



Wolfgang Lehner

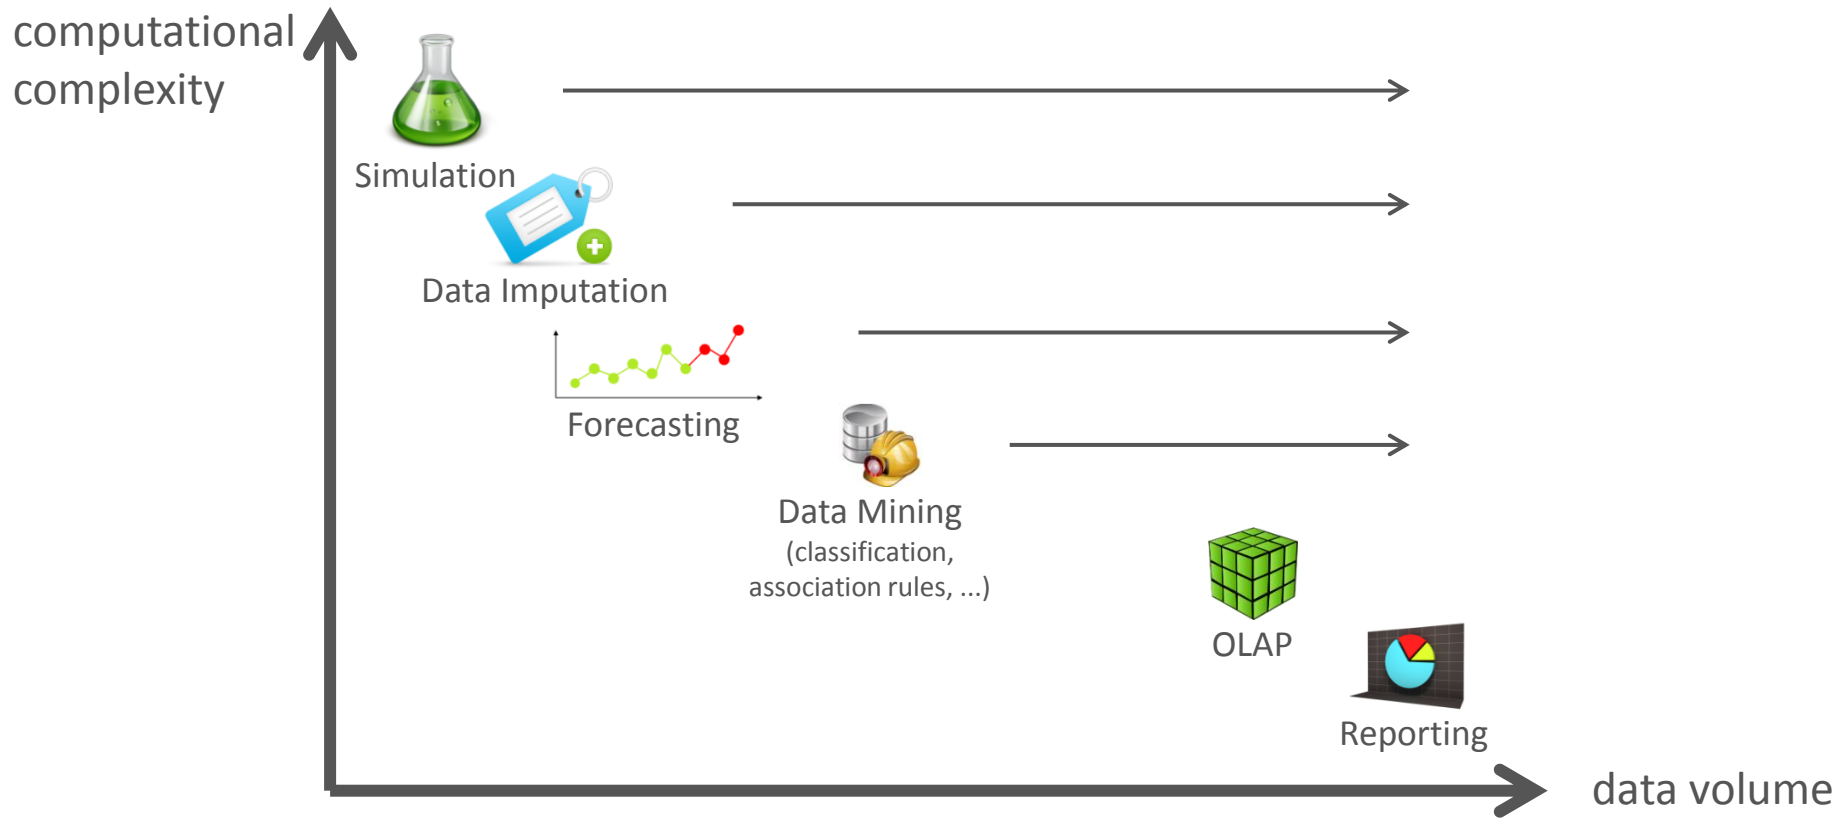
Forecasting and Data Imputation Strategies in Database Systems

11.07.2013

> Novel type of applications !!!



data crunching meets *number crunching*



> The NetFlix Competition





Netflix' star rating system helps determine personalized movie recommendations. Now the company is looking to outside developers to improve those recommendations.

BUSINESS

The \$1 Million Netflix Challenge

FRIDAY, OCTOBER 6, 2006 | BY KATE GREENE

VP Jim Bennett discusses how recommendation systems suggest your next movie and the challenges of building a better one.

[E-mail](#) [Audio](#) [Print](#)

Earlier this week, Netflix, the online movie rental service, announced it will award **\$1 million** to anyone who can come up with an algorithm that improves the accuracy of its movie recommendation service.

In doing so, the company is putting out a call to researchers who specialize in machine learning—the type of artificial intelligence used to build systems that recommend music, books, and movies. The entrant who can increase the accuracy of the Netflix recommendation system, which is called Cinematch, by 10 percent by 2011 will win the prize.

Recommendation systems such as those used by Netflix, Amazon, and other Web retailers are based on the principle that if two people enjoy the same product, they're likely to have other favorites in common too.

But behind this simple premise is a complex algorithm that incorporates millions of ratings, tens of thousands of items, and ever-changing relationships between user preferences.



BellKor's Pragmatic Chaos

> The NetFlix Competition (3)



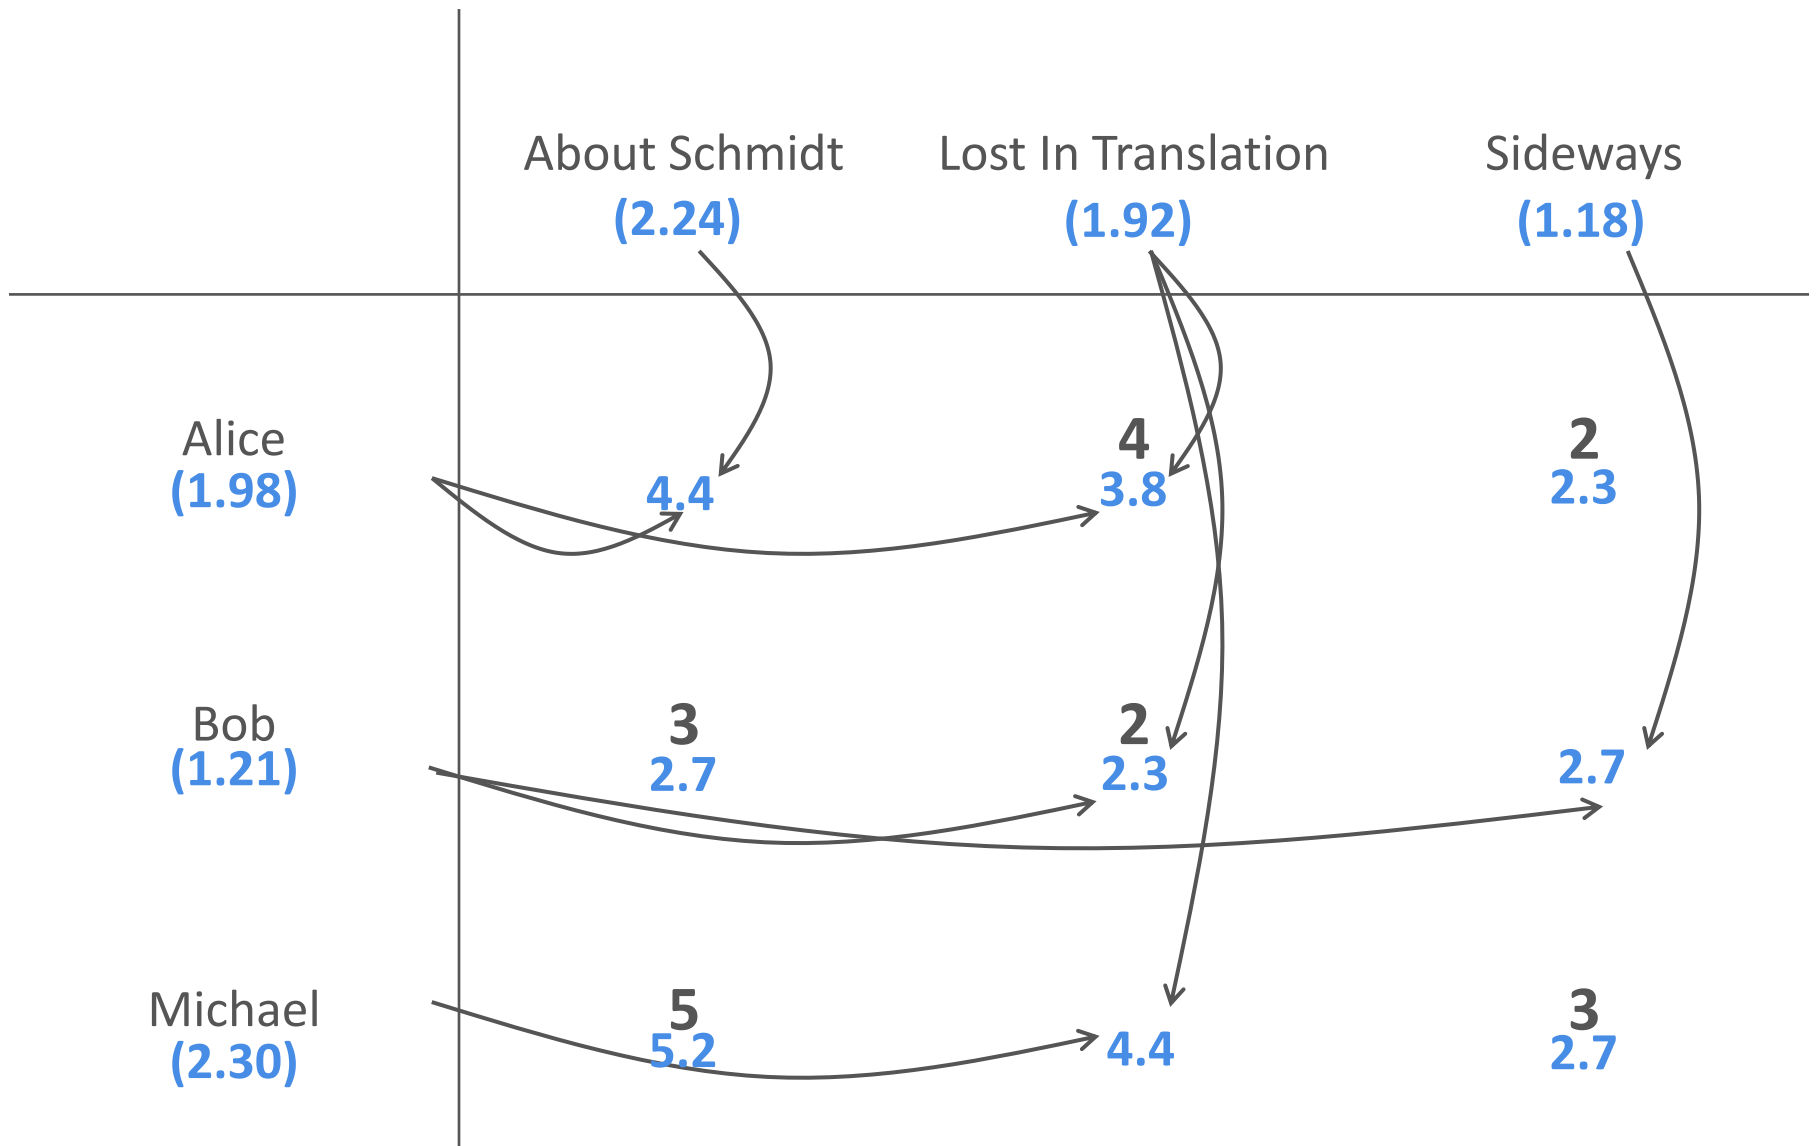
$$\hat{r}_{ui} = b_{ui} + |\mathbf{N}(u)|^{-\frac{1}{2}} \sum_{j \in \mathbf{N}(u)} e^{-\beta_u \cdot |t_{ui} - t_{uj}|} c_{ij} +$$

$$|\mathbf{R}(u)|^{-\frac{1}{2}} \sum_{j \in \mathbf{R}(u)} e^{-\beta_u \cdot |t_{ui} - t_{uj}|} ((r_{uj} - \tilde{b}_{uj}) w_{ij}) +$$

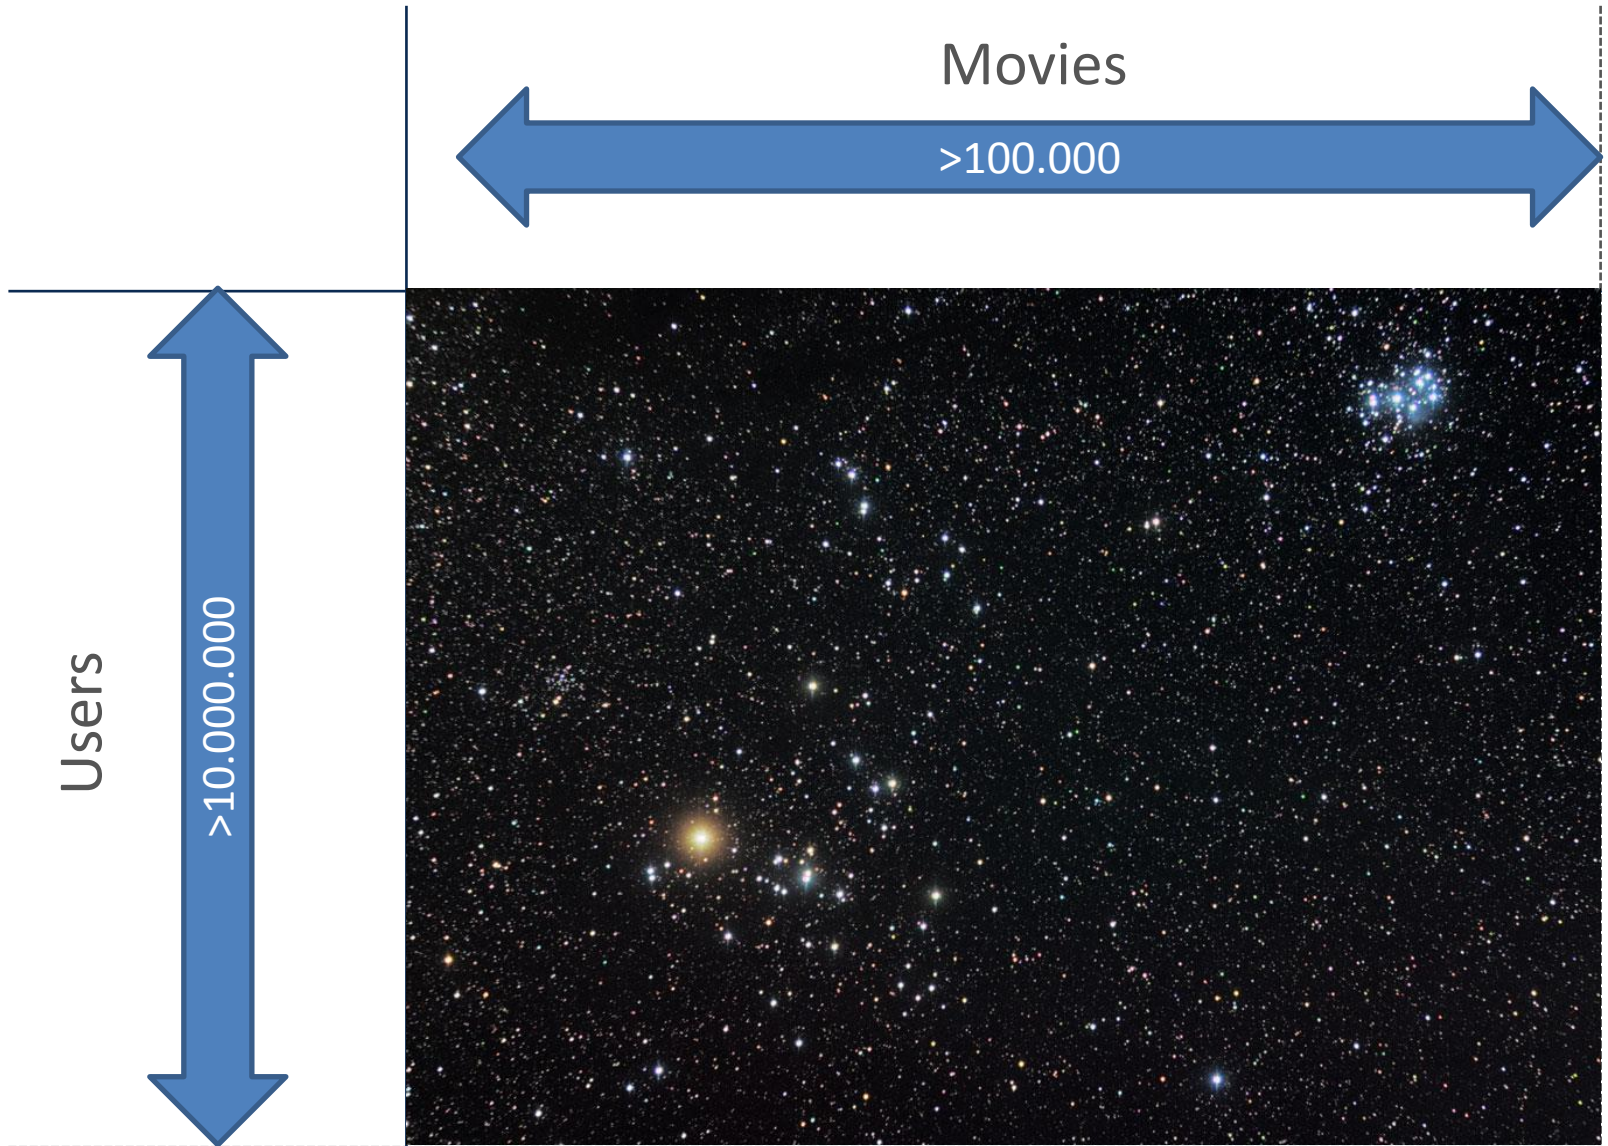
$$\sum_{j \in \mathbf{R}(u)} e^{-\gamma_u \cdot |t_{ui} - t_{uj}|} ((r_{uj} - \tilde{b}_{uj}) d_{ij}).$$



> The NetFlix Competition (4)



> The NetFlix Competition (6)



> A simple experiment ...



> ... our test object



color code := user rating

different movies

different users





Phase 1: drop 75% of all pixels






Phase 2: Random permutation of rows and columns



> The Experiment ...



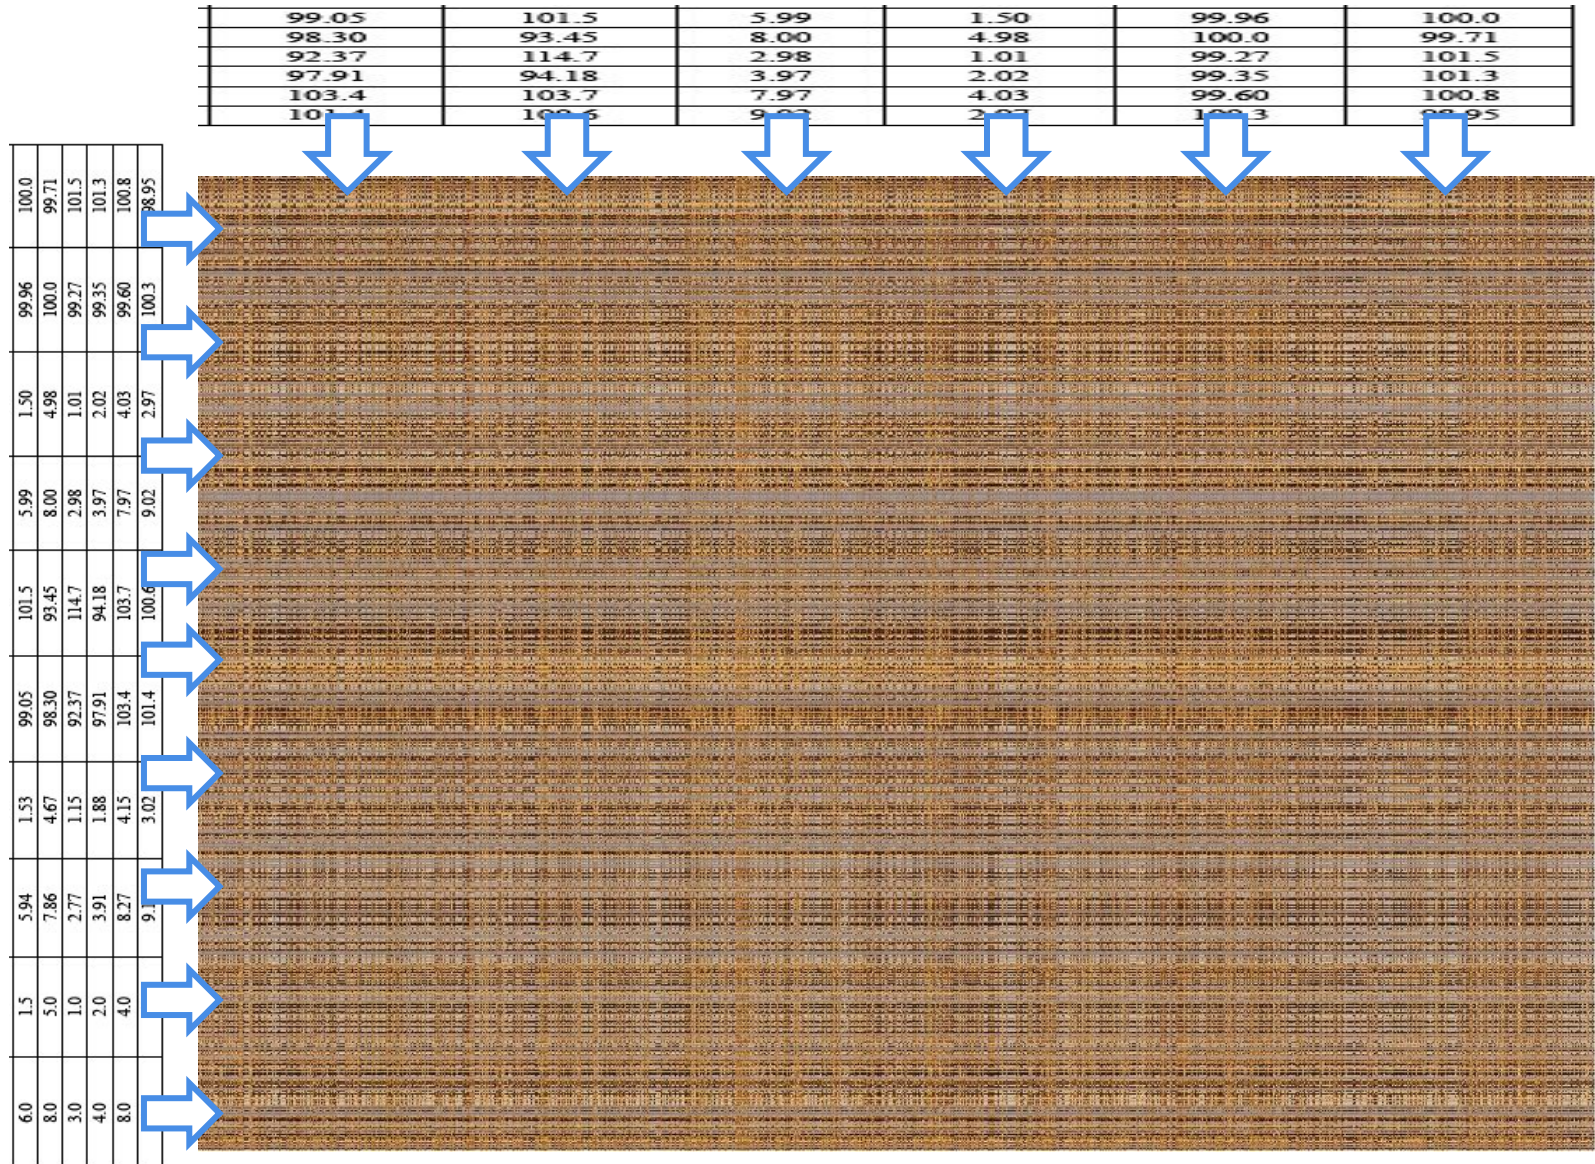
Phase 3: Determine the latent factors



> The Experiment ...



Phase 4: Reconstruction








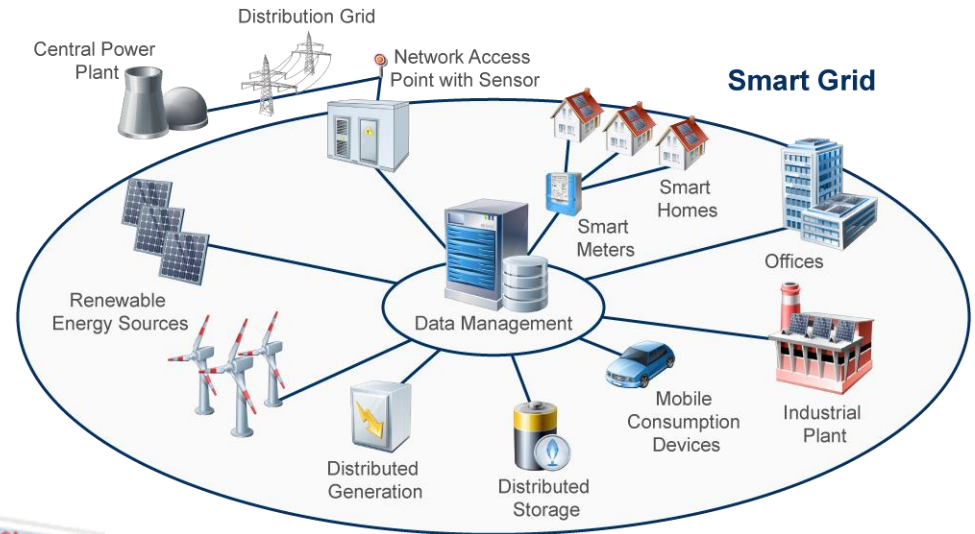

Phase 5: Final Result Generation



> The Experiment ...



> Time Series Forecasting



➔ Multi-Dimensional Time Series Data



“It’s tough to make predictions,
especially about the future.”
-- Mark Twain

Given

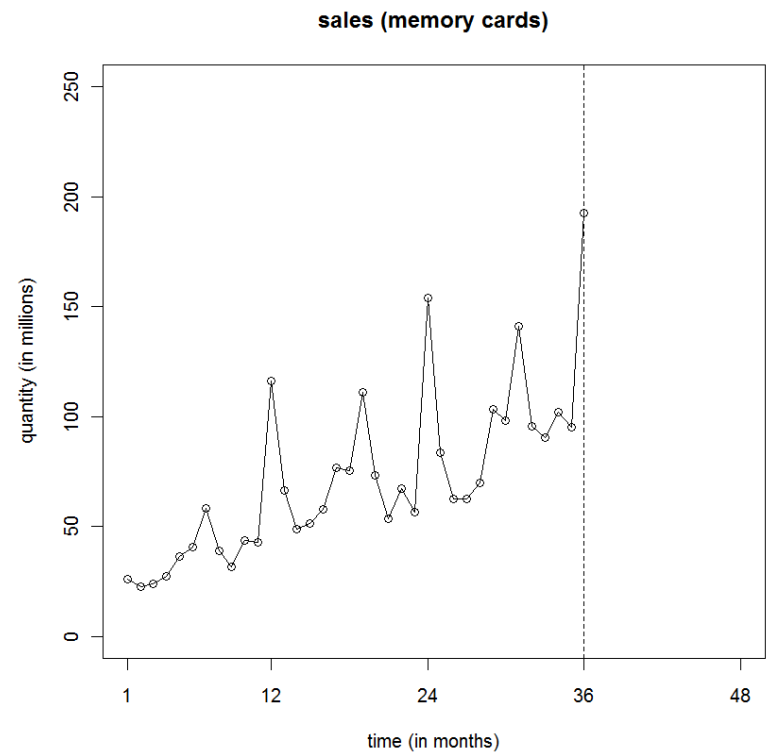
- Time series with numerical values as training data

Goal

- Predict future values for arbitrary future point in times (forecast horizon)
- Include trend and seasonality

Applications

- Planning of sales and budget
- Price development
- Inventory, manufacturing
- Climate, weather, environment
- Economic indicators
- Stocks





Diverse Time Series Data

- Find forecast model that minimizes the forecast error

Example: Exponential Smoothing Framework

- 30 variants (additional additive or multiplicative errors)

Trend/Season	No Season	Add. Season	Mult. Season
No Trend			
Add. Trend			
Add. Damp. Trend			
Mult. Trend			
Mult. Damp. Trend			



Mathematical Foundations

In-DBMS Time Series Forecasting

- Flash-Forward Query Project



Query Processing and Optimization

Model Configuration Advisor

- Forecasting the Data Cube
- Selection of Model Configurations

Notification-based Forecast Queries

Beyond Forecast Models

- Model Refinement



Mathematical Foundations

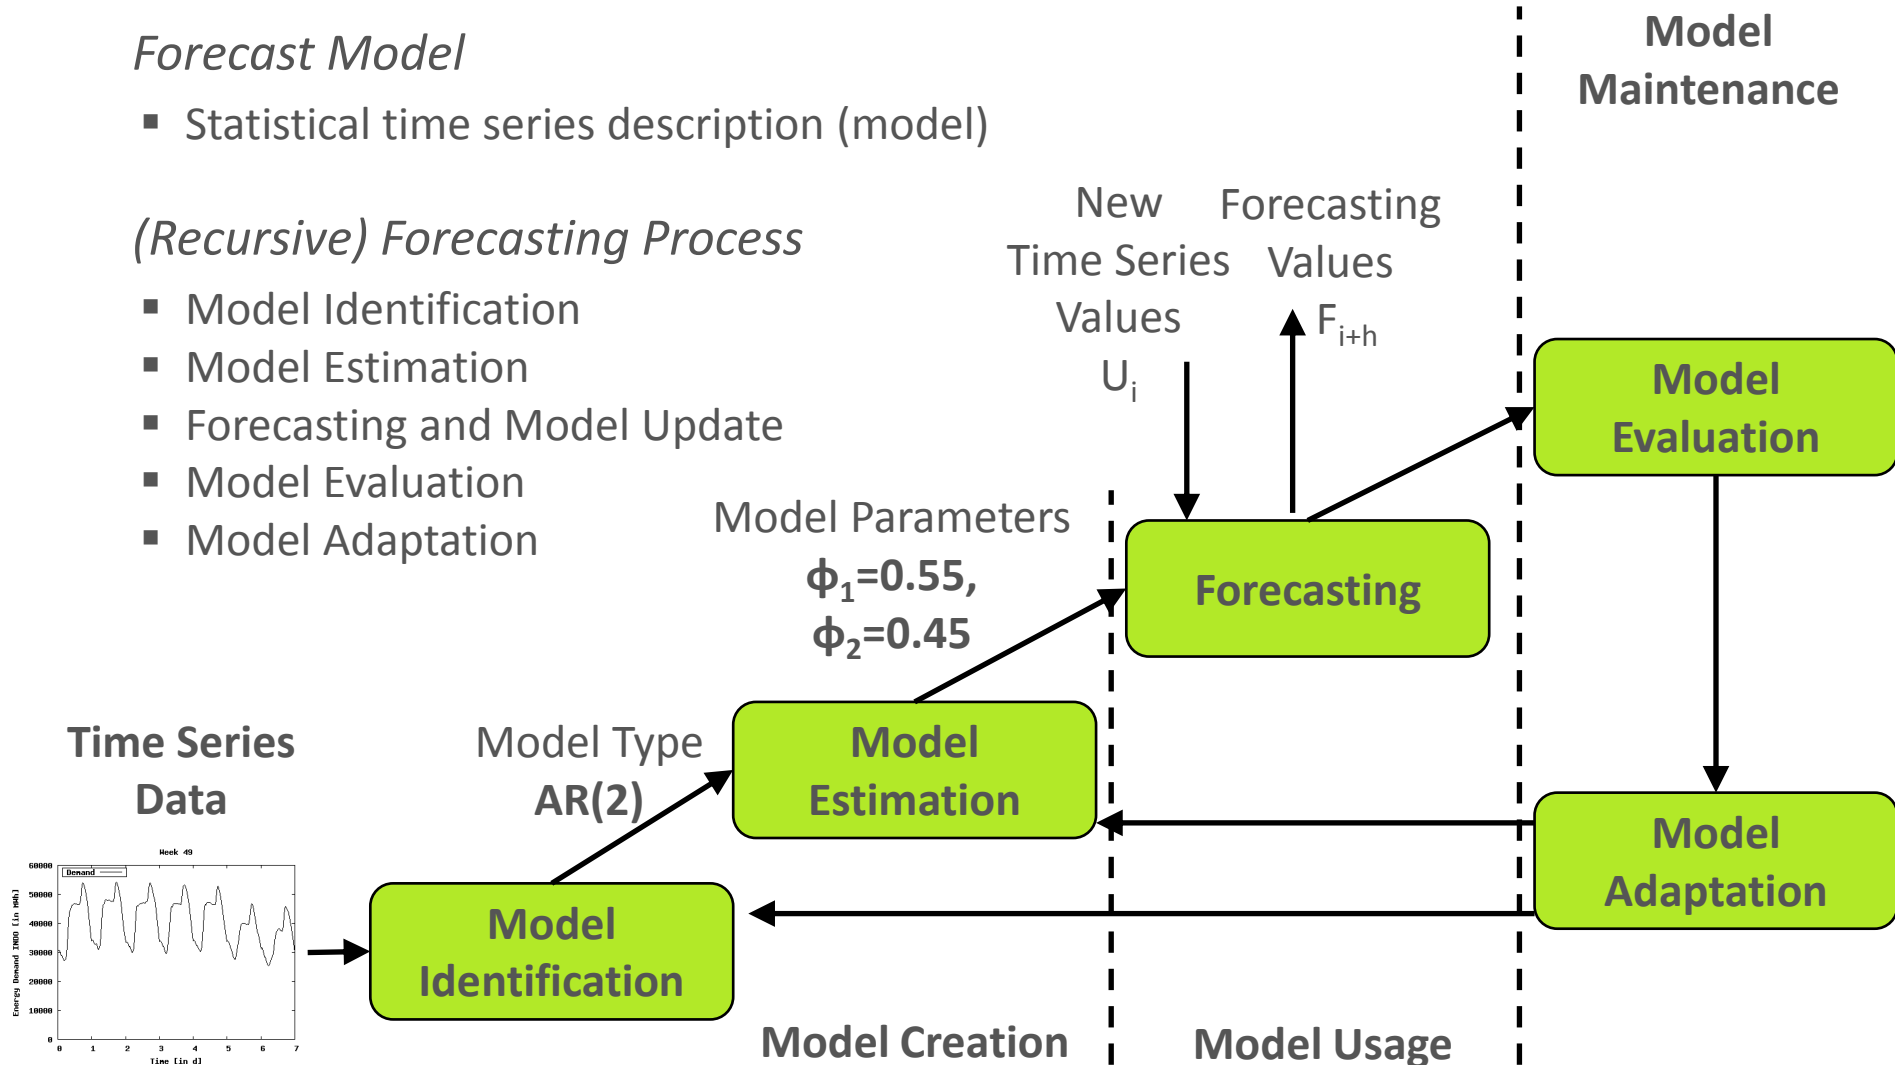


Forecast Model

- Statistical time series description (model)

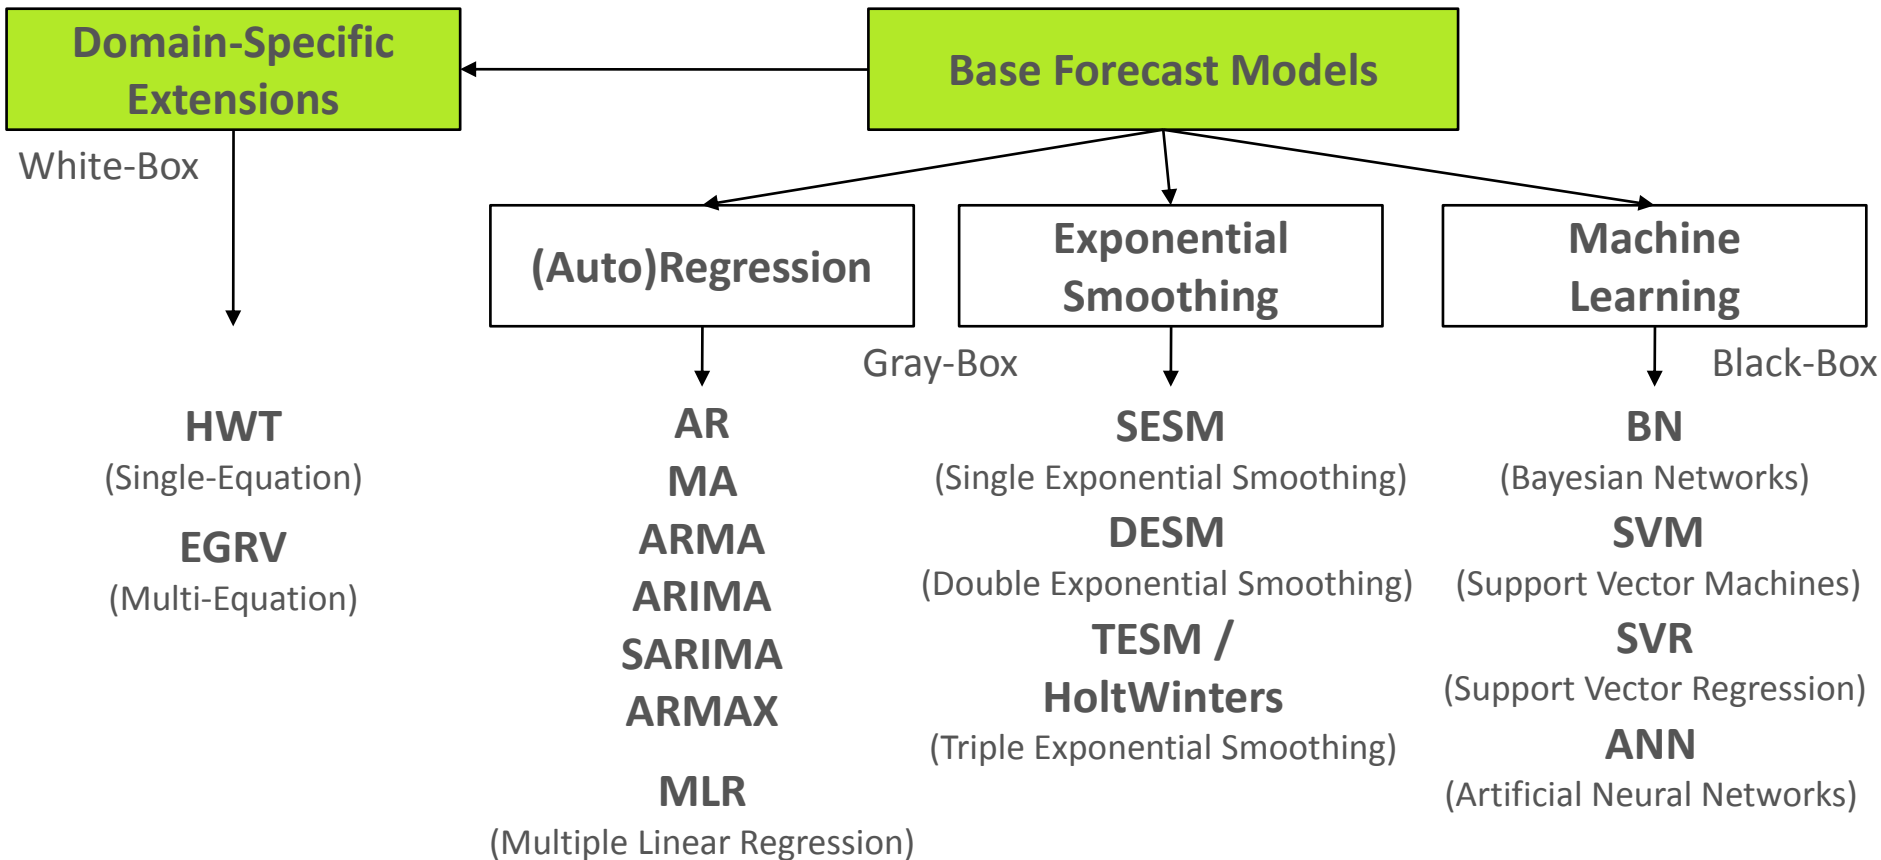

(Recursive) Forecasting Process

- Model Identification
- Model Estimation
- Forecasting and Model Update
- Model Evaluation
- Model Adaptation





Forecast Model Types / Classes



> Model Estimation



Problem

- Instantiate a forecast model w.r.t. meta model and training data set

Example




- Forecast Model Type **AR(2)**:

$$\hat{y}_t = \phi_1 \cdot y_{t-1} + \phi_2 \cdot y_{t-2}$$

- Error Metric: **MSE**

$$\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$


- Parameter Estimator **L-BFGS-B**






Forecasting

- Use the estimated forecast model $\hat{y}_t = 0.56 \cdot y_{t-1} + 0.23 \cdot y_{t-2}$
- Create h forecast values (forecast horizon)
- Update model state for new measurements (e.g., exponential smoothing)



Example Forecast (EGRV)



- $\text{SMAPE}_{\text{vshort}} = 0.0021$
- $\text{SMAPE}_{\text{long}} = 0.0755$





Model Evaluation

- Goal: Trigger model adaptation only if necessary
- Fixed Interval Techniques (# updates, time interval)
- Continuous Evaluation Techniques (threshold, on-demand)



Model Adaptation

- Goal: Adapt the forecast model to the changed time series (if necessary)
- Model Re-Identification
- Model Re-Estimation (old model as start point)



In-DBMS Time Series Forecasting





Magnetic

- „Attract data and practitioners“
- Use all available data sources independent of their quality



Agile

- „Rapid iteration: ingest, analyze, productionalize“
- Continuous and rapid evolution of physical and logical structures
- ELT (Extraction, Loading, Transformation)

Deep

- „Sophisticated analytics in Big Data“
- Extended algorithmic runtime environment
- Ad-hoc advanced analytics and statistics


1. mad skills

92 up,

To be able to do/perform amazing/unexpected things

I gots me mad skills, yo.


To be said after performing an extraordinary feat.







Integration of advanced analytics into scalable database management systems

- Traditionally forecasting is performed manually in external statistical systems
- Support of transparent and automatic in-DBMS forecasting




Traditional

Analysis
(e.g., R, SPSS)




In-DBMS Forecasting





Model-based Prediction

- We employ (classical) time series models (e.g. exponential smoothing, ARIMA)
- Used to
 - Explain time series from it's history
 - And—possibly—from exogenous inputs

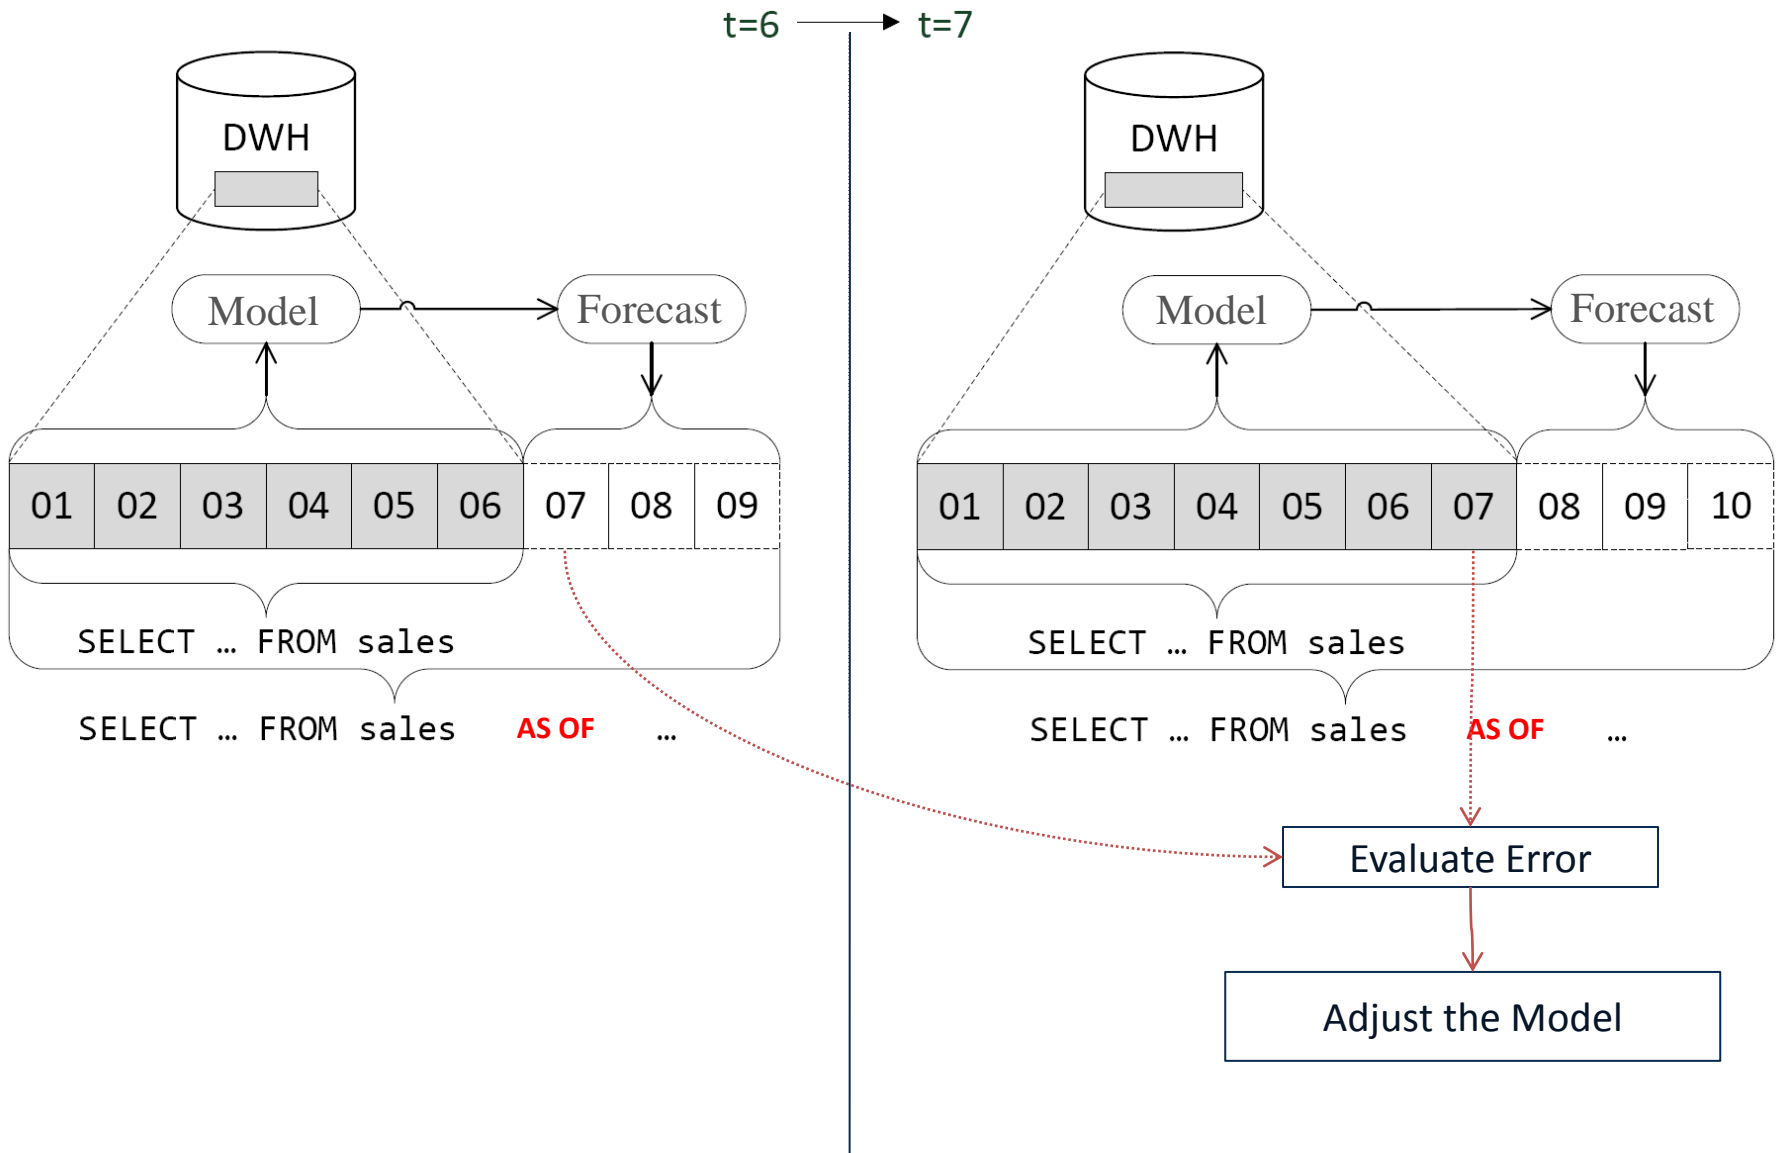


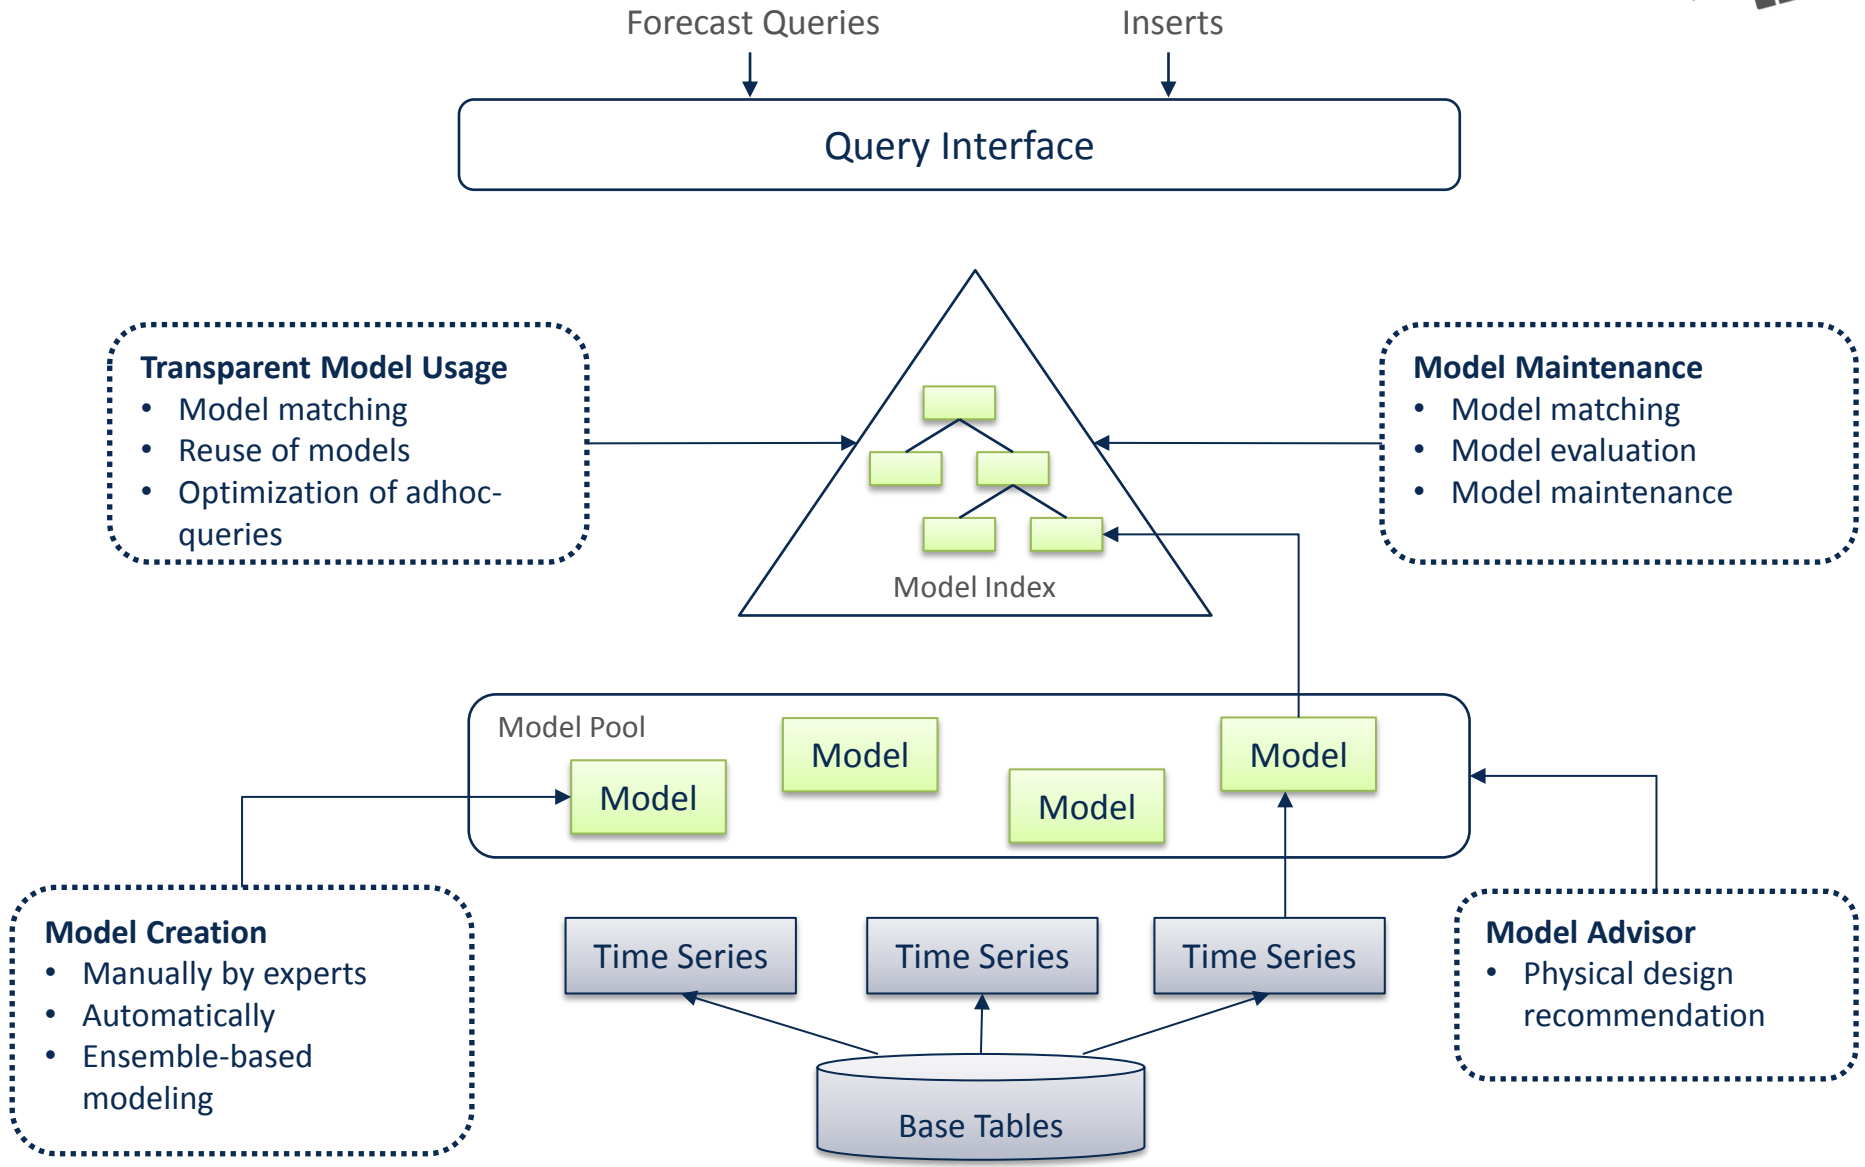
Related Work

- Customized functions with proprietary languages
 - SQL Server 2012: ARIMA, autoregressive trees
 - Oracle: exponential smoothing, non-linear regression
- Bi-directional communication
 - Reuse existing statistical tools (e.g. R)
 - SAP HANA, Oracle, IBM Netezza
- Model-based views

Key Elements

- Declarative querying
- Automatic model creation
- Automatic model maintenance
- Forecast model advisor







F2DB

Connect Disconnect Reconnect Load History Save History Reset Quit

Console Forecast Query Execution Plan Model Index Configuration Maintenance System Catalog History Preferences About

```
SELECT s_time,SUM(s_amount) measure
FROM sales
GROUP BY s_time
ORDER BY s_time
ASOF '2011-01-01'
ALGORITHM HOLTWINTERS
PARAMETERS (season=12)
```

Table: sales, Time Column: s_time, Measure Column: s_amount

Where ... explicit explicit

Load SQL Stmt... Save SQL Stmt... OptimProperties Execute!

10 ASOF 10 ASOF 2011-01-01 Use AutoArima Algorithm HOLTWINTERS With P.I. Storage

Parameters (season=12)

Elapsed Time:110.163907ms

Connection to 'MyConnection_db4711' established.



Console Forecast Query Execution Plan Model Index Configuration Maintenance System Catalog History Preferences About

sales greedy Alpha: 1 Create ModelIndex Explain! Equal node width

```

    graph TD
      root[root] --> Netherlands[Netherlands]
      root --> Portugal[Portugal]
      Netherlands --> MC17[17 MemoryCard]
      Netherlands --> PT18[18 PlasmaTV]
      Netherlands --> T19[19 [Total]]
      Portugal --> AD20[20 AudioDevice]
      Portugal --> MC21[21 MemoryCard]
      Portugal --> PT22[22 PlasmaTV]
  
```

ModelSelect

- Root
 - Belgium
 - England
 - Germany
 - Italy
 - Netherlands
 - Portugal
 - AudioDevice
 - amt on s_time
 - M m20
 - D m39
 - MemoryCard
 - PlasmaTV

ModelInfo

model name	m20
model type	hwmodel
error	'sse'
has_season	1
has_trend	1
season	12
training data type	select s_time, sum(s_amour model

of forecasts: 10

Query Model

Create Model

Copy Query to TSPlot

60.5 %

Connection to 'MyConnection_db4711' established.



Query Processing and Optimization



SQL Forecast Query

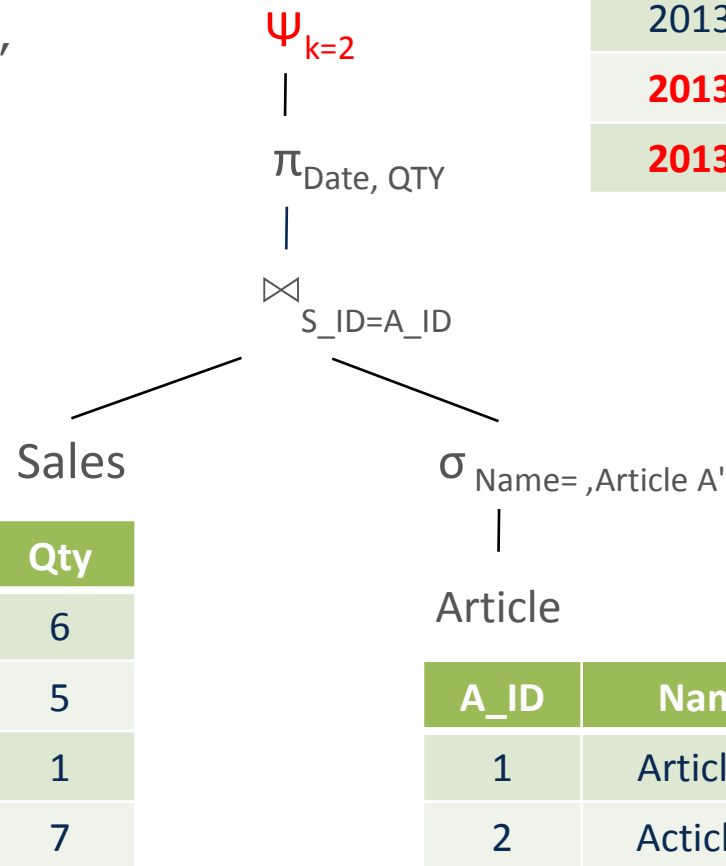
```
SELECT Date, Quantity
FROM Sales, Article
WHERE S_ID = A_ID AND
      Name = 'Article A'
AS OF 2013-05-02
```

Query Plan

- Forecast operator Ψ

Date	Qty
2013-04-28	6
2013-04-29	5
2013-04-30	7
2013-05-01	6
2013-05-02	6.5

...	S_ID	Date	Qty
...	1	2013-04-28	6
...	1	2013-04-29	5
...	2	2013-04-30	1
...	1	2013-04-30	7

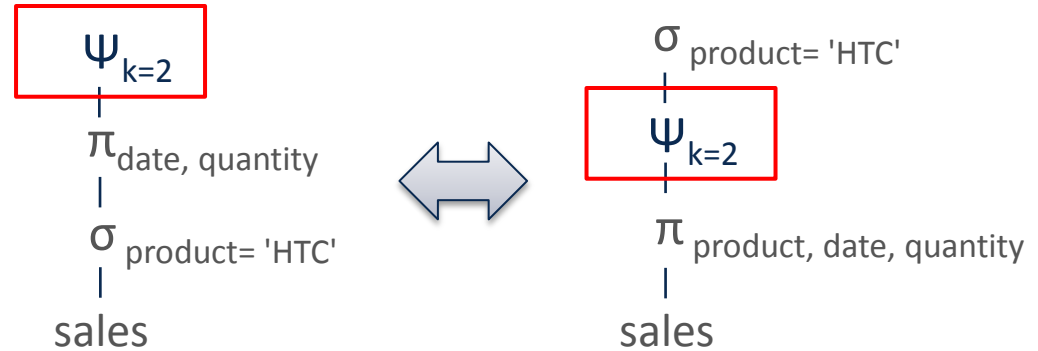


A_ID	Name	...
1	Article A	...
2	Article B	...



Selection invariance

- Condition filtering out whole clusters

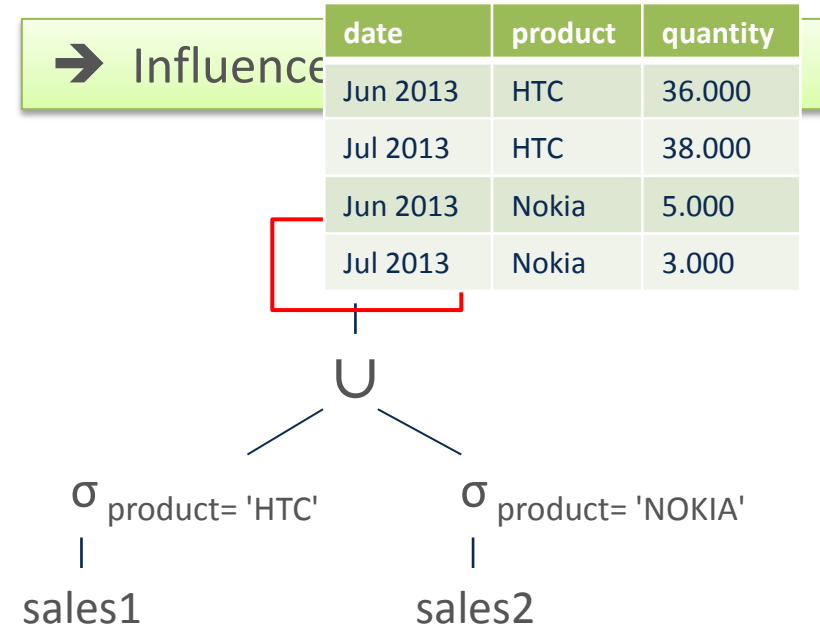



Projection invariance

- If neither time, measure or cluster attribute

Union invariance

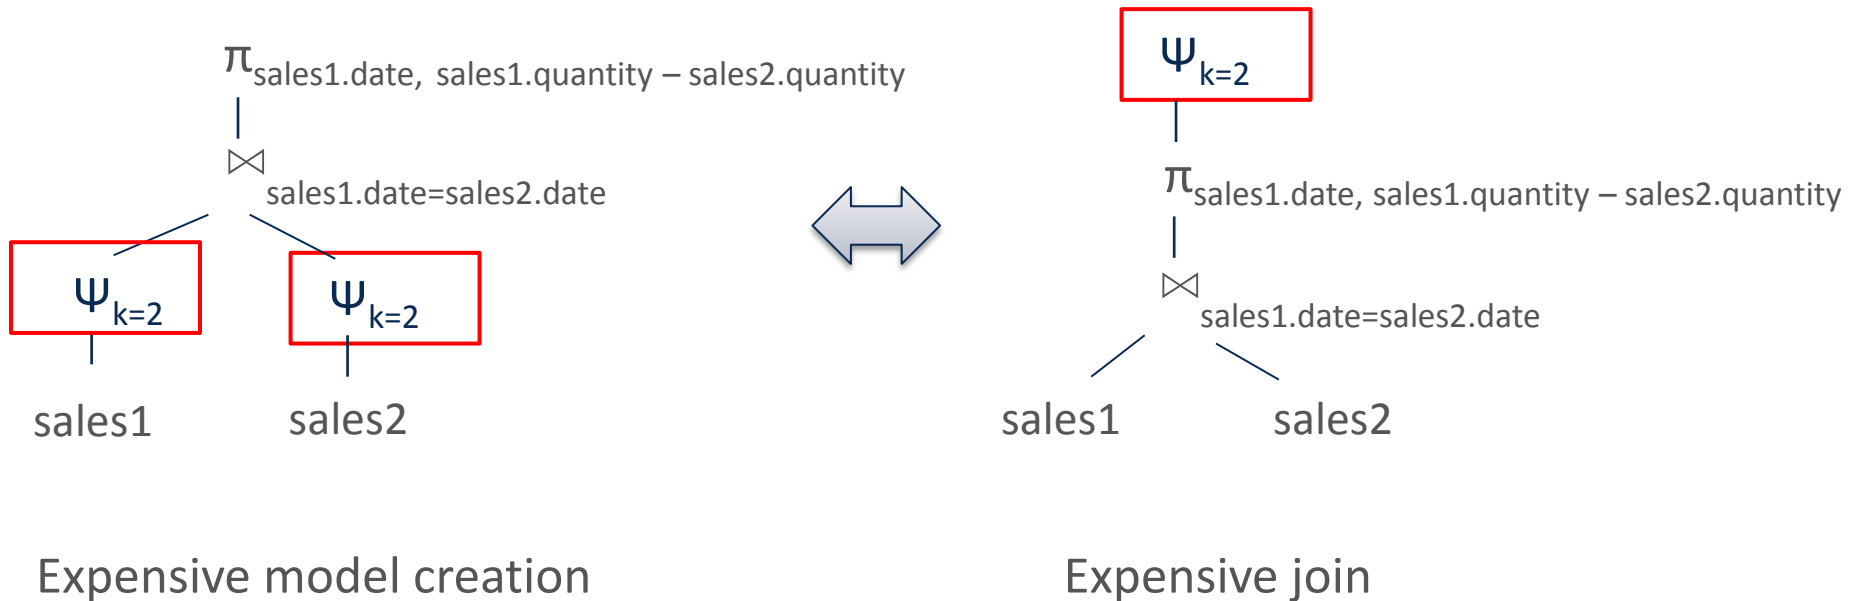
- On different relations





Restructuring might also influence accuracy

- E.g. join, aggregation




→ Cost model required: accuracy and runtime

> Example: Aggregation Queries

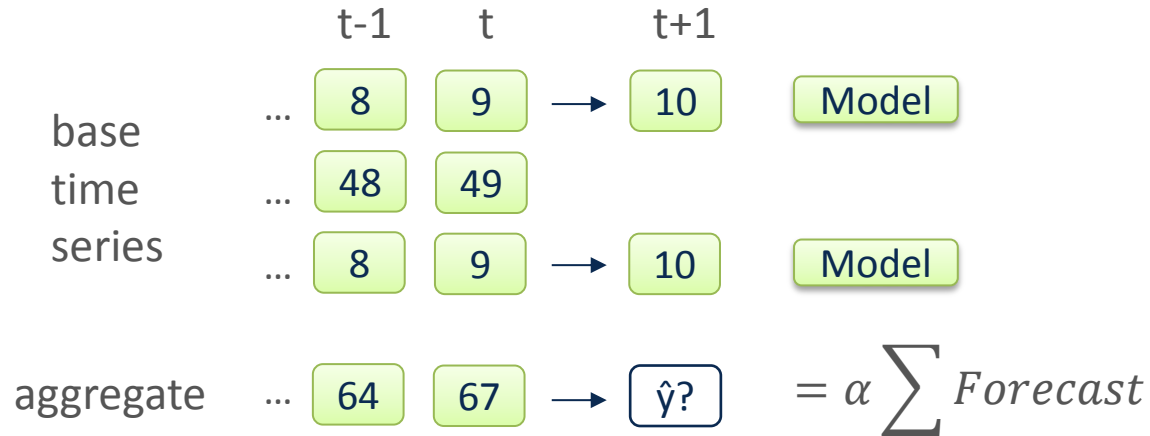


How to compute aggregation forecast queries?



	Runtime	Accuracy
	↑	↓
	↓	↑
	→	↑

> Example: Aggregation Queries



Uniform Estimation $\alpha = \frac{\# \text{ base series}}{\# \text{ base series in sample}} \quad \frac{3}{2} \cdot 20 = 30$


Estimation with Historical Ratios

$ratio = \frac{\text{base series}}{\text{aggregate series}} \quad \Rightarrow \quad \alpha = \frac{1}{\sum ratios} \quad \frac{1}{\frac{9}{67} + \frac{9}{67}} = 3.7 \cdot 20 = 74$




Calculation of Historical Ratios

- Different approaches possible
 - Simple averages
 - Lagged proportions
 - ...
- Seasonality of data is important



Combined strategy: mixture of past ratios and ratio one season ago



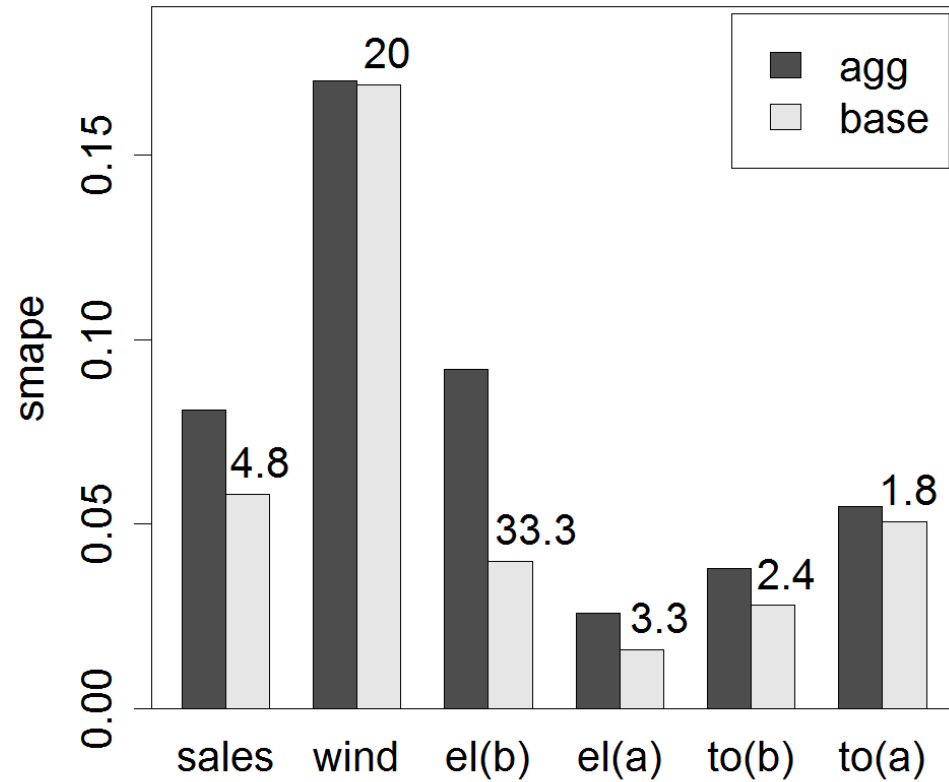

> Example: Aggregation Queries



agg




base






Aggregation - Sales






Model Configuration Advisor

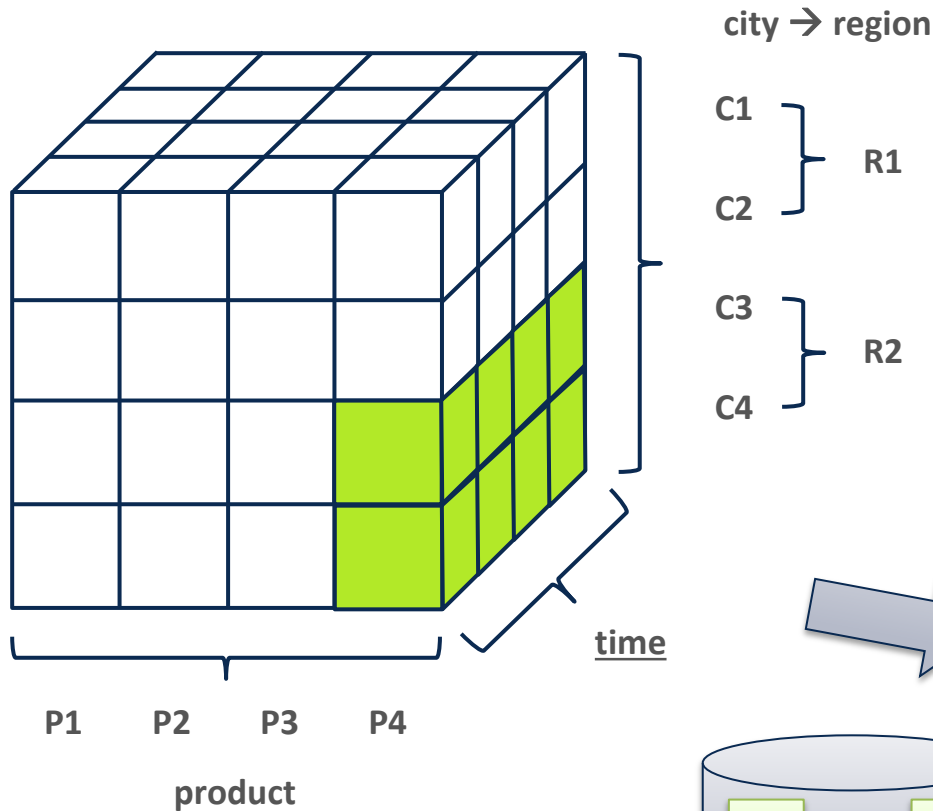
> Forecasting the Data Cube



```
SELECT time, measure  
FROM facts  
WHERE product = P4  
AND city = C4  
AS OF now() + 1 day
```





> Forecasting the Data Cube



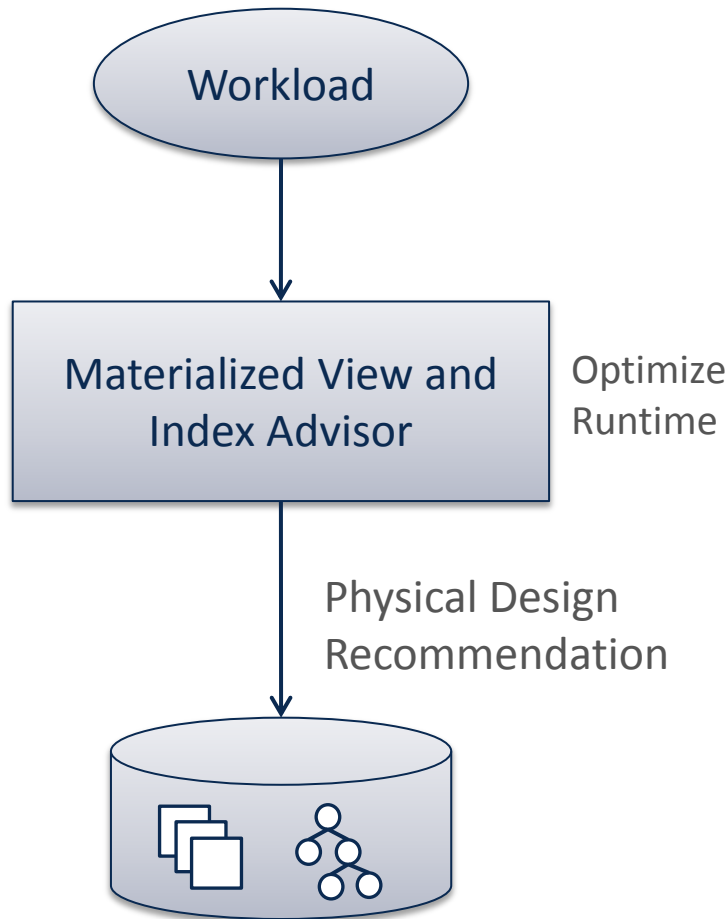
```

SELECT  time, SUM(measure)
FROM    facts
WHERE   product = P4
AND     region = R2
GROUP BY time
AS OF  now() + 1 day
    
```

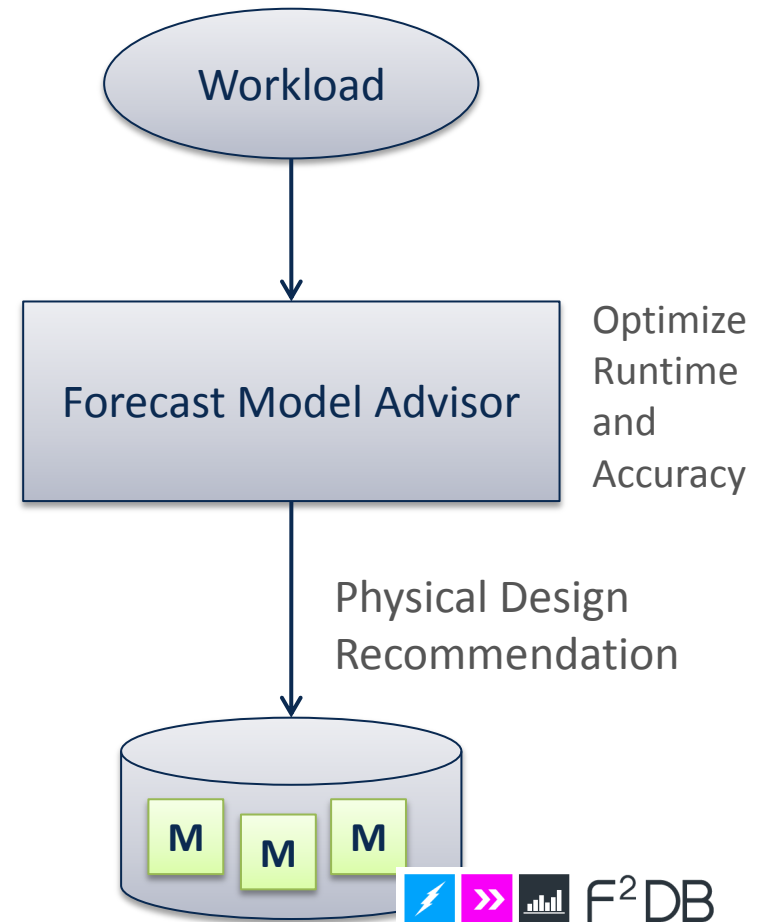




Traditional Database Systems

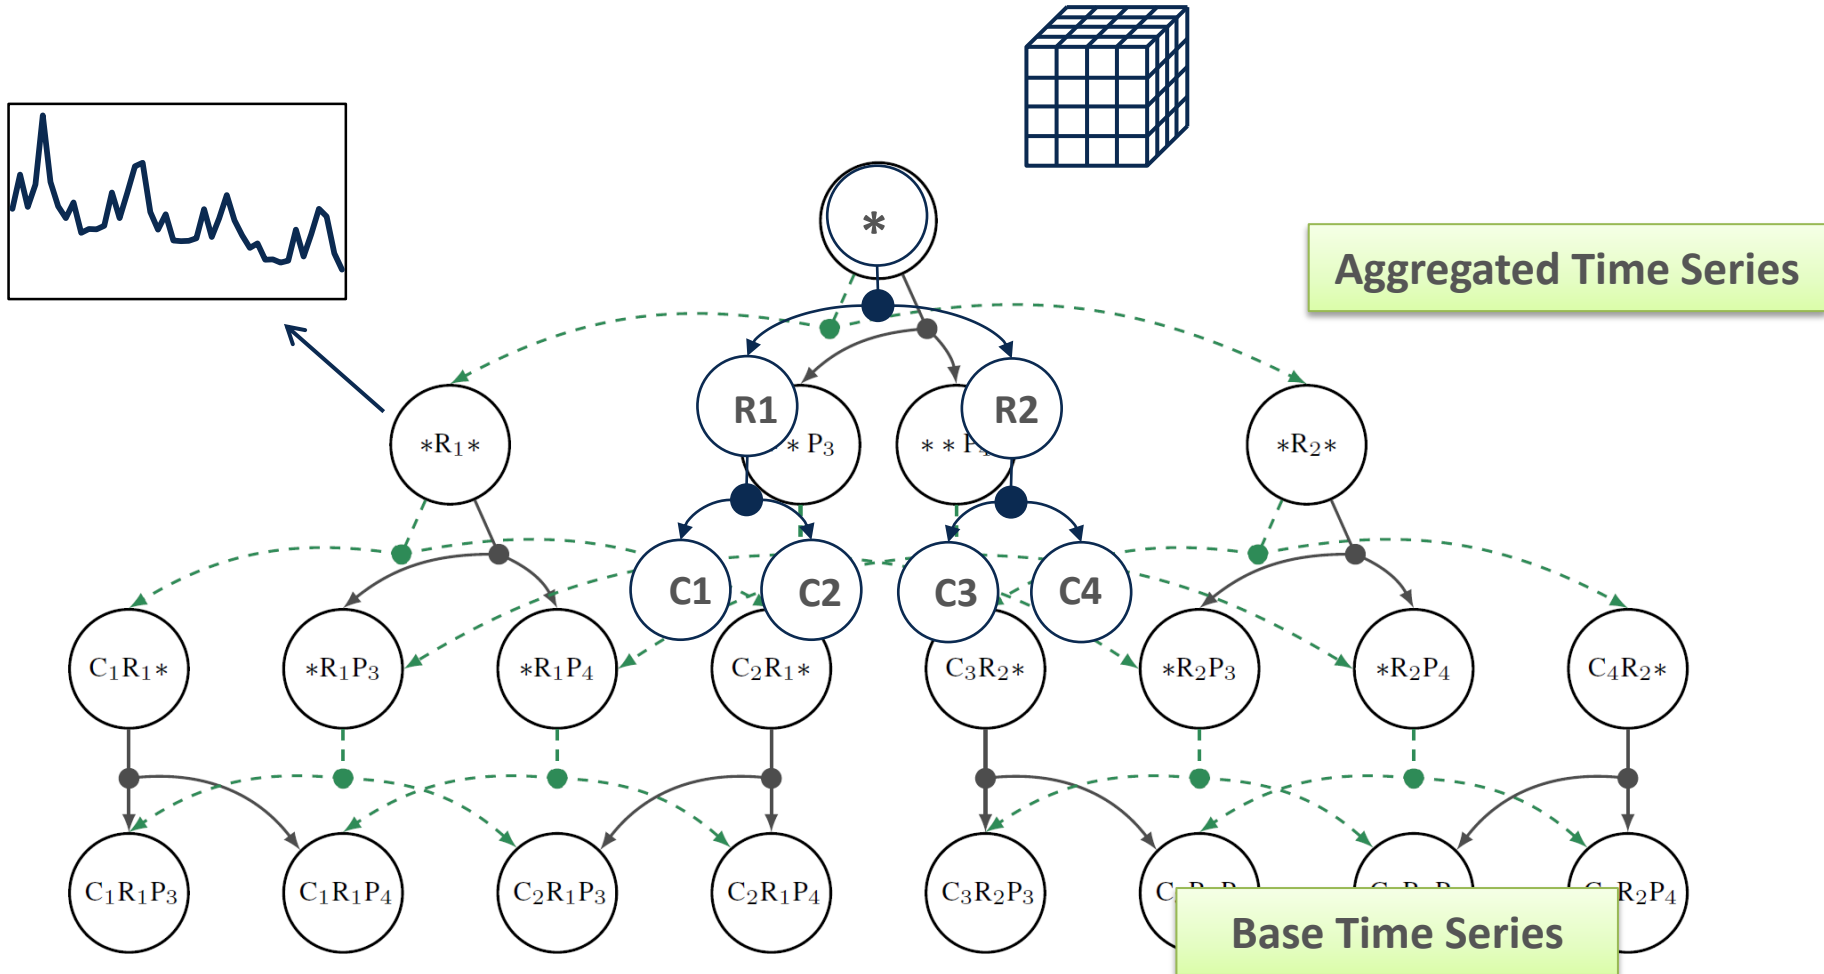



Statistical Database Systems






Conceptually, we organize the aggregation possibilities as a directed time series hyper graph








Each node (or time series) may be associated with a forecast model



Time series methods

- Exponential Smoothing
- ARIMA

A query describes one or several nodes in the hyper graph







Forecast values of a node can be computed by any nodes in the graph

Derivation weight k


- Based on history of source and target time series




Direct



Aggregation




Disaggregation






Model configuration

Assignment of models



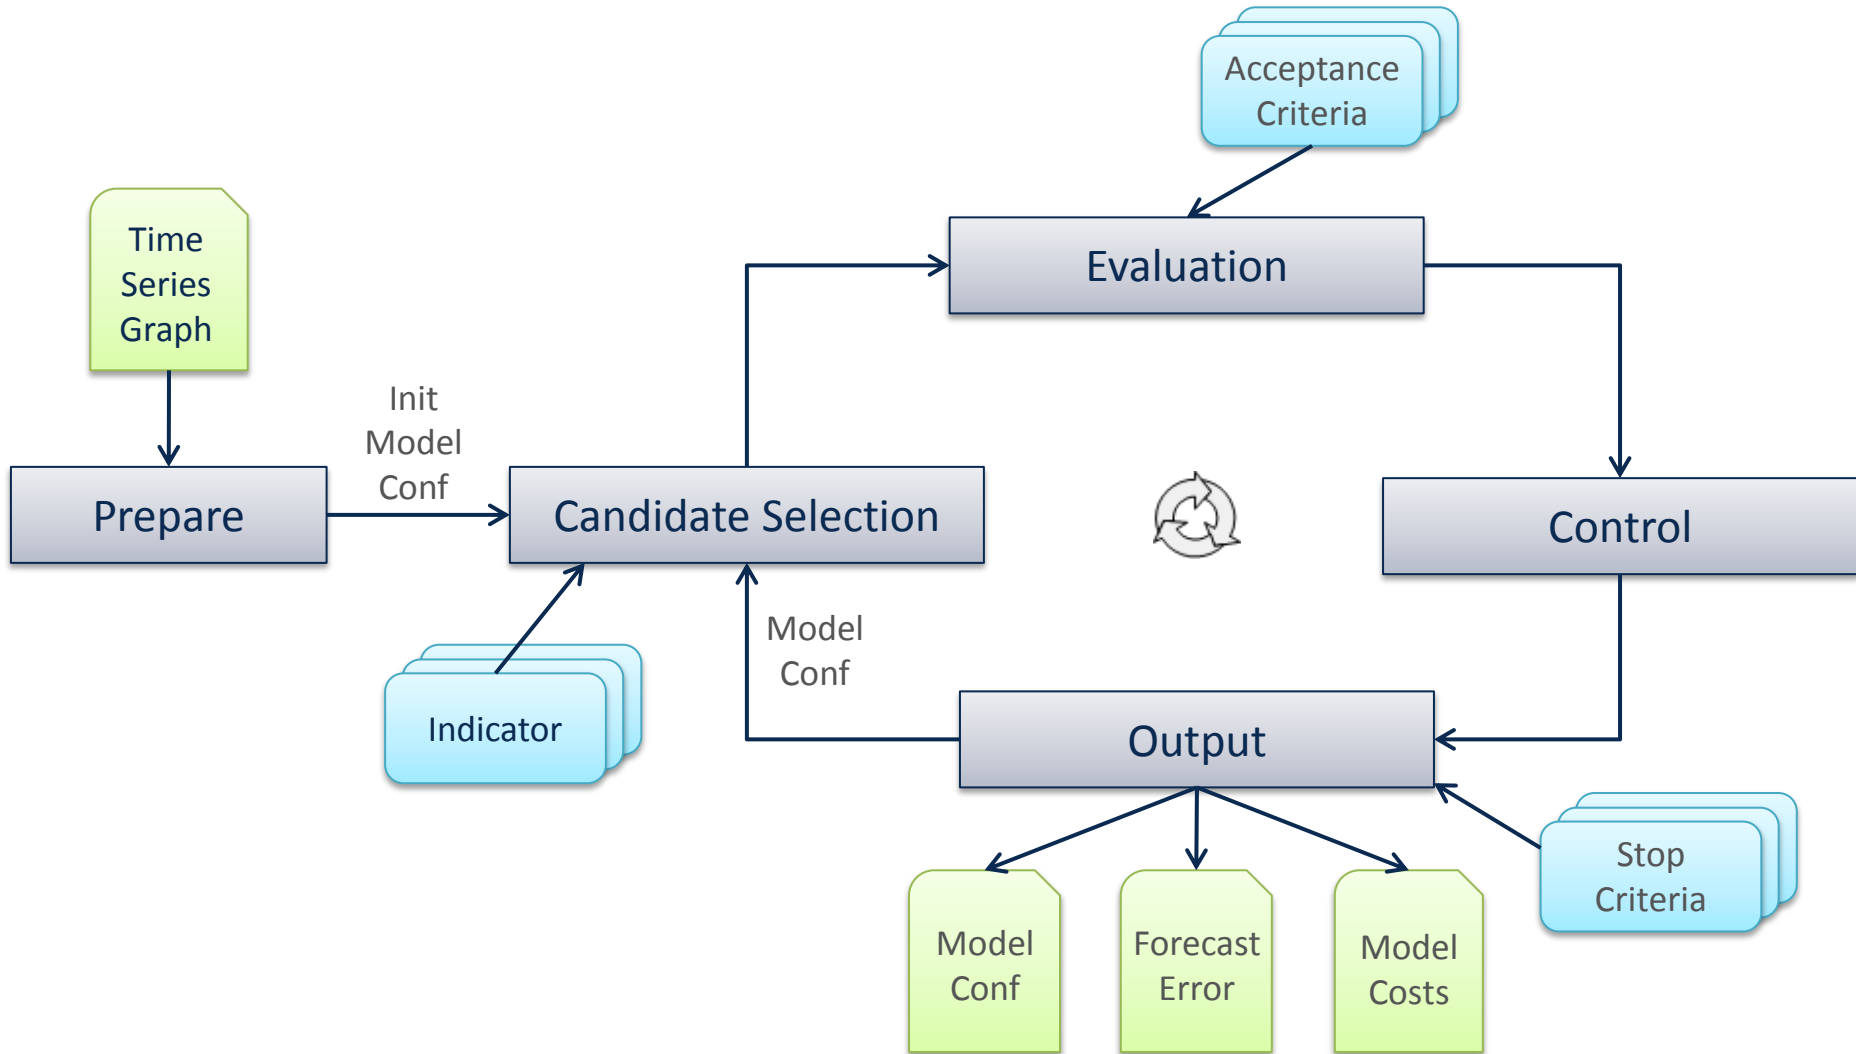
+



Assignment of derivation schemes

Configuration evaluation

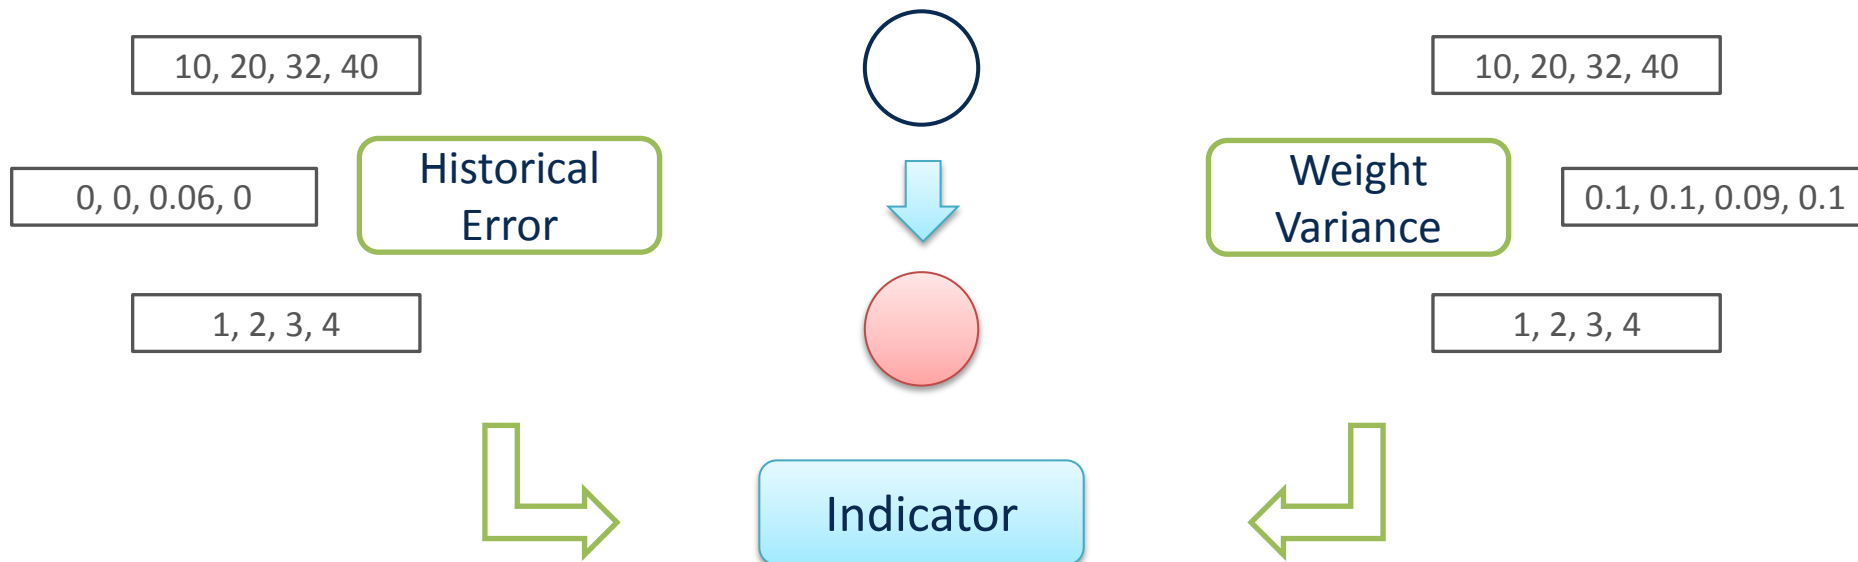






*Heuristically indicate the expected benefit of a model at a node
(without building the model)*

- Focus on time series relationships
- Measure to specify the derivation error between two nodes
- Low indicator value \rightarrow low error (*good derivation*)
- High indicator value \rightarrow high error (*poor derivation*)



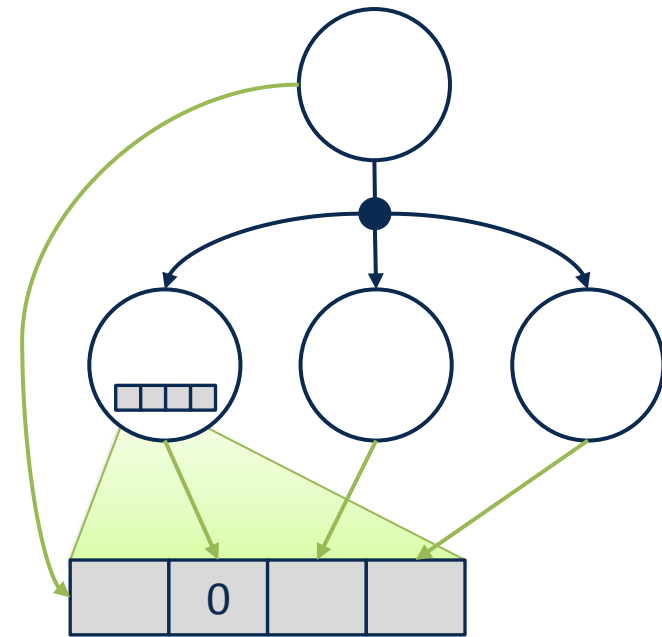
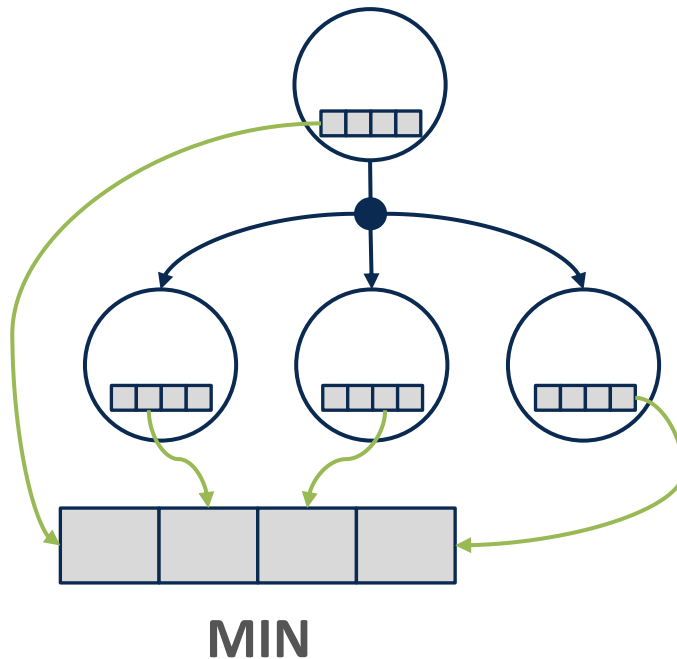


Local indicator arrays


→ *Derivation errors of one node*

Global indicator array


→ *Minimum over all local arrays*




> Selection of Model Configurations



Select start configuration




> Selection of Model Configurations




Select start configuration

Candidate selection

- Preselection




> Selection of Model Configurations




Select start configuration

Candidate selection

- Preselection



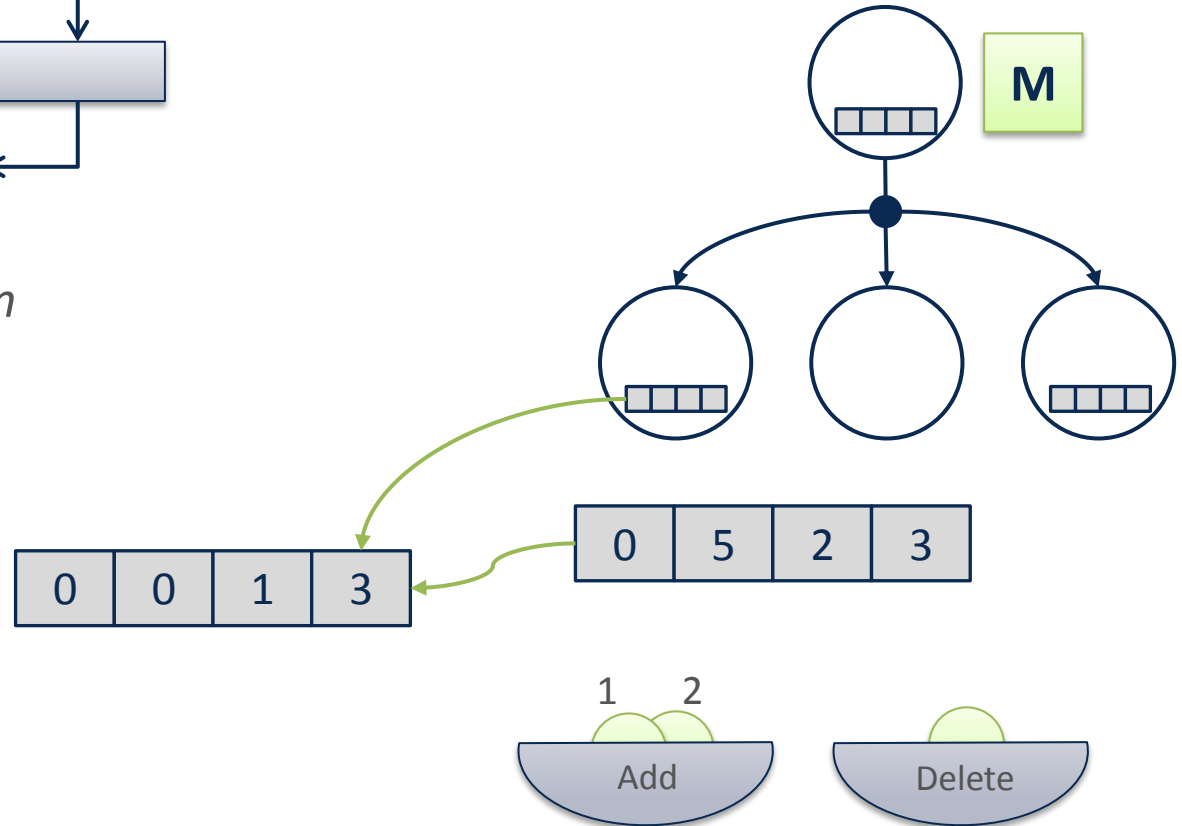
> Selection of Model Configurations




Select start configuration

Candidate selection

- Preselection
- Ranking



> Selection of Model Configurations



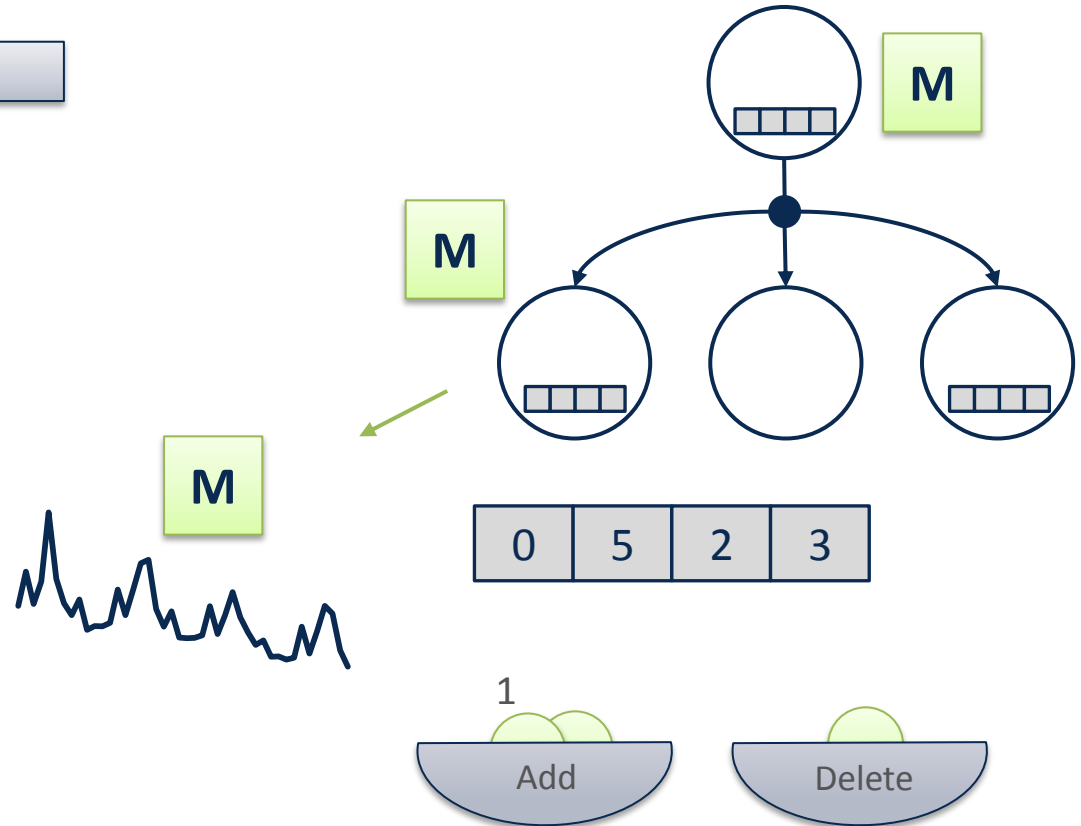
Select start configuration

Candidate selection


- Preselection
- Ranking

Evaluation

- Model creation
- Acceptance



> Selection of Model Configurations




Select start configuration

Candidate selection


- Preselection
- Ranking

Evaluation

- Model creation
- Acceptance



> Selection of Model Configurations




Select start configuration

Candidate selection


- Preselection
- Ranking

Evaluation

- Model creation
- Acceptance



> Selection of Model Configurations



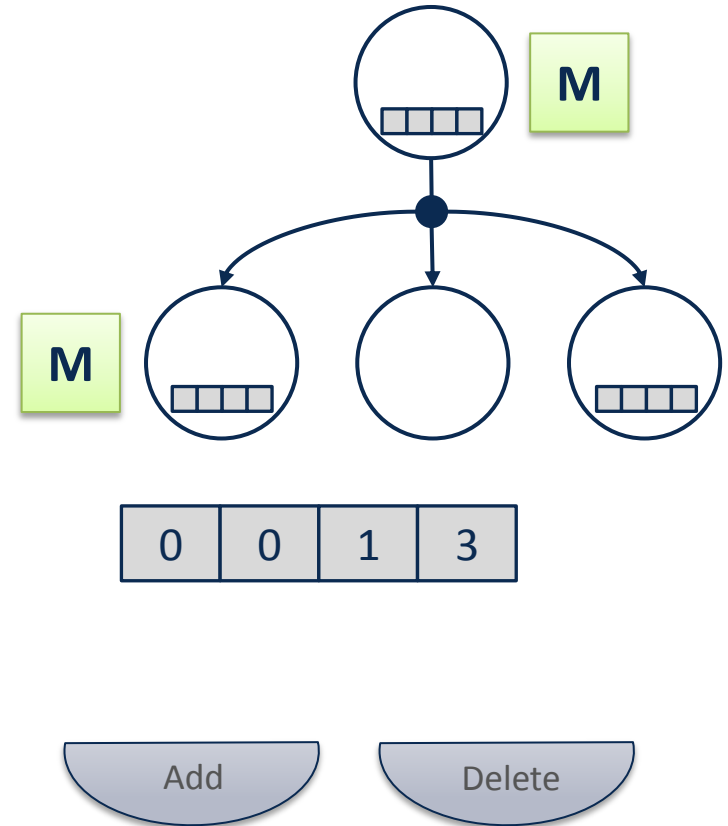
Select start configuration

Candidate selection

- Preselection
- Ranking

Evaluation

- Model creation
- Acceptance

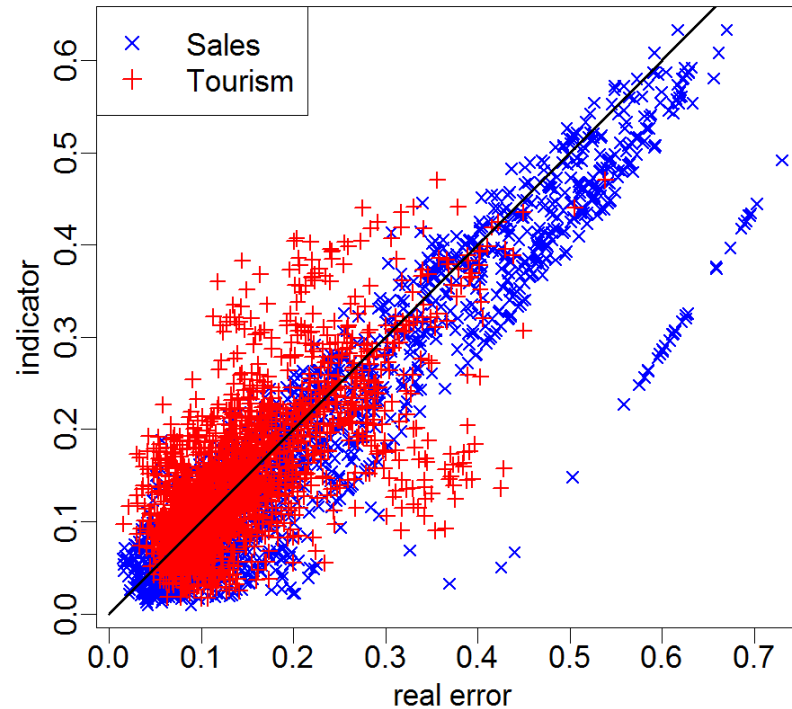


Model costs: 2

Forecast error: 10 %




Correlation between indicators and real forecast error






Tourism




Sales



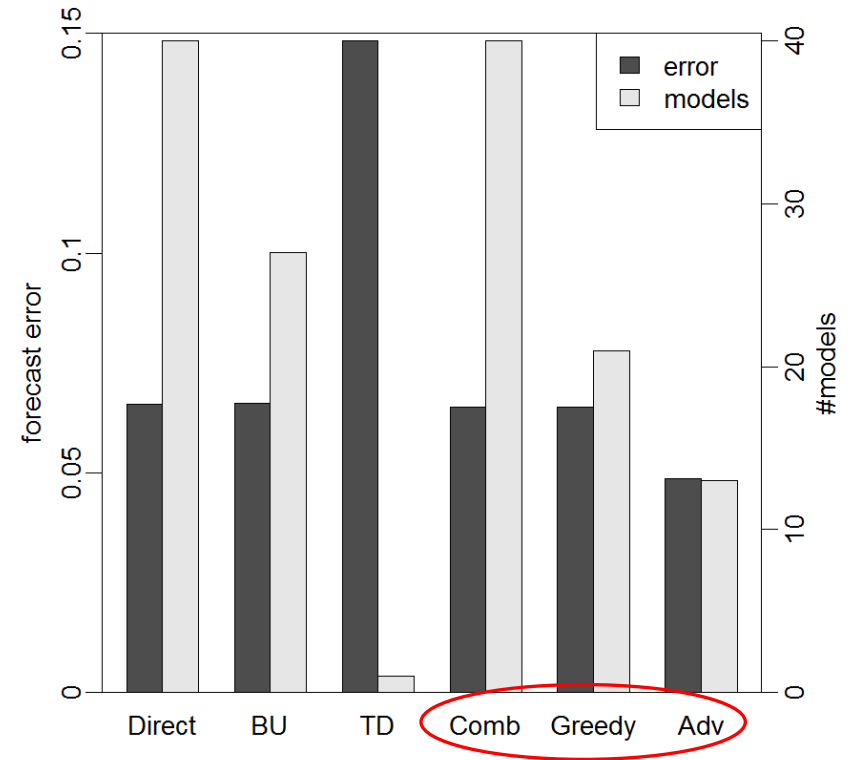
Static approaches – data independent



Tourism




Sales



Dynamic approaches – empirical selection



Scalability





Subscription-Based Forecast Queries




Energy data management systems

- Provide stable energy supply while including larger amounts of renewable energy
- **Continuously** require forecasts of energy demand and supply




Subscription-based forecast queries





When to notify subscriber?

After each new
real value ...



- Many notifications
- **High subscriber costs**

As less as possible ...



- Long messages
- Low accuracy
 - Resend messages
- **High subscriber costs**

Optimal ...




- Reduce number of notifications and notification length



Parameters

- Time series description
- Minimum continuous forecast horizon
- Accuracy threshold

```
SELECT datetime, production
FROM ts_powerproduction
WHERE type = „wind“
FORECAST 3 hours
THRESHOLD 0.1
```






Horizon Violation

- When? Subscriber has less than minimum horizon
- What? Send missing values + horizon extension

Threshold Violation

- When? Threshold is violated
- What? Resend all values

```
SELECT    datetime, production
FROM      ts_powerproduction
WHERE     type = „wind“
FORECAST  3 hours
THRESHOLD 0.1
```




threshold violation ↑
horizon violation ↑
first notification ↑



Processing costs of the subscriber

- Analytically known or learned function
- Depend on the forecast horizon
- Complete costs F_C
 - Complete restart of processing
- Incremental costs F_I
 - Processing of additional values



threshold violation ↑

horizon violation ↑

first notification ↑



Assume we know ...

- The sequence of threshold + horizon violations
- The subscriber cost functions F_C and F_I




Subscriber costs over subscription lifetime

- Sum over all F_C and F_I

Optimization Goal

- Find forecast horizon that minimizes total costs
- Depends on subscriber cost function
- Depends on forecast accuracy



How to get threshold violations?




Analyze past to predict future



Core Idea: Calculate best forecast horizon using our cost model on the time series history


Offline – Static

- One forecast horizon over whole lifetime




Offline – Dynamic

- Adapts to periodic changes of time series accuracy
- Sequence of forecast horizons for time slices



Online

- Adapts to arbitrary changes in data or cost functions
- Continuously adapts forecast horizon





Real-world energy demand and supply data sets

- National Energy Demand
- Household Energy Demand
- National Wind Supply

Subscriber cost functions

- Synthetic linear function
- Real world cost function (obtained from MIRABEL)


Forecast Methods

- Tailor-made for short-term energy forecasting
- Extension of exponential smoothing




Comparison of Approaches

- Fixed subscription parameters and linear cost function



Evaluation of Time Slice Approach







Influence of subscriber threshold

- Relationship between number of notifications, subscriber costs and runtime
- Real-world cost function


number of notifications



subscriber costs




runtime





Cost Model Validation

- Estimated costs vs. real costs
- Real-world cost function
- Queries with increasing complexity (Q1, Q2, Q3)





Beyond Forecast Models

Towards a Model-based Database System



Forecast

- Base approach
- Predict missing data from history
- Requires no known data

Impute

- Exploit local patterns
- Infer missing data from similar units
- Requires adequate set of known data

Refine

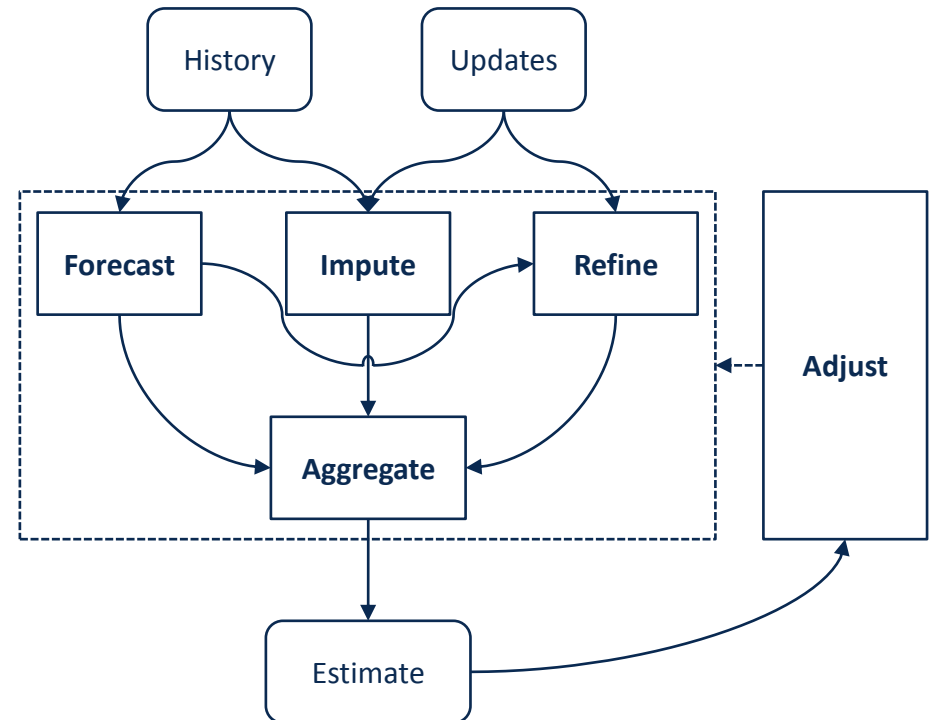
- Detect local and global shifts
- Infer error → yields new forecasts
- Requires few known data

Aggregate

- Calculate aggregate (e.g. report)

Adjust

- Maintain models and synopses
- Optimize accuracy of estimates




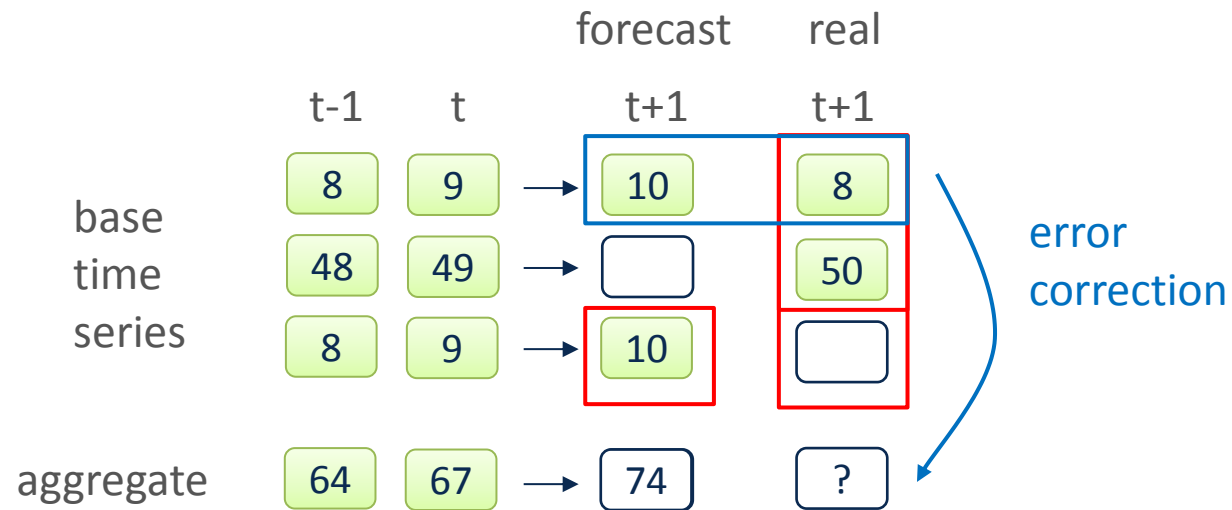


How to include new real data?

Data delivery may be late

- There might be missing data for the last period
- Reports still have to be generated
- Estimate missing data





Refine I $\alpha \sum Real$

Refine II $\alpha \sum (Real + Forecast)$

Refine III $\alpha \sum (Real + Forecast) - \beta \sum Error$



Refine III: Estimation of forecast errors

	forecast	real	error
	t+1	t+1	
base time series	10	8	8 - 10
	10		?
	50	50	50 - 50
aggregate	74	?	

Case 1

- Forecast and real value
- Calculate true model error

Case 2/3

- No real value
- No error calculation


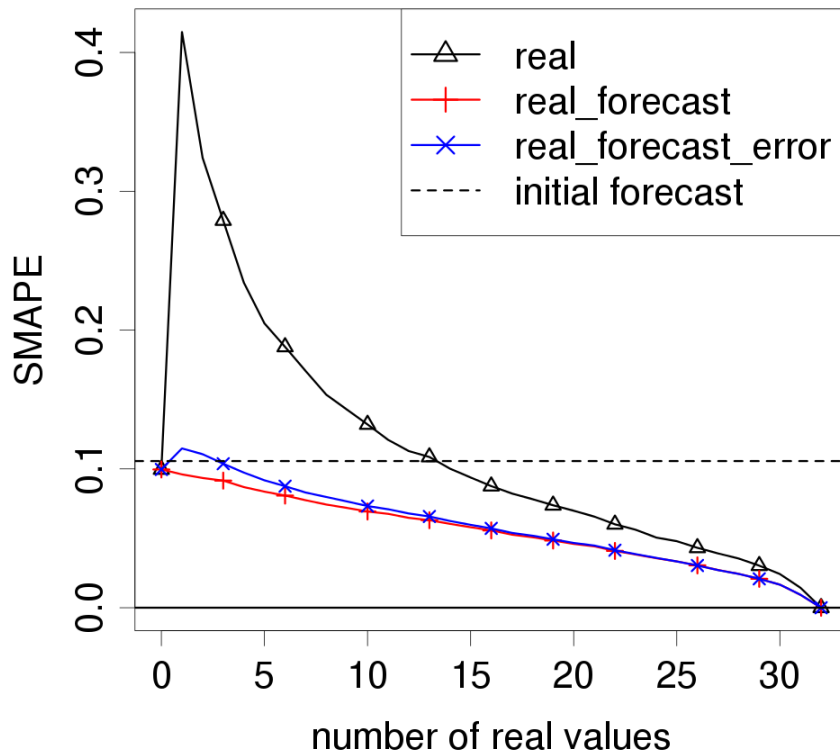
Case 4


- No forecast value
- Estimate model error



Refine – Sales

Refine – Wind production





FFQ project

- Provide forecasting as 1st class citizen within a database system
- Preserve logical and physical data independence (e.g. transparent model usage, transparent model maintenance, and model creation)
- Extend traditional processing and optimization techniques
- Apply concept of traditional index advisors to forecast models

Towards a model-based database system

- Data is increasingly inconsistent, incomplete and imprecise
- Extend concept of models to other use cases (missing data, uncertain data, data compression ...)



Wolfgang Lehner

Forecasting and Data Imputation Strategies in Database Systems

11.07.2013