Metanome

HPI Hasso Plattner Institut
IT Systems Engineering | Universität Potsdam

# Data Profiling with Metanome

## Exercise: UCCs

Thorsten Papenbrock
PhD Candidate
Hasso-Plattner-Institute

# Agenda

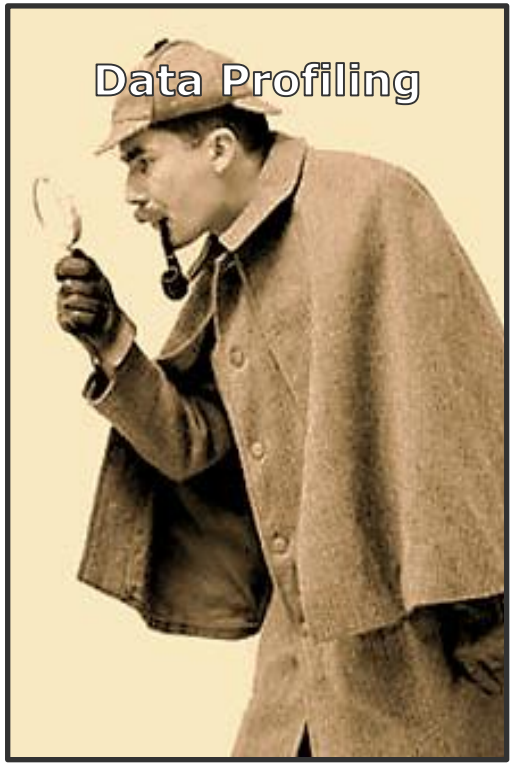


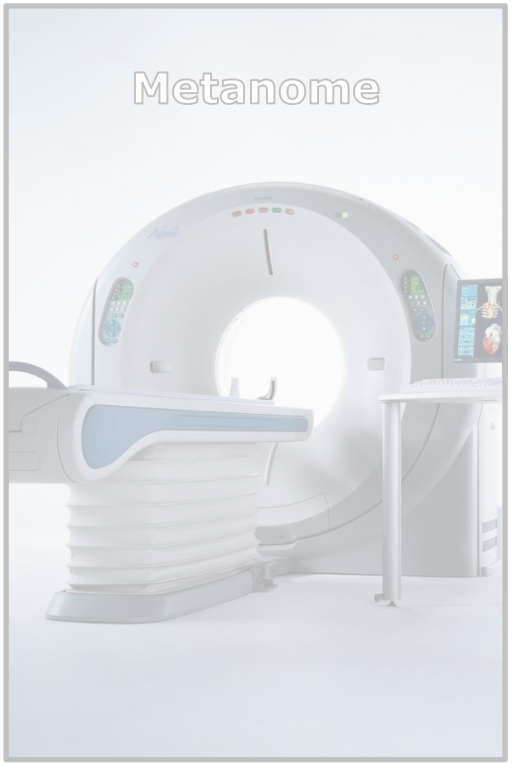



**Data Profiling with Metanome**

ThorstenPapenbrock,
PhD Candidate,
9th of October, 2014

# Data Profiling
## The need to know your data

# Data Profiling
# The meaning of "knowing data"

Hasso Plattner Institut

**range** min = 0.382, max = 11.209

**aggregation** sum = 173, avg = 21.625

**distribution**

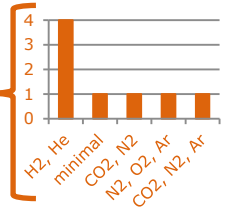


**format**

| Name | Type | Equatorial diameter | Mass | Orbital radius | Orbital period | Rotation period | Confirmed moons | Rings | Atmosphere |
|---|---|---|---|---|---|---|---|---|---|
| Mercury | Terrestrial | 0.382 | 0.06 | 0.47 | 0.24 | 58.64 | 0 | no | minimal |
| Venus | Terrestrial | 0.949 | 0.82 | 0.72 | 0.62 | −243.02 | 0 | no | $CO_2$, $N_2$ |
| Earth | Terrestrial | .0 | 1.( | | 0 | 1.00 | 1 | no | $N_2$, $O_2$, Ar |
| Mars | Terrestrial | 53 | 0. | | 2 | 1.03 | 2 | no | $CO_2$, $N_2$, Ar |
| Jupiter | Giant | .2 | 7.8 | 5 | 0 | 11.86 | 67 | yes | $H_2$, He |
| Saturn | Giant | 9.449 | 95.2 | 9.54 | 29.46 | 0.43 | 62 | yes | $H_2$, He |
| Uranus | Giant | 4.007 | 14.6 | 19.22 | 84.01 | −0.72 | 27 | yes | $H_2$, He |
| Neptune | Giant | 3.883 | 17.2 | 30.06 | 164.8 | 0.67 | 14 | yes | $H_2$, He |

**size**

SIMPLE

CHAR(32) CHAR(16) FLOAT FLOAT FLOAT FLOAT FLOAT INTEGER BOOLEAN VARCHAR

**data types**

Chart **4**

# Data Profiling
# The meaning of “knowing data”

**inter table dependencies**
Planet.Name ⊆ Moon.Planet

**relationships**
Mass ~ Confirmed Moons

**keys**
unique

| Name | Type | Equatorial diameter | Mass | Orbital radius | Orbital period | Rotation period | Confirmed moons | Rings | Atmosphere |
|---|---|---|---|---|---|---|---|---|---|
| Mercury | Terrestrial | 0.382 | 0.06 | 0.47 | 0.24 | 58.64 | 0 | no | minimal |
| Venus | Terrestrial | 0.949 | 0.82 | 0.72 | 0.62 | −243.02 | 0 | no | $CO_2$, $N_2$ |
| Earth | Ter | | 0 | 1.0 | 0 | 1 | | no | $N_2$, $O_2$, Ar |
| Mars | Te strial | 0.5 | | 1.5 | .88 | 1 | | no | $CO_2$, $N_2$, Ar |
| Jupiter | | | 3 80 | 5.2 | 11.8 | 0 | 6 | yes | $H_2$, He |
| Saturn | Giant | 9.449 | 95.2 | 9.54 | 29.46 | 0.43 | 62 | yes | $H_2$, He |
| Uranus | Giant | 4.007 | 14.6 | 19.22 | 84.01 | −0.72 | 27 | yes | $H_2$, He |
| Neptune | Giant | 3.883 | 17.2 | 30.06 | 164.8 | 0.67 | 14 | yes | $H_2$, He |

COMPLEX

Equatorial diameter x Mass > 0

**rules**

Atmosphere → Rings

**intra table dependencies**

# Agenda


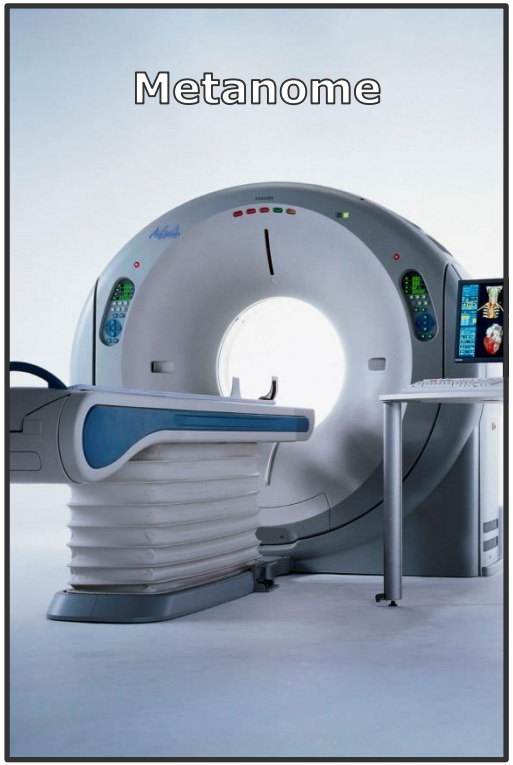



**Data Profiling with Metanome**

# Metanome
## An extensible architecture

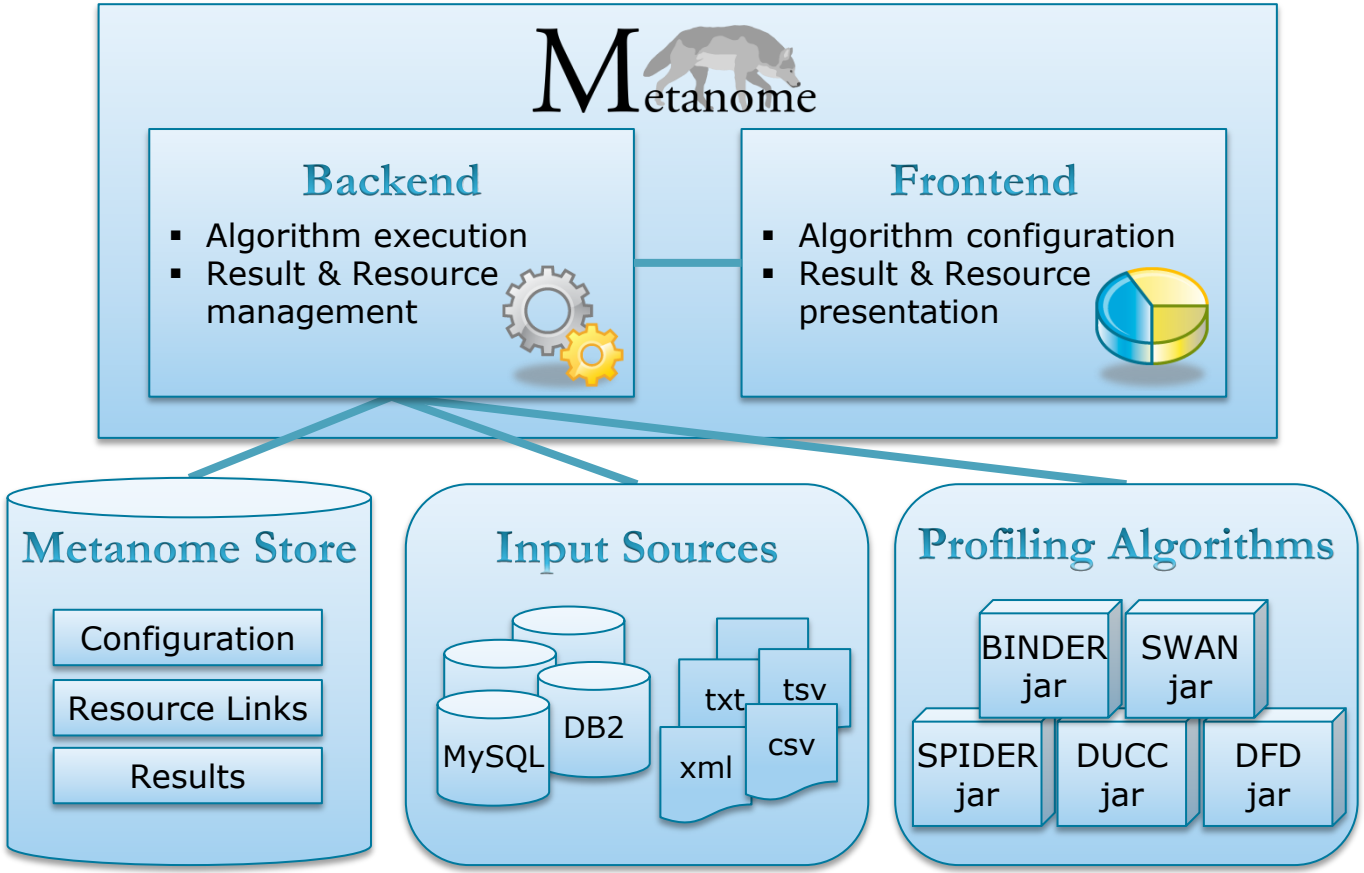

# Metanome
# The typical workflow

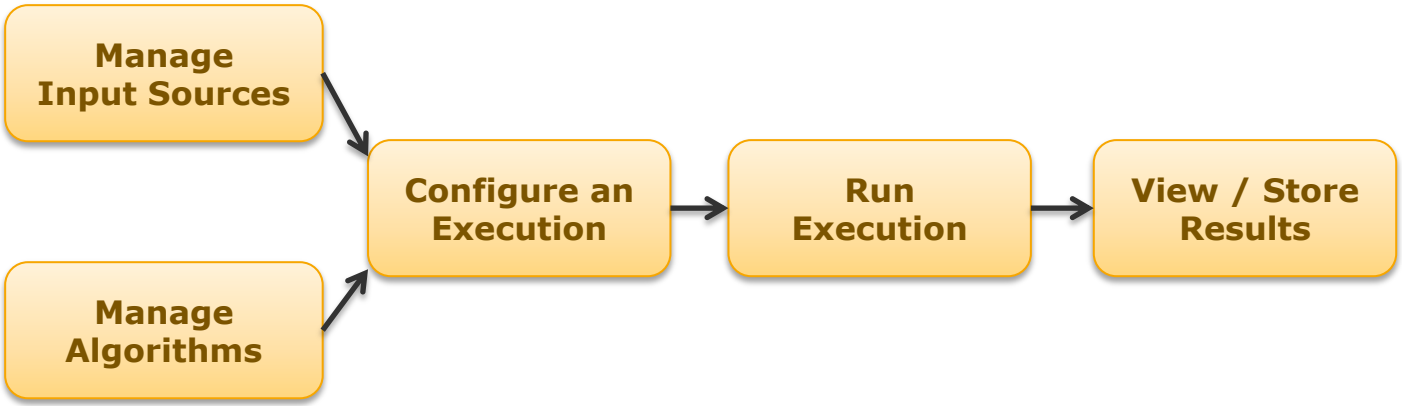

# Demo

Hasso Plattner Institut

Metanome

**Data Profiling with Metanome**

ThorstenPapenbrock,
PhD Candidate,
9th of October, 2014

# Metanome
# The exercise: Build your own profiling tool

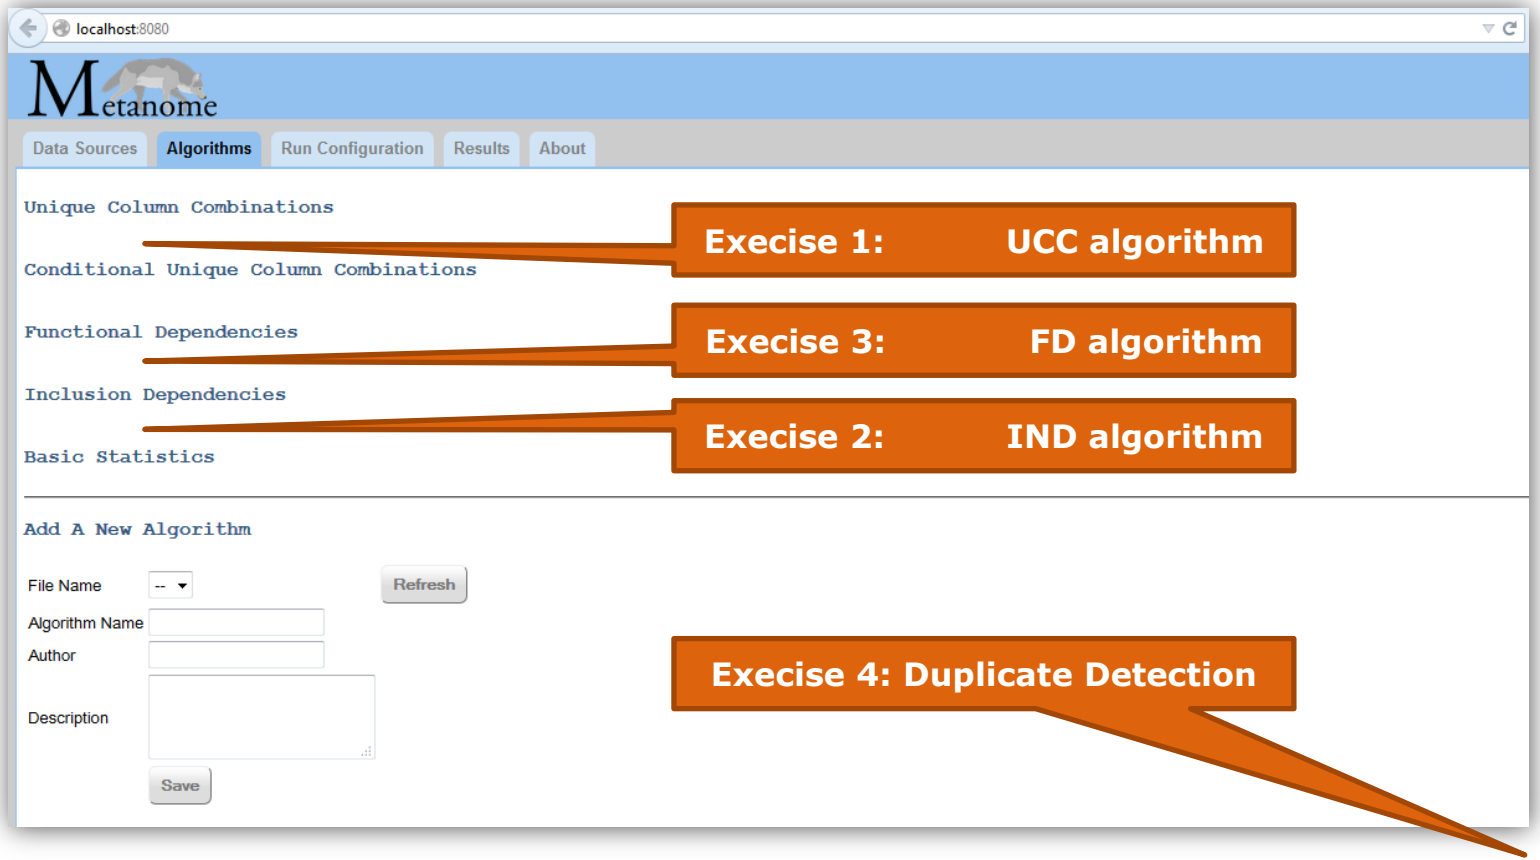

# Agenda


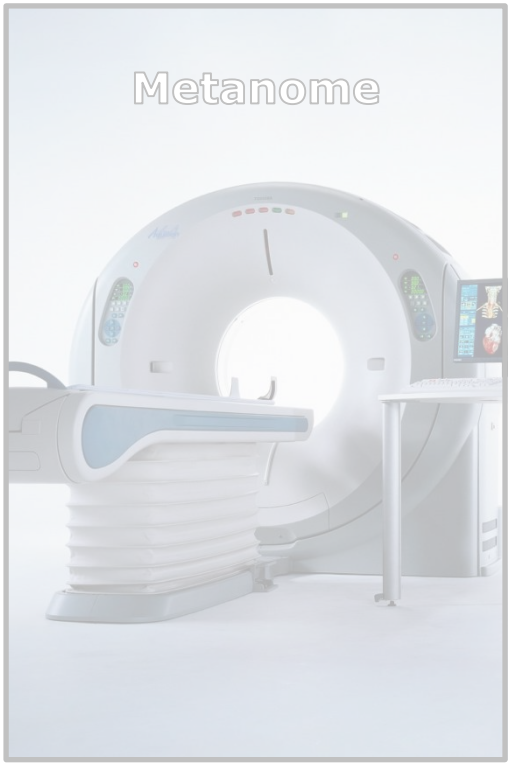


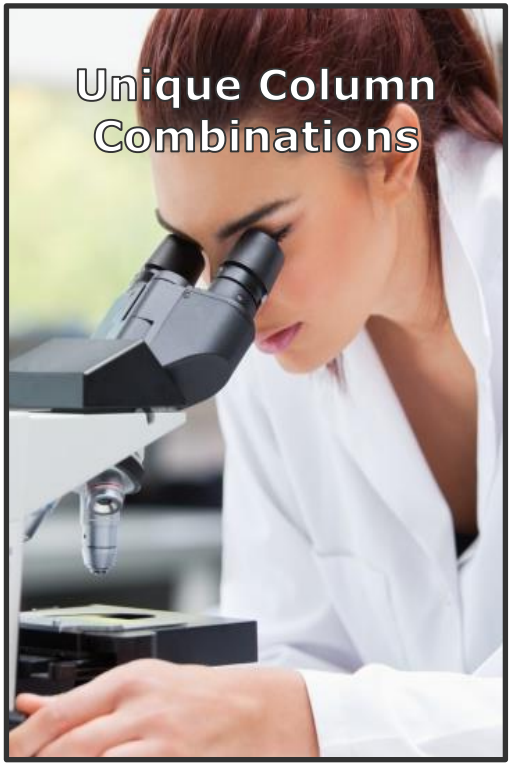

# Advanced Profiling
# Three important metadata

Hasso Plattner Institut

## Unique Column Combinations

**Key candidates**

| Name | Type | Equatorial diameter | Mass |
|---|---|---|---|
| Mercury | Terrestrial | 0.382 | 0.06 |
| Venus | Terrestrial | 0.949 | 0.82 |
| Earth | Terrestrial | 1.000 | 1.00 |
| Mars | Terrestrial | 0.532 | 0.11 |
| Jupiter | Giant | 11.209 | 317.8 |
| Saturn | Giant | 9.449 | 95.2 |
| Uranus | Giant | 4.007 | 14.6 |
| ... | ... | ... | ... |

**|Name|**

## Inclusion Dependencies

**Foreign key candidates**

| Sign | Domicile |
|---|---|
| Aries | Mars |
| Taurus | Venus |
| Gemini | Mercury |
| Cancer | Moon |
| Leo | Sun |
| Virgo | Mercury |
| Libra | Venus |
| Scorpio | Pluto |
| Sagittarius | Jupiter |
| Capricorn | Saturn |
| Aquarius | Uranus |
| ... | ... |

| Name | Type |
|---|---|
| Mercury | Terrestrial |
| Venus | Terrestrial |
| Earth | Terrestrial |
| Mars | Terrestrial |
| Jupiter | Giant |
| Saturn | Giant |
| Uranus | Giant |
| ... | ... |

**Domicile ⊆ Name**

## Functional Dependencies

**Normalization criterion**

| Name | Atmosphere | Rings |
|---|---|---|
| Mercury | minimal | no |
| Venus | $CO_2$, $N_2$ | no |
| Earth | $N_2$, $O_2$, Ar | no |
| Mars | $CO_2$, $N_2$, Ar | no |
| Jupiter | $H_2$, He | yes |
| Saturn | $H_2$, He | yes |
| Uranus | $H_2$, He | yes |
| ... | ... | ... |

**Atmosphere → Rings**

**Data Profiling with Metanome**

ThorstenPapenbrock,
PhD Candidate,
9th of October, 2014

# Advanced Profiling
# Discovery of key candidates

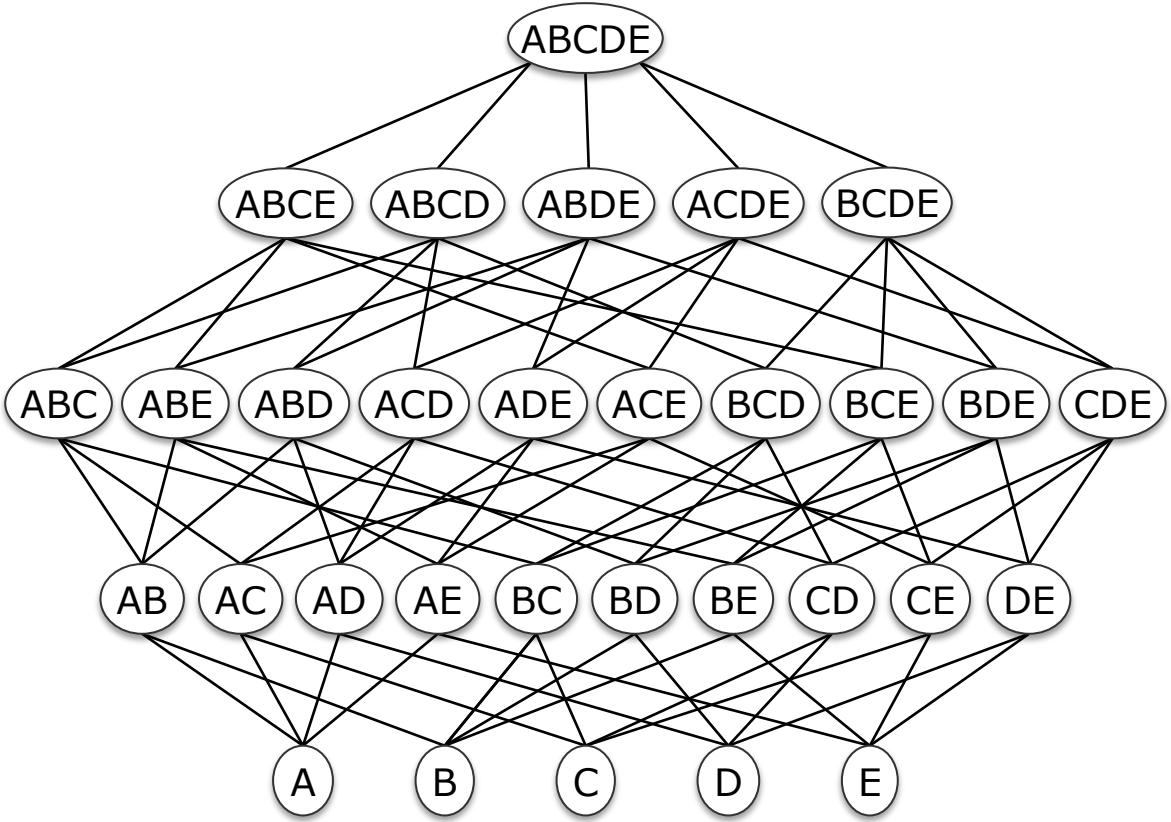

# Advanced Profiling
## Discovery of key candidates

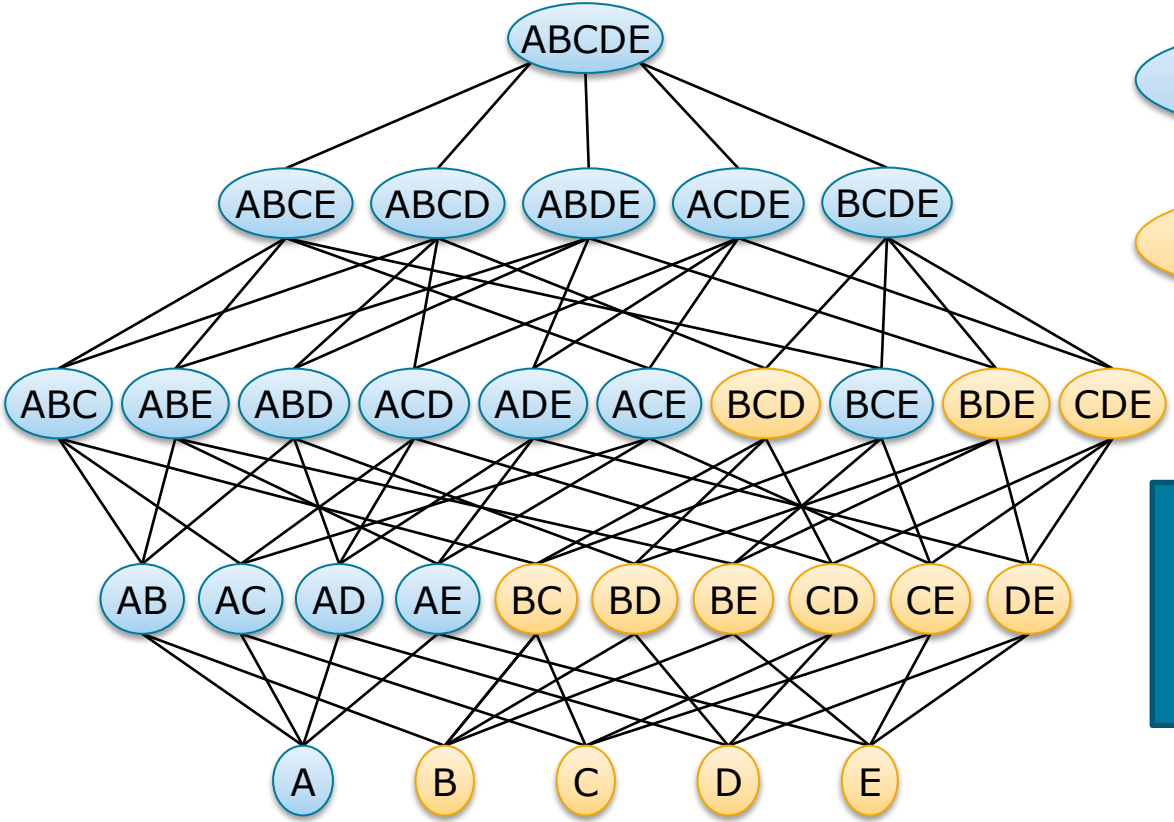


**Unique**
Column Combination

**Non-Unique**
Column Combination

**Complexity:** $O(2^n-1)$
for n attributes

**Example:**
10 attr ~ 1,023 checks
30 attr ~ 1,073,741,823 checks

# Discovery of key candidates – DUCC

**Random Walk**

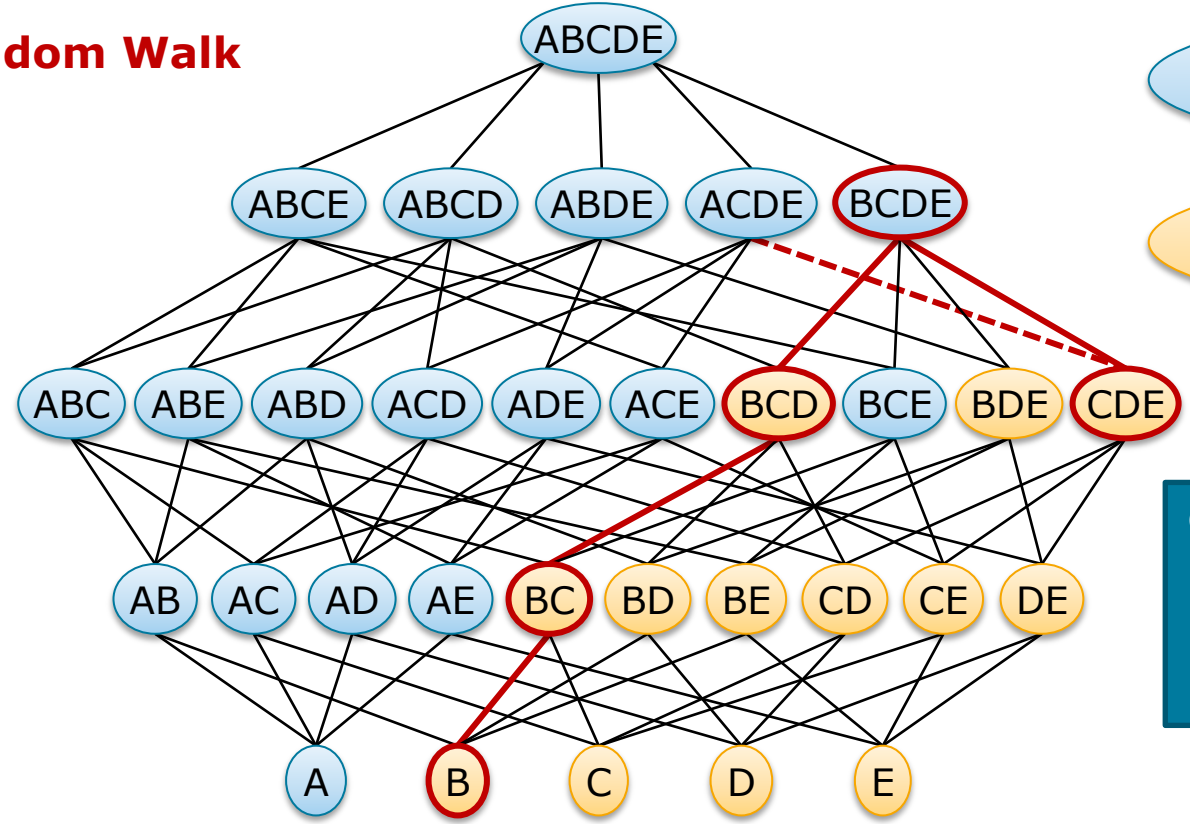


**Unique**
Column Combination

**Non-Unique**
Column Combination

**Complexity:** $O(2^n-1)$
for n attributes

**Example:**
10 attr ~ 1,023 checks
30 attr ~ 1,073,741,823 checks

# Discovery of key candidates – DUCC

## Column Scalability

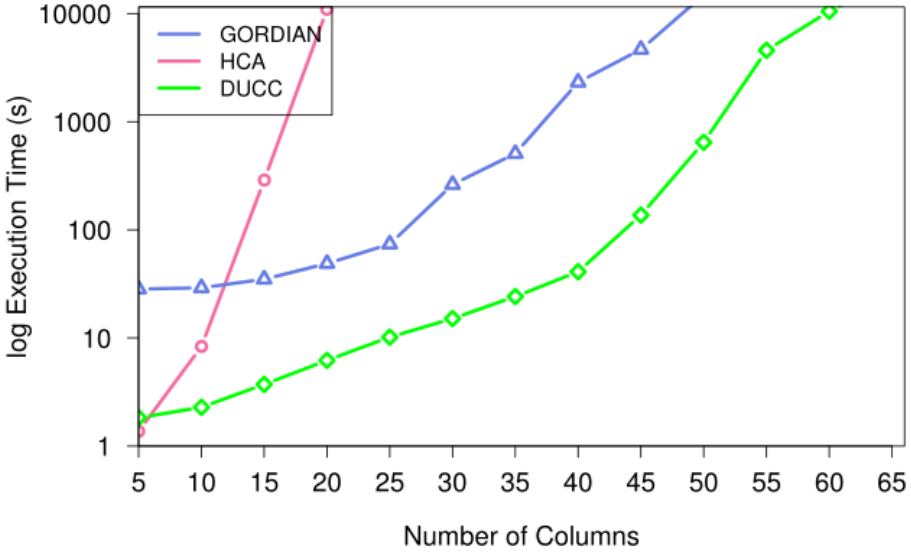


NCVoter

## Row Scalability

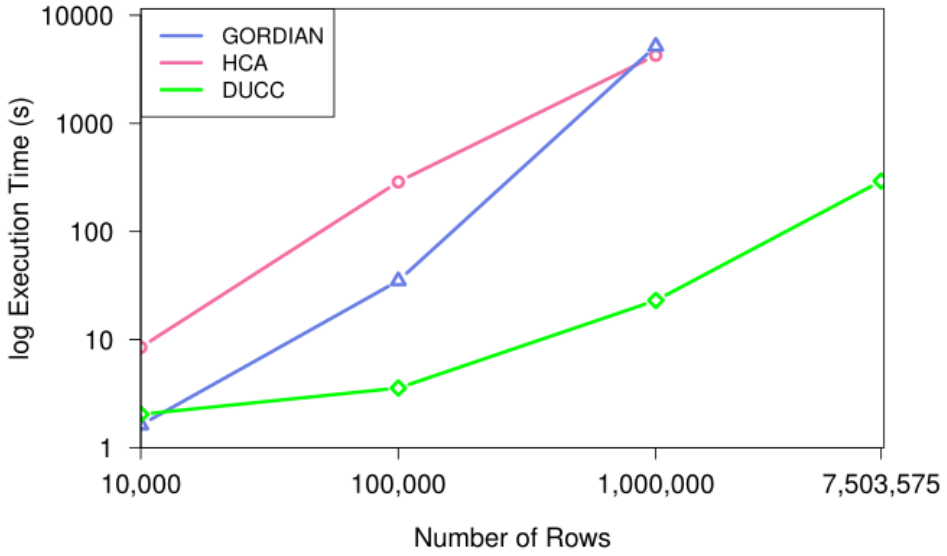


NCVoter

# Advanced Profiling
# Discovery of key candidates – Execise

## Exercise 1
## Unique Column Combinations

- Deadline: **Monday, 10.11.14**
- The admission to the exam requires *all* exercises to be solved.
- The exercises should be solved in teams of two students.
- The Metanome project is available at GitHub:
  `https://github.com/HPI-Information-Systems/Metanome`
- The datasets and supplemental material can be found at network drive S:
  `\\fs3\bbs\DPDC`
- The submission system can be found at:
  `https://www.dcl.hpi.uni-potsdam.de/submit/`
- To solve an exercise, please submit a zip file containing the following items:
  - **<algorithm_name>.jar**: An executable Metanome algorithm.
  - **<algorithm_name>.zip**: The algorithm's source code (maven project).
  - **<algorithm_name>_docu.pdf**: Short documentation of the algorithm.
  - **<algorithm_name>_pres.pptx/ppt/pdf**: Two slides presentation of the algorithm.

# Advanced Profiling
# Discovery of key candidates – Execise

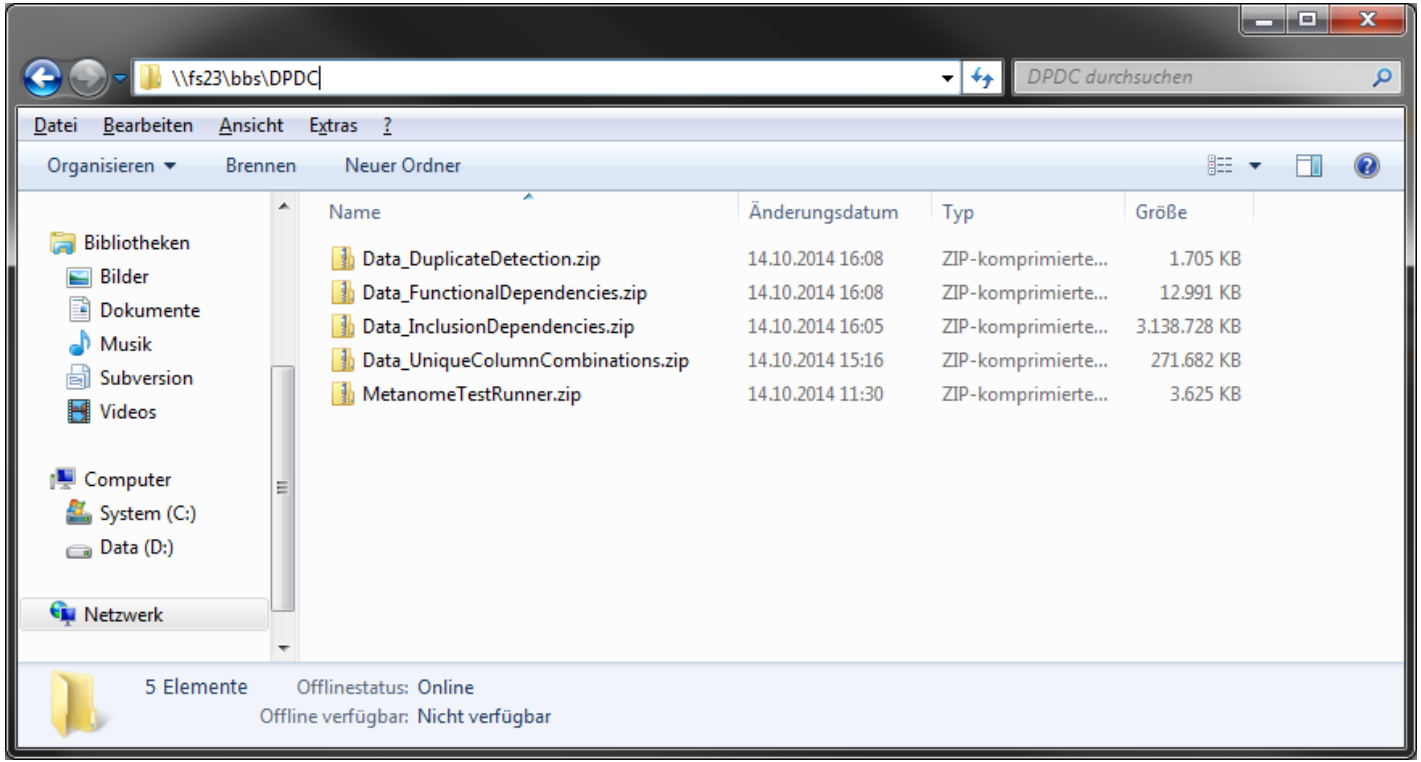

# Advanced Profiling
# Discovery of key candidates – Execise

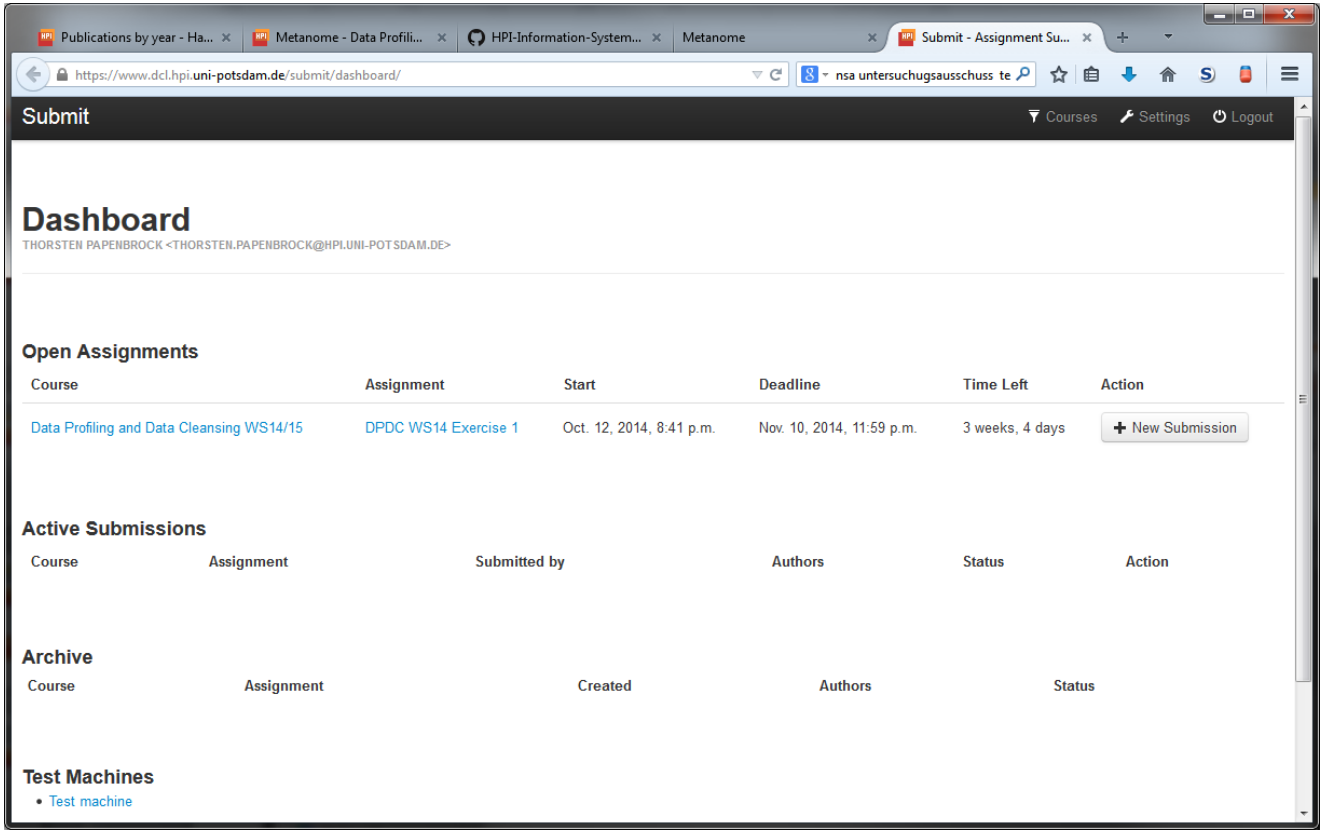

# Advanced Profiling
# Discovery of key candidates – Execise

## Task 1: Unique Column Combinations - A discovery algorithm

Write an algorithm that discovers *all* unique column combinations on the given datasets. The rules for your implementation are as follows:

a) The algorithm discovers *exact* results, so no approximate or fuzzy results are allowed.

b) The algorithm is not allowed to use parallelization.

c) The algorithm implements the Metanome interface and is compatible with Metanome.

d) For the NULL semantic, assume NULL $\neq$ NULL.

You can re-implement an existing algorithm from literature or find your own algorithm. To test and evaluate the algorithm, use the datasets provided on the network share. Your algorithm should at least be able to process the WDC_planets dataset! Before submitting your algorithm, check that it correctly executes within Metanome!

*BONUS TASK: If you like to dive deeper, you can try to parallelize your discovery algorithm. What effect does parallelization have on your pruning strategies? You can also try to make the search approximate or conditional. Which changes make your algorithm faster and which make it slower? If you made changes to your main algorithm, please submit them as a separate algorithm, e.g. <algorithm_name>_conditional.jar*

**Template Algorithm:**
https://github.com/HPI-Information-Systems/Metanome

# Advanced Profiling
## How to start

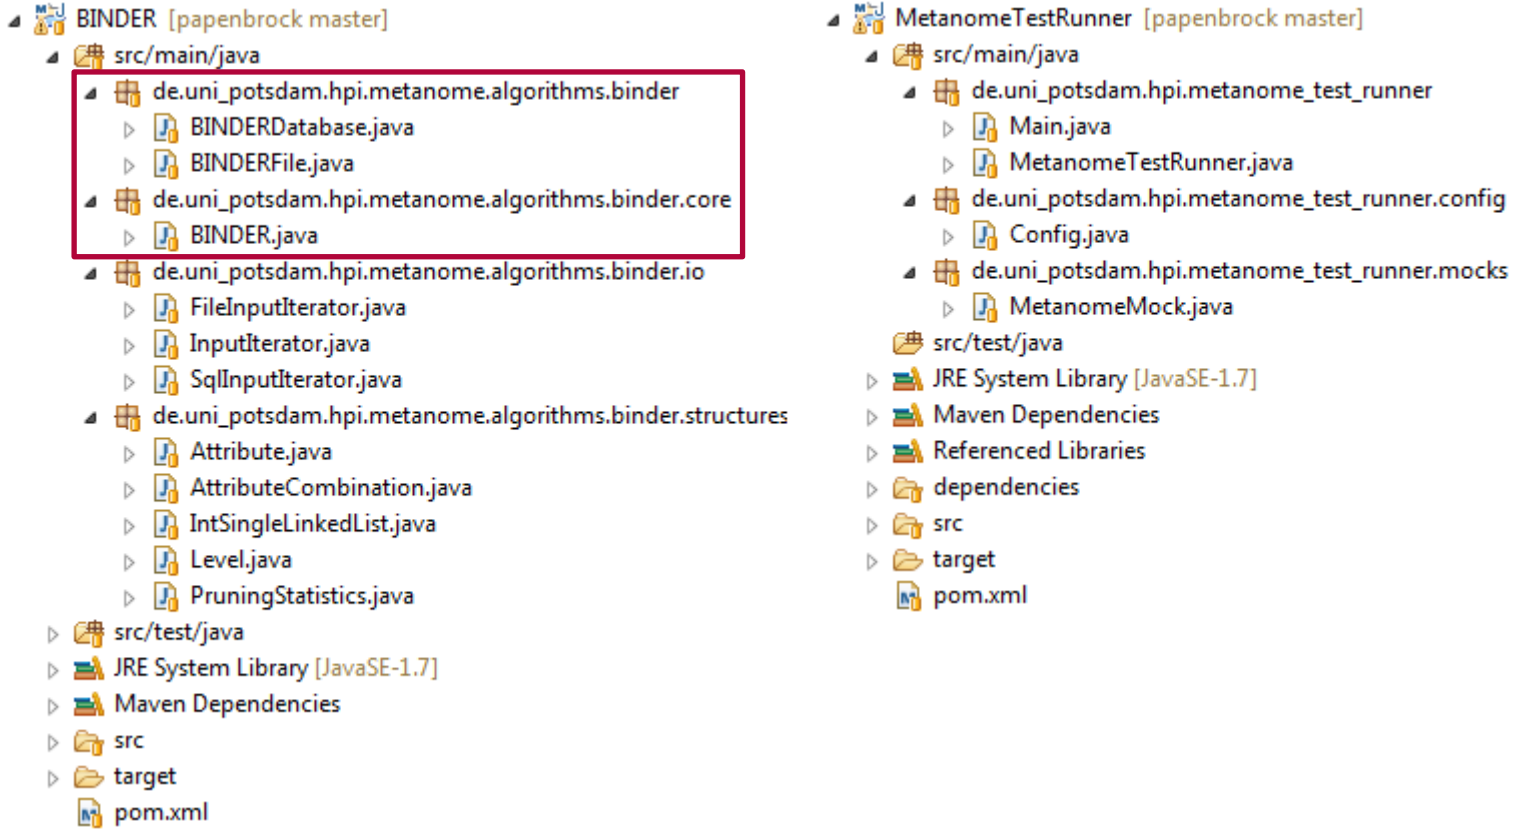

# Discovery of key candidates – Execise

## Task 2: Documentation

Write a short (max one A4 page) documentation for your algorithm describing the algorithm that you implemented:

a) Describe the algorithm's basic idea. How does the algorithm cope with the complexity of the given task?

b) If you used an algorithm from literature, provide a reference to the according publication.

c) If you came up with an own approach, provide one or two arguments why it is or could be better than related algorithms.

d) If you algorithm implements an adaption or optimization of existing approaches, describe these briefly.

e) State if your algorithm (jar) should be published as Metanome algorithm providing you as the authors.

f) *If you solved a bonus task, please discuss your findings here as well.*

In the same document, answer the following questions:

a) How many unique column combinations did your algorithm find on the provided datasets?

b) How long did the discovery take on each dataset and what machine did you use?

c) Did you discover any limitations of your approach (e.g. runtime or memory consumption) that made computing a certain dataset impossible?

d) What is the conceptual differences between NULL $\neq$ NULL and NULL = NULL and how

# Discovery of key candidates – Execise

## Task 3: Presentation

Prepare two slides for a short, 5 min presentation of your algorithm in the lecture. One slide about the algorithm and one slide about its performance. Here are some ideas for these slides:

a) Explain the idea of your approach and why you think it works well in comparison to others.

b) How did your algorithm perform on the given datasets?

c) Did you make any interesting observations?

d) What lets your algorithm crash?

e) How hard was it to implement your algorithm?

f) *This is also a good time to introduce an optimization or adaption of your algorithm.*

Note that each team will present its work once!

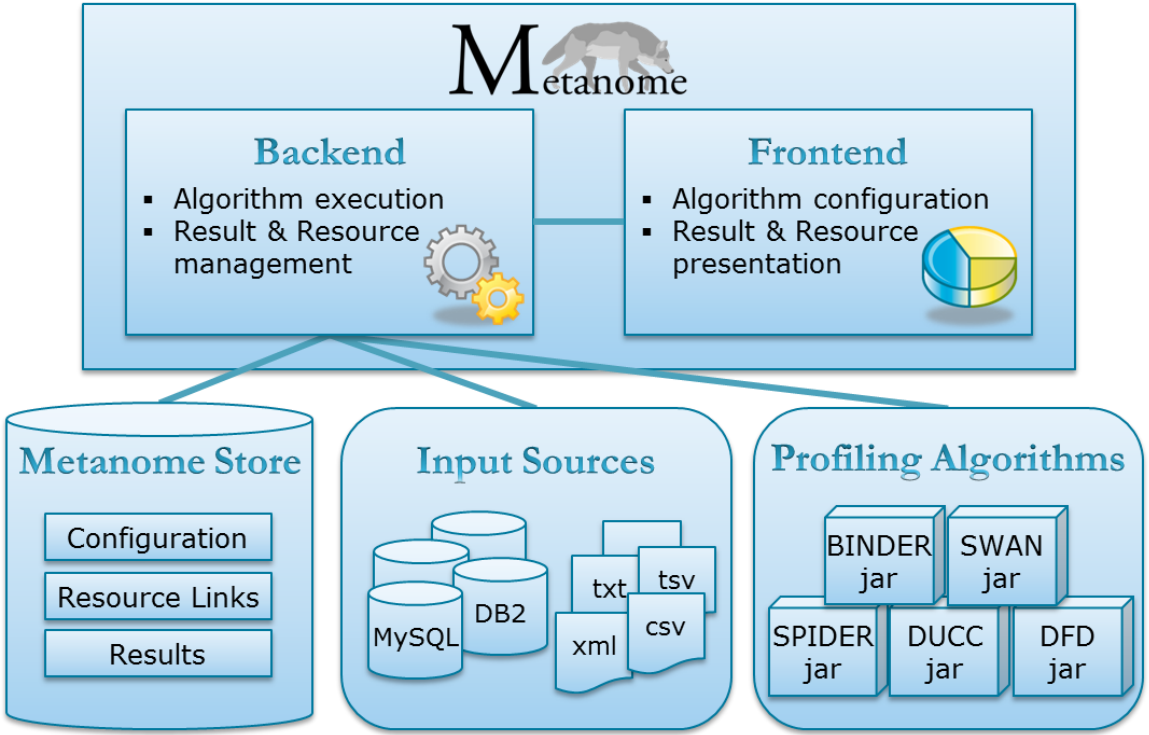


HPI Hasso Plattner Institut
IT Systems Engineering | Universität Potsdam

# Data Profiling with Metanome

*www.metanome.de*

Thorsten Papenbrock
PhD Candidate
Hasso-Plattner-Institute