

A Layered Aggregate Engine for Analytics Workloads

Maximilian Schleich
University of Oxford

Dan Olteanu
University of Oxford

Mahmoud Abo Khamis
relationalAI

Hung Q. Ngo
relationalAI

XuanLong Nguyen
University of Michigan

ABSTRACT

This paper introduces LMFAO (Layered Multiple Functional Aggregate Optimization), an in-memory optimization and execution engine for batches of aggregates over the input database. The primary motivation for this work stems from the observation that for a variety of analytics over databases, their data-intensive tasks can be decomposed into group-by aggregates over the join of the input database relations. We exemplify the versatility and competitiveness of LMFAO for a handful of widely used analytics: learning ridge linear regression, classification trees, regression trees, and the structure of Bayesian networks using Chow-Liu trees; and data cubes used for exploration in data warehousing.

LMFAO consists of several layers of logical and code optimizations that systematically exploit sharing of computation, parallelism, and code specialization.

We conducted two types of performance benchmarks. In experiments with four datasets, LMFAO outperforms by several orders of magnitude on one hand, a commercial database system and MonetDB for computing batches of aggregates, and on the other hand, TensorFlow, Scikit, R, and AC/DC for learning a variety of models over databases.

*Aggregation is the aspirin to all problems.
– contemporary Greek philosopher*

ACM Reference Format:

Maximilian Schleich, Dan Olteanu, Mahmoud Abo Khamis, Hung Q. Ngo, and XuanLong Nguyen. 2019. A Layered Aggregate Engine for Analytics Workloads. In *2019 International Conference on Management of Data (SIGMOD '19)*, June 30–July 5, 2019, Amsterdam, Netherlands. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3299869.3324961>

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.
SIGMOD '19, June 30–July 5, 2019, Amsterdam, Netherlands
© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5643-5/19/06...\$15.00
<https://doi.org/10.1145/3299869.3324961>

1 INTRODUCTION

This work has its root in two observations. First, the majority of practical analytics tasks involve relational data, with the banking or retail domains exceeding 80% [27]. Second, for a variety of such analytics tasks, their data-intensive computation can be reformulated as batches of group-by aggregates over the join of the database relations [4, 46].

We introduce LMFAO (Layered Multiple Functional Aggregate Optimization), an in-memory optimization and execution engine for batches of aggregates over relational data. We exemplify the versatility and competitiveness of LMFAO for a handful of widely used analytics: learning ridge linear regression, classification trees, regression trees, and the structure of Bayesian networks using Chow-Liu trees; and data cubes used for exploration in data warehousing.

Query processing lies at the core of database research, with four decades of innovation and engineering on query engines for relational databases. Without doubt, the efficient computation of a handful of group-by aggregates over a join is well-supported by mature academic and commercial systems and also widely researched. There is relatively less development for large batches of such queries, with initial work in the context of data cubes [22, 24, 37] and SQL-aware data mining systems [12, 13] from two decades ago.

We show that by designing for the workload required by analytics tasks, LMFAO can outperform general-purpose mature database systems such as PostgreSQL, MonetDB, and a commercial database system by orders of magnitude. This is not only a matter of query optimization, but also of execution. Aspects of LMFAO's optimized execution for query batches can be cast in SQL and fed to a database system. Such SQL queries capture decomposition of aggregates into components that can be pushed past joins and shared across aggregates, and as such they may create additional intermediate aggregates. This poses scalability problems to these systems due to, e.g., design limitations such as the maximum number of columns or lack of efficient query batch processing, and led to larger compute times than for the plain unoptimized queries. This hints at LMFAO's distinct design that departs from mainstream query processing.

The performance advantage brought by LMFAO's design becomes even more apparent for the end-to-end applications.

For the aforementioned use cases in machine learning, the application layer takes relatively insignificant time as it offloads all data-intensive computation to LMFAO. LMFAO computes from the input database sufficient statistics whose size ranges from hundreds of KBs to tens of MBs (Table 2) and that are used for learning regression and classification models. Mainstream solutions, e.g., MADlib [25], R [45], Scikit-learn [42], and TensorFlow [1], either take orders of magnitude more time than LMFAO to train the same model or do not work due to various design limitations. These solutions use data systems to materialize the training dataset, which is defined by a feature extraction query over a database of multiple relations, and ML libraries to learn models over this dataset. We confirm experimentally that the main bottleneck of these solutions is this materialization: The training datasets can be an order of magnitude larger than the input databases used to create them (Table 1). In addition to being expected to work on much larger inputs, the ML libraries are less scalable than the data systems. These solutions also inherit the limitations of both of their underlying systems, e.g., the maximum data frame size in R and number of columns in PostgreSQL are much less than typical database sizes and respectively number of model features.

1.1 Problem Statement

LMFAO evaluates batches of queries of the following form:

```
SELECT  F1, ..., Ff, SUM(α1), ..., SUM(αℓ)
FROM    R1 NATURAL JOIN ... NATURAL JOIN Rm
GROUP BY F1, ..., Ff;
```

The user-defined aggregate functions (UDAFs), or simply aggregates, $\alpha_1, \dots, \alpha_\ell$ can be sums of products of functions:

$$\forall i \in [\ell] : \quad \alpha_i = \sum_{j \in [s_i]} \prod_{k \in [p_{ij}]} f_{ijk}, \text{ where } s_i, p_{ij} \in \mathbb{N}$$

We next give examples of such aggregates. To express count and sum aggregates, i.e., $\text{SUM}(1)$ and $\text{SUM}(X_1)$ for some attribute X_1 , we take $s_i = p_{ij} = 1$ and then f_{i11} is the constant function $f_{i11}() = 1$ and respectively the identity function $f_{i11}(X_1) = X_1$. To encode a selection condition X_1 op t that defines a decision tree node, where X_1 is an attribute, op is a binary operator, and t is a value in the domain of X_1 , we use the Kronecker delta $f_{i11}(X_1) = \mathbf{1}_{X_1 \text{ op } t}$, which evaluates to 1 if the condition is true and 0 otherwise. A further example is given by $s_i = n, p_{ij} = 2$, and for $j \in [n]$ the constant functions $f_{ij1}() = \theta_j$ and the identity functions $f_{ij2}(X_j) = X_j$. Then, α_i is $\sum_{j \in [n]} \theta_j \cdot X_j$ and captures the linear regression function with parameters θ_j and features X_j . A final example is that of an exponential n -ary function $f_{i11}(X_1, \dots, X_n) = e^{\sum_{j \in [n]} \theta_j \cdot X_j}$, which is used for logistic regression.

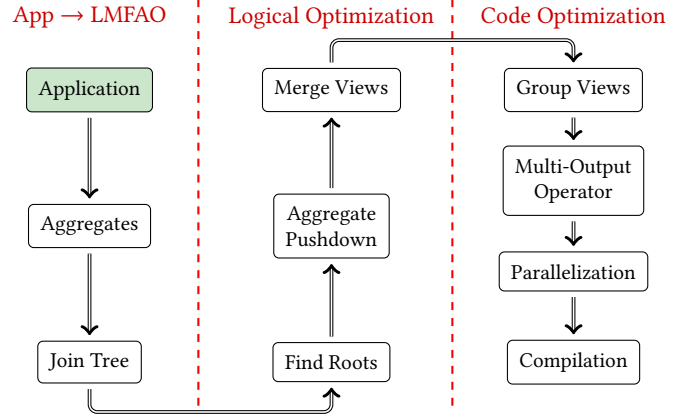


Figure 1: The Optimization Layers of LMFAO.

Applications, such as those in Section 2, may generate batches of tens to thousands of aggregates (Table 2). They share the same join of database relations and possibly of relations defined by queries of the same above form.

1.2 The Layers of LMFAO

To evaluate aggregate batches, LMFAO employs a host of techniques, either novel or adaptations of known concepts to our specific workload. The layered architecture of LMFAO is given in Figure 1 and highlighted next. Section 3 expands on the key design choices behind LMFAO.

The **Join Tree** layer takes as input the batch of aggregates, the database schema, and cardinality constraints (e.g., sizes of relations and attribute domains) and produces one join tree that is used to compute all aggregates. This step uses state-of-the-art techniques¹ [3].

The **Find Roots** layer is novel and affects the design of all subsequent layers. By default, LMFAO computes each group-by aggregate in one bottom-up pass over the join tree, by decomposing the aggregate into views computed along each edge in the join tree. We allow for different traversals of the join tree: different aggregates may be computed over the same join tree rooted at different nodes. This can reduce the overall compute time for the batch as it can reduce the number of views and increase the sharing of their computation. In our experiments, the use of multiple roots for the computation of aggregate batches led to an up to 5× speedup.

LMFAO uses *directional views* to support different traversals of the join tree: For each edge between two nodes, there may be views flowing in both directions. Directional views are similar in spirit with messages in the message passing algorithm used for inference in graphical models [41]. Figure 3 (middle) depicts directional views along a join tree.

¹For cyclic queries, we first compute a hypertree decomposition and materialize its bags (cycles) to obtain a join tree.

In the **Aggregate Pushdown** layer, each aggregate is decomposed into one directional view per edge in the join tree. Its view at an edge going out of a node n computes the aggregate when restricted to the subtree rooted at n and is defined over the join of the views at the incoming edges of n and of the relation at n . The directions of these views are from the leaves to the root of the aggregate. The rationale for this decomposition is twofold. First, it partially pushes the aggregates past joins (represented by edges in the tree), as in prior work on eager computation of one aggregate [55] and its generalization to factorized databases [8]. Second, it allows for sharing common views across the aggregates.

The **Merge Views** layer consolidates the views generated in the previous layer. There are three types of merging possible for views along the same edge in the join tree, depending on whether they have the same: group-by attributes; aggregates; and body. Views with the same direction are defined over the same subtree of the join tree. We first identify identical views constructed for different aggregates and only keep one copy. We then merge the views with the same group-by attributes and body but different aggregates. We finally merge views with the same group-by attributes and different bodies. This consolidation is beneficial. For instance, there are 814 aggregates to compute for learning a linear regression model over the join of five relations in our Retailer dataset. This amounts to $814 \text{ aggregates} \times 4 \text{ edges} = 3,256$ views, which are consolidated into 34 views that have between themselves 1,468 aggregates.

The previous three layers are concerned with logical transformations of view expressions. The remaining four layers consider optimizations not expressible at the syntactic level.

In the **Group Views** layer, we group the views going out of the same node possibly along different edges such that there is no dependency between them. No dependency means that they can be evaluated together once the incoming views used in their joins are computed. The views in a group do not necessarily have the same group-by attributes, so a view group has multiple outputs. To continue our example, the remaining 34 views are clustered into 7 groups.

The view group is a computational unit in LMFAO. At the **Multi-Output Operator** layer, each view group at a node is computed in one pass over the relation at that node, with lookups using the join keys into the incoming views to fetch aggregates needed for the computation of the views in the group. This is yet another instance of sharing in LMFAO: The computation of different views share the scan of the relation at the node. This is particularly beneficial for fact relations in star and snowflake schemas as they can be very large (as it is the case for Inventory in Retailer and Sales in Favorita datasets). This scan sees the relation organized logically as a trie: first grouped by one attribute, then by the next in the

context of values for the first, and so on. This trie iteration of a relation is reminiscent of factorized databases [9] and LeapFrog TrieJoin [53] and can visit up to three times less values than a standard row-based scan for our datasets. In our experiments, this layer brought up to $2\times$ extra speedup.

The **Parallelization** layer addresses task and domain parallelism. LMFAO parallelizes the computation of multi-output operators for view groups that do not depend on each other. For this, it computes the dependency graph of the view groups. LMFAO partitions the large (fact) input relations and allocates a thread per partition to compute the multi-output operator on that partition. This layer brought $2 - 3\times$ extra speedup on a machine with four vCPUs (AWS d2.xlarge).

Finally, the **Compilation** layer generates C++ code for the parallel execution of multi-output operators that is specialized to the join tree and database schema, with separate code for each view group and also for general tasks such as data loading. The specialized code chunks are compiled in parallel. LMFAO uses an intermediate representation for this code and applies various optimizations, including: inlining function calls; organization of aggregate computation to minimize the number of updates to each aggregate; organization of the aggregates for each view in a contiguous fixed-size array and ordered to allow sequential read/write; reuse of arithmetic operations, such as repeating multiplication of two entries in the aggregate array; and synthesis of loops from long sequences of lockstep aggregate computations. The latter two optimizations are enabled by sorted input relations and views that are accessed in lockstep. The organization of aggregates allows us to manage them in batch. This is reminiscent of vectorization [58], now applied to aggregates instead of data records. LMFAO generates C++ code out of the optimized intermediate representation. This layer brought up to $15\times$ extra speedup.

An additional complexity is brought by applications that require the computation of UDAFs that change between iterations depending on the outcome of computation. This is the case for decision trees whose nodes are iteratively constructed in the context of conditions that are selected based on the data. The application tags these functions as *dynamic* to instruct LMFAO to avoid inlining their calls and instead generate separate code that is compiled between iterations and linked dynamically.

1.3 Contributions

To sum up, the contributions of this work are as follows:

1. We introduce LMFAO, a principled layered approach to computing large batches of group-by aggregates over joins. Its layers encompass several stages of logical and code optimization that come with novel contributions as well as adaptations of known techniques to a novel setting. The

novel contributions are on: using different traversals of the same join tree to solve many aggregates with different group-by clauses; synthesizing directional views out of large sets of views representing components of aggregate queries; and the multi-output operator for computing groups of directional views using one pass over the input relations. It adapts compilation techniques to generate specialized code for the parallel computation of multi-output operators for aggregate queries with static and dynamic user-defined functions.

2. We show the versatility of LMFAO for a range of analytics applications built on top of it.

3. We implemented LMFAO in C++ and conducted two kinds of performance benchmarks: The computation of batch aggregates and of the complete applications requiring these aggregates. In experiments with four datasets, LMFAO outperforms by several orders of magnitude on one hand, PostgreSQL, MonetDB and a commercial DBMS for computing batches of aggregates, and, on the other hand, TensorFlow, Scikit, R, and AC/DC for learning models over databases.

2 APPLICATIONS

LMFAO encompasses a unified formulation and processing of core data processing tasks in database, data mining, and machine learning problems. We exemplify with a small sample of such problems: data cubes; gradients and covariance matrices used for linear regression, polynomial regression, factorization machines; classification and regression trees; and mutual information of pairwise variables used for learning the structure of Bayesian networks.

We next introduce a compact query syntax and use it to formulate the above-mentioned data processing tasks.

Query Language. We are given a database D of m (materialized) relations R_1, \dots, R_m over relation schemas $\omega_{R_1}, \dots, \omega_{R_m}$. For convenience, we see relation schemas, which are lists of attributes, also as sets. The list of attributes in the database is denoted by $\mathbf{X} = \bigcup_{r \in [m]} \omega_{R_r} = (X_1, \dots, X_n, X_{n+1})$.

We would like to compute a set of group-by aggregates over the natural join of these relations. This join may represent the training dataset for machine learning models, the input for multi-dimensional data cubes, or the joint probability distribution to be approximated by a Bayesian network.

We use the following query formulation, which is more compact than the SQL form from Section 1.1:

$$Q(F_1, \dots, F_f; \alpha_1, \dots, \alpha_\ell) += R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (1)$$

In the head of Q , the group-by attributes F_1, \dots, F_f are separated from the aggregate functions by semicolon; we omit the semicolon if there are no group-by attributes. The aggregate functions are as defined in Section 1.1. We use $+=$ to capture the SUM over each aggregate function. In the query body, we make explicit the attributes of each relation for

a clearer understanding of the definitions of the aggregate functions. By definition, there is a functional dependency $F_1, \dots, F_f \rightarrow \alpha_1, \dots, \alpha_\ell$.

Our queries generalize FAQ-SS [6] and MPF (Marginalize a Product Function) [7] by allowing tuples of arbitrary UDAFs.

Ridge Linear Regression. Assuming one parameter θ_j per attribute (feature) X_j , the linear regression model is given by:

$$LR(X_1, \dots, X_n) = \sum_{j \in [n]} \theta_j \cdot X_j$$

In practice, features may be defined by attributes in both the input relations and results of queries over these relations.

We assume without loss of generality that (1) X_1 only takes value 1 and then θ_1 is the so-called intercept and (2) X_{n+1} is the label and has a corresponding new parameter $\theta_{n+1} = -1$.

The error of the model is given by an objective function that is the sum of the least squares loss function and of the penalty term that is the ℓ_2 norm of the parameter vector θ :

$$\begin{aligned} J(\theta) &= \frac{1}{2|D|} \sum_{\mathbf{X} \in D} \left(\sum_{j \in [n]} \theta_j \cdot X_j - X_{n+1} \right)^2 + \frac{\lambda}{2} \|\theta\|^2 \\ &= \frac{1}{2|D|} \sum_{\mathbf{X} \in D} \left(\sum_{j \in [n+1]} \theta_j \cdot X_j \right)^2 + \frac{\lambda}{2} \|\theta\|^2 \end{aligned}$$

We optimize the model using batch gradient descent (BGD), which updates the parameters in the direction of the gradient vector $\nabla J(\theta)$ of $J(\theta)$ using a step size s :

repeat until convergence:

$$\forall k \in [n] : \theta_k := \theta_k - s \cdot \nabla_k J(\theta)$$

$$:= \theta_k - s \cdot \left(\frac{1}{|D|} \sum_{\mathbf{X} \in D} \left(\sum_{j \in [n+1]} \theta_j \cdot X_j \right) \cdot X_k + \lambda \theta_k \right)$$

The above update relies on the aggregates for the size of the dataset D and the product of X_k with the inner product $\langle \theta, \mathbf{X} \rangle = \sum_{j \in [n+1]} \theta_j \cdot X_j$. There are two ways to express these aggregates. The common approach, which we call the gradient vector, is to compute this inner product and then, for each gradient k , multiply it with the corresponding X_k . This requires recomputation for each new vector of parameters. The second approach [46] is to rewrite $\sum_{\mathbf{X} \in D} (\sum_{j \in [n+1]} \theta_j \cdot X_j) \cdot X_k$ as $\sum_{j \in [n+1]} \theta_j \cdot \sum_{\mathbf{X} \in D} (X_j \cdot X_k)$ and compute the non-centered covariance matrix (the covar matrix hereafter).

The covar matrix accounts for all pairwise multiplications $X_j \cdot X_k$. Each entry can be computed as aggregate query:

$$Covar_{j,k}(X_j \cdot X_k) += R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}). \quad (2)$$

Categorical attributes are one-hot encoded in a linear regression model. In our formalism, such attributes become group-by attributes. If only X_j is categorical, we get:

$$Covar_{j,k}(X_j; X_k) += R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (3)$$

If both X_j and X_k are categorical, we get instead:

$$Covar_{j,k}(X_j, X_k; 1) += R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (4)$$

The computation of covar does not depend on the parameters θ , and can be done once for all BGD iterations.

Higher-degree Regression Models. A polynomial regression models of degree d is defined as follows:

$$PR_d(X_1, \dots, X_n) = \sum_{(a_1, \dots, a_n) \in \mathbf{A}} \theta_{(a_1, \dots, a_n)} \cdot \prod_{j=1}^n X_j^{a_j}, \text{ where} \\ \mathbf{A} = \{(a_1, \dots, a_n) \mid a_1 + \dots + a_n \leq d, \forall j \in [n] : a_j \in \mathbb{N}\}$$

The covar matrix for PR_d has the following aggregates in the gradient of the square loss function:

$$\text{Covar}_{(a_1, \dots, a_{n+1})}(X_1^{a_1} \dots X_{n+1}^{a_{n+1}}) = R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \\ \forall (a_1, \dots, a_{n+1}) : \sum_{j=1}^{n+1} a_j \leq 2d, a_{n+1} \leq 1, \forall j \in [n+1] : a_j \in \mathbb{N} \quad (5)$$

A similar generalization works for factorization machines [5, 40]. Categorical attributes can be accommodated as for linear regression and then each categorical attribute X_j with exponent $a_j > 0$ becomes a group-by attribute.

Data Cubes. Data cubes [22] are popular in data warehousing scenarios. For a set $S_k \subseteq \mathbf{X}$ of k attributes or dimensions, a k -dimensional data cube is a shorthand for the union of 2^k cube aggregates with the same aggregation function α over the same (measure) attribute out of v attributes. We define one aggregate for each of the 2^k possible subsets of S_k :

$$\forall F_i \subseteq S_k : \text{Cube}_i(F_i; \alpha_1, \dots, \alpha_v) = R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (6)$$

The cube aggregates have a similar structure with covar matrices for polynomial regression models. They both represent sets of group-by aggregates over the same join. However, the two constructs compute different aggregates and use different data representations. Whereas all cube aggregates use the same measure aggregation, the covar aggregates sum over different attributes. Data cubes are represented as tables in 1NF using a special ALL value standing for a set of values, whereas the covar matrices for regression models are matrices whose entries are the regression aggregates whose outputs have varying sizes and arities.

A polynomial regression model of degree d (PR_d) over categorical features given by k attributes ($k \geq 2d$) requires regression aggregates whose group-by clauses are over all subsets of size at most $2d$ of the set of k attributes. In contrast, a $2d$ -dimensional data cube for a given set of $2d$ (dimension) attributes defines aggregates whose group-by clauses are over all subsets of the $2d$ attributes. The set of group-by clauses used by the aggregates for PR_d is captured by all $2d$ -dimensional data cubes constructed using the k attributes.

Mutual Information. The mutual information of two distinct discrete random variables X_i and X_j is a measure of their mutual dependence and determines how similar the joint distribution is to the factored marginal distribution. In our database setting, we capture the distributions of two attributes

X_i and X_j using the following count queries that group by any subset of $\{X_i, X_j\}$ (thus expressible as a 2-dimensional data cube with a count measure):

$$\forall S \subseteq \{i, j\} : Q_S((X_k)_{k \in S}; 1) = R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (7)$$

The mutual information of X_i and X_j is the given by the following query with a 4-ary aggregate function f over the aggregates of the queries Q_S defined above:

$$MI(f(\alpha, \beta, \gamma, \delta)) = Q_{\emptyset}(\alpha), Q_{\{i\}}(X_i; \beta), Q_{\{j\}}(X_j; \gamma), Q_{\{i,j\}}(X_i, X_j; \delta) \\ f(\alpha, \beta, \gamma, \delta) = \frac{\delta}{\alpha} \cdot \log \left(\frac{\alpha \cdot \delta}{\beta \cdot \gamma} \right)$$

Mutual information has many applications as it is used: as cost function in learning decision trees; in determining the similarity of two different clusterings of a dataset; as criterion for feature selection; in learning the structure of Bayesian networks. The Chow-Liu algorithm [16] constructs an optimal tree-shaped Bayesian network T with one node for each input attribute in the set \mathbf{X} . It proceeds in rounds and in each round it adds to T an edge (X_i, X_j) between the nodes X_i and X_j such that the mutual information of X_i and X_j is maximal among all pairs of attributes not chosen yet.

Classification and Regression Trees. Decision trees are popular machine learning models that use trees with inner nodes representing conditional control statements to model decisions and their consequences. Leaf nodes represent predictions for the label. If the label is continuous, we learn a regression tree and the prediction is the average of the label values in the fragment of the training dataset that satisfies all control statements on the root to leaf path. If the label is categorical, the tree is a classification tree, and the prediction is the most likely category for the label in the dataset fragment. Figure 2 shows an example of a regression tree.

The CART algorithm [11] constructs the tree one node at a time. Given an input dataset D , CART repeatedly finds a condition $X_j \text{ op } t$ on one of the attributes X_1, \dots, X_n of D that splits D so that a given cost function over the label X_{n+1} is minimized. For categorical attributes (e.g., city), t may be a set of categories and op denotes inclusion. For continuous attributes (e.g., age), t is a real number and op is inequality. Once this condition is found, a new node $X_j \text{ op } t$ is constructed and the algorithm proceeds recursively to construct the subtree rooted at this node for the dataset representing the fragment of D satisfying the conditions at the new node and at its ancestors.

Practical implementations of CART compute at each node the cost for 20-100 conditions per continuous attribute and for categorical attributes the best subset of categories is chosen based on the cost of splitting on each individual category.

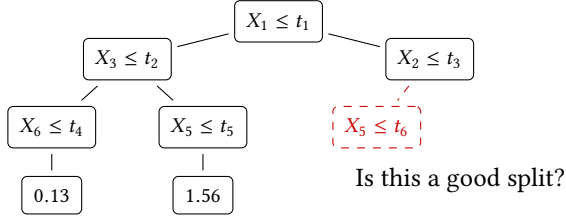


Figure 2: Example of a regression tree. Classification trees replace numerical leaves by categorical values.

For regression trees, the cost is given by the variance:

$$\text{variance} = \sum_{X \in D_i} X_{n+1}^2 - \frac{1}{|D_i|} \left(\sum_{X \in D_i} X_{n+1} \right)^2$$

computed over the fragment D_i of the dataset D . For the tree depicted in Figure 2, $D_i = \sigma_{X_1 \geq t_1 \wedge X_2 \leq t_3 \wedge X_5 \leq t_6}(D)$, where $X_5 \leq t_6$ is the new condition for which we compute the cost of the split in the context of the conjunction of conditions $X_1 \geq t_1 \wedge X_2 \leq t_3$ representing its ancestors in the tree. The computation of this cost needs the aggregates $\text{COUNT}()$, $\text{SUM}(X_{n+1})$, and $\text{SUM}(X_{n+1}^2)$ over D_i :

$$RT(1 \cdot \alpha, X_{n+1} \cdot \alpha, X_{n+1}^2 \cdot \alpha) += R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (8)$$

where $\alpha = \mathbf{1}_{X_1 \geq t_1} \cdot \mathbf{1}_{X_2 \leq t_3} \cdot \mathbf{1}_{X_5 \leq t_6}$

The product aggregate α evaluates to 1 whenever all conditions in the subscript are satisfied and to 0 otherwise.

For classification trees, the label X_{n+1} has a set $\text{Dom}(X_{n+1})$ of categories. The cost is given by the entropy or Gini index:

$$\text{entropy} = - \sum_{k \in \text{Dom}(X_{n+1})} \pi_k \log(\pi_k) \quad \text{gini} = 1 - \sum_{k \in \text{Dom}(X_{n+1})} \pi_k^2$$

The aggregates π_k for $k \in \text{Dom}(X_{n+1}) = \{k_1, \dots, k_p\}$ compute the frequencies of each category k for the label X_{n+1} in the dataset D_i , i.e., for category k this frequency is the fraction of the tuples in D_i where $X_{n+1} = k$ and of all tuples in D_i : $\frac{1}{|D_i|} \sum_{X \in D_i} \mathbf{1}_{X_{n+1}=k}$. These frequencies can all be computed with the following two aggregate queries:

$$CT(X_{n+1}; \alpha) += R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (9)$$

$$CT(\alpha) += R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m}) \quad (10)$$

Applications need a large number of aggregates. The number of aggregates in a batch is a function of the number n of attributes in the database: $\frac{1}{2}(n+1)(n+2)$ for linear regression; $\frac{1}{2} \left[\binom{n+d}{d}^2 + \binom{n+d}{d} \right]$ for polynomial regression of degree d ; $2^d v$ for d -dimensional data cubes with v measures; $\frac{n(n-1)}{2}$ for Chow-Liu trees with n nodes; and $dn(p+1)c$ for classification/regression trees with d nodes where c conditions are tried per attribute and the response has p categories in case of classification tree; the formula for regression tree is obtained with $p = 2$. Table 2 gives the number of aggregates for these applications and our four datasets, whose details are in Table 1. This number ranges from tens to tens of thousands.

Further Applications. Virtually any in-database machine learning setting can benefit from an efficient processor for aggregate batches over joins. Although not reported in this work, we also investigated SVM, k-means clustering, and low-rank models such as quadratically regularized PCA and Non-Negative Matrix Factorization, as well as linear algebra operations such as QR and SVD decompositions of matrices defined by the natural join of database relations. All these applications decompose into batches of aggregates of a similar form to those mentioned here.

3 THE LMFAO ENGINE

In this section we discuss key design choices for LMFAO and motivate them using examples. Section 1 already provided an overview of its layers that are depicted in Figure 1.

3.1 Aggregates over Join Trees

LMFAO evaluates a batch of aggregate queries of the form (1) over a join tree of the database schema, or equivalently of the natural join of the database relations. We next recall the notion of join trees and exemplify the evaluation of aggregates over join trees by decomposing them into views.

The *join tree* of the natural join of the database relations $R_1(\omega_{R_1}), \dots, R_m(\omega_{R_m})$ is an undirected tree T such that [3]:

- The set of nodes of T is $\{R_1, \dots, R_m\}$.
- For every pair of nodes R_i and R_j , their common attributes are in the schema of every node R_k along the distinct path from R_i to R_j , i.e., $\omega_{R_i} \cap \omega_{R_j} \subseteq \omega_{R_k}$.

Figure 3 shows a possible join tree for the natural join of the six relations in the Favorita dataset [18] (details are given in Appendix A). Instead of showing join attributes on the edges, we underline them in the schema (left) to avoid clutter.

Acyclic joins always admit join trees. Arbitrary joins are transformed into acyclic ones by means of hypertree decompositions and materialization of their nodes (called bags) using worst-case optimal join algorithms [33, 53]. We next exemplify the computation of aggregates over a join tree [6, 8].

Example 3.1. Let us compute the sum of the product of two aggregate functions $f(\text{units})$ and $g(\text{price})$ over the natural join of the Favorita relations:

$$Q_1(f(\text{units}) \cdot g(\text{price})) += S(\omega_S), T(\omega_T), R(\omega_R), O(\omega_O), H(\omega_H), I(\omega_I)$$

We abbreviated the names of the Favorita relations as highlighted in Figure 3. The aggregate functions f and g are over the attributes units in Sales and price in Oil. We can rewrite Q_1 to push these functions down to the relations and also

Sales: date, store, item, units, promo
Holidays: date, htype, locale, transferred
StoRes: store, city, state, stype, cluster
Items: item, family, class, perishable
Transactions: date, store, txns
Oil: date, price

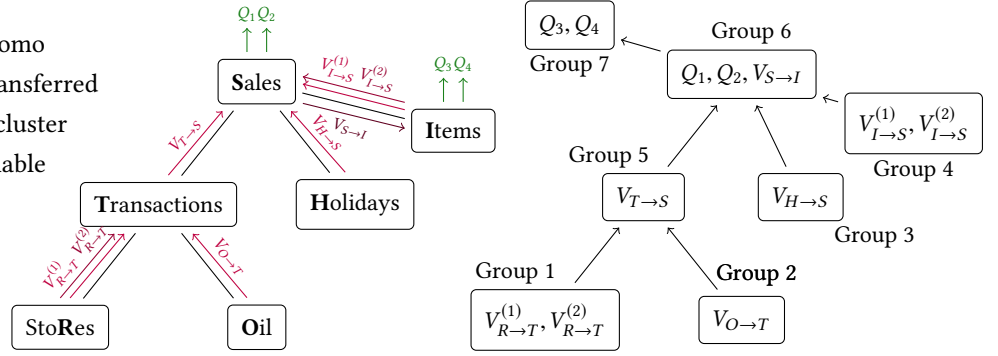


Figure 3: (left) The schema for the Favorita dataset. (middle) A join tree for this schema with directional views and four queries, partitioned in 7 groups. (right) The dependency graph of the view groups.

follow the structure of the join tree in Figure 3:

$$\begin{aligned}
 V_O(\text{date}; g(\text{price})) &+= O(\text{date}, \text{price}) \\
 V_R(\text{store}; 1) &+= R(\text{store}, \text{city}, \text{state}, \text{stype}, \text{cluster}) \\
 V_T(\text{date}, \text{store}; c \cdot p) &+= T(\text{date}, \text{store}, t), V_R(\text{store}; c), V_O(\text{date}; p) \\
 V_H(\text{date}; 1) &+= H(\text{date}, \text{htype}, \text{locale}, \text{transferred}) \\
 V_I(\text{item}; 1) &+= I(\text{item}, \text{family}, \text{class}, \text{perishable}) \\
 Q_1(f(\text{units}) \cdot p \cdot c_1 \cdot c_2) &+= V_T(\text{date}, \text{store}; p), V_I(\text{item}; c_2), \\
 &\quad V_H(\text{date}; c_1), S(\text{date}, \text{store}, \text{item}, \text{units})
 \end{aligned}$$

Except for S and O , which have attributes in the aggregate functions, we only need to count the number of tuples with the same join key in each of the other relations. \square

The computation of several aggregates over the same join tree may share views between themselves.

Example 3.2. Consider now $Q_2(\text{family}; g(\text{price}))$ over the same join. This query reports the sum of $g(\text{price})$ for each item family. We can rewrite it similarly to Q_1 in Example 3.1:

$$\begin{aligned}
 V_I'(\text{family}, \text{item}; 1) &+= I(\text{item}, \text{family}, \text{class}, \text{perishable}) \\
 Q_2(\text{family}; p \cdot c_1 \cdot c_2) &+= V_T(\text{date}, \text{store}; p), V_I'(\text{family}, \text{item}; c_2), \\
 &\quad V_H(\text{date}; c_1), S(\text{date}, \text{store}, \text{item}, \text{units})
 \end{aligned}$$

We can share the views V_T , and thus its underlying views V_O and V_R , and V_H between Q_1 and Q_2 . \square

3.2 Directional Views over Join Trees

An alternative evaluation for Q_2 in Example 3.2 would not create the view V_I' and instead create a view $V_S(\text{item}; p)$ over the subtree rooted at Sales and then join it with Items in the body of Q_2 . This effectively means that we use the same join tree but rooted at different nodes: Sales for Q_1 and Items for Q_2 . This also means that the edge between Sales and Item has two views, yet with different direction.

To accommodate this evaluation approach, we introduce *directional views*: These are queries of the form (1) where we

also specify their direction. They flow along an edge from a source node to a neighboring target node and are computed using the relation at the source node and some of its incoming views. Examples 3.1 and 3.2 showed views whose directions are always towards the root Sales. The direction of V_I is $I \rightarrow S$ and of V_S for the alternative evaluation of Q_2 is $S \rightarrow I$.

Consider a join tree T with root S and children C_1, \dots, C_k , where child C_i is the root of a subtree T_i in T . We use ω_{C_i} and ω_{T_i} to denote the schema of the relation C_i in T and respectively the union of the schemas of all relations in T_i .

We decompose a query $Q(F; \alpha)$ with group-by attributes F and aggregate function α as follows:

$$Q(F; \alpha) += S(\omega_S), V_{C_1 \rightarrow S}(F_1; \alpha_1), \dots, V_{C_k \rightarrow S}(F_k; \alpha_k)$$

The view $V_{C_i \rightarrow S}(F_i; \alpha_i)$ for a child C_i of S is defined as the “projection” of Q onto T_i as follows. Its group-by attributes are $F_i = (F \cap \omega_{T_i}) \cup (\omega_S \cap \omega_{C_i})$; here, $\omega_S \cap \omega_{C_i}$ are the attributes shared between S and a child C_i and $F \cap \omega_{T_i}$ are the group-by attributes from F present in T_i . Its body is the natural join of the relations in T_i . If all attributes of α are (are not) in ω_{T_i} , then $\alpha_i = \alpha$ (respectively $\alpha_i = 1$). Otherwise, T_i has some of the attributes required to compute α , in which case we add them as group-by attributes, i.e., $F_i := F_i \cup (\omega_{T_i} \cap \omega_\alpha)$, and use the aggregate $\alpha_i = 1$ to count. We can now decompose the views $V_{C_i \rightarrow S}$ recursively as explained for Q .

Using different roots for different queries may lower the overall complexity of evaluating a batch of aggregates. At the same time, we would like to share computation as much as possible, which is intuitively maximized if all queries are computed at the same root. We next discuss our solution to this tension between complexity and sharing.

3.3 Each Aggregate to Its Own Root

We next exemplify the advantage of evaluating a batch of queries, which are common in linear regression and mutual information settings where all attributes are categorical, at

different roots in the join tree and then explain how to find a root for a given aggregate in a batch of aggregates.

Example 3.3. Consider the following count queries over the join of relations $S_k(X_k, X_{k+1})$ of size N , $\forall k \in [n-1]$:

$$\forall i \in [n] : Q_i(X_i; 1) \text{ += } S_1(X_1, X_2), \dots, S_{n-1}(X_{n-1}, X_n).$$

We first explain how to compute these n queries by decomposing them into directional views that are over the join tree $S_1 - S_2 - \dots - S_{n-1}$ with root S_1 and have the same direction along this path towards the root.

For simplicity, we denote by L_k^i the view constructed for Q_i with direction from S_k to S_{k-1} . The views are defined as follows, with the addition of $L_n^n(X_n, X_n; 1) \text{ += } \text{Dom}(X_n)$ that associates each value in the domain of X_n with 1.

$$\begin{aligned} \forall k \in [i-1] : L_k^i(X_k, X_i; c) & \text{ += } S_k(X_k, X_{k+1}), L_{k+1}^i(X_{k+1}, X_i; c) \\ \forall i \in [n-1] : L_i^i(X_i, X_i; c) & \text{ += } L_{i+1}^{i+1}(X_{i+1}, X_i; c) \\ \forall i \in [n-1] : Q_i(X_i; c) & \text{ += } L_1^i(X_1, X_i; c) \end{aligned}$$

The above decomposition proceeds as follows. Q_n counts the number of occurrences of each value for X_n in the join. We start with 1, as provided by L_n^n , and progress to the left towards the root S_1 . The view L_{n-1}^n computes the counts for X_n in the context of each value for X_{n-1} as obtained from S_{n-1} . We need to keep the values for X_{n-1} to connect with S_{n-2} . Eventually, we reach L_1^n that gives the counts for X_n in the context of X_1 , and we sum them over the values of X_1 . The same idea applies to any Q_i with one additional optimization: Instead of starting from the leaf S_{n-1} , we can jump-start at S_{i-1} and reuse the computation of the counts for X_{i+1} in the context of X_i as provided by L_{i+1}^{i+1} . We need $O(n^2)$ many views and those of them that have group-by attributes from two different relations take $O(N^2)$ time.

We can lower this complexity to $O(N)$ by using different roots for different queries. We show the effect of using the root S_i for query Q_i . For each query Q_i , we construct two directional views: view R_i from S_{i-1} to S_i (i.e., from left to right) and view L_i from S_{i+1} to S_i (i.e., from right to left). The counts for X_i values are the products of the counts in the left view L_i and the right view R_i :

$$\begin{aligned} \forall 1 \leq i < n : L_i(X_i; c) & \text{ += } S_i(X_i, X_{i+1}), L_{i+1}(X_{i+1}; c) \\ \forall 1 < i \leq n : R_i(X_i; c) & \text{ += } S_{i-1}(X_{i-1}, X_i), R_{i-1}(X_{i-1}; c) \\ \forall 1 \leq i \leq n : Q_i(X_i; c_1 \cdot c_2) & \text{ += } L_i(X_i; c_1), R_i(X_i; c_2) \end{aligned}$$

We also use two trivial views $R_1(X_1; 1) \text{ += } \text{Dom}(X_1)$ and $L_n(X_n; 1) \text{ += } \text{Dom}(X_n)$. Note how the left view L_i is expressed using the left view L_{i+1} coming from the node S_{i+1} below S_i . Similarly for the right views. Each of the $2n$ views takes linear time. Moreover, they share much more computation among them than the views L_k^i used in the first scenario.

The second approach that chooses the root S_i for query Q_i can also be used for queries over all pairs of attributes:

$$\forall i, j \in [n] : Q_{i,j}(X_i, X_j; 1) \text{ += } S_1(X_1, X_2), \dots, S_{n-1}(X_{n-1}, X_n).$$

Each of these n^2 queries takes time $O(N)$ for $|i-j| \leq 1$ and $O(N^2)$ otherwise. At each node S_i , we compute a left view $L_{i,j}$, for any $i < j \leq n$, that counts the number of tuples for (X_i, X_j) over the path $S_i - \dots - S_{n-1}$. Then, the overall count c in $Q_{i,j}(X_i, X_j; c)$ is computed as the product of the count for X_i given by the right view R_i and the count for (X_i, X_j) given by the left view $L_{i,j}$ ($\forall 1 \leq i < j < n$):

$$\begin{aligned} L_{i,j}(X_i, X_j; c) & \text{ += } S_i(X_i, X_{i+1}), L_{i+1,j}(X_{i+1}, X_j; c) \\ Q_{i,j}(X_i, X_j; c_1 \cdot c_2) & \text{ += } L_{i,j}(X_i, X_j; c_1), R_i(X_i; c_2) \end{aligned}$$

The trivial views $\forall i \in [n] : L_{i,i}(X_i, X_i; 1) \text{ += } \text{Dom}(X_i)$ assign a count of 1 to each value of X_i . \square

LMFAO chooses the root in a join tree for each query in a batch using a simple and effective approximation for the problem of minimizing the overall size of the views used to compute the entire batch. For each query Q in the batch, we assign a weight to each relation R in the join tree that is the fraction of the number of group-by attributes of Q in R ; if Q has no group-by attribute, then any relation is a possible root and we assign to each relation the same weight that is an equal fraction of the number of relations. At the end of this weight assignment phase, each relation will be assigned some weight. We assign roots in the reverse order of their weights. A relation with the largest weight is then assigned as root to all queries that considered it as possible root. We break ties by choosing a relation with the largest size. The rationale for this choice is threefold. The choice for the largest relation avoids the creation of possibly large views over it. If the root for Q has no group-by attribute of Q , then we will create views carrying around values for these attributes, so larger views. A root with a large weight ensures that many views share the same direction towards it, so their computation may be shared and they may be merged or grouped (as explained in the next sections).

3.4 Merging and Grouping Views

The views generated for a batch of aggregates can be consolidated or *merged* if they have in common: (1) only the group-by attributes and direction, (2) also the body, and (3) also the aggregates. The common case (3), which is also the most restrictive one, has been seen in Example 3.2: The same view is created for several queries, in which case we only need to compute it once. Case (2) concerns views with the same group-by attributes and join but different aggregates. Such views are merged into a single one that keeps the same group-by attributes and join but merges the lists of aggregates. Case (1) is the most general form of merging supported by LMFAO and consolidates the views into a new view that is a join of these views on their (same) group-by attributes. The reason why this merging is sound is twofold. First, these views are over the same join, so they have the same set of

tuples over their group-by attributes. Second, the aggregates are functionally determined by the group-by attributes.

Example 3.4. We continue Examples 3.1 and 3.2 and add a third count query $Q_3(\text{family}; h(\text{txns}, \text{city}))$ over the same join body as Q_1 and Q_2 . This is decomposed into the following views over the same join tree rooted at Sales:

$$\begin{aligned} V'_O(\text{date}; 1) &+= O(\text{date}, \text{price}) \\ V'_R(\text{store}, \text{city}; 1) &+= R(\text{store}, \text{city}, \text{state}, \text{stype}, \text{cluster}) \\ V'_T(\text{date}, \text{store}; h(\text{txns}, \text{city}) \cdot c_1 \cdot c_2) &+= V'_R(\text{store}, \text{city}; c_1), \\ &T(\text{date}, \text{store}, \text{txns}), V'_O(\text{date}; c_2) \\ Q_3(\text{family}; c_1 \cdot c_2 \cdot c_3) &+= S(\text{date}, \text{store}, \text{item}, \text{units}), V_H(\text{date}; c_2), \\ &V'_T(\text{date}, \text{store}; c_1), V'_I(\text{family}, \text{item}; c_2) \end{aligned}$$

where V_H is shared with Q_1 and Q_2 ; V'_I is shared with Q_2 .

The view $V_O(\text{date}; g(\text{price}))$ for Q_1 can be merged with $V'_O(\text{date}; 1)$ for Q_3 into a new view W_O since they have the same group-by attributes and body:

$$W_O(\text{date}; g(\text{price}), 1) += O(\text{date}, \text{price})$$

Both views V_T for Q_1 and V'_T for Q_3 are now defined over W_O instead of the views V_O and V'_O . Views V_T and V'_T have the same group by attributes and direction, but different bodies (one joins over V_R and the other over V'_R). We can merge them in W_T following Case (1):

$$W_T(\text{date}, \text{store}; a_1, a_2) += V_T(\text{date}, \text{store}; a_1), V'_T(\text{date}, \text{store}; a_2) \quad \square$$

Besides merging, *grouping* is another way of clustering the views that can share computation: We form groups of views that go out of the same node, regardless of their group-by attributes and bodies. We group the views as follows. We compute a topological order of these views: If a view V_1 uses a view V_2 in its body, i.e., it depends directly on it, then V_1 appears after V_2 in this order. We then traverse this order and create a group with all views such that (1) no view in the group depends on another view, and (2) all views within the group go out of the same relation in the join tree.

Figure 3(center) shows a scenario with directional views and four queries along the edges of our Favorita join tree. Their grouping is shown in Figure 3(right).

In the next section, we show how to compute all views within a group in one scan over their common relation.

3.5 Multi-Output Operator

The view group is a computational unit in LMFAO. We introduce a new operator that can compute all views in a group in one scan of their common input relation. Since this operator outputs the results for several views, we call it the *multi-output operator*, or MOO for short.

One source of complexity in MOO is that the views in the group are defined over different incoming views. While scanning the common relation, MOO looks up into the incoming views to fetch aggregates needed for the computation of the

views in the group. A second challenge is to update the aggregates of each view in the group as soon as possible and with minimal number of computation steps.

MOO has four steps: (1) Compute the order of join attributes of the common relation; Register (2) incoming and outgoing views and (3) aggregate functions to attributes in the attribute order; (4) Generate specialized code to compute view group. We next present and exemplify each of these steps. For this, we consider the following group of three views with the common relation Sales (S):

$$\begin{aligned} Q_4(f(\text{units}) \cdot \alpha_1 \cdot \alpha_4 \cdot \alpha_8) &+= S(\text{item}, \text{date}, \text{store}, \text{units}, \text{promo}), \\ &V_T(\text{date}, \text{store}; \alpha_8), V_H(\text{date}; \alpha_4), V_I(\text{item}; \alpha_1) \\ Q_5(\text{store}; g(\text{item}) \cdot h(\text{date}, \text{family}) \cdot \alpha_4 \cdot \alpha_8 \cdot \alpha_{13}) &+= V_H(\text{date}; \alpha_4), \\ &V_T(\text{date}, \text{store}; \alpha_8), V'_I(\text{item}, \text{family}; \alpha_{13}), \\ &S(\text{item}, \text{date}, \text{store}, \text{units}, \text{promo}) \\ Q_6(\text{item}; g(\text{item}) \cdot f(\text{units}) \cdot \alpha_1 \cdot \alpha_4 \cdot \alpha_8) &+= V_T(\text{date}, \text{store}; \alpha_8), \\ &S(\text{item}, \text{date}, \text{store}, \text{units}, \text{promo}), V_H(\text{date}; \alpha_4), V_I(\text{item}; \alpha_1) \end{aligned}$$

Join attribute order. The scan uses a total order on the join attributes of the relation S and sees S logically as a (partial) trie, grouped by the first join attribute and so on until the last join attribute. The leaves of the tries are relations over the remaining non-join attributes. This order can be computed offline. To avoid exploring all possible permutations of the join attributes, we proceed with the following approximation. We first compute the domain size for each join attribute in S , i.e., the number of its distinct values. We choose the order that is the increasing order in the domain sizes of these attributes: item – date – store. We sort S in this order.

MOO uses a multi-way nested-loops join over the relation and the incoming views, with one loop per join attribute. It sees the incoming and outgoing views as functions that, for a given tuple over the group-by attributes, look up the corresponding aggregate value. The aggregates to compute are also functions, in particular sums of products of functions that are UDAFs or lookups into incoming views. For instance, the aggregate of Q_4 is the product $f(\text{units}) \cdot \alpha_1 \cdot \alpha_4 \cdot \alpha_8$, where the last three components are provided by lookups in incoming views: $\alpha_1 = V_I(i)$ for the aggregate α_1 in the view V_I , where the group-by attribute item is set to i ; similarly for $\alpha_4 = V_H(d)$ and $\alpha_8 = V_T(d, s)$.

View registration. Each (incoming or outgoing) view V is registered at the lowest attribute in the order that is a group-by attribute of V . The reason is that at this attribute, all of the join attributes that are group-by attributes of V are fixed to constants and we can construct the tuples over its group-by attributes. The outgoing views without group-by attributes are registered outside the join attributes, as they are computed outside the outermost loop. Figure 4 depicts the registration of views in our example (left).

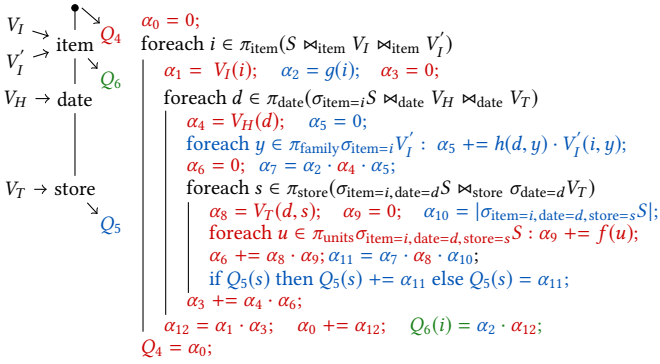


Figure 4: Multi-output operator for executing Q_4 , Q_5 , and Q_6 in the example of Section 3.5.

Aggregate function registration. Let d_Q be the depth in the attribute order where we registered an outgoing view Q . We discuss the registration of a product p of aggregate functions in Q . We decompose p into minimal partial products, such that no pair of functions from different partial products depend on each other. Two functions depend on each other if they have non-join attributes in the same relation or view: In Q_5 , $h(\text{date}, \text{family})$ and $V'_I(\text{item}, \text{family})$ depend on each other, because they share the non-join attribute family. Dependent functions need to be evaluated together in loops over the distinct values of the non-join attributes in the context of the values of the join attributes. The reason for non-join attributes is that the join attributes are fixed by the nested-loops join. The evaluation for each partial product is to be performed at the attribute of largest depth that is a parameter of any of the dependent functions.

If several functions in p are registered at the same depth d , we multiply their values. This is the partial product p_d of functions in p that can be computed at depth d . In order to obtain the final product p , we combine the partial products that were computed at each depth d as follows. If $d < d_Q$, we register at d one intermediate aggregate a_d that is the product of p_d and the intermediate aggregate a_{d-1} that is computed at depth $d - 1$. This intermediate aggregate is computed *before* we proceed to the next attribute in the order. If $d > d_Q$, we register at d a running sum r_d over the product of p_d and the running sum r_{d+1} . The running sum is computed *after* we return from the next attribute in the order. If $d = d_Q$, the product p corresponds to the product of p_d , the intermediate aggregate a_{d-1} , and the running sum r_{d+1} . The product p is computed *after* we return from depth $d + 1$ in the order, and then added to the tuple over the group-by attributes of Q .

Example 3.5. Figure 4 depicts the computation of Q_4 , Q_5 , and Q_6 . We register the components of the aggregate in Q_4 depending on the group-by attributes of their respective views: $f(\text{units})$ at store since the non-join attribute units is accessible once the join attributes are fixed in S ; α_8 also

at store; α_4 at date; α_1 at item. The function $f(\text{units})$ has a special treatment, since units is not a join attribute. Within the context in relation S of an item i , date d , and store s , we iterate over the qualifying tuples in S and accumulate in the local variable α_9 the sum over all values $f(u)$ for each value u for units. Once we computed locally the values for the component aggregates, we combine them with a minimal number of computation steps. We use local variables for running sums of multiplications of these values. As soon as the aggregates $f(\text{units})$ and α_8 are computed within the context of an item i , date d , and store s , we add their multiplication to a local variable α_6 ; this variable is initialized to 0 outside the loop over stores and its content is accumulated in α_3 right after the same loop. This accumulation is also used for the loops over dates and then items. Since Q_4 has no group-by attributes, its result is the scalar representing the aforementioned accumulation: $Q_4 = \alpha_0$.

Q_5 and Q_6 are treated similarly to Q_4 , with the difference that they have group-by attributes. We insert tuples in Q_5 within the loop over stores and update the aggregate value for a given store if the same store occurs under different (item, date) pairs. Q_6 reuses the aggregates α_{12} computed for Q_4 and α_2 computed for Q_5 . The tuples for Q_6 are constructed in the order of the items enumerated in the outermost loop. \square

Code Generation. Instead of interpreting and registering the views and their aggregates at runtime, we designed MOO to generate succinct and efficient C++ code for the shared computation of many aggregates in a view group. This code looks very similar to that in Figure 4 and features code specialization and optimization. Here are examples of code optimization already present in the code in Figure 4. The local variable α_{10} stores the size of a fragment of S . Since S is an array and sorted by item, date, and store, this fragment is a contiguous range whose size can be provided right away without having to enumerate over it. We do not allocate local variables if there is no need, e.g., the view lookup in α_5 . MOO also distinguishes the different requirements for data structures representing the results of Q_5 and Q_6 . Since we iterate over distinct items and Q_6 has one tuple per item, we can store Q_6 as a vector where each new item value is appended. In contrast, MOO may encounter the same store under different (item, date) pairs and therefore stores Q_5 as a hashmap to support efficient out-of-order updates. Further optimizations are highlighted in Appendix C.

4 EXPERIMENTS

We conducted two types of performance benchmarks on four datasets: (1) the computation of batches of aggregates in LMFAO, MonetDB, and DBX (a commercial DBMS); and (2) the training of machine learning models in LMFAO, TensorFlow, MADlib, and AC/DC. The outcome of these experiments

	R	F	Y	T
Tuples in Database	87M	125M	8.7M	30M
Size of Database	1.5GB	2.5GB	0.2GB	3.4GB
Tuples in Join Result	86M	127M	360M	28M
Size of Join Result	18GB	7GB	40GB	9GB
Relations	5	6	5	10
Attributes	43	18	37	85
Categorical Attributes	5	15	11	26

Table 1: Characteristics of used datasets: Retailer (R), Favorita (F), Yelp (Y), and TPC-DS (T).

is twofold: (1) Current database systems cannot efficiently compute large batches of aggregates as required by a variety of analytics workloads; (2) Scalability challenges faced by state-of-the-art machine learning systems can be mitigated by a combination of database systems techniques.

Datasets We consider four datasets: (1) *Retailer* [4] is used by a large US retailer for forecasting user demands and sales; (2) *Favorita* [18] is a public real dataset that is also used for retail forecasting; (3) *Yelp* is based on the public Yelp Dataset Challenge [56] and contains information about review ratings that users give to businesses; (4) *TPC-DS* [38] (scale factor 10, excerpt) is a synthetic dataset designed for decision support applications. The structure and size of these datasets are common in retail and advertising, where data is generated by sales transactions or click streams. Retailer and TPC-DS have a snowflake schema, Favorita has a star schema. They have a large fact table in the center and several smaller dimension tables. Yelp also has a star schema, but contains many-to-many joins, which increases the size of the join result significantly compared to the input database.

Table 1 provides key statistics for each dataset. Appendix A gives a detailed description of each dataset, including the schema and join tree used for our experiments.

Appendix B details the experimental setup for each workload and the limitations of the competing systems.

4.1 Computing Batches of Aggregates

We compute the batches of aggregates for the following workloads and each of the four datasets: (1) the covar matrix; (2) a single node in a regression tree; (3) the mutual information of all pairwise combinations of discrete attributes; and (4) a data cube. For each workload and dataset, Table 2 details how many aggregates, views, and groups are computed. It also gives the size on disk of the aggregates. This is a strong indicator of the running time; except for data cubes, these sizes are much smaller than for the underlying join.

Competitors We benchmark LMFAO against MonetDB 1.1 [26] and DBX. We also benchmarked PostgreSQL (PSQL) 11.1, but it was consistently slower than DBX and MonetDB. We attempted to benchmark against EmptyHeaded [2], but it failed to compute our workloads (cf. Appendix B).

Takeaways Table 3 presents the performance of each system for the four workloads and for the count query, which is used to assess the performance of many queries relative to this simple query. LMFAO consistently outperforms both DBX and MonetDB on all experiments, with a speedup of up to three orders of magnitude. The reason is as follows. Whereas DBX and MonetDB compute each individual query efficiently, they do not share computation across them. In contrast, LMFAO clusters the batch of queries into a few groups that are computed together in a single pass over the fact table and at most two passes over the smaller dimension tables. This is particularly beneficial for Retailer and Favorita, which have large fact tables with few attributes and most aggregates are computed over the dimension tables. The fact table in TPC-DS has many more attributes and LMFAO needs more time to compute over it. This explains the relatively lower performance improvement for TPC-DS. For Yelp, LMFAO’s decomposition of aggregates into views avoids the materialization of the many-to-many joins.

The performance gap is particularly large for regression tree nodes, which, in contrast to the other workloads, do not require the computation of queries with group-by attributes from different relations. LMFAO merges all aggregates into few views and shares their computation over input relation.

We use the count query to show how much computation is shared in LMFAO. For instance, the covar matrix for Retailer has 814 aggregates. Without sharing, the performance would be at least $814\times$ that of the count query, or 6510 seconds. The performance of LMFAO, however, is $55\times$ better!

Performance Breakdown of LMFAO Optimizations Figure 5 shows that the combination of several optimizations explains the superior performance of LMFAO for computing the covar matrix. Most notably, the compilation improves the performance up to $15\times$ for Yelp. Multi-output and multiple roots together further improve the performance by roughly $4\times$ to $7\times$ for each dataset. Parallelization with four cores further improves the performance by up to $3\times$.

Compilation Overhead The compilation overhead of LMFAO depends on the workload. When compiled with g++6.4.0 and eight cores, it ranges from 2 seconds for data cubes to 50 seconds for mutual information aggregate batch over TPC-DS. This overhead is independent from aggregate computation. If necessary, it could be reduced by using LLVM code generation and compilation [39].

4.2 Training Models

We report the end-to-end performance of LMFAO for learning three machine learning models: (1) ridge linear regression model; (2) regression tree; and (3) classification tree. Models (1) and (2) are computed over Retailer and Favorita, and used to predict the number of inventory units and respectively

	Retailer				Favorita				Yelp				TPC-DS			
	A+I	V	G	Size	A+I	V	G	Size	A+I	V	G	Size	A+I	V	G	Size
CM	814 + 654	34	7	0.1	140 + 46	125	9	0.9	730 + 309	99	8	1795	3061 + 590	286	14	577
RT	3141 + 16	19	9	0.1	270 + 20	26	11	0.1	1392 + 16	22	9	39	4299 + 138	52	17	0.2
MI	56 + 22	78	8	76	106 + 35	141	9	10	172 + 64	236	9	1759	301 + 95	396	15	55
DC	40 + 8	12	5	3944	40 + 7	13	6	5463	40 + 7	13	5	1876	40 + 12	17	10	3794

Table 2: Number of application aggregates (A), additional intermediate aggregates synthesised by LMFAO (I), views (V), and groups of views (G) for each dataset and aggregate batches: covar matrix (CM), regression tree node (RT), mutual information (MI), and data cube (DC). The size on disk of the final aggregates is given in MB.

Aggregate batch		Retailer		Favorita		Yelp		TPC-DS	
Count	LMFAO	0.80	1.00×	0.97	1.00×	0.68	1.00×	5.01	1.00×
	DBX	2.38	2.98×	4.04	4.15×	2.53	3.72×	2.84	0.57×
	MonetDB	3.75	4.70×	8.11	8.32×	4.37	6.44×	2.84	0.57×
Covar Matrix	LMFAO	11.87	1.00×	38.11	1.00×	108.81	1.00×	274.55	1.00×
	DBX	2,647.36	223.10×	773.46	20.30×	2,971.88	27.31×	9,454.31	34.44×
	MonetDB	3,081.02	259.64×	1,354.47	35.54×	5,840.18	53.67×	9,234.01	33.63×
Regression Tree Node	LMFAO	1.80	1.00×	3.49	1.00×	8.83	1.00×	105.66	1.00×
	DBX	3,134.67	1,739.55×	431.11	123.58×	2,409.59	272.97×	2,480.49	23.48×
	MonetDB	3,395.00	1,884.02×	674.06	193.23×	13,489.20	1,528.11×	3,085.60	29.20×
Mutual Information	LMFAO	30.05	1.00×	111.68	1.00×	345.35	1.00×	252.96	1.00×
	DBX	178.03	5.92×	596.01	5.34×	794.00	2.30×	1,002.84	3.96×
	MonetDB	297.30	9.89×	1,088.31	9.74×	1,952.02	5.65×	1,032.17	4.08×
Data Cube	LMFAO	15.47	1.00×	22.85	1.00×	23.75	1.00×	15.65	1.00×
	DBX	100.08	6.47×	273.10	11.95×	156.67	6.60×	66.12	4.23×
	MonetDB	111.08	7.18×	561.03	24.55×	260.39	10.96×	74.38	4.75×

Table 3: Time performance (seconds) for computing various batches of aggregates using LMFAO, MonetDB, and DBX and the relative speedup of LMFAO.

number of units sold. Model (3) is learned over TPC-DS and used to predict whether a customer is a preferred customer, as proposed in the Relational Dataset Repository [36]. To assess the accuracy of the models, we separate out a test dataset. The training dataset for each model is defined by the natural join of the remaining tuples in the input database.

Competitors We benchmarked LMFAO against several analytics tools commonly used in data science: TensorFlow 1.12 (compiled with AVX optimization enabled) [1], MADlib 1.8 [25], R [45], scikit-learn 0.20 [42], and Python StatsModels [52]. The latter three fail to compute the models either due to internal design limitations or out-of-memory error. TensorFlow mitigates this issue by using an iterator interface that only loads a small batch of tuples at a time. MADlib is an in-database analytics engine integrated with PSQL.

We also compared against AC/DC [4] which learns linear regression models over databases. Apart of AC/DC, all other systems require the full materialization of the training dataset. In addition, TensorFlow requires a random shuffling of the training dataset for linear regression models. We used PSQL to compute the join and the shuffling steps.

Takeaways Tables 4 and 5 give the performance of the systems. LMFAO is able to compute all models orders-of-magnitude faster than the competitors. For Retailer and Favorita, LMFAO learns the model over the input database even faster than it takes PSQL to compute the join. This is because LMFAO avoids the materialization of the large training dataset and works directly on the input database: For Retailer, the former is 10× larger than the latter (Table 1). Furthermore, for linear regression, the convergence step takes as input the result of the aggregate batch, which is again at least an order of magnitude smaller than the input database.

LMFAO learns the linear regression models with the same accuracy as the closed-form solution computed by MADlib yet in a fraction of the time it takes MADlib. Tensorflow takes order of magnitudes longer for one epoch (one pass over the training dataset). The model that Tensorflow learns for Favorita also has comparable accuracy to the closed-form solution, but for Retailer the root-mean-square-error of the model is 600× higher. TensorFlow would require more epochs to converge to the solution of LMFAO. LMFAO also outperforms the specialized AC/DC engine by up to 18×.

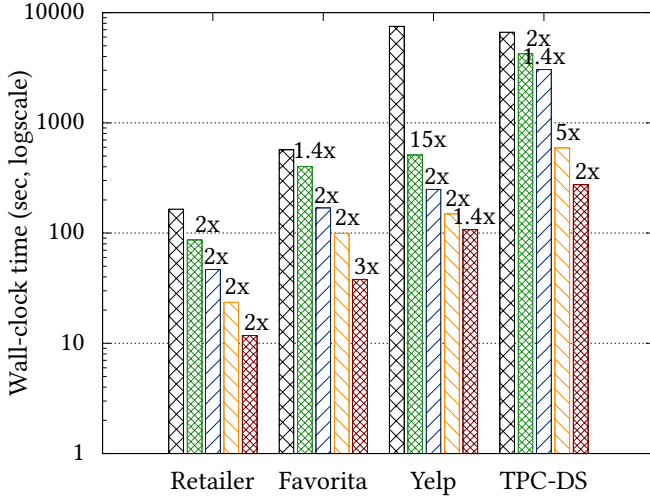


Figure 5: Performance impact of optimization layers in LMFAO for computing the covar matrix. The left-most bar corresponds to no optimization. Optimizations are enabled one at a time (left to right): compilation, multi-output, multi-root, and parallelization with 4 threads. Bars are annotated by relative speedup over the previous optimizations.

		Retailer	Favorita
Join	PSQL	152.06	129.32
Join Shuffle	PSQL	5,488.73	1,720.02
Join Export	PSQL	351.76	241.03
Linear Regression			
TensorFlow	(1 epoch)	7,249.58	4,812.01
MADlib		5,423.05	19,445.58
AC/DC		110.88	364.17
LMFAO		6.08	21.23
Regression Trees			
TensorFlow	(1 node)	8,234.30	20,842.98
MADlib	(max 31 nodes)	10,309.89	11,334.87
LMFAO	(31 nodes)	29.46	26.53

Table 4: Time performance (seconds) for learning linear regression models and regression trees over Favorita and Retailer.

LMFAO learns decision trees with the same accuracy orders of magnitudes faster than MADlib. TensorFlow times out after 12 hours in all cases; we show the time to compute the tree root as indication of the overall runtime.

5 RELATED WORK

LMFAO builds on a vast literature of database research. We cited highly relevant work in previous sections. We next mention further connections to work on sharing computation and data systems for learning models. LMFAO computes a batch of group-by aggregates over the same joins without

		TPC-DS
Join	PSQL	219.04
Join Export	PSQL	350.02
Classification Trees		
TensorFlow	(1 node)	10,336.61
MADlib	(max 31 nodes)	43,617.05
LMFAO	(31 nodes)	593.71

Table 5: Time performance (seconds) for learning classification trees over TPC-DS.

materializing these joins, in the spirit of ad-hoc mining [12], eager aggregation [55], and factorized databases [8].

Sharing Computation Prior techniques for data cubes use a lattice of sub-queries to capture sharing across the group-by aggregates defining data cubes [24, 37]. Which cells to materialize in a data cube is decided based on space or user-specified constraints [24, 37]. More recent work revisited shared workload optimization for queries with hash joins and shared scans and proposes an algorithm that, given a set of statements and their relative frequency in the workload, outputs a global plan over shared operators [21]. Data Canopy is a library of frequently used statistics in the form of aggregates that can speed up repeating requests for the same statistics. It is concerned with how to decompose, represent, and access such statistics in an efficient manner [54].

Multi-Query Optimization (MQO) [47] is concerned with identifying common subexpressions across a set of queries with the purpose of avoiding redundant computation. One of the three types of view merging in LMFAO is also concerned with the same goal, though for directional views with group-by aggregates. LMFAO’s view merging proved useful in case of very many and similar views, such as for the applications detailed in Section 2. An alternative type of MQO is concerned with caching intermediate query results, such as in the MonetDB system that we used in experiments.

Learning over Multi-Relational Data There are *structure-agnostic* and *structure-aware* solutions depending on whether they exploit the structure in the data or not.

The *structure-agnostic* solutions are by far the most common. They first construct the training dataset using a data system capable of computing queries and then learn the model over the materialized training dataset using an ML library or statistical package. The first step is performed in Python Pandas, R dplyr, or database systems such as PostgreSQL and SparkSQL [57]. The second step commonly uses scikit-learn [42], Python StatsModels [52], TensorFlow [1], R [45], MLib [35], SystemML [10], or XGBoost [15]. Although one could combine any data system and ML library, working solutions feature combinations that avoid the expensive data export/import at the interface between the two systems, e.g., MLib over SparkSQL, the Python packages over Pandas, R over dplyr, and MADlib [25] over PostgreSQL.

MADlib, Bismarck [19], and GLADE PF-OLA [44] define ML tasks as user-defined aggregate functions (UDAFs). Although UDAFs share the same execution space with the query computing the training dataset, they are treated as black boxes and executed after the training dataset is materialized.

A disadvantage of two-step solutions is the required materialization of the training dataset that may be much larger than the input data (cf. Table 1). This is exacerbated by the stark asymmetry between the two steps: Whereas data systems tend to scale to large datasets, this is not the case for ML libraries. Yet, the two-step solutions expect by design that the ML libraries work on even larger inputs than the data systems! A further disadvantage is that these solutions inherit the limitations of both underlying systems. For instance, the R data frame can host at most 2^{31} values, which makes it impossible to learn models over large datasets, even if data systems can process them. Database systems can only handle up to a few thousand columns per relation, which is usually smaller than the number of features of the model.

The *structure-aware* solutions tightly integrate the dataset construction and the learning steps, and allow the second step to exploit the relational structure in the input data. There is typically one unified execution plan for both the feature extraction query and the subsequent learning task, with sub-components of the latter possibly pushed past the joins in the former. This plan works directly on the input data and computes sufficient statistics of much smaller size than of the training dataset (cf. Table 2). Our system LMFAO is a prime example of this class. It builds on F [46] and AC/DC [4]. F supports linear regression models. AC/DC generalizes F to non-linear models, categorical features, and model reparameterization under functional dependencies. A key aspect that sets apart F, AC/DC, and LMFAO from all other efforts is the use of execution plans for the mixed workload of queries and learning whose complexity may be asymptotically lower even than that of the materialization step. In particular, this line of work shows that all machine learning approaches that require as input the materialization of the result of the feature extraction query may be asymptotically suboptimal.

Further examples in this category are: Orion [30] and Hamlet [31], which support generalized linear models and Naïve Bayes classification; recent efforts on scaling linear algebra using existing distributed database systems [32]; the declarative language BUDS [20], whose compiler can perform deep optimizations of the user’s program; and Morpheus [14]. Morpheus factorizes the computation of linear algebra operators summation, matrix-multiplication, pseudo-inverse, and element-wise operations over training datasets defined by key-foreign key star or chain joins. It represents the training dataset as a normalized matrix, which is a triple of the fact table, a list of dimension tables, and a list of indicator matrices that encode the join between the fact table and each

dimension table. Morpheus provides operator rewritings that exploit the relational structure by pushing computation past joins to the smaller dimension tables. Initial implementations of Morpheus are built on top of the R and Python numpy linear algebra packages. Morpheus only supports key-foreign key star or chain joins and models that are expressible in linear algebra. In contrast, LMFAO supports arbitrary joins and rich aggregates that can capture computation needed by a large heterogeneous set of models beyond those expressible in linear algebra, including, e.g., decision trees.

Optimizations in ML packages Most ML libraries exploit sparsity in the form of zero-values (due to missing values or one-hot encoding), yet are not structure-aware. LMFAO exploits a more powerful form of sparsity that is prevalent in training datasets defined by joins of multiple relations: This is the join factorization that avoids the repeated representation of and computation over arbitrarily-sized data blocks. LMFAO’s code optimizations aim specifically at generating succinct and efficient C++ code for the shared computation of many aggregates over the join of a large table and several views represented as ordered vectors or hashmaps. The layout of the generated code is important: how to decompose the aggregates, when to initialize and update them, how to share partial computation across many aggregates with different group-by and UDAFs (Section 3.5). Lower-level optimizations (Appendix C) are generic and adapted to our workload, e.g., how to manage large amounts of aggregates and how to update them in sequence. LMFAO’s multi-aggregate optimizations are absent in ML and linear algebra packages. We next highlight some of their code optimizations. BLAS and LAPACK provide cache-efficient block matrix operations. Eigen [23] supports both dense and sparse matrices, fuses operators to avoid intermediate results, and couples loop unrolling with SIMD vectorization. SPOOF [17] translates linear algebra operations into sum-product form and detects opportunities for aggregate pushdown and operator fusion. LGen [50] uses compilation to generate efficient basic linear algebra operators for small dense, symmetric, or triangular matrices by employing loop fusion, loop tiling, and vectorization. TACO [29] can generate compound linear algebra operations on both dense and sparse matrices. LMFAO can also learn decision trees, which cannot be expressed in linear algebra. XGBoost [15] is a gradient boosting library that uses decision trees as base learners. It represents the training dataset in a compressed sparse columnar (CSC) format, which is partitioned into blocks that are optimized for cache access, in-memory computation, parallelization, and can be stored on disk for out-of-core learning. LMFAO may also benefit from a combination of value-based compression and factorized representation of the training dataset, as well as from an out-of-core learning mechanism.

ACKNOWLEDGMENTS

Olteanu has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 682588. He also acknowledges a Google research award and an Infor research gift. Nguyen gratefully acknowledges support from NSF grants CAREER DMS-1351362 and CNS-1409303, Adobe Research and Toyota Research, and a Margaret and Herman Sokol Faculty Award. Schleich acknowledges the AWS Proof of Concept Program.

REFERENCES

- [1] Martín Abadi, Paul Barham, Jianmin Chen, and et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*. 265–283.
- [2] Christopher R. Aberger, Susan Tu, Kunle Olukotun, and Christopher Ré. 2016. EmptyHeaded: A Relational Engine for Graph Processing. In *SIGMOD*. 431–446.
- [3] S. Abiteboul, R. Hull, and V. Vianu. 1995. *Foundations of Databases*. Addison-Wesley.
- [4] 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.
- [5] Mahmoud Abo Khamis, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. 2018. In-Database Learning with Sparse Tensors. In *PODS*. 325–340.
- [6] Mahmoud Abo Khamis, Hung Q. Ngo, and Atri Rudra. 2016. FAQ: Questions Asked Frequently. In *PODS*. 13–28.
- [7] S. M. Aji and R. J. McEliece. 2006. The Generalized Distributive Law. *IEEE Trans. Inf. Theor.* 46, 2 (2006), 325–343.
- [8] Nurzhan Bakibayev, Tomás Kociský, Dan Olteanu, and Jakub Závodný. 2013. Aggregation and Ordering in Factorised Databases. *PVLDB* 6, 14 (2013), 1990–2001.
- [9] Nurzhan Bakibayev, Dan Olteanu, and Jakub Závodný. 2012. FDB: A Query Engine for Factorised Relational Databases. *PVLDB* 5, 11 (2012), 1232–1243.
- [10] Matthias Boehm et al. 2016. SystemML: Declarative Machine Learning on Spark. *PVLDB* 9, 13 (2016), 1425–1436.
- [11] L. Breiman, J. Friedman, R. Olshen, and C. Stone. 1984. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA.
- [12] Surajit Chaudhuri. 1998. Data Mining and Database Systems: Where is the Intersection? *IEEE Data Eng. Bull.* 21, 1 (1998), 4–8.
- [13] Surajit Chaudhuri, Usama M. Fayyad, and Jeff Bernhardt. 1999. Scalable Classification over SQL Databases. In *ICDE*. 470–479.
- [14] Lingjiao Chen, Arun Kumar, Jeffrey F. Naughton, and Jignesh M. Patel. 2017. Towards Linear Algebra over Normalized Data. *PVLDB* 10, 11 (2017), 1214–1225.
- [15] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *KDD*. 785–794.
- [16] C. Chow and C. Liu. 2006. Approximating Discrete Probability Distributions with Dependence Trees. *IEEE Trans. Inf. Theor.* 14, 3 (2006), 462–467.
- [17] Tarek Elgamal, Shangyu Luo, Matthias Boehm, Alexandre V. Efimievski, Shirish Tatikonda, Berthold Reinwald, and Prithviraj Sen. 2017. SPOOF: Sum-Product Optimization and Operator Fusion for Large-Scale Machine Learning. In *CIDR*.
- [18] Corporacion Favorita. 2017. Corp. Favorita Grocery Sales Forecasting: Can you accurately predict sales for a large grocery chain? (2017). <https://www.kaggle.com/c/favorita-grocery-sales-forecasting/>
- [19] Xixuan Feng, Arun Kumar, Benjamin Recht, and Christopher Ré. 2012. Towards a unified architecture for in-RDBMS analytics. In *SIGMOD*. 325–336.
- [20] Zekai J. Gao, Shangyu Luo, Luis Leopoldo Perez, and Chris Jermaine. 2017. The BUDS Language for Distributed Bayesian Machine Learning. In *SIGMOD*. 961–976.
- [21] Georgios Giannikis, Darko Makreshanski, Gustavo Alonso, and Donald Kossmann. 2014. Shared Workload Optimization. *PVLDB* 7, 6 (2014), 429–440.
- [22] Jim Gray, Adam Bosworth, Andrew Layman, and Hamid Pirahesh. 1996. Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Total. In *ICDE*. 152–159.
- [23] Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. (2010). <http://eigen.tuxfamily.org>
- [24] Venky Harinarayan, Anand Rajaraman, and Jeffrey D. Ullman. 1996. Implementing Data Cubes Efficiently. In *SIGMOD*. 205–216.
- [25] Joseph M. Hellerstein, Christopher Ré, Florian Schoppmann, Daisy Zhe Wang, Eugene Fratkin, Aleksander Gorajek, Kee Siong Ng, Caleb Welton, Xixuan Feng, Kun Li, and Arun Kumar. 2012. The MADlib Analytics Library or MAD Skills, the SQL. *PVLDB* 5, 12 (2012), 1700–1711.
- [26] Stratos Idreos, Fabian Groffen, Niels Nes, Stefan Manegold, K. Sjoerd Mullender, and Martin L. Kersten. 2012. MonetDB: Two Decades of Research in Column-oriented Database Architectures. *IEEE Data Eng. Bull.* 35, 1 (2012), 40–45.
- [27] Kaggle. 2018. Kaggle ML & DB Survey. (2018). <https://www.kaggle.com/kaggle/kaggle-survey-2018>
- [28] Timo Kersten, Viktor Leis, Alfons Kemper, Thomas Neumann, Andrew Pavlo, and Peter Boncz. 2018. Everything You Always Wanted to Know About Compiled and Vectorized Queries but Were Afraid to Ask. *PVLDB* 11, 13 (2018), 2209–2222.
- [29] Fredrik Kjolstad, Shoaib Kamil, Stephen Chou, David Lugato, and Saman Amarasinghe. 2017. The Tensor Algebra Compiler. *OOPSLA* 1, Article 77 (2017), 77:1–77:29 pages.
- [30] Arun Kumar, Jeffrey F. Naughton, and Jignesh M. Patel. 2015. Learning Generalized Linear Models Over Normalized Data. In *SIGMOD*. 1969–1984.
- [31] Arun Kumar, Jeffrey F. Naughton, Jignesh M. Patel, and Xiaojin Zhu. 2016. To Join or Not to Join?: Thinking Twice about Joins before Feature Selection. In *SIGMOD*. 19–34.
- [32] Shangyu Luo, Zekai J. Gao, Michael N. Gubanov, Luis Leopoldo Perez, and Christopher M. Jermaine. 2018. Scalable Linear Algebra on a Relational Database System. *SIGMOD Rec.* 47, 1 (2018), 24–31.
- [33] Daniel Marx. 2010. Approximating Fractional Hypertree Width. *ACM Trans. Algorithms* 6, 2, Article 29 (April 2010), 17 pages.
- [34] H. Brendan McMahan and et al. 2013. Ad Click Prediction: A View from the Trenches. In *KDD*. 1222–1230.
- [35] Xiangrui Meng, Joseph Bradley, et al. 2016. MLlib: Machine Learning in Apache Spark. *J. Mach. Learn. Res.* 17, 1 (2016), 1235–1241.
- [36] Jan Motl and Oliver Schulte. 2015. The CTU Prague Relational Learning Repository. (2015). [arXiv:cs.LG/1511.03086](https://arxiv.org/abs/1511.03086)
- [37] Inderpal Singh Mumick, Dallen Quass, and Barinderpal Singh Mumick. 1997. Maintenance of Data Cubes and Summary Tables in a Warehouse. In *SIGMOD*. 100–111.
- [38] Raghunath Othayoth Nambiar and Meikel Poess. 2006. The Making of TPC-DS. In *PVLDB*. 1049–1058.
- [39] Thomas Neumann. 2011. Efficiently Compiling Efficient Query Plans for Modern Hardware. *PVLDB* 4, 9 (2011), 539–550.
- [40] Dan Olteanu and Maximilian Schleich. 2016. Factorized Databases. *SIGMOD Rec.* 45, 2 (2016), 5–16.
- [41] Judea Pearl. 1982. Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach. In *AAAI*. 133–136.
- [42] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, and et al. 2011. Scikit-learn: Machine Learning in Python. *J. Machine Learning Research* 12 (2011), 2825–2830.

- [43] Holger Pirk, Oscar Moll, Matei Zaharia, and Samuel Madden. 2016. Voodoo - A Vector Algebra for Portable Database Performance on Modern Hardware. *PVLDB* 9 (2016), 1707–1718.
- [44] Chengjie Qin and Florin Rusu. 2015. Speculative Approximations for Terascale Distributed Gradient Descent Optimization. In *DanaC*. 1:1–1:10.
- [45] R Core Team. 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Stat. Comp., www.r-project.org.
- [46] Maximilian Schleich, Dan Olteanu, and Radu Ciucanu. 2016. Learning Linear Regression Models over Factorized Joins. In *SIGMOD*. 3–18.
- [47] Timos K. Sellis. 1988. Multiple-query Optimization. *ACM Trans. Database Syst.* 13, 1 (1988), 23–52.
- [48] 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), 45 pages.
- [49] Amir Shaikhha, Yannis Klonatos, Lionel Parreaux, Lewis Brown, Mohammad Dashti, and Christoph Koch. 2016. How to Architect a Query Compiler. In *SIGMOD*. 1907–1922.
- [50] Daniele G. Spampinato and Markus Püschel. 2016. A basic linear algebra compiler for structured matrices. In *CGO*. 117–127.
- [51] Ruby Y. Tahboub, Grégory M. Essertel, and Tiark Rompf. 2018. How to Architect a Query Compiler, Revisited. In *SIGMOD*. 307–322.
- [52] The StatsModels development team. 2012. StatsModels: Statistics in Python, <http://statsmodels.sourceforge.net>. (2012).
- [53] Todd L. Veldhuizen. 2014. Triejoin: A Simple, Worst-Case Optimal Join Algorithm. In *ICDT*. 96–106.
- [54] Abdul Wasay, Xinding Wei, Niv Dayan, and Stratos Idreos. 2017. Data Canopy: Accelerating Exploratory Statistical Analysis. In *SIGMOD*. 557–572.
- [55] Weipeng P. Yan and Per-Åke Larson. 1995. Eager Aggregation and Lazy Aggregation. In *VLDB*. 345–357.
- [56] Yelp. 2017. Yelp Dataset Challenge. (2017). <https://www.yelp.com/dataset/challenge/>
- [57] Matei Zaharia, Mosharaf Chowdhury, et al. 2012. Resilient Distributed Datasets: A Fault-tolerant Abstraction for In-memory Cluster Computing. In *NSDI*. 2–2.
- [58] Marcin Zukowski, Mark van de Wiel, and Peter Boncz. 2012. Vector-wise: A Vectorized Analytical DBMS. In *ICDE*. 1349–1350.

A DATASETS

Figure 6 gives the join trees for the four datasets used in the experiments in Section 4.

Retailer has five relations: *Inventory* stores the number of inventory units for each date, store, and stock keeping unit (sku); *Location* keeps for each store: its zipcode, the distance to competitors, the store type; *Census* provides 14 attributes that describe the demographics of each zipcode, including the population or median age; *Weather* stores statistics about the weather condition for each date and store, including the temperature or if it rained; *Items* keeps the price, category, subcategory, and category cluster of each sku.

Favorita has six relations. Its schema is given in Figure 3. *Sales* stores the number of units sold for each store, date, and item, and whether the item was on promotion; *Items* provides information about the skus, such as the item class and price; *Stores* keeps information on stores, like the city they are located in; *Transactions* stores the number of transactions for

each date and store; *Oil* provides the oil price for each date; and *Holiday* indicates whether a given date is a holiday.

Yelp has five relations: *Review* gives the rating and date for each review by users of businesses; *User* keeps information about the users, including how many reviews they made, or when they joined; *Business* provides information about the business, e.g., their location and average rating; *Category* and *Attribute* keep the categories, e.g., Restaurant, and respectively attributes, e.g., open late, of the businesses. A business can have many attributes and categories.

TPC-DS [38] is an excerpt of the snowflake query with the Store Sales fact table and scale factor 10. We consider the ten relations and schema shown in Figure 6(d). We modified the generated relations by (1) turning strings into integer ids, (2) populating null values, and (3) dropping attributes that are not relevant for our analytics workloads, e.g. street name or categorical attributes with only one category. We provide further details on the modifications and the scripts to do them on our website: <https://github.com/fdbresearch/fbench/tree/master/data/tpc-ds>.

Test Data In order to assess the accuracy of a learned linear regression model, we separate test data for each dataset that the model is not trained over. The test data constitutes the sales in the last month in the dataset, for Retailer and Favorita, and the last 15 days for TPC-DS. This simulates the realistic usecase where the ML model predicts future sales.

B EXPERIMENTAL EVALUATION

Experimental Setup Since DBX is only available in the cloud, we run all experiments in Section 4.1 on a dedicated AWS d2.xlarge instance with Ubuntu 18.04 and four vCPUs.

The experiments in Section 4.2 are performed on an Intel(R) Core(TM) i7-4770 3.40GHz/64bit/32GB with Linux 3.13.0/g++6.4.0 and eight cores.

We used the O3 compiler optimization flag and report wall-clock times by running each system once and then reporting the average of four subsequent runs with warm cache. We do not report the times to load the database into memory. All relations are given sorted by their join attributes.

Setup for Aggregate Computation The covar matrix and regression tree node are computed over all attributes in Yelp, all but join keys in Retailer and TPC-DS, and all but date and item in Favorita. We compute all pairwise mutual information aggregates over nine attributes for Retailer, 15 for Favorita, 11 for Yelp, and 19 for TPC-DS. These attributes include all categorical and some discrete continuous attributes in each dataset. For data cubes, we used three dimensions and five measures for all experiments. We provide DBX and MonetDB with the same list of queries as LMFAO, which may have multiple aggregates in the head.

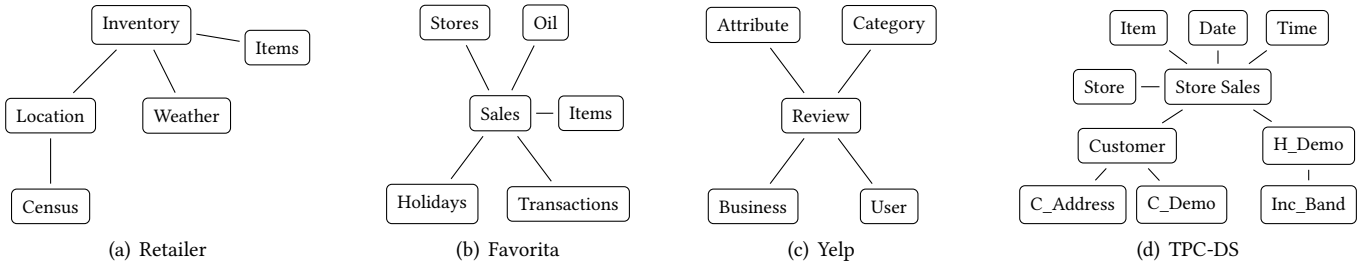


Figure 6: Join Trees used in experiments for the Retailer, Favorita, Yelp and TPC-DS datasets.

In Figure 5, the baseline is computed with AC/DC [4], which is an (imperfect) proxy for computing the covar matrix in an interpreted version of LMFAO without optimizations.

Setup for Model Training We learn regression and regression tree models over all attributes but join keys for Retailer, and all but date and item for Favorita. For TPC-DS, we learn classification trees over all attributes but join keys.

LMFAO and AC/DC first compute the covar matrix and then optimize the parameters over it using gradient descent with Armijo backtracking line search and Barzilai-Borwein step size [4]. MADlib computes the closed form solution of the model with ordinary least squares over the non-materialized view of the training dataset. (OLS is the fastest approach supported by MADlib for this problem.) We evaluate the accuracy of the model by computing the root-mean-square-error over the test dataset and ensure that it is the same for LMFAO’s model and MADlib’s closed form solution.

TensorFlow requires as input the materialized training dataset shuffled in random order. TensorFlow fails to shuffle the entire dataset and runs out-of-memory during learning when the entire dataset is represented and shuffled in-memory with Python Pandas. We therefore materialize and shuffle the datasets in PSQL, and use TensorFlow’s iterator interface to load the dataset. The model is then learned with the default settings of the built-in LinearRegressor Estimator, which optimizes the parameters with a variant of stochastic gradient descent called FTRL [34]. We used a batch size of 500K for learning, because this gave us the best performance/accuracy tradeoff amongst all batch sizes we considered. We could not set TensorFlow to run until convergence easily, so we computed the time it takes for one epoch (one pass over the training data) and compared the accuracy of the obtained model with the closed form solution.

All systems learn the decision trees with the CART algorithm [11]. As cost function, we use the variance for regression trees, and the Gini index for classification trees. The maximum depth of the tree is 4, so that it can have at most 31 nodes. Continuous attributes are bucketized into 20 buckets.

We verify that LMFAO learns decision trees that have the same accuracy as the decision trees learned in MADlib.

We use TensorFlow’s built-in BoostedTrees Estimator with a batch size of 1M to learn decision trees. Larger batch sizes cause either out-of-memory errors or a lot of memory swaps, which significantly degrades the performance. For continuous attributes, TensorFlow requires the buckets as input, and we provide it with the same buckets as LMFAO.

Finally, we use PSQL to compute, shuffle, and export the join results. We tuned PSQL for in-memory processing by setting its working memory to 28GB, shared buffers to 128MB, and turning off the parameters fsync, synchronous commit, full page writes, and bgwriter LRU maxpages.

Benchmarking against EmptyHeaded We attempted to benchmark against EmptyHeaded [2], which computes single aggregates over join trees. The implementation, however, requires an extensive preprocessing of the dataset to turn the relations into the input format that EmptyHeaded requires. This preprocessing step introduces significant overhead, which, when applied to our datasets, blows up the size of the data to the extent that it no longer fits into memory. For instance, the Inventory relation in the Retailer dataset (2GB) is blown up to more than 300GB during preprocessing. Our observation is that EmptyHeaded has difficulty preprocessing relations whose arity is beyond 2. We were therefore unable to compare against EmptyHeaded.

Limitations of ML Competitors We detail here further limitations of the systems we encountered while preparing the experiments. (1) The iterator interface of TensorFlow is both a blessing, because it allows TensorFlow to compute models over large datasets, but also a curse, because of its overhead and poor performance for learning models over large datasets. In particular, it needs to repeatedly load, parse and cast the batches of tuples. (2) R can load at most 2.5 billion values, which is less than the training datasets require. (3) Scikit-learn and StatsModels succeed in loading the training dataset, but runs out of memory during the one-hot encoding. (4) Scikit-learn and StatsModels require the all values have the same type, so they go for the most general type: Floats.

```

aggregate[435] += aggregatesV3[0];
...
aggregate[444] += aggregatesV3[9];
aggregate[445] += aggregate[0]*aggregatesV3[0];
...
aggregate[473] += aggregate[28]*aggregatesV3[0];
aggregate[474] += aggregate[0]*aggregatesV3[1];
...
aggregate[476] += aggregate[0]*aggregatesV3[3];

```

(a) Aggregates are stored and accessed consecutively in fixed size array.

```

for (size_t i = 0; i < 10; ++i)
    aggregate[435+i] += aggregatesV3[i];
for (size_t i = 0; i < 29; ++i)
    aggregate[445+i] += aggregate[i]*aggregatesV3[0];
for (size_t i = 0; i < 3; ++i)
    aggregate[474+i] += aggregate[0]*aggregatesV3[i];

```

(b) Updates to consecutive aggregates are fused into tight loops.

Figure 7: Snippet of code generated by LMFAO that shows how aggregates are stored and computed.

This can add significant overhead and is one of the reasons why the Python variants run out of memory.

C LMFAO COMPILATION

Recent work uses code compilation to reduce the interpretation overhead of query evaluation [28, 39, 43, 48, 49, 51]. The compilation approach in LMFAO is closest in spirit to DBLAB [49], which advocates for the use of intermediate representations (IR) to enable code optimizations that cannot be achieved by conventional query optimizers or query compilation techniques without IRs.

The various optimization layers of LMFAO can be viewed as optimizations over the following increasingly granular IRs: (1) the join tree; (2) orders of join attributes; and (3) the multi-output operator which registers the computation of aggregates to specific attributes in the attribute order.

LMFAO relies on these IRs to identify optimizations that are not available in conventional query processing. For instance, the join tree is used to identify groups of views that become the main computational unit in LMFAO (c.f. Section 3.4) and the data then moves between different view groups. This departs from standard query processing that computes queries by pipelining tuples between relational operators in a query plan. The three IRs exploit information about the workload at compile time to specialize the generated code. We next explain some of these optimizations; others have been already mentioned at the end of Section 3.5.

Code Specialization The database catalog and join tree provide a lot of statistics that LMFAO exploits to generate specific data structures to represent the relations. For instance, given its size and schema, a relation is represented as

a fixed size array of tuples that are represented with specialized C++ structs that have the exact type for each attribute. The attribute order provides for each join the join attribute and the views that are joined over. LMFAO uses this information to generate specialized code that computes these joins without dynamic casting and function calls to an iterator.

Fixed size arrays The registration of aggregates to the attribute order allows us to derive at compile time how many aggregates are computed in a group and the order in which they are accessed. LMFAO uses this information during the code generation to generate fixed size arrays that store all aggregates consecutively, in an order that allows for sequential reads and writes. Figure 7(a) presents a snippet of generated code that computes partial aggregates for the covar matrix. Each aggregate array is accessed sequentially, which improves the cache locality of the generated code.

Loop Synthesis The sequential access to the array of aggregates further allows us to compress long sequences of arithmetic operations over aggregates addressed in lockstep into tight loops, as shown in Figure 7(b). This optimization allows the compiler to vectorize the computation of the loop and reduces the amount of code to compile.

Avoid repeated computation of UDAFs As presented in Section 3.5, the MOO decomposes the computation of aggregates over the attribute order. This allows LMFAO to register functions to the lowest possible node in the attribute order. The effect of this is that we evaluate each function only when necessary, and we minimize the number of updates to each aggregate. For instance, in Figure 4 the function $g(item)$ is evaluated only once per item value, and not repeatedly for different dates and stores that are joined with this item value.

Inlining Function Calls LMFAO knows which UDAFs are computed at compile time, and can thus inline them during code generation. For instance, LMFAO generates the following code snippet for a product of three functions that constitutes one decision tree node aggregate over Retailer:

```

aggregate[96] =(t.avghhi <= 52775 ? 1.0 : 0.0)
                * (t.area_sq_ft > 93580 ? 1.0 : 0.0)
                * (t.distance_comp <= 5.36 ? 1.0 : 0.0);

```

An interpreted version of LMFAO would make one function call for each term in the product to compute this aggregate.

Dynamic Functions Some workloads require repeated computation of slightly different aggregates. To learn decision trees, for instance, we repeatedly compute the same set of aggregates, where the only difference is one additional threshold function per node. To avoid recompiling the entire generated code for each decision tree node, we generate dynamic functions in a separate C++ file with a few lines of code. This file can be recompiled efficiently and dynamically loaded into a running instance of LMFAO to recompute the aggregates.