + and × to \*.

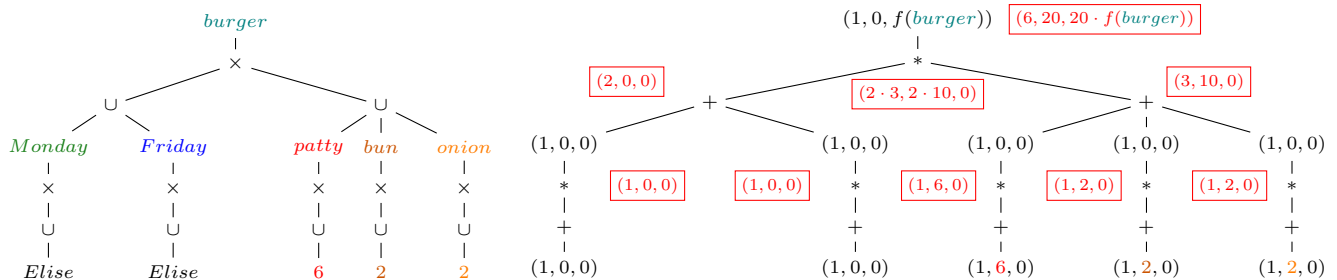


Figure 10: Left: Factorised join. Right: Computing SUM(1), SUM(price), SUM(price \* dish) using a ring whose elements are triples of numbers, one per aggregate. The sum (+) and product (\*) operations are now defined over triples. This specialised ring enables shared computation across the aggregates.

## 5. HIGHLIGHTS

We exemplify some of the query computation and maintenance techniques mentioned in Sections 3 and 4.

### 5.1 Factorised computation

*Factorisation is a normal form for the query result that comes with time and space improvement.*

Consider the relations and their natural join in Figure 7. The relations, including the join, can be interpreted as relational algebra expressions over union and Cartesian product: Each tuple is a product of values and a relation is a union of such tuples. The set of possible relations together with these two operations form a semiring<sup>4</sup>.

As highlighted in Figure 7, the same value may occur several times in a relation. Using the distributivity and commutativity laws of the semiring, we may factor out occurrences of the same value or even entire expressions so as to achieve a small representation. Figure 8 depicts a factorisation of the join. There are several algebraically equivalent factorisations. The one we depict has fewer data values than the input relations. In contrast, the number of values in the non-factorised join can be much larger.

The factorised join is modelled on an order of the variables (or attributes) in the join result. This variable order dictates the compression factor of the factorisation. It exploits the join dependencies in the join result to avoid redundancy in the representation. In this sense, the factorisation can be thought of as a normal form for the join result (in analogy with the normalisation commonly applied to the input relations). To avoid redundancy, the factorisation represents separately conditionally independent information. For instance, the days are independent of the items *given* a dish. Indeed, the latter and the former are in different relations, which are joined via dish. We express graphically this conditional independence by having these variables in different branches under dish. Not every conditional independence can be captured by branching. For instance, price is independent of dish given item, yet all three variables are along the same path in the order. To signal this independence, we explicitly adorn each variable in the order by the set of its ancestors on which it depends. For instance, price is adorned with item only and not also with dish.

The structure of the variable order is mirrored in the factorised join. We first group by dish; for each dish, we represent separately the days and customers from the items and their prices. Furthermore, since price is independent from dish given item, we cache the price for a specific item so as to avoid its repetition under several dishes.

By avoiding data repetitions, factorised representations can be much more succinct than their equivalent tabular representation. The variable order can also effectively guide us to compute the factorisation directly from the input relations and in time proportional to its size.

We can compute aggregates of the form given in Section 2 directly on the factorised join. Moreover, the operators for join and aggregates can be fused so that we do not need to first compute the factorised join. Figure 9 shows how to compute a count and a sum-product aggregate used for

<sup>4</sup>Constraint: Only relations with the same schema can be added and relations with disjoint schemas can be multiplied.

the covariance matrix. They are computed in one bottom-up traversal of the factorised join using a different semiring. For the count, we use the semiring of the natural numbers and map each value to 1. For the sum-product aggregate, we use the same semiring mapping, with the exception that the price values are kept (since we want to sum over them) while the dish values are mapped to numbers (or become keys, if we were to group by dish).

### 5.2 Sum-product abstraction

*Ring operations can be redefined to capture the shared computation of aggregates in a batch.*

We continue the example in Section 5.1 and consider a strategy that shares the computation of several aggregates. This is motivated by the large number of similar aggregates needed to train models, cf. Section 2. In particular, we would like to design a ring that captures this shared computation.

Figure 10 gives a factorised join (a fragment of the one in Figure 8 for simplicity) and a bottom-up run over it annotated with triples of values corresponding to local computations for the aggregates  $\text{SUM}(1)$ ,  $\text{SUM}(\text{price})$ ,  $\text{SUM}(\text{price} * \text{dish})$ . We can observe that our aggregates are of increasing complexity: the first does not have a variable, the second has one variable, and the third has two variables. Such aggregates are common in covariance matrices.

We proceed with our computation as in Figure 9, yet now with all three aggregates together. There are two instances where computation sharing occurs. At the product node under the root: The count from the left branch is used to scale the sum of prices from the right. At the root node, the second aggregate is used to compute the third aggregate.

A natural question is whether we can generalise this to an entire covariance matrix. For a set of numerical features  $x_1, \dots, x_n$ , this matrix is defined by the aggregates  $\text{SUM}(1)$ ,  $\text{SUM}(x_i)$ , and  $\text{SUM}(x_i * x_j)$  for  $i, j \in [n]$ . Following our example, we can use their partial computation as follows:

- $\text{SUM}(1)$  can be reused for all  $\text{SUM}(x_i)$  and  $\text{SUM}(x_i * x_j)$ .
- $\text{SUM}(x_i)$  can be reused for all  $\text{SUM}(x_i * x_j)$ .

The ring  $(\mathcal{R}, +, *, \mathbf{0}, \mathbf{1})$  over triples of aggregates  $(c, \mathbf{s}, \mathbf{Q})$  captures this shared computation:

$$\begin{aligned} (c_1, \mathbf{s}_1, \mathbf{Q}_1) + (c_2, \mathbf{s}_2, \mathbf{Q}_2) &= (c_1 + c_2, \mathbf{s}_1 + \mathbf{s}_2, \mathbf{Q}_1 + \mathbf{Q}_2) \\ (c_1, \mathbf{s}_1, \mathbf{Q}_1) * (c_2, \mathbf{s}_2, \mathbf{Q}_2) &= (c_1 \cdot c_2, c_2 \cdot \mathbf{s}_1 + c_1 \cdot \mathbf{s}_2, c_2 \cdot \mathbf{Q}_1 + \\ &\quad c_1 \cdot \mathbf{Q}_2 + \mathbf{s}_1 \mathbf{s}_2^T + \mathbf{s}_2 \mathbf{s}_1^T) \\ \mathbf{0} &= (0, \mathbf{0}_{n \times 1}, \mathbf{0}_{n \times n}) \\ \mathbf{1} &= (1, \mathbf{0}_{n \times 1}, \mathbf{0}_{n \times n}) \end{aligned}$$

A ring element consists of a scalar, a vector, and a matrix:

$$\left( \square, \begin{array}{|c|} \hline \color{red}{\blacksquare} \\ \hline \end{array}, \begin{array}{|c|c|} \hline \color{blue}{\blacksquare} & \color{blue}{\blacksquare} \\ \hline \end{array} \right)$$

$\underbrace{\hspace{1.5cm}}_{\text{SUM}(1)} \quad \underbrace{\hspace{1.5cm}}_{\text{SUM}(x_i)} \quad \underbrace{\hspace{1.5cm}}_{\text{SUM}(x_i * x_j)}$

In the definition of the product operation, the scalars  $c_i$  are used to scale the vectors  $\mathbf{s}_i$  and matrices  $\mathbf{Q}_i$ , while the vectors  $\mathbf{s}_i$  are used to compute the matrices  $\mathbf{Q}_i$  ( $i \in [2]$ ).

Shared computation across the aggregates as emulated by the above ring is partly responsible for the performance benefits shown in Figure 4 for LMFAO [65] and F-IVM [53].

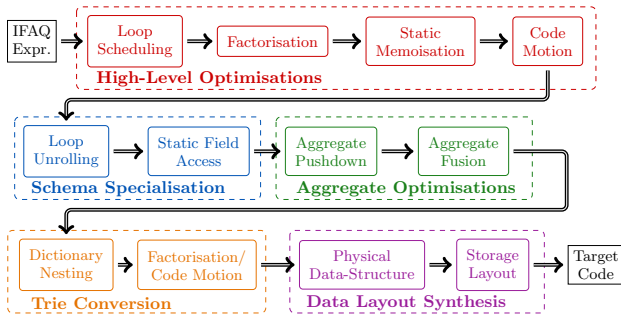


Figure 11: Transformation steps in IFAQ.

### 5.3 Multi-stage compilation

IFAQ can automatically synthesise and optimise aggregates from ML+DB workloads.

We exemplify IFAQ [70] using a simplified gradient descent program that learns the parameters of a linear regression model over the join  $\mathbf{Q} = S \bowtie R \bowtie I$  of relations  $S(i, s, u)$ ,  $R(s, c)$ , and  $I(i, p)$ <sup>5</sup>. The model predicts the response  $u$  given the set of numerical features  $\mathbf{F} = \{i, s, c, p\}$ .

IFAQ takes the gradient descent program through the sequence of equivalence-preserving rule-based transformations shown in Figure 11. We start with the program:

```

let  $\mathbf{F} = \{i, s, p, c\}$  in
 $\theta \leftarrow \theta_0$ 
while( not converged ) {
 $\theta = \lambda_{f_1 \in \mathbf{F}} \left( \theta(f_1) - \sum_{\mathbf{x} \in \text{sup}(\mathbf{Q})} \mathbf{Q}(\mathbf{x}) * \left( \sum_{f_2 \in \mathbf{F}} \theta(f_2) * \mathbf{x}(f_2) \right) * \mathbf{x}(f_1) \right)$ 
}
 $\theta$ 

```

IFAQ supports collections via dictionaries, which map keys to values that can be simple data types or again dictionaries. Sets are dictionaries, where we only show the keys. Example of dictionaries are:  $\text{sup}(\mathbf{Q})$  mapping of tuples in the join  $\mathbf{Q}$  to feature vectors; the feature set  $x$  mapping features ( $i$ ,  $s$ ,  $p$ , and  $c$ ) to values; the corresponding parameter set  $\theta$ . IFAQ has three iteration constructs:  $\lambda_{e \in \text{set}} f(e)$  constructs a dictionary that maps each key  $e$  to  $f(e)$ ; the stateful computation  $\sum_{e \in \text{set}} f(e)$  sums  $f(e)$  over the elements in the set; and the **while** construct is used for general iterations.

Without loss of generality, we left out the convergence condition and the step size and assume that the response  $u$  is also included in  $\mathbf{x}$  ( $\theta(u) = 1$  need not be updated).

**High-level optimisations.** We first normalise the gradient by bringing  $\mathbf{Q}(x)$  and  $\mathbf{x}(f_1)$  inside the innermost summation, swap the two summations so that the inner one iterates over larger dictionaries than the outer one, and factor out  $\theta(f_2)$  to use less arithmetic operations. The update step becomes:

$$\theta = \lambda_{f_1 \in \mathbf{F}} \left( \theta(f_1) - \sum_{f_2 \in \mathbf{F}} \theta(f_2) * \left( \sum_{\mathbf{x} \in \text{sup}(\mathbf{Q})} \mathbf{Q}(\mathbf{x}) * \mathbf{x}(f_2) * \mathbf{x}(f_1) \right) \right)$$

Since the data-intensive computation over  $\mathbf{Q}$  in the above expression is independent of the two loops over  $\mathbf{F}$ , we can

<sup>5</sup>The relations can be for instance `Sales(item, store, units)`, `StoRes(store, city)`, and `Items(item, price)`.

hoist it outside the convergence loop. To achieve this, we first use memoisation to synthesise the dictionary  $\mathbf{M}$ :

$$\text{let } \mathbf{M} = \lambda_{f_1 \in \mathbf{F}} \lambda_{f_2 \in \mathbf{F}} \sum_{\mathbf{x} \in \text{sup}(\mathbf{Q})} \mathbf{Q}(\mathbf{x}) * \mathbf{x}(f_2) * \mathbf{x}(f_1) \text{ in}$$

$$\theta = \lambda_{f_1 \in \mathbf{F}} \left( \theta(f_1) - \sum_{f_2 \in \mathbf{F}} \theta(f_2) * \mathbf{M}(f_1)(f_2) \right)$$

The dictionary  $\mathbf{M}$  encodes the (non-centred) covariance matrix. Next, the code motion transformation moves the computation of  $\mathbf{M}$  outside the **while** convergence loop:

```

let  $\mathbf{F} = \{i, s, p, c\}$  in
let  $\mathbf{M} = \lambda_{f_1 \in \mathbf{F}} \lambda_{f_2 \in \mathbf{F}} \sum_{\mathbf{x} \in \text{sup}(\mathbf{Q})} \mathbf{Q}(\mathbf{x}) * \mathbf{x}(f_2) * \mathbf{x}(f_1)$  in
 $\theta \leftarrow \theta_0$ 
while( not converged ) {
 $\theta = \lambda_{f_1 \in \mathbf{F}} \left( \theta(f_1) - \sum_{f_2 \in \mathbf{F}} \theta(f_2) * \mathbf{M}(f_1)(f_2) \right)$ 
}
 $\theta$ 

```

**Schema specialisation.** Since the set  $\mathbf{F}$  of features is known at compile time, we can unroll the loops over  $\mathbf{F}$  in the definitions for  $\mathbf{M}$  and  $\theta$ . We show below the expanded definition of  $\mathbf{M}$ :

$$\mathbf{M} = \left\{ c \rightarrow \left\{ \dots, p \rightarrow \sum_{\mathbf{x} \in \text{sup}(\mathbf{Q})} \mathbf{Q}(\mathbf{x}) * \mathbf{x}(c) * \mathbf{x}(p), \dots \right\}, \dots \right\}$$

We next enforce static field access wherever possible. This means that we convert dictionaries over  $\mathbf{F}$  into records so that the dynamic accesses to the dictionary now become static accesses. For instance, the dictionary  $\mathbf{x}$  becomes the record  $x$  and then the dynamic access  $\mathbf{x}(c)$  becomes the static access  $x.c$ . The program becomes:

```

let  $M = \left\{ c = \left\{ \dots, p = \sum_{x \in \text{sup}(\mathbf{Q})} \mathbf{Q}(x) * x.c * x.p, \dots \right\}, \dots \right\}$ 
in
 $\theta \leftarrow \theta_0$ 
while( not converged ) {
 $\theta = \left\{ c = \theta.c - \left( \dots + \theta.c * M.c.c + \theta.p * M.c.p \dots \right), \dots \right\}$ 
}
 $\theta$ 

```

**Aggregate optimisations.** The next set of transformations focus on database optimisations. In particular, we extract the aggregates from  $M$ , push them past the joins in  $\mathbf{Q}$ , and share them across several entries in  $M$ . IFAQ identifies all aggregates in  $M$  and pulls them out of  $M$ :

```

let  $M_{cc} = \sum_{x \in \text{sup}(\mathbf{Q})} \mathbf{Q}(x) * x.c * x.c$  in
let  $M_{cp} = \sum_{x \in \text{sup}(\mathbf{Q})} \mathbf{Q}(x) * x.c * x.p$  in
let  $M = \left\{ c = \left\{ \dots, c = M_{cc}, p = M_{cp}, \dots \right\}, \dots \right\}$  in ...

```

We next unfold the definition of the join  $\mathbf{Q}$ . Recall that we do not want to materialise this join for runtime performance reasons as argued in Section 1. IFAQ represents relations as

dictionaries mapping tuples to their multiplicities. The join query  $\mathbf{Q}$  can be expressed in IFAQ as follows:

$$\mathbf{Q} = \sum_{x_s \in \text{sup}(\mathbf{S})} \sum_{x_r \in \text{sup}(\mathbf{R})} \sum_{x_i \in \text{sup}(\mathbf{I})} \left( \begin{aligned} &\text{let } k = \{i = x_s.i, s = x_s.s, c = x_r.c, p = x_i.p\} \text{ in} \\ &\{k \rightarrow \mathbf{S}(x_s) * \mathbf{R}(x_r) * \mathbf{I}(x_i) * (x_s.i == x_i.i) * (x_s.s == x_r.s)\} \end{aligned} \right)$$

By inlining  $\mathbf{Q}$  in  $M_{cp}$ , we obtain:

$$\begin{aligned} M_{cp} &= \sum_{x_s \in \text{sup}(\mathbf{S})} \sum_{x_r \in \text{sup}(\mathbf{R})} \sum_{x_i \in \text{sup}(\mathbf{I})} \left( \begin{aligned} &\text{let } k = \{i = x_s.i, s = x_s.s, c = x_r.c, p = x_i.p\} \text{ in} \\ &\text{let } \mathbf{Q} = \{k \rightarrow \mathbf{S}(x_s) * \mathbf{R}(x_r) * \mathbf{I}(x_i) * \\ &\quad (x_s.i == x_i.i) * (x_s.s == x_r.s)\} \text{ in} \\ &\sum_{x \in \text{sup}(\mathbf{Q})} \mathbf{Q}(x) * x.c * x.p \end{aligned} \right) \end{aligned}$$

Since  $\text{sup}(\mathbf{Q}) = \{k\}$  is a singleton set, we can replace the last line above by  $\mathbf{Q}(k) * x.c * x.p$ . Further inlining field accesses and retrieving the value of a singleton dictionary yields the following expression:

$$M_{cp} = \sum_{x_s \in \text{sup}(\mathbf{S})} \sum_{x_r \in \text{sup}(\mathbf{R})} \sum_{x_i \in \text{sup}(\mathbf{I})} \left( \begin{aligned} &\mathbf{S}(x_s) * \mathbf{R}(x_r) * \mathbf{I}(x_i) * (x_s.i == x_i.i) * (x_s.s == x_r.s) \\ &* x.c * x.p \end{aligned} \right)$$

We can now leverage the distributivity of multiplication over addition to factorise the expression:

$$M_{cp} = \sum_{x_s \in \text{sup}(\mathbf{S})} \mathbf{S}(x_s) * \sum_{x_r \in \text{sup}(\mathbf{R})} \mathbf{R}(x_r) * (x_s.s == x_r.s) * x_r.c * \sum_{x_i \in \text{sup}(\mathbf{I})} \mathbf{I}(x_i) * (x_s.i == x_i.i) * x_i.p$$

By static memoisation and loop-invariant code motion, we partially push the aggregates past the joins. We effect this using the dictionaries  $\mathbf{V}_R$  and  $\mathbf{V}_I$  for the partial aggregates:

$$\begin{aligned} \text{let } \mathbf{V}_R &= \sum_{x_r \in \text{sup}(\mathbf{R})} \mathbf{R}(x_r) * \{\{s = x_r.s\} \rightarrow x_r.c\} \text{ in} \\ \text{let } \mathbf{V}_I &= \sum_{x_i \in \text{sup}(\mathbf{I})} \mathbf{I}(x_i) * \{\{i = x_i.i\} \rightarrow x_i.p\} \text{ in} \\ \text{let } M_{cp} &= \sum_{x_s \in \text{sup}(\mathbf{S})} \mathbf{S}(x_s) * \mathbf{V}_R(\{s = x_s\}) * \mathbf{V}_I(\{i = x_s.i\}) \end{aligned}$$

We can obtain  $M_{cc}$  similarly:

$$\begin{aligned} \text{let } \mathbf{V}'_R &= \sum_{x_r \in \text{sup}(\mathbf{R})} \mathbf{R}(x_r) * \{\{s = x_r.s\} \rightarrow x_r.c * x_r.c\} \text{ in} \\ \text{let } \mathbf{V}'_I &= \sum_{x_i \in \text{sup}(\mathbf{I})} \mathbf{I}(x_i) * \{\{i = x_i.i\} \rightarrow 1\} \text{ in} \\ \text{let } M_{cc} &= \sum_{x_s \in \text{sup}(\mathbf{S})} \mathbf{S}(x_s) * \mathbf{V}'_R(\{s = x_s\}) * \mathbf{V}'_I(\{i = x_s.i\}) \end{aligned}$$

The expressions for  $M_{cp}$  and  $M_{cc}$  are similar. To benefit from shared computation, we fuse them:

$$\begin{aligned} \text{let } \mathbf{W}_R &= \sum_{x_r \in \text{sup}(\mathbf{R})} \mathbf{R}(x_r) * \{\{s = x_r.s\} \rightarrow \\ &\quad \{v_r = x_r.c, v'_R = x_r.c * x_r.c\}\} \text{ in} \\ \text{let } \mathbf{W}_I &= \sum_{x_i \in \text{sup}(\mathbf{I})} \mathbf{I}(x_i) * \{\{i = x_i.i\} \rightarrow \\ &\quad \{v_I = x_i.p, v'_I = 1\}\} \text{ in} \\ \text{let } M_{cc,cp} &= \sum_{x_s \in \text{sup}(\mathbf{S})} \mathbf{S}(x_s) * \left( \begin{aligned} &\text{let } w_R = \mathbf{W}_R(\{s = x_s.s\}), w_I = \mathbf{W}_I(\{i = x_s.i\}) \text{ in} \\ &\{m_{cp} = w_R.v_R * w_I.v_I, m_{cc} = w_R.v'_R * w_I.v'_I\} \end{aligned} \right) \end{aligned}$$

Then  $M_{cc} = M_{cc,cp}.m_{cc}$  and  $M_{cp} = M_{cc,cp}.m_{cp}$ . **Triе conversion.** By representing  $\mathbf{S}$  as a nested dictionary, first grouped by  $x_s$  and then by  $x_i$ , IFAQ factors out the multiplication of the elements in  $w_R$  with those in  $w_I$  before the enumeration of the  $x_i$  values. This further reduces the number of arithmetic operations needed to compute  $M_{cc,cp}$ . **Data layout.** As data structures for dictionaries, IFAQ supports hash tables, balanced-trees, and sorted dictionaries. Each of them show advantages for different workloads.

## 6. REFLECTIONS

The work overviewed in this paper follows a three-step recipe for efficient learning over relational data:

1. Turn the learning problem into a database problem.
2. Exploit the problem structure to lower the complexity.
3. Generate optimised code to lower the constant factors.

This recipe proved surprisingly effective for a good number of machine learning models and also reinvigorated the fundamental problems of query evaluation and optimisation. Further principled work on the theory and systems of structure-aware learning approaches is needed. What are the limits of structure-aware learning? What other classes of machine learning models can benefit from it? What other types of structure are relevant to the learning task? A major target remains the development of data systems that can natively run database and machine learning workloads as one. A pragmatic step towards this goal is to integrate database techniques for feature extraction queries, such as those described in this paper, in popular tools used by data scientists, e.g., TensorFlow, scikit-learn, R, and jupyter notebooks. A second major target is to translate the rich knowledge readily available in curated relational databases into the accuracy of the trained models.

It may be a painstaking endeavour to understand and translate success stories on machine learning models and optimisation to the database coordinate system, though this may reveal deeper connections between the two fields and also draw them closer. Database research may also strengthen the machine learning research landscape with its rigour in theoretical and system development, benchmarking, and reproducibility. Designing and teaching courses that present a more integrated and unified view of Computer Science fields, such as databases and machine learning, may equip both the lecturers and their students with the tools necessary to advance research and practice at the interface between seemingly disparate fields.

## 7. REFERENCES

- [1] M. Abadi, P. Barham, J. Chen, and et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, pages 265–283, 2016.
- [2] C. R. Aberger, A. Lamb, S. Tu, A. Nötzli, K. Olukotun, and C. Ré. Emptyheaded: A relational engine for graph processing. *ACM Trans. Database Syst.*, 42(4):20:1–20:44, 2017.
- [3] S. Abiteboul and et al. Research directions for principles of data management (dagstuhl perspectives workshop 16151). *Dagstuhl Manifestos*, 7(1):1–29, 2018.
- [4] M. Abo Khamis, R. R. Curtin, B. Moseley, H. Q. Ngo, X. Nguyen, D. Olteanu, and M. Schleich. On functional aggregate queries with additive inequalities. In *PODS*, pages 414–431, 2019.
- [5] M. Abo Khamis, H. Q. Ngo, X. Nguyen, D. Olteanu, and M. Schleich. AC/DC: In-database learning thunderstruck. In *DEEM*, pages 8:1–8:10, 2018.
- [6] M. Abo Khamis, H. Q. Ngo, X. Nguyen, D. Olteanu, and M. Schleich. In-database learning with sparse tensors. In *PODS*, pages 325–340, 2018.
- [7] M. Abo Khamis, H. Q. Ngo, X. Nguyen, D. Olteanu, and M. Schleich. Learning models over relational data using sparse tensors and functional dependencies. *ACM Trans. Database Syst.*, 45(2), 2020.
- [8] M. Abo Khamis, H. Q. Ngo, D. Olteanu, and D. Suciu. Boolean tensor decomposition for conjunctive queries with negation. In *ICDT*, pages 21:1–21:19, 2019.
- [9] M. Abo Khamis, H. Q. Ngo, and A. Rudra. FAQ: questions asked frequently. In *PODS*, pages 13–28, 2016.
- [10] M. Abo Khamis, H. Q. Ngo, and D. Suciu. What do shannon-type inequalities, submodular width, and disjunctive datalog have to do with one another? In *PODS*, pages 429–444, 2017.
- [11] M. Aref, B. ten Cate, T. J. Green, B. Kimelfeld, D. Olteanu, E. Pasalic, T. L. Veldhuizen, and G. Washburn. Design and implementation of the LogicBlox system. In *SIGMOD*, pages 1371–1382, 2015.
- [12] A. Atserias, M. Grohe, and D. Marx. Size bounds and query plans for relational joins. In *FOCS*, pages 739–748, 2008.
- [13] P. Avgustinov, O. de Moor, M. P. Jones, and M. Schäfer. QL: Object-oriented Queries on Relational Data. In *ECOOOP*, pages 2:1–2:25, 2016.
- [14] N. Bakibayev, T. Kociský, D. Olteanu, and J. Závodný. Aggregation and ordering in factorised databases. *PVLDB*, 6(14):1990–2001, 2013.
- [15] N. Bakibayev, D. Olteanu, and J. Závodný. FDB: A query engine for factorised relational databases. *PVLDB*, 5(11):1232–1243, 2012.
- [16] P. Barceló, E. V. Kostylev, M. Monet, J. Pérez, J. L. Reutter, and J. P. Silva. The expressive power of graph neural networks as a query language. *SIGMOD Rec.*, 49(2), 2020. To appear.
- [17] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA, 1984.
- [18] S. Chaudhuri. Data mining and database systems: Where is the intersection? *IEEE Data Eng. Bull.*, 21(1):4–8, 1998.
- [19] T. Chen and C. Guestrin. Xgboost: A scalable tree boosting system. In *KDD*, pages 785–794, 2016.
- [20] G. Chichirim. Structured Sum-Product Networks, May 2020. BSc thesis, University of Oxford.
- [21] A. A. Chien. Owning computing’s environmental impact. *Commun. ACM*, 62(3):5, Feb. 2019.
- [22] R. R. Curtin and et al. mlpack 3: A fast, flexible machine learning library. *J. Open Source Soft.*, 3:726, 2018.
- [23] R. R. Curtin, B. Moseley, H. Q. Ngo, X. Nguyen, D. Olteanu, and M. Schleich. Rk-means: Fast clustering for relational data. In *AISTATS*, volume 108 of *Proc. ML Res.*, pages 2742–2752, 2020.
- [24] A. Darwiche. Three modern roles for logic in AI. In *PODS*, pages 229–243, 2020.
- [25] I. Elghandour, A. Kara, D. Olteanu, and S. Vansummeren. Incremental techniques for large-scale dynamic query processing. In *CIKM*, pages 2297–2298, 2018. Tutorial.
- [26] X. Feng, A. Kumar, B. Recht, and C. Ré. Towards a unified architecture for in-rdbms analytics. In *SIGMOD*, pages 325–336, 2012.
- [27] T. J. Green and V. Tannen. The semiring framework for database provenance. In *PODS*, pages 93–99, 2017.
- [28] M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017.
- [29] M. Grohe and D. Marx. Constraint solving via fractional edge covers. In *SODA*, pages 289–298, 2006.
- [30] T. Grust. Accelerating XPath location steps. In *SIGMOD*, pages 109–120, 2002.
- [31] P. J. Haas and J. M. Hellerstein. Ripple joins for online aggregation. In *SIGMOD*, pages 287–298, 1999.
- [32] J. M. Hellerstein and et al. The MADlib analytics library or MAD skills, the SQL. *PVLDB*, 5(12):1700–1711, 2012.
- [33] G.-R. Inelus. Quadratically Regularised Principal Component Analysis over multi-relational databases, 2019. MSc thesis, University of Oxford.
- [34] Kaggle. The State of Data Science and Machine Learning, 2017. <https://www.kaggle.com/surveys/2017>.
- [35] A. Kara, H. Q. Ngo, M. Nikolic, D. Olteanu, and H. Zhang. Counting triangles under updates in worst-case optimal time. In *ICDT*, pages 4:1–4:18, 2019.
- [36] A. Kara, H. Q. Ngo, M. Nikolic, D. Olteanu, and H. Zhang. Maintaining triangle queries under updates. *ACM Trans. Database Syst.*, Accepted April 2020.
- [37] A. Kara, M. Nikolic, D. Olteanu, and H. Zhang. Trade-offs in static and dynamic evaluation of hierarchical queries. In *PODS*, pages 375–392, 2020.
- [38] C. Koch. Incremental query evaluation in a ring of databases. In *PODS*, pages 87–98, 2010.
- [39] C. Koch, Y. Ahmad, O. Kennedy, M. Nikolic, A. Nötzli, D. Lupei, and A. Shaikhha. Dbtoaster: higher-order delta processing for dynamic, frequently fresh views. *VLDB J.*, 23(2):253–278, 2014.
- [40] A. Kumar, M. Boehm, and J. Yang. Data management in machine learning: Challenges, techniques, and



- systems. In *SIGMOD*, pages 1717–1722, 2017.
- [41] A. Kumar, J. F. Naughton, and J. M. Patel. Learning generalized linear models over normalized data. In *SIGMOD*, pages 1969–1984, 2015.
- [42] F. Li, B. Wu, K. Yi, and Z. Zhao. Wander join and XDB: online aggregation via random walks. *ACM Trans. Database Syst.*, 44(1):2:1–2:41, 2019.
- [43] D. Marx. Approximating fractional hypertree width. *ACM Trans. Alg.*, 6(2):29:1–29:17, 2010.
- [44] D. Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM*, 60(6):42:1–42:51, 2013.
- [45] W. McKinney. Pandas: a foundational python library for data analysis and statistics. *Python for High Performance and Scientific Computing*, 14, 2011.
- [46] F. McSherry, M. Isard, and D. G. Murray. Scalability! but at what cost? In *HotOS*, 2015.
- [47] X. Meng, J. Bradley, et al. MLlib: Machine learning in Apache Spark. *J. Mach. Learn. Res.*, 17(1):1235–1241, 2016.
- [48] K. P. Murphy. *Machine learning: a probabilistic perspective*. MIT Press, Cambridge, Mass., 2013.
- [49] T. Neumann. Efficiently Compiling Efficient Query Plans for Modern Hardware. *PVLDB*, 4(9):539–550, 2011.
- [50] H. Q. Ngo, E. Porat, C. Ré, and A. Rudra. Worst-case optimal join algorithms. In *PODS*, pages 37–48, 2012.
- [51] H. Q. Ngo, C. Ré, and A. Rudra. Skew strikes back: New developments in the theory of join algorithms. In *SIGMOD Rec.*, pages 5–16, 2013.
- [52] D. T. Nguyen, M. Aref, M. Bravenboer, G. Kollias, H. Q. Ngo, C. Ré, and A. Rudra. Join processing for graph patterns: An old dog with new tricks. *CoRR*, abs/1503.04169, 2015.
- [53] M. Nikolic and D. Olteanu. Incremental view maintenance with triple lock factorization benefits. In *SIGMOD*, pages 365–380, 2018.
- [54] M. Nikolic, H. Zhang, A. Kara, and D. Olteanu. F-IVM: learning over fast-evolving relational data. In *SIGMOD*, pages 2773–2776, 2020.
- [55] D. Olteanu and M. Schleich. F: Regression models over factorized views. *PVLDB*, 9(13):1573–1576, 2016.
- [56] D. Olteanu and M. Schleich. Factorized databases. *SIGMOD Rec.*, 45(2):5–16, 2016.
- [57] D. Olteanu and J. Závodný. Size bounds for factorised representations of query results. *ACM Trans. Database Syst.*, 40(1):2, 2015.
- [58] Y. Park, J. Qing, X. Shen, and B. Mozafari. Blinkml: Efficient maximum likelihood estimation with probabilistic guarantees. In *SIGMOD*, pages 1135–1152, 2019.
- [59] F. Pedregosa, G. Varoquaux, A. Gramfort, and et al. Scikit-learn: Machine learning in python. *J. Machine Learning Research*, 12:2825–2830, 2011.
- [60] C. Perlich and F. J. Provost. Distribution-based aggregation for relational learning with identifier attributes. *Mach. Learn.*, 62(1-2):65–105, 2006.
- [61] N. Polyzotis, S. Roy, S. E. Whang, and M. Zinkevich. Data management challenges in production machine learning. In *SIGMOD*, pages 1723–1726, 2017.
- [62] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Stat. Comp., [www.r-project.org](http://www.r-project.org), 2013.
- [63] S. Rendle. Factorization machines with libFM. *ACM Trans. Intell. Syst. Technol.*, 3(3):57:1–57:22, 2012.
- [64] M. Schleich and D. Olteanu. LMFAO: An engine for batches of group-by aggregates. *PVLDB*, 13(12):2945–2948, 2020.
- [65] M. Schleich, D. Olteanu, M. Abo Khamis, H. Q. Ngo, and X. Nguyen. A layered aggregate engine for analytics workloads. In *SIGMOD*, pages 1642–1659, 2019.
- [66] M. Schleich, D. Olteanu, and R. Ciucanu. Learning linear regression models over factorized joins. In *SIGMOD*, pages 3–18, 2016.
- [67] M. Schleich, D. Olteanu, M. A. Khamis, H. Q. Ngo, and X. Nguyen. Learning models over relational data: A brief tutorial. In *SUM*, volume 11940, pages 423–432, 2019.
- [68] A. Shaikhha, Y. Klonatos, and C. Koch. Building efficient query engines in a high-level language. *ACM Trans. Database Syst.*, 43(1):4:1–4:45, 2018.
- [69] A. Shaikhha, Y. Klonatos, L. Parreaux, L. Brown, M. Dashti, and C. Koch. How to architect a query compiler. In *SIGMOD*, pages 1907–1922, 2016.
- [70] A. Shaikhha, M. Schleich, A. Ghita, and D. Olteanu. Multi-layer optimizations for end-to-end data analytics. In *CGO*, pages 145–157, 2020.
- [71] D. Suci, D. Olteanu, C. Ré, and C. Koch. *Probabilistic Databases*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2011.
- [72] The StatsModels development team. StatsModels: Statistics in Python, <http://statsmodels.sourceforge.net>, 2012.
- [73] B. van Geffen. QR decomposition of normalised relational data, 2018. MSc thesis, University of Oxford.
- [74] T. L. Veldhuizen. Leapfrog triejoin: A simple, worst-case optimal join algorithm. In *ICDT*, pages 96–106, 2014.
- [75] G. Xiao, D. Calvanese, R. Kontchakov, D. Lembo, A. Poggi, R. Rosati, and M. Zakharyashev. Ontology-based data access: A survey. In *IJCAI*, pages 5511–5519, 2018.
- [76] M. Zaharia, M. Chowdhury, et al. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *NSDI*, pages 2–2, 2012.