Learning Models over Relational Databases fdbresearch.github.io relational.ai

Dan Olteanu

Oxford & relationalAl

ICDT & EDBT Keynote Lisbon, March 2019



Acknowledgments

FDB team, in particular:







Max



Milos



Ahmet



Fabian

relationalAl team, in particular:



Mahmoud



Hung



Long

ML: The Next Big Opportunity

ML is emerging as general purpose technology

• Just as computing became general purpose 70 years ago

A core ability of intelligence is the ability to predict

• Turn information you have into information you need

The quality of the prediction is increasing as the cost per prediction is decreasing

Most Enterprises Rely on Relational Data for Their ML Models



Retail: 86% relational Insurance: 83% relational Marketing: 82% relational Financial: 77% relational

> Source: The State of Data Science & Machine Learning 2017, Kaggle, October 2017 (based on 2017 Kaggle survey of 16,000 ML practitioners)

Relational Model: The Jewel in the Data Management Crown

Last decades have witnessed massive adoption of the Relational Model

Many human hours invested in building relational models

Relational databases are rich with knowledge of the underlying domains

Availability of curated data made it possible to learn from the past and to predict the future for both humans (BI) and machines (ML)



Current State of Affairs in Building Predictive Models



Current State of Affairs in Building Predictive Models





THROWS AWAY

the relational structure

that can help build

FASTER

BETTER MODELS

Data matrix Features



Learning over Relational Databases: Revisit from First Principles



Structure-agnostic learning requires:

- 1. Materialisation of the query result (Recomputation in case of data updates)
- 2. Data export from DBMS and import into ML tool
- 3. One/multi-hot encoding of categorical variables



Structure-agnostic learning requires:

- Materialisation of the query result (Recomputation in case of data updates)
- 2. Data export from DBMS and import into ML tool
- 3. One/multi-hot encoding of categorical variables

All these steps are very expensive and unnecessary!



Structure-aware learning avoids the three expensive steps.

Structure-aware Learning FASTER even than Feature Extraction Query!

Example: A Retailer Use Case



| Relation | Cardinality | Arity (Keys+Values) | File Size (CSV) |
|--------------|-------------|---------------------|-----------------|
| Inventory | 84,055,817 | 3 + 1 | 2 GB |
| Items | 5,618 | 1 + 4 | 129 KB |
| Stores | 1,317 | 1 + 14 | 139 KB |
| Demographics | 1,302 | 1 + 15 | 161 KB |
| Weather | 1,159,457 | 2 + 6 | 33 MB |
| Join | 84,055,817 | 3 + 41 | 23GB |

Train a linear regression model to predict inventory given all features

| PostgreSQL+TensorFlow | | |
|-----------------------|----------------|------------|
| | Time | Size (CSV) |
| Database | _ | 2.1 GB |
| Join | 152.06 secs | 23 GB |
| Export | 351.76 secs | 23 GB |
| Shuffling | 5,488.73 secs | 23 GB |
| Query batch | - | - |
| Grad Descent | 7,249.58 secs | - |
| Total time | 13,242.13 secs | |

Train a linear regression model to predict inventory given all features

| | PostgreSQL+TensorFlow | | Our approach (SIGMOD'19) | | |
|--------------|-----------------------|------------|--------------------------|------------|--|
| | Time | Size (CSV) | Time | Size (CSV) | |
| Database | _ | 2.1 GB | _ | 2.1 GB | |
| Join | 152.06 secs | 23 GB | _ | - | |
| Export | 351.76 secs | 23 GB | _ | - | |
| Shuffling | 5,488.73 secs | 23 GB | _ | - | |
| Query batch | - | - | 6.08 secs | 37 KB | |
| Grad Descent | 7,249.58 secs | - | 0.05 secs | - | |
| Total time | 13,242.13 secs | | 6.13 secs | | |

 $2,160 \times$ faster while $600 \times$ more accurate (RMSE on 2% test data) TensorFlow trains one model. Our approach takes < 0.1 sec for any extra model over a subset of the given feature set.

9/44

Similar behaviour (or outright failure) for more:

- datasets: Favorita, TPC-DS, Yelp, Housing
- systems:
 - used in industry: R, scikit-learn, Python StatsModels, mlpack, XGBoost, MADlib
 - academic prototypes: Morpheus, libFM
- models: decision trees, factorization machines, k-means, ...

This is to be contrasted with the scalability of DBMSs!

How to Achieve This Performance Improvement?

Idea 1: Turn the Learning Problem into a Database Problem



Through DB Glasses, Everything is a Batch of Queries



11/44

Models under Consideration

So far:

- Polynomial regression
- Factorization machines
- Classification/regression trees
- Mutual information
- Chow Liu trees
- k-means clustering
- k-nearest neighbours
- PCA
- SVM

On-going:

- Boosting regression trees
- AdaBoost
- Sum-product networks
- Random forests
- Logistic regression
- Linear algebra:
 - QR decomposition
 - SVD
 - low-rank matrix factorisation

Models under Consideration

So far:

- Polynomial regression
- Factorization machines
- Classification/regression trees
- Mutual information
- Chow Liu trees
- k-means clustering
- k-nearest neighbours
- PCA
- SVM

On-going:

- Boosting regression trees
- AdaBoost
- Sum-product networks
- Random forests
- Logistic regression
- Linear algebra:
 - QR decomposition
 - SVD
 - low-rank matrix factorisation

All these cases can benefit from structure-aware computation

Linear function: $f_{\theta}(\mathbf{x}) = \langle \theta, \mathbf{x} \rangle = \theta_0 x_0 + \theta_1 x_1 + \dots$

- Training dataset D defined by feature extraction query and consists of tuples (x, y) of feature vector x and response y
- Parameters heta obtained by minimising the objective function:

$$J(\boldsymbol{\theta}) = \underbrace{\frac{1}{2|D|} \sum_{(\mathbf{x}, y) \in D} (\langle \boldsymbol{\theta}, \mathbf{x} \rangle - y)^2}_{(\mathbf{x}, y) \in D} + \underbrace{\frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2}_{\ell_2 - \text{regulariser}}$$

We can solve $\theta^* := \arg \min_{\theta} J(\theta)$ by repeatedly updating θ in the direction of the gradient until convergence:

$$\boldsymbol{\theta} := \boldsymbol{\theta} - \boldsymbol{\alpha} \cdot \boldsymbol{\nabla} \boldsymbol{J}(\boldsymbol{\theta}).$$

Model reformulation idea: Decouple

- data-dependent (\mathbf{x}, y) computation from
- data-independent (θ) computation

in the formulations of the objective $J(\theta)$ and its gradient $\nabla J(\theta)$.

From Optimisation to Query Batch

$$J(\boldsymbol{\theta}) = \frac{1}{2|D|} \sum_{(\mathbf{x},y)\in D} \left(\langle \boldsymbol{\theta}, \mathbf{x} \rangle - y \right)^2 + \frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2$$

$$=rac{1}{2}oldsymbol{ heta}^{ op}oldsymbol{\Sigma}oldsymbol{ heta}-\langleoldsymbol{ heta},oldsymbol{c}
angle+rac{s_Y}{2}$$

$$\nabla J(\theta) = \Sigma \theta - \mathbf{c} + \lambda \theta$$
, where

From Optimisation to Query Batch

$$J(\boldsymbol{\theta}) = \frac{1}{2|D|} \sum_{(\mathbf{x},y)\in D} \left(\langle \boldsymbol{\theta}, \mathbf{x} \rangle - y \right)^2 + \frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2$$

$$=\frac{1}{2}\boldsymbol{\theta}^{\top}\boldsymbol{\Sigma}\boldsymbol{\theta}-\langle\boldsymbol{\theta},\mathbf{c}\rangle+\frac{s_{Y}}{2}$$

$$\nabla J(\theta) = \Sigma \theta - \mathbf{c} + \lambda \theta$$
, where

matrix $\Sigma = (\sigma_{ij})_{i,j \in [n]}$, vector $\mathbf{c} = (c_i)_{i \in [n]}$, and scalar s_Y are:

$$\boldsymbol{\sigma}_{ij} = \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} \mathbf{x}_i \mathbf{x}_j^\top \qquad \mathbf{c}_i = \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} y \cdot \mathbf{x}_i \qquad s_Y = \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} y^2$$

15/44

Queries for
$$\sigma_{ij} = \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} \mathbf{x}_i \mathbf{x}_j^{\top}$$
 (w/o factor $\frac{1}{|D|}$):

Queries for
$$\sigma_{ij} = \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} \mathbf{x}_i \mathbf{x}_j^{\top}$$
 (w/o factor $\frac{1}{|D|}$):

• x_i, x_j continuous

SELECT SUM $(x_i * x_j)$ FROM D;

where D is the feature extraction query over the input DB.

Queries for
$$\sigma_{ij} = \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} \mathbf{x}_i \mathbf{x}_j^{\top}$$
 (w/o factor $\frac{1}{|D|}$):

• x_i, x_j continuous

SELECT SUM $(x_i * x_j)$ FROM D;

x_i categorical, x_j continuous
 SELECT x_i, SUM(x_i) FROM D GROUP BY x_i;

where D is the feature extraction query over the input DB.

Queries for
$$\sigma_{ij} = \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} \mathbf{x}_i \mathbf{x}_j^{\top}$$
 (w/o factor $\frac{1}{|D|}$):

• x_i, x_j continuous

SELECT SUM $(x_i * x_j)$ FROM D;

- x_i categorical, x_j continuous
 SELECT x_i, SUM(x_j) FROM D GROUP BY x_i;
- x_i, x_j categorical

SELECT x_i , x_j , SUM(1) FROM D GROUP BY x_i , x_j ;

where D is the feature extraction query over the input DB.

Query batch sizes in practice:

| Application/Dataset | Retailer | Favorita | Yelp | TPC-DS |
|--------------------------|----------|----------|-------|--------|
| Covariance Matrix | 937 | 157 | 730 | 3,299 |
| Decision Tree (one node) | 3,150 | 273 | 1,392 | 4,299 |
| k-means | 44 | 19 | 38 | 92 |

Query batch sizes in practice:

| Application/Dataset | Retailer | Favorita | Yelp | TPC-DS |
|--------------------------|----------|----------|-------|--------|
| Covariance Matrix | 937 | 157 | 730 | 3,299 |
| Decision Tree (one node) | 3,150 | 273 | 1,392 | 4,299 |
| <i>k</i> -means | 44 | 19 | 38 | 92 |

Queries in a batch:

- Same aggregates but over different attributes
- Expressed over the same join of the database relations

AMPLE opportunities for sharing computation in a batch.

Natural Attempt: Use Existing DB Technology to Compute the Query Batch

Existing DBMSs are NOT Designed for Query Batches



C = Covariance Matrix; R = Regression Tree Node; AWS d2.xlarge (4 vCPUs, 32GB)

Existing DSMSs are NOT Designed for Query Batches

Task: Maintain the covariance matrix over Retailer

- Round-robin insertions in all relations
- All maintenance strategies implemented in DBToaster



Azure DS14, Intel Xeon, 2.40GHz, 112GB, 1 thread; one hour timeout

Idea 2: Exploit the Problem Structure to Lower the Complexity



Structure-aware Tools of a Database Researcher

Algebraic structure: (semi)rings $(\mathcal{R}, +, *, 0, 1)$

• Distributivity law \rightarrow Factorisation

Factorised Databases

[VLDB'12+'13]

Factorised Machine Learning

[SIGMOD'16,VLDB'16,SIGREC'16,DEEM'18,PODS'18+'19]
Algebraic structure: (semi)rings $(\mathcal{R}, +, *, 0, 1)$

- Distributivity law \rightarrow Factorisation
 - Factorised Databases

[VLDB'12+'13]

Factorised Machine Learning

[SIGMOD'16,VLDB'16,SIGREC'16,DEEM'18,PODS'18+'19]

• Additive inverse \rightarrow Uniform treatment of updates

Factorised Incremental Maintenance

[SIGMOD'18]

Algebraic structure: (semi)rings $(\mathcal{R}, +, *, 0, 1)$

- Distributivity law \rightarrow Factorisation
 - Factorised Databases

[VLDB'12+'13]

Factorised Machine Learning

[SIGMOD'16,VLDB'16,SIGREC'16,DEEM'18,PODS'18+'19]

• Additive inverse \rightarrow Uniform treatment of updates

Factorised Incremental Maintenance

[SIGMOD'18]

 Sum-Product abstraction → Same processing for distinct tasks DB queries, Covariance matrix, PGM inference, Matrix chain multiplication [SIGMOD'18]

Combinatorial structure: query width and data degree measures

• Width measure w for FEQ \rightarrow Low complexity $\tilde{O}(N^w)$

factorisation width \geq fractional hypertree width \geq sharp-submodular width worst-case optimal size and time for factorised joins

[ICDT'12+'18,TODS'15,PODS'19]

Combinatorial structure: query width and data degree measures

• Width measure w for FEQ \rightarrow Low complexity $\tilde{O}(N^w)$

factorisation width \geq fractional hypertree width \geq sharp-submodular width worst-case optimal size and time for factorised joins

[ICDT'12+'18,TODS'15,PODS'19]

• Degree \rightarrow Adaptive processing depending on high/low degrees

worst-case optimal incremental maintenance for triangle count [ICDT'19a] evaluation of queries with negated relations of bounded degree [ICDT'19b]

Combinatorial structure: query width and data degree measures

• Width measure w for FEQ \rightarrow Low complexity $\tilde{O}(N^w)$

factorisation width \geq fractional hypertree width \geq sharp-submodular width worst-case optimal size and time for factorised joins

[ICDT'12+'18,TODS'15,PODS'19]

• Degree \rightarrow Adaptive processing depending on high/low degrees

worst-case optimal incremental maintenance for triangle count[ICDT'19a]evaluation of queries with negated relations of bounded degree[ICDT'19b]

 Functional dependencies → Learn simpler, equivalent models reparameterisation of polynomial regression models and factorisation machines [PODS'18,TODS'19]

Statistical structure: sampling for joins and models, coresets

• Sampling through joins: ripple/wander joins [SIGMOD'99+'16]



• Sampling specific to classes of models

[SIGMOD'19]



• Succinct approximate data representations

coresets for k-means with constant-factor approximation

22/44[submission'19]

Case in Point (1): Factorised Query Computation

 \Rightarrow

Exponential Time/Size Improvement

Example: Factorised Query Computation

| Orders (O for short) | | | Dish (D for short) | | Items (I for short) | |
|----------------------|----------|--------|--------------------|---------|---------------------|-------|
| customer | day dish | | dish | item | item | price |
| Elise | Monday | burger | burger | patty | patty | 6 |
| Elise | Friday | burger | burger | onion | onion | 2 |
| Steve | Friday | hotdog | burger | bun | bun | 2 |
| Joe | Friday | hotdog | hotdog | bun | sausage | 4 |
| | | | hotdog | onion | | |
| | | | hotdog | sausage | | |

Example: Factorised Query Computation

| Orders (O for short) | | | Dish (D | Dish (D for short) | | Items (I for short) | |
|----------------------|--------|----------|---------|--------------------|---------|---------------------|--|
| customer | day | day dish | | item | item | price | |
| Elise | Monday | burger | burger | patty | patty | 6 | |
| Elise | Friday | burger | burger | onion | onion | 2 | |
| Steve | Friday | hotdog | burger | bun | bun | 2 | |
| Joe | Friday | hotdog | hotdog | bun | sausage | 4 | |
| | | | hotdog | onion | | | |
| | | | hotdog | sausage | | | |

Consider the natural join of the above relations:

O(customer, day, dish), D(dish, item), I(item, price)

| customer | day | dish | item | price |
|----------|--------|--------|-------|-------|
| Elise | Monday | burger | patty | 6 |
| Elise | Monday | burger | onion | 2 |
| Elise | Monday | burger | bun | 2 |
| Elise | Friday | burger | patty | 6 |
| Elise | Friday | burger | onion | 2 |
| Elise | Friday | burger | bun | 2 |
| | | | | |

23/44

Example: Factorised Query Computation

| O(customer, day, dish), D(dish, item), I(item, price) | | | | | | | |
|---|--------|--------|-------|-------|--|--|--|
| customer | day | dish | item | price | | | |
| Elise | Monday | burger | patty | 6 | | | |
| Elise | Monday | burger | onion | 2 | | | |
| Elise | Monday | burger | bun | 2 | | | |
| Elise | Friday | burger | patty | 6 | | | |
| Elise | Friday | burger | onion | 2 | | | |
| Elise | Friday | burger | bun | 2 | | | |
| | | | | | | | |

An algebraic encoding uses product (×), union (U), and values:

| Elise | \times | Monday | \times | burger | \times | patty | \times | 6 | U | |
|-------|----------|--------|----------|--------|----------|-------|----------|---|--------------|-------|
| Elise | × | Monday | × | burger | × | onion | × | 2 | U | |
| Elise | × | Monday | × | burger | × | bun | × | 2 | U | |
| Elise | × | Friday | × | burger | × | patty | × | 6 | U | |
| Elise | × | Friday | × | burger | × | onion | × | 2 | U | |
| Elise | × | Friday | × | burger | × | bun | × | 2 | $\cup \dots$ | 24/44 |

Factorised Join



There are several algebraically equivalent factorised joins defined by distributivity of product over union and their commutativity.

.. Now with Further Compression



Observation:

- price is under item, which is under dish, but only depends on item,
- .. so the same price appears under an item *regardless* of the dish.

Idea: Cache price for a specific item and avoid repetition!

Factorising the Computation of Aggregates (1/2)



COUNT(*) computed in one pass over the factorisation:

- $\bullet \ \text{values} \mapsto 1\text{,}$
- $\cup \mapsto +$, $\times \mapsto *$.

Factorising the Computation of Aggregates (1/2)



COUNT(*) computed in one pass over the factorisation:

- $\bullet \ \text{values} \mapsto 1\text{,}$
- $\cup \mapsto +$, $\times \mapsto *$.

Factorising the Computation of Aggregates (2/2)



SUM(dish * price) computed in one pass over the factorisation:

- Assume there is a function f that turns dish into numbers.
- All values except for dish & price \mapsto 1,
- $\cup \mapsto +$, $\times \mapsto *$.

Factorising the Computation of Aggregates (2/2)



SUM(dish * price) computed in one pass over the factorisation:

- Assume there is a function *f* that turns dish into numbers.
- All values except for dish & price \mapsto 1,

•
$$\cup \mapsto +$$
, $\times \mapsto *$.

Case in Point (2): Sum-Product Ring Abstraction

 \Rightarrow

Sharing Aggregate Computation

Shared Computation of Several Aggregates (1/2)



Ring for computing SUM(1), SUM(price), SUM(price * dish):

- Elements = triples of numbers, one per aggregate
- Sum (+) and product (*) now defined over triples They enable shared computation across the aggregates

Shared Computation of Several Aggregates (2/2)



Ring for computing SUM(1), SUM(price), SUM(price * dish):

- Elements = triples of numbers, one per aggregate
- Sum (+) and product (*) now defined over triples They enable shared computation across the aggregates

Ring Generalisation for the Entire Covariance Matrix

Ring $(\mathcal{R}, +, *, \mathbf{0}, \mathbf{1})$ over triples of aggregates $(c, \mathbf{s}, \mathbf{Q}) \in \mathcal{R}$:

$$\begin{pmatrix} \Box , \bullet , \bullet , \bullet \\ \mathsf{SUM}(1) & \mathsf{SUM}(x_i) & \mathsf{SUM}(x_i^*x_j) \end{pmatrix}$$

$$(c_{1}, \mathbf{s}_{1}, \mathbf{Q}_{1}) + (c_{2}, \mathbf{s}_{2}, \mathbf{Q}_{2}) = (c_{1} + c_{2}, \mathbf{s}_{1} + \mathbf{s}_{2}, \mathbf{Q}_{1} + \mathbf{Q}_{2})$$

$$(c_{1}, \mathbf{s}_{1}, \mathbf{Q}_{1}) * (c_{2}, \mathbf{s}_{2}, \mathbf{Q}_{2}) = (c_{1} \cdot c_{2}, c_{2} \cdot \mathbf{s}_{1} + c_{1} \cdot \mathbf{s}_{2},$$

$$c_{2} \cdot \mathbf{Q}_{1} + c_{1} \cdot \mathbf{Q}_{2} + \mathbf{s}_{1}\mathbf{s}_{2}^{T} + \mathbf{s}_{2}\mathbf{s}_{1}^{T})$$

$$\mathbf{0} = (0, \mathbf{0}_{n \times 1}, \mathbf{0}_{n \times n})$$

$$\mathbf{1} = (1, \mathbf{0}_{n \times 1}, \mathbf{0}_{n \times n})$$

- SUM(1) reused for all SUM(x_i) and SUM($x_i * x_j$)
- $SUM(x_i)$ reused for all $SUM(x_i * x_j)$

Idea 3: Lower the Constant Factors



Engineering Tools of a Database Researcher

1. Specialisation for workload and data

Generate code specific to the query batch and dataset Improve cache locality for hot data path

2. Sharing low-level data access

Aggregates decomposed into views over join tree Share data access across views with different output schemas

3. Parallelisation: multi-core (SIMD & distribution to come)

Task and domain parallelism

[DEEM'18,SIGMOD'19]

Engineering Tools of a Database Researcher





Added optimisations for covariance matrix computation:

specialisation \rightarrow + sharing \rightarrow + parallelization

AWS d2.xlarge (4 vCPUs, 32GB)

Case in Point (3): Code Optimisations

 \Rightarrow

Non-trivial Speedup

Sharing Computation for a Query Batch in Favorita

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



Sharing Computation for a Query Batch in Favorita

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



Aggregates to compute over the join of relations:

 $\begin{array}{ll} Q_1: \ \text{SUM} \ (f(\text{units})) \\ Q_2: \ \text{SUM} \ (g(\text{item}) \cdot h(\textit{date},\textit{family})) & \quad \text{GROUP BY store} \\ Q_3: \ \text{SUM} \ (f(\text{units}) \cdot g(\text{item})) & \quad \text{GROUP BY family} \end{array}$

Shared Execution Plan for a Query Batch



- For each query, decide its output (root) node in the join tree
- Break down each query into directional views over the join tree
- Merge/share/group views for different queries
- Computational unit: group of views

Parallelisation: Dependency Graph of View Groups



- Task parallelism: Evaluate independent groups in parallel
- Domain parallelism: Partition the large relation used for each group



Traverse Sales as a trie following an order of its join variables



Lookup into incoming views, e.g., $V_l^{(1)}$, as early as possible

39/44

$$\begin{array}{ll} V_{l}^{(1)} & \alpha_{4} = 0; \\ \text{foreach } i \in \pi_{\text{item}}(S \Join_{\text{item}} V_{l}^{(1)} \bowtie_{\text{item}} V_{l}^{(2)}) \\ & \alpha_{l}^{(1)} = V_{l}^{(1)}(i) \\ & \alpha_{3} = 0; \\ \text{foreach } d \in \pi_{\text{date}}(\sigma_{\text{item}=i}S \bowtie_{\text{date}} V_{H} \bowtie_{\text{date}} V_{T}) \\ & \alpha_{H} = V_{H}(d); \\ & \alpha_{2} = 0; \\ \text{foreach } s \in \pi_{\text{store}}(\sigma_{\text{item}=i,\text{date}=d}S \bowtie_{\text{store}} \sigma_{\text{date}=d}V_{T}) \\ & \alpha_{T} = V_{T}(d,s); \quad \alpha_{1} = 0; \\ & \text{foreach } u \in \pi_{\text{units}}\sigma_{\text{item}=i,\text{date}=d,\text{store}=s}S : \alpha_{1} + = f(u); \\ & \alpha_{2} + = \alpha_{1} \cdot \alpha_{T}; \\ & \alpha_{3} + = \alpha_{2} \cdot \alpha_{H}; \\ & \alpha_{4} + = \alpha_{3} \cdot \alpha_{l}^{(1)} \\ & Q_{1} = \alpha_{4}; \end{array}$$

Insert code for partial aggregates as early as possible

39/44

```
\alpha_{4} = 0;
V_{I}^{(1)} \stackrel{\text{def}}{\to} \text{item} \qquad \text{foreach } i \in \pi_{\text{item}}(S \Join_{\text{item}} V_{I}^{(1)} \bowtie_{\text{item}} V_{I}^{(2)}) \\ \left| \begin{array}{c} \alpha_{I}^{(1)} = V_{I}^{(1)}(i); \end{array} \right|
                                           \alpha_3 = 0:
  V_H \rightarrow date
                                           foreach d \in \pi_{date}(\sigma_{item=i}S \bowtie_{date} V_H \bowtie_{date} V_T)
                                               \alpha_H = V_H(d);
                                                 \alpha_2 = 0:
   V_T \rightarrow \text{store}
                                                  foreach s \in \pi_{\text{store}}(\sigma_{\text{item}=i,\text{date}=d}S \bowtie_{\text{store}} \sigma_{\text{date}=d}V_T)
                                                       \alpha_T = V_T(d,s); \quad \alpha_1 = 0;
                                                      foreach u \in \pi_{\text{units}} \sigma_{\text{item}=i,\text{date}=d,\text{store}=s} S : \alpha_1 + = f(u);
                                                      \alpha_2 += \alpha_1 \cdot \alpha_T;
                                           \alpha_3 += \alpha_2 \cdot \alpha_H;
\alpha_4 += \alpha_3 \cdot \alpha_1^{(1)};
                                        Q_1 = \alpha_4:
Q_1: SUM (f(units))
```

$$V_{I}^{(1)} \qquad \alpha_{4} = 0;$$
foreach $i \in \pi_{item}(S \bowtie_{item} V_{I}^{(1)} \bowtie_{item} V_{I}^{(2)})$

$$\downarrow \alpha_{I}^{(1)} = V_{I}^{(1)}(i);$$

$$\alpha_{5} = g(i);$$

$$\alpha_{3} = 0;$$
foreach $d \in \pi_{date}(\sigma_{item=i}S \bowtie_{date} V_{H} \bowtie_{date} V_{T})$

$$\downarrow \alpha_{H} = V_{H}(d);$$

$$\alpha_{2} = 0;$$
foreach $s \in \pi_{store}(\sigma_{item=i,date=d}S \bowtie_{store} \sigma_{date=d}V_{T})$

$$\downarrow \alpha_{T} = V_{T}(d, s); \quad \alpha_{1} = 0;$$
foreach $u \in \pi_{units}\sigma_{item=i,date=d,store=s}S : \alpha_{1} + = f(u);$

$$\alpha_{3} + = \alpha_{2} \cdot \alpha_{H};$$

$$\alpha_{4} + = \alpha_{3} \cdot \alpha_{I}^{(1)}; \quad V_{S \to I}(i) = \alpha_{3} \cdot \alpha_{5};$$

$$Q_{1} = \alpha_{4};$$

$$V_{S \to I}: \text{ SUM } (f(\text{units}) \cdot g(\text{item})) \quad \text{GROUP BY item}$$

$$V_{I}^{(1)} \qquad \alpha_{4} = 0;$$

foreach $i \in \pi_{item}(S \bowtie_{item} V_{I}^{(1)} \bowtie_{item} V_{I}^{(2)})$

$$\alpha_{I}^{(1)} = V_{I}^{(1)}(i);$$

$$\alpha_{5} = g(i);$$

$$\alpha_{3} = 0;$$

foreach $d \in \pi_{date}(\sigma_{item=i}S \bowtie_{date} V_{H} \bowtie_{date} V_{T})$

$$\alpha_{H} = V_{H}(d); \quad \alpha_{6} = 0;$$

foreach $y \in \pi_{family}\sigma_{item=i}V_{I}^{(2)} : \alpha_{6} += h(d, y) \cdot V_{I}^{(2)}(i, y);$

$$\alpha_{2} = 0; \quad \alpha_{7} = \alpha_{6} \cdot \alpha_{5} \cdot \alpha_{H};$$

foreach $s \in \pi_{store}(\sigma_{item=i,date=d}S \bowtie_{store} \sigma_{date=d}V_{T})$

$$\begin{vmatrix} \alpha_{T} = V_{T}(d, s); & \alpha_{1} = 0; & \alpha_{8} = |\sigma_{item=i,date=d,store=s}S|;$$

foreach $u \in \pi_{units}\sigma_{item=i,date=d,store=s}S : \alpha_{1} += f(u);$

$$\alpha_{2} += \alpha_{1} \cdot \alpha_{T};$$

if $Q_{2}(s)$ then $Q_{2}(s) += \alpha_{7} \cdot \alpha_{8} \cdot \alpha_{T}$ else $Q_{2}(s) = \alpha_{7} \cdot \alpha_{8} \cdot \alpha_{T};$

$$\alpha_{4} += \alpha_{3} \cdot \alpha_{I}^{(1)}; \quad V_{S \rightarrow I}(i) = \alpha_{3} \cdot \alpha_{5};$$

 $Q_{1} = \alpha_{4};$
 $Q_{2}:$ SUM $(g(\text{item}) \cdot h(date, family))$ GROUP BY store

40/44

Three-step recipe for efficient machine learning over databases:

- 1. Turn the learning problem into a database problem
- 2. Exploit the problem structure to lower the complexity
- 3. Specialise and optimise the code to lower the constant factors

Q.E.D.

Call to arms

We need more sustained work on theory and systems for

Structure-aware Approaches to Data Analytics
References

[ICDT'12] Dan Olteanu and Jakub Závodný. *Factorised Representations of Query Results: Size Bounds and Readability.* In ICDT 2012

[VLDB'12] Nurzhan Bakibayev, Dan Olteanu, and Jakub Závodný. FDB: A Query Engine for Factorised Relational Databases. In VLDB 2012

[VLDB'13] Nurzhan Bakibayev, Tomáš Kočiský, Dan Olteanu, and Jakub Závodný. *Aggregation and Ordering in Factorised Databases*. In VLDB 2013

[TODS'15] Dan Olteanu and Jakub Závodný. *Size Bounds for Factorised Representations of Query Results.* In TODS Vol. 40(1):2 2015

[SIGMOD'16] Maximilian Schleich, Dan Olteanu, and Radu Ciucanu. *Learning Linear Regression Models over Factorized Joins*. In SIGMOD 2016

[SIGREC'16] Dan Olteanu and Maximilian Schleich. *Factorized Databases*. In SIGMOD Record Vol 45 (2) 2016

[VLDB'16] Dan Olteanu and Maximilian Schleich. F: Regression Models over Factorized Views. In VLDB 2016

[DEEM'18] Mahmoud Abo Khamis, Hung Q. Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. *AC/DC: In-Database Learning Thunderstruck.* In DEEM@SIGMOD 2018

[ICDT'18] Ahmet Kara and Dan Olteanu. *Covers of Query Results.* In ICDT 2018

[PODS'18] Mahmoud Abo Khamis, Hung Ngo, XuanLong Nguyen, Dan Olteanu, and Maximilian Schleich. *In-Database Learning with Sparse Tensors*. In PODS 2018

[SIGMOD'18] Milos Nikolic and Dan Olteanu. *Incremental View Maintenance with Triple Lock Factorisation Benefits*. In SIGMOD 2018

[ICDT'19a] Ahmet Kara, Hung Q. Ngo, Miklos Nikolic, Dan Olteanu and Haozhe Zhang. *Counting Triangles under Updates in Worst-Case Optimal Time*. In ICDT 2019

[ICDT'19b] Mahmoud Abo Khamis, Hung Ngo, Dan Olteanu and Dan Suciu. Boolean Tensor Decomposition for Conjunctive Queries with Negation. In ICDT 2019

[PODS'19] Mahmoud Abo Khamis, Ryan Curtin, Ben Moseley, Hung Ngo, Long Nguyen, Dan Olteanu, and Maximilian Schleich. *On Functional Aggregate Queries with Additive Inequalities.* In PODS 2019

[SIGMOD'19] Maximilian Schleich, Dan Olteanu, Mahmoud Abo Khamis, Hung Ngo, and Long Nguyen. *A Layered Aggregate Engine for Analytics Workloads.* In SIGMOD 2019