

Data Curation for ML: Toward a Principled Approach

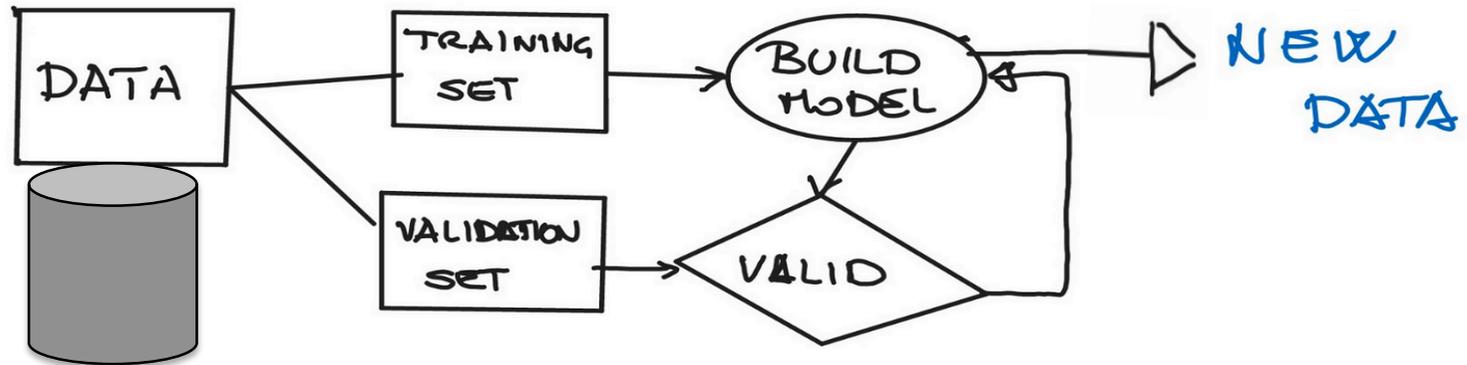
Laure Berti-Equille

laure.berti@ird.fr

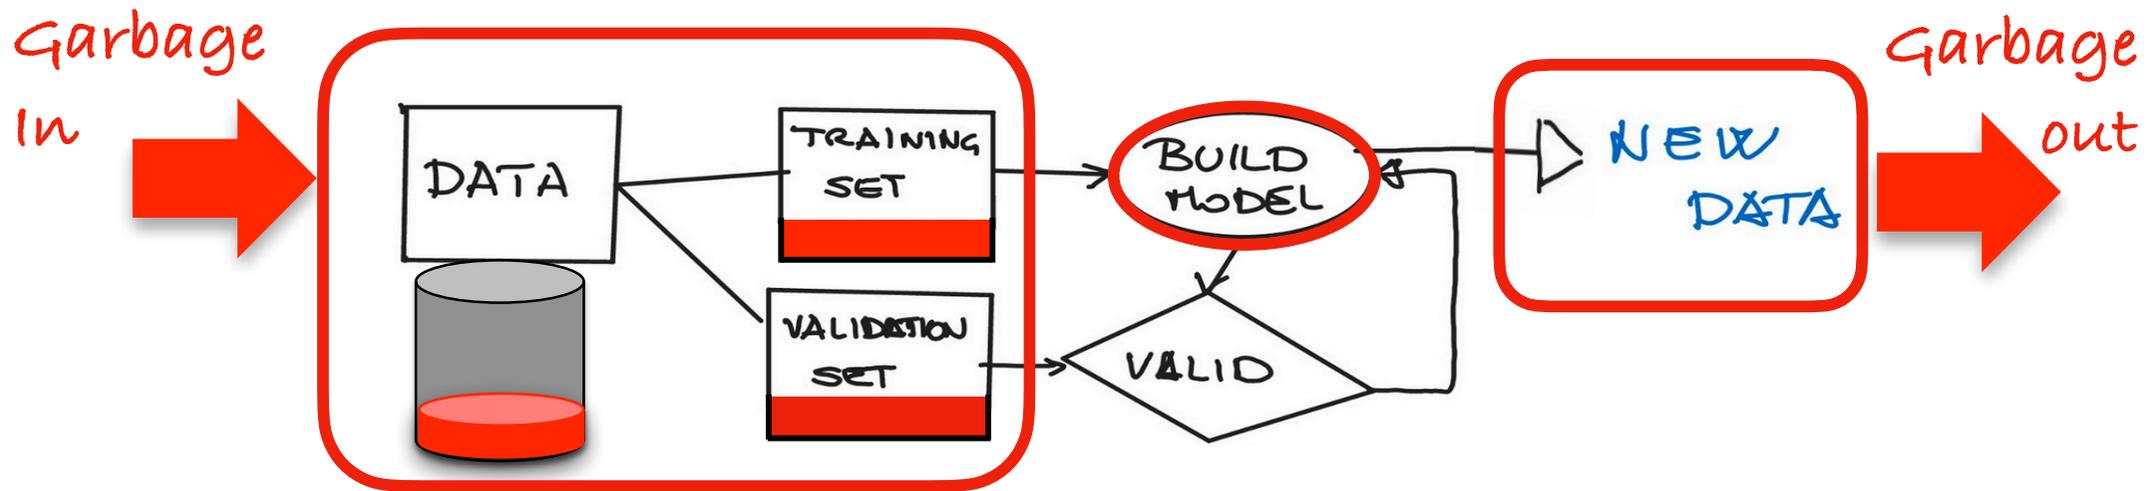


Espace-Dev, IRD, Univ Montpellier, Univ Guyane, Univ La Réunion,
Univ Antilles, Univ Nouvelle Calédonie, Montpellier France

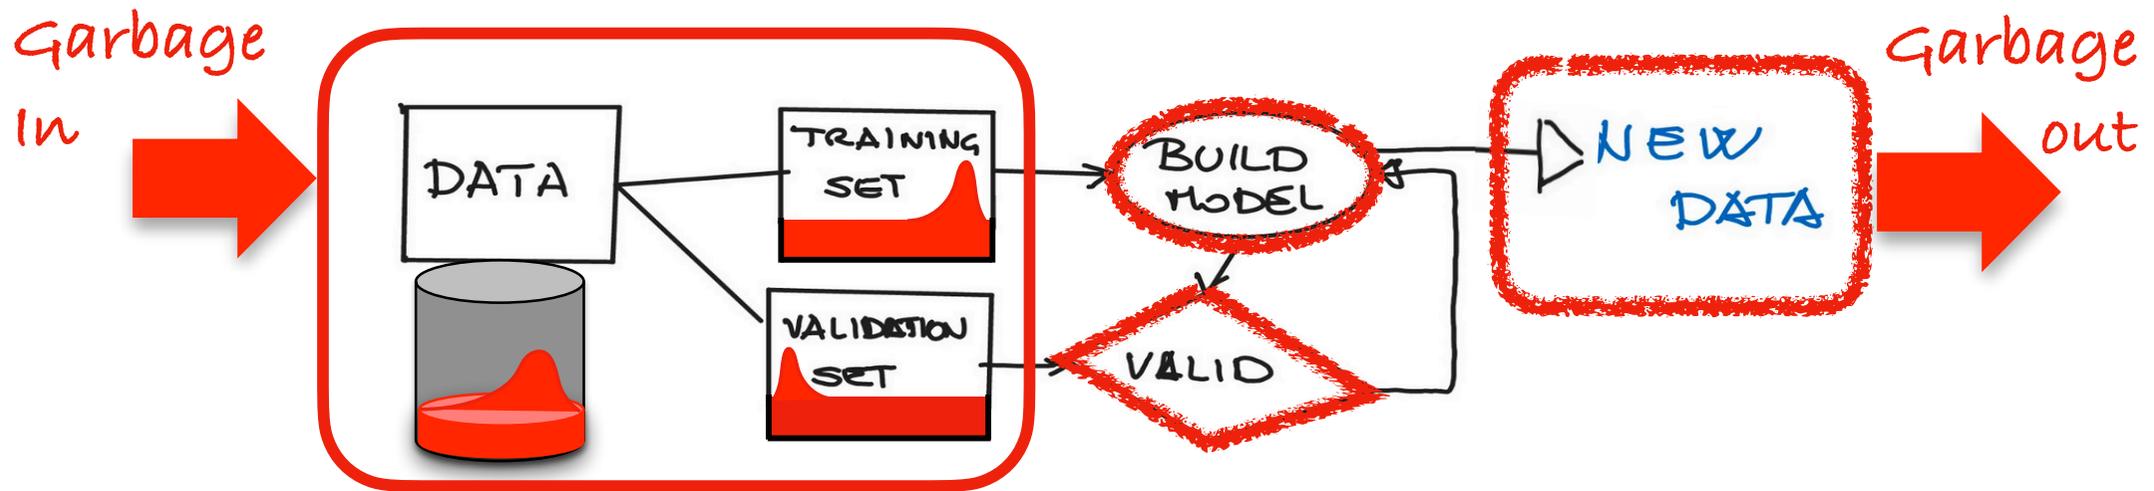
Learning from dirty data is risky



Learning from dirty data is risky



Learning from dirty data is risky



Glitch types and distributions can be very different in the datasets used for training, testing, and validation and they affect accuracy of ML models in different ways.

Two complementary approaches



INTERVENE

How to efficiently fix the data:

- ◆ Detect the anomalies
- ◆ Correct them with minimal cost (domain expert intervention, time, external master data, etc.)
- ◆ Select the repair/preparation strategies that will maximize the ML result quality



MITIGATE

How to reduce the impact of dirty data:

- ◆ Robustify the ML algorithms and apply ML ensembling strategies
- ◆ Use AutoML to find optimal parameter setting
- ◆ Select portions of the data and/or augment the data

Two complementary approaches



INTERVENE

How to efficiently fix the data:

- ◆ Detect the anomalies
- ◆ Correct them with minimal cost (domain expert intervention, time, external master data, etc.)
- ◆ Select the repair/preparation strategies that will maximize the ML result quality



MITIGATE

How to reduce the impact of dirty data:

- ◆ Robustify the ML algorithms and apply ML ensembling strategies
- ◆ Use AutoML to find optimal parameter setting
- ◆ Select portions of the data and/or augment the data

Outline

1. **Detection of data quality problems**
Profiling data quality
2. **Data cleaning**
Leveraging the patterns of glitches
3. **Data preparation strategies for ML**
Learning to clean and prepare the data

Outline

- 1. Detection of data quality problems**
Profiling data quality
- 2. Data cleaning**
Leveraging the patterns of glitches
- 3. Data preparation strategies**
Learning to clean and prepare the data

Data Quality Problems

DATA TYPES

0101010101

ACACGTGT

John Doe

High
Medium
Low

RELATIONSHIPS

- Structural (record)
- Sequential
- Graph-based
- Temporal
- Spatial
- Spatio-Temporal

DATA QUALITY PROBLEMS

TYPE

Missing Data
Anomalous Data
Duplicate Data
Inconsistent Data
Obsolete Data
Incorrect data

CARDINALITY

Single-Point
Collection

DETECTION MODE

Model-based
Data distribution-based
Constraint-based
Pattern-based

Existing approaches for detecting/fixing DQ problems



Declarative

- Data debugging
- Checking data assertions
- Transform



ML-based

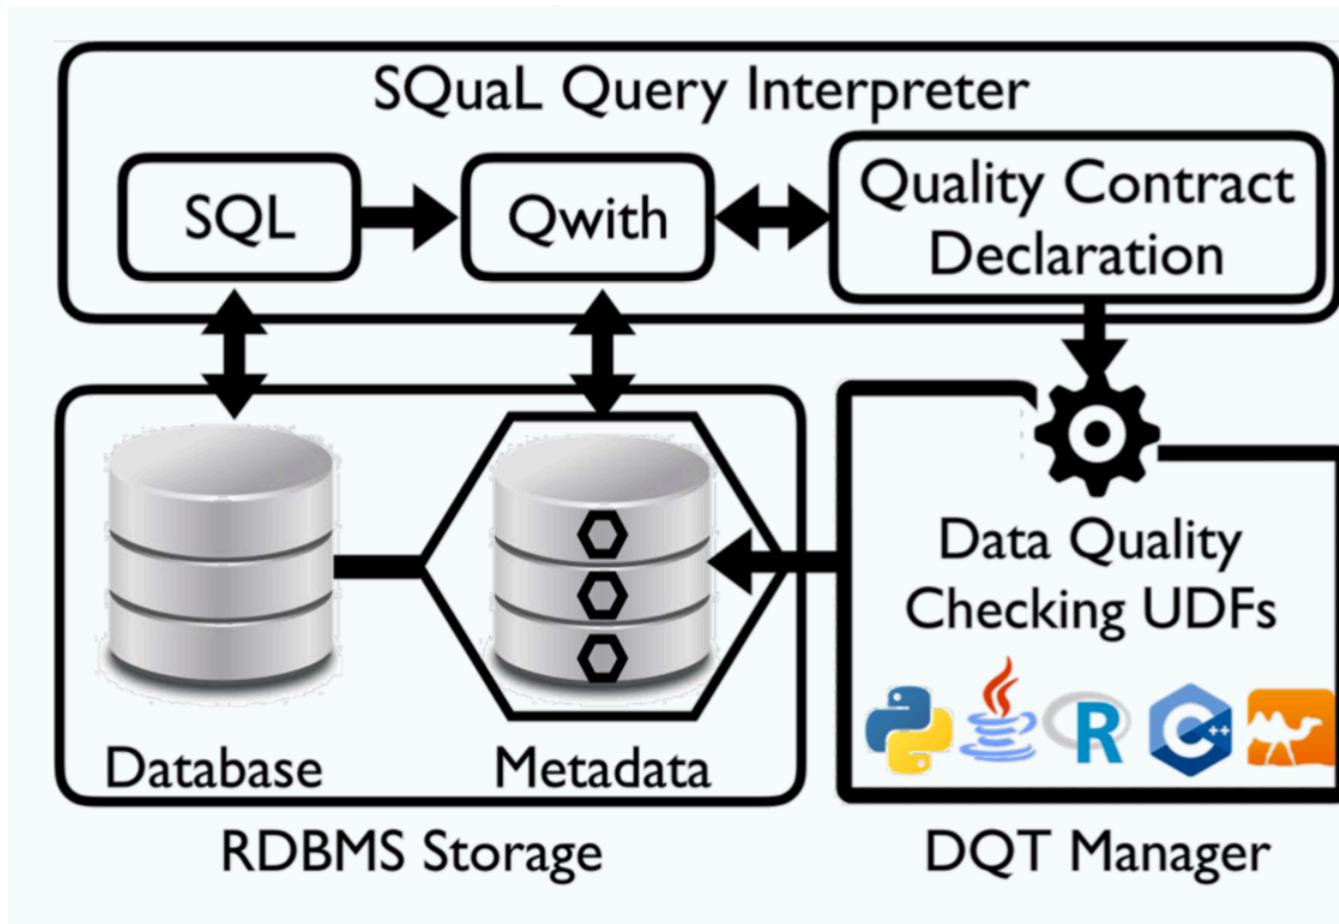
Learn from clean data and replace

Declarative Approaches

Checking data assertions and transform

- ◆ **Deequ** [Schelter et al., VLDB 2018] requires cloud infrastructure and manual integration into training and serving systems; dependent on Apache Spark
- ◆ **TensorFlow Data Validation** (TFDV) [Caveness et al., SIGMOD 2020] integrated with Google TFX difficult to use outside of these platforms
- ◆ Lightweight Python-based approaches like **great_expectations** (<https://greatexpectations.io>) or **hooqu** (<https://github.com/mfcabrera/hooqu>) not integrated with the ML development process

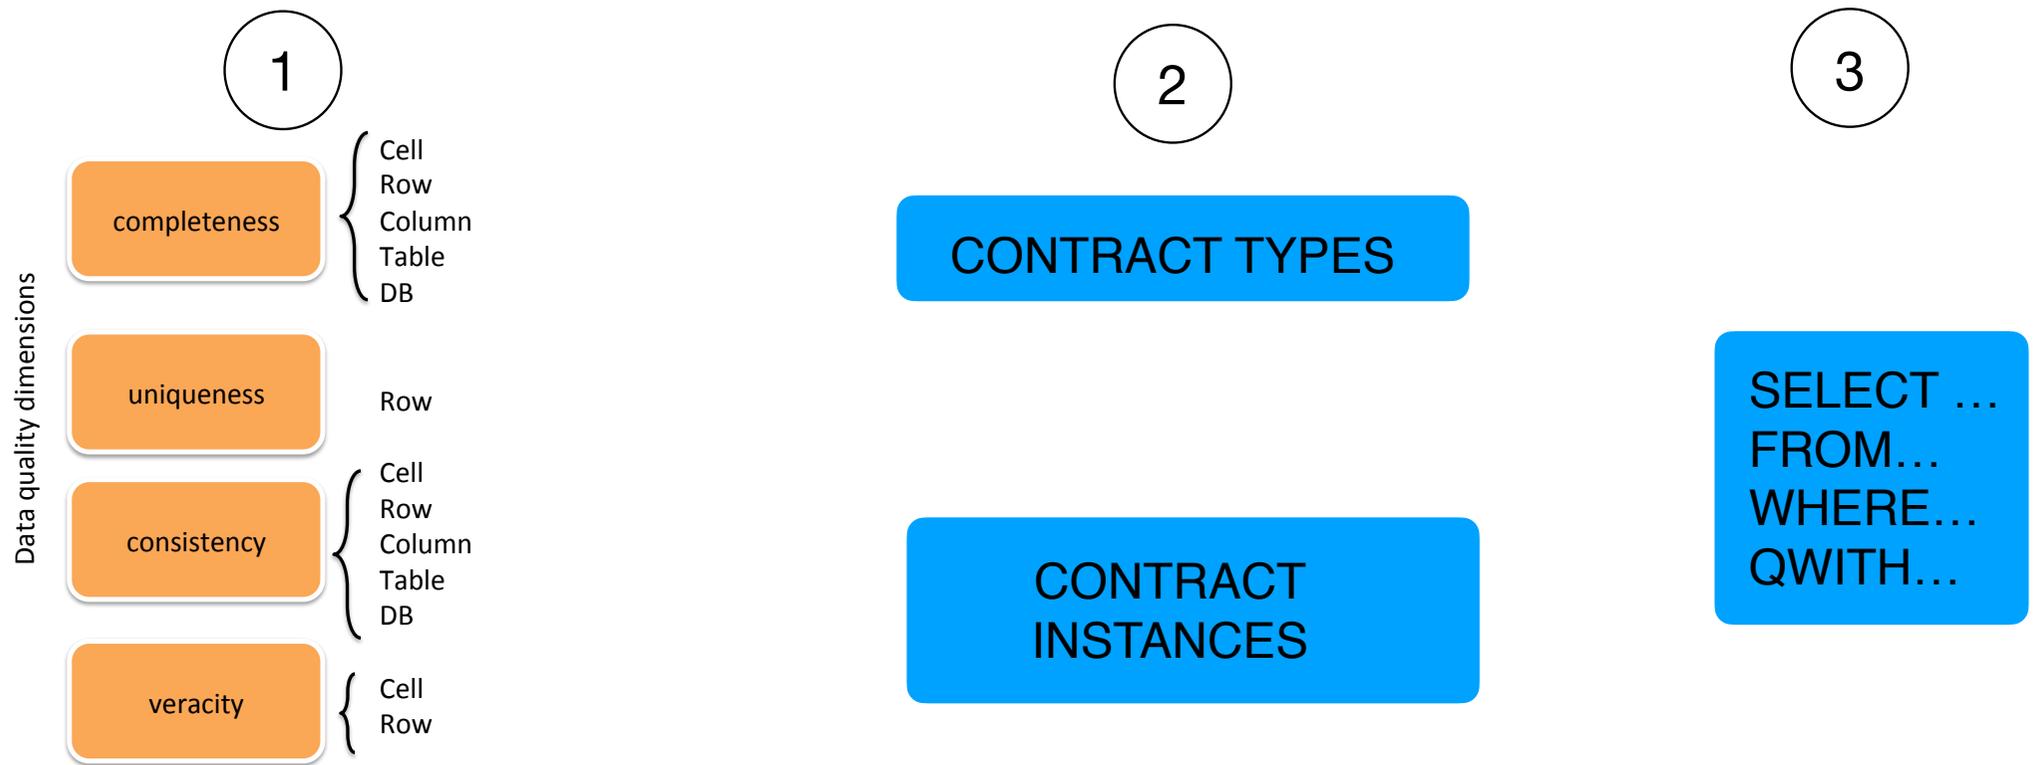
Declarative data profiling with MeSQual



<https://github.com/ucomignani/MeSQual>

MeSQual Key Concepts

Flexible declarative data quality profiling with UDFs



Procedural approach with UDFs

Declarative approach

Extended query

MeSQual Examples

DECLARATION

```
CREATE CONTRACTTYPE StatTests (
  autocorrelation BY FUNCTION 'durbinWatsonTest.py' LANGUAGE PYTHON,
  multicollinearity BY FUNCTION 'varInflationFactor.py' LANGUAGE PYTHON,
  heteroscedasticity BY FUNCTION 'BreuschPaganTest.py' LANGUAGE PYTHON,
  KMerrorNormality BY FUNCTION 'KolmogorovSmirnov.py' LANGUAGE PYTHON,
  SWerrorNormality BY FUNCTION 'ShapiroWilkTest.py' LANGUAGE PYTHON);
```

```
CREATE CONTRACT RegressionAssumptions (
  StatTests.autocorrelation > 0
  AND StatTests.autocorrelation < 4
  AND StatTests.multicollinearity <= 4
  AND StatTests.heteroscedasticity < 0.05
  AND StatTests.SWerrorNormality < 0.05);
```

```
CREATE CONTRACTTYPE CheckQDB (
  completeness BY FUNCTION 'completeness.py' LANGUAGE PYTHON,
  uniqueness BY FUNCTION 'uniqueness.py' LANGUAGE PYTHON,
  consistency BY FUNCTION 'consistency.py' LANGUAGE PYTHON,
  outlyingness BY FUNCTION 'outlyingness.py' LANGUAGE PYTHON);
```

```
CREATE CONTRACT CheckBeforeAnalysis (
  RegressionAssumptions
  AND CheckQDB.consistency > 0.9
  AND CheckQDB.outlyingness < 0.2);
```

MANIPULATION

AoT	{ SELECT * FROM ChicagoDataset } QWITH CheckQDB.completeness> 0.95;
	{ SELECT * FROM ChicagoDataset } QWITH CheckBeforeAnalysis AND RegressionAssumptions;
	{ SELECT timestamp, node_id,value_raw,valuehrf FROM ChicagoDataset WHERE ChicagoDataset.sensor = 'o3' } QWITH CheckBeforeAnalysis AND CheckQDB.completeness> 0.95;
MIMIC-III	{ SELECT * FROM Admissions } QWITH CheckQDB.completeness> 0.95;
	{ SELECT * FROM Admissions WHERE Admissions.insurance = 'Private' } QWITH CheckBeforeAnalysis AND CheckQDB.completeness> 0.95;
	{ SELECT gender, dob, admittime FROM Admissions INNER JOIN (SELECT * FROM Patients WHERE dob < '2090-12-12 00:00:00' QWITH CheckQDB.completeness> 0.95) as Pat ON Admissions.subject_id=Pat.subject_id; } QWITH CheckQDB.completeness> 0.95;

MeSQual GUI

SQual Query A

```
{ SELECT timestamp, node_id,value_raw,valuehrf
FROM ChicagoDataset
WHERE ChicagoDataset.sensor = 'o3'
}
QWITH CheckBeforeAnalysis AND CheckQDB.completeness > 0.95
```

Run

Query Results B

timestamp	node_id	value_raw	value_hrf
2019/11/18 12:55:07	001e061146cb	-629.00	0.00
2019/11/18 12:55:06	001e06117b41	970.00	0.00
2019/11/18 12:55:02	001e0610ee43	1.55 K	0.00
2019/11/18 12:54:59	001e061183f3	1.83 K	0.00
2019/11/18 12:54:54	001e061144be	1.67 K	0.00
2019/11/18 12:54:52	001e0610f6db	436.00	0.00

Tables C

Test	CONTRACTTYPE
Test	CONTRACT
Test	ChicagoDataset

Contracts

contractName	constraintOperator	dimensionName	comparedValue
CheckBeforeAnalysis	CONTRACT	RegressionAssumptions	-1.00
CheckBeforeAnalysis	LESSER	outlyingness	0.20
CheckBeforeAnalysis	GREATER	consistency	0.90

Contract Types

contractTypeName	dimensionName	language	functionPath
CheckQDB	outlyingness	PYTHON	outlyingness.py
CheckQDB	consistency	PYTHON	consistency.py
CheckQDB	uniqueness	PYTHON	uniqueness.py

Queries

queryId	query
c21418d8-5e3c-4814-9556-5e0a7196b502	{ SELECT timestamp, node_id,value_raw,valuehrf FROM ChicagoDataset WHERE ChicagoDataset.sensor = 'o3' } QWITH CheckBeforeAnalysis AND CheckQDB.completeness > 0.95

Data Quality Checks D

Check	Value
completeness db	0.95
consistency db	0.90
completeness value_raw	0.80
completeness value_hrf	0.90
consistency value_raw	0.90
consistency value_hrf	0.90
heteroscedasticity value_hrf	0.10

Monitoring E

Query & Check Logs F

queryId	queryTimestamp	contractType	dimensionName	comparedValue	operator	elementGranularity	elementId	udfResult	violation
c21418d8-5e3c-4814-9556-5e0a7196b502	2019-11-26 15:00:00.000	CheckQDB	consistency	0.90	>	att	value_hrf	0.90	1.00
c21418d8-5e3c-4814-9556-5e0a7196b502	2019-11-26 15:00:00.000	StatTests	heteroscedasticity	0.05	<	att	value_hrf	0.10	1.00
c21418d8-5e3c-4814-9556-5e0a7196b502	2019-11-26 15:00:00.000	CheckQDB	completeness	0.95	>	att	value_raw	0.80	1.00

ML-based Approaches

Learn from clean data and replace/repair

- Pattern enforcement
 - Syntactic patterns (date formatting)
 - Semantic patterns (name/address)
- Value update to satisfy a set of rules, constraints, FDs, CFDs, Denial Constraints (DCs), Matching Dependencies (MDs) with minimal number of changes.
- Value replacement
- Entity resolution

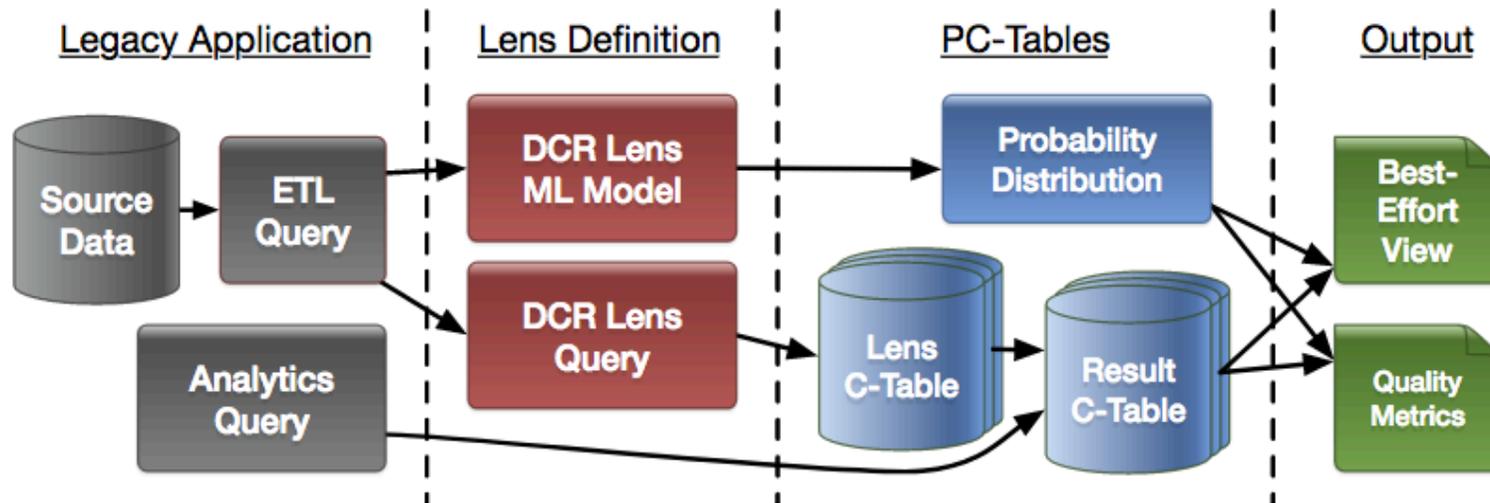
EXAMPLES

- ◆ SCARE: Scalable Automatic Repair
- ◆ On-demand ETL [Yang et al., VLDB'15]
- ◆ ActiveClean [Krishnan et al., VLDB'16]
- ◆ HoloClean [Rekatsinas et al., VLDB 2017]
- ◆ Deep learning for Entity Resolution
- ◆ Transformers for data prep

On-demand ETL with Lenses

[Yang et al., VLDB'15]

Specification of Lens with classifiers from the massive online analysis (MOA) framework for Domain Constraint Repair (DCR).



```
CREATE LENS SaneProduct AS SELECT * FROM Product
  USING DOMAIN_REPAIR( category string NOT NULL,
                       brand   string NOT NULL );
```

id	name	brand	category
P123	Apple 6s, White	$Var('X', R1)$	phone
P124	Apple 5s, Black	$Var('X', R2)$	phone
P125	Samsung Note2	Samsung	phone
P2345	Sony 60 inches	$Var('X', R4)$	$Var('Y', R4)$
P34234	Dell, Intel 4 core	Dell	laptop
P34235	HP, AMD 2 core	HP	laptop

HoloClean

[Rekatsinas et al., VLDB 2017]

<https://github.com/HoloClean/HoloClean>

HoloClean generates a factor graph capturing co-occurrences, correlations based on a set of constraints and external evidences. It uses SGD to learn parameters and infer the marginal distribution of unknown variables with Gibbs sampling.

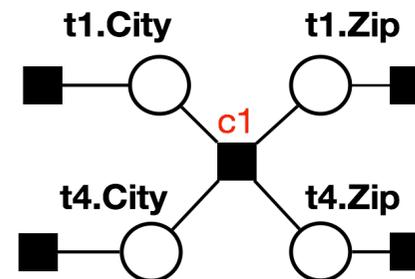
Each cell is a random variable

	Address	City	State	Zip
t1	3465 S Morgan ST	Chicago	IL	60608
t2	3465 S Morgan ST	Chicago	IL	60609
t3	3465 S Morgan ST	Chicago	IL	60609
t4	3465 S Morgan ST	Chicago	IL	60608

Value co-occurrences capture data statistics

Constraints introduce correlations c1: Zip → City

"Address= 3465 S Morgan St"



○ : Unknown (to be inferred) RV
 ■ : Factor (encodes correlations)

Denial constraints:

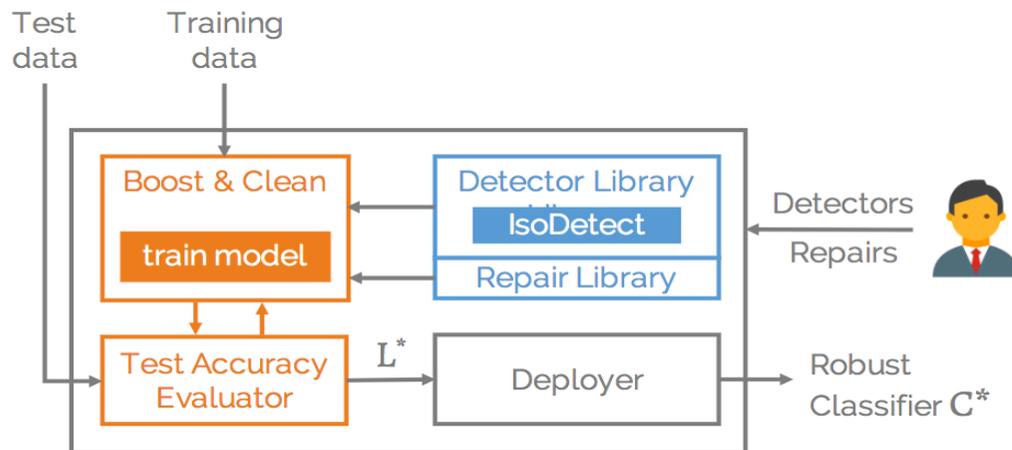
$$\forall t_1, t_2 \in D : \neg(t_1[Zip] = t_2[Zip] \wedge t_1[City] \neq t_2[City])$$

$$\forall t_1, t_2 \in D : \neg(t_1[Zip] = t_2[Zip] \wedge t_1[State] \neq t_2[State])$$

BoostClean

[Krishnan et al., 2017]

BoostClean selects an ensemble of methods (statistical and logic rules) for error detection and for repair combinations using statistical boosting.

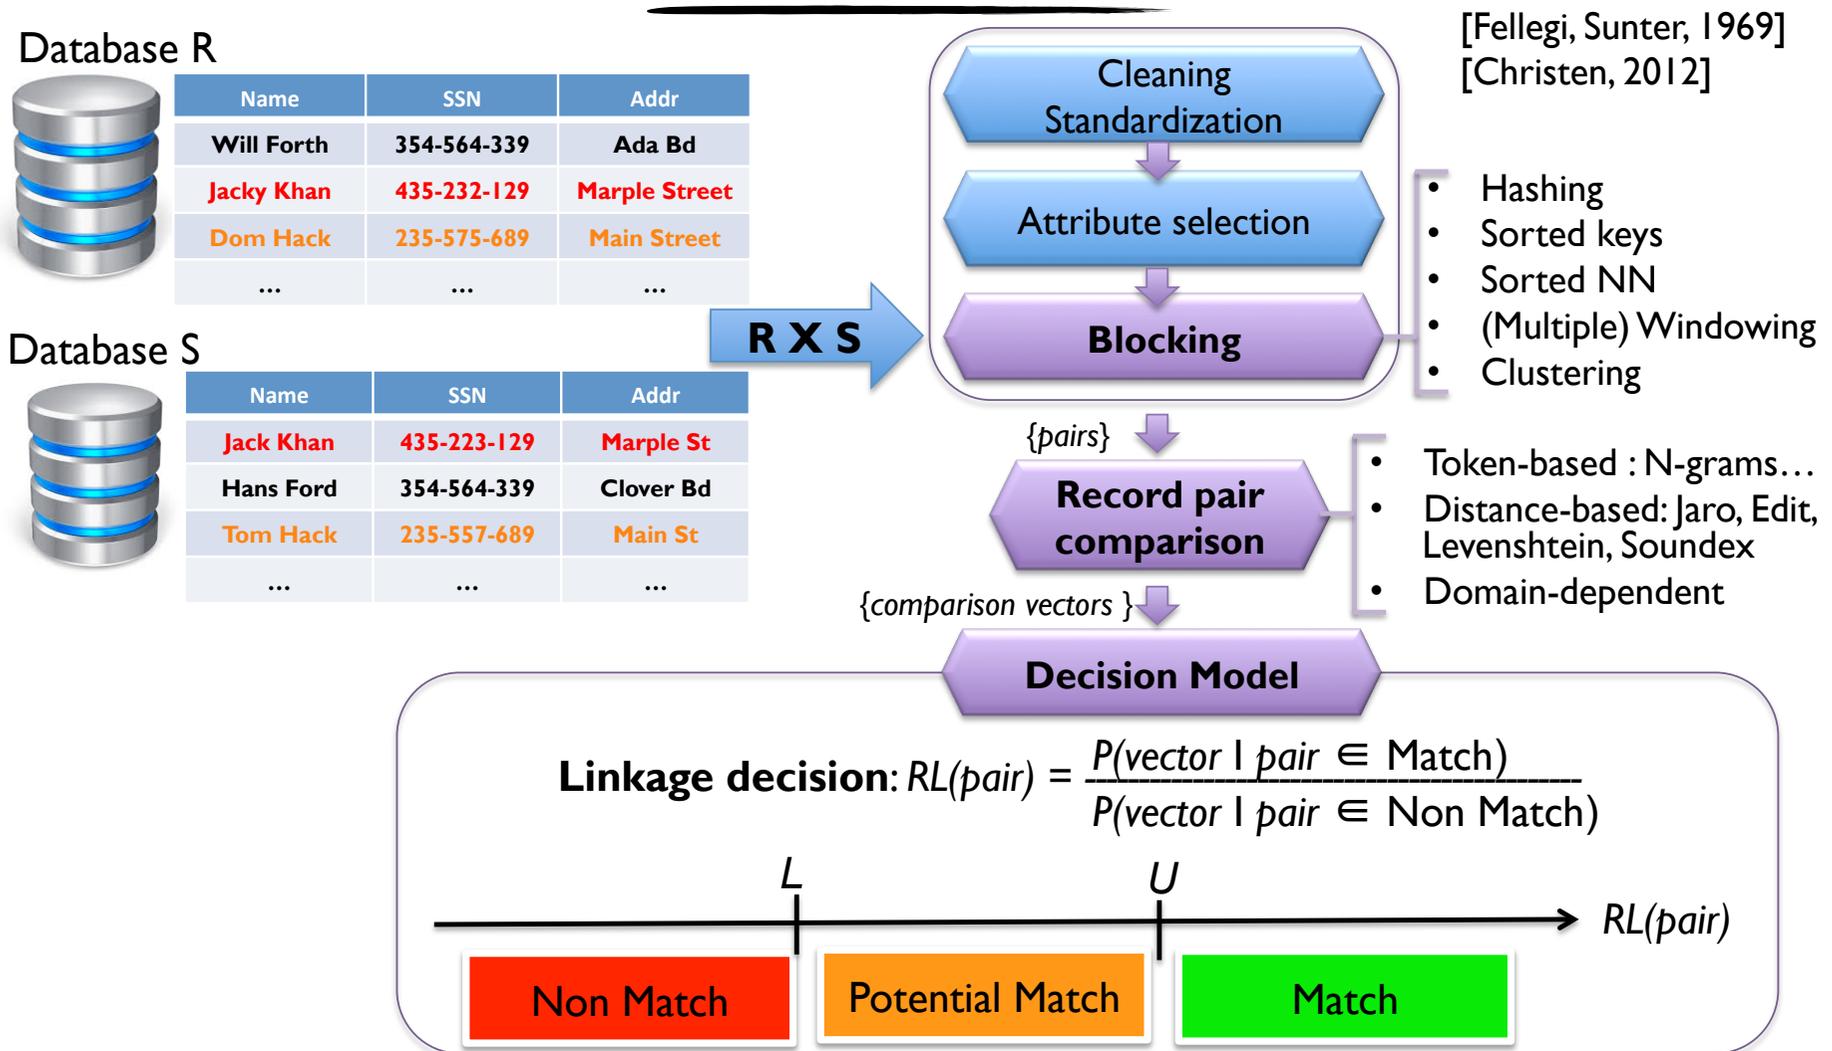


Algorithm 1: Boost-and-Clean Algorithm

Data: (X, Y)

- 1 Initialize $W_i^{(1)} = \frac{1}{N}$
- 2 \mathcal{L} generates a set of classifiers $\mathcal{C} \{C^{(0)}, C^{(1)}, \dots, C^{(k)}\}$ where $C^{(0)}$ is the base classifier and $C^{(1)}, \dots, C^{(k)}$ are derived from the cleaning operations.
- 3 **for** $t \in [1, T]$ **do**
- 4 $C_t = \text{Find } C_t \in \mathcal{C}$ that maximizes the weighted accuracy on the test set. $\epsilon_t = \text{Calculate weighted classification error on the test set } \alpha_t = \ln\left(\frac{1-\epsilon_t}{\epsilon_t}\right)$
 $W_i^{(t+1)} \propto W_i^{(t)} e^{-\alpha_t y_i C_t(x_i)}$: down-weight correct predictions, up-weight incorrectly predictions.
- 5 **return** $C(x) = \text{sign}\left(\sum_t^T \alpha_t C_t(x)\right)$

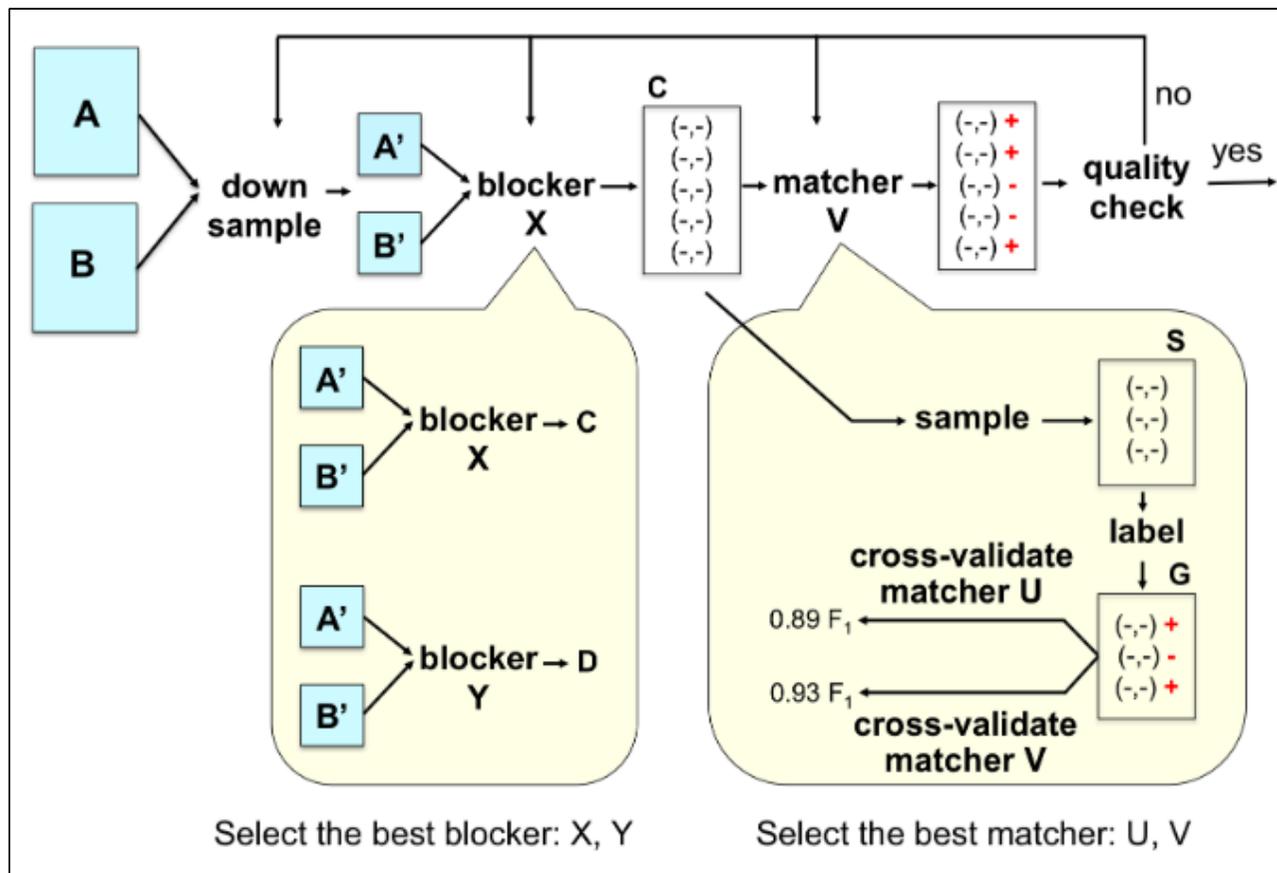
Record Linkage (RL): Generic Workflow



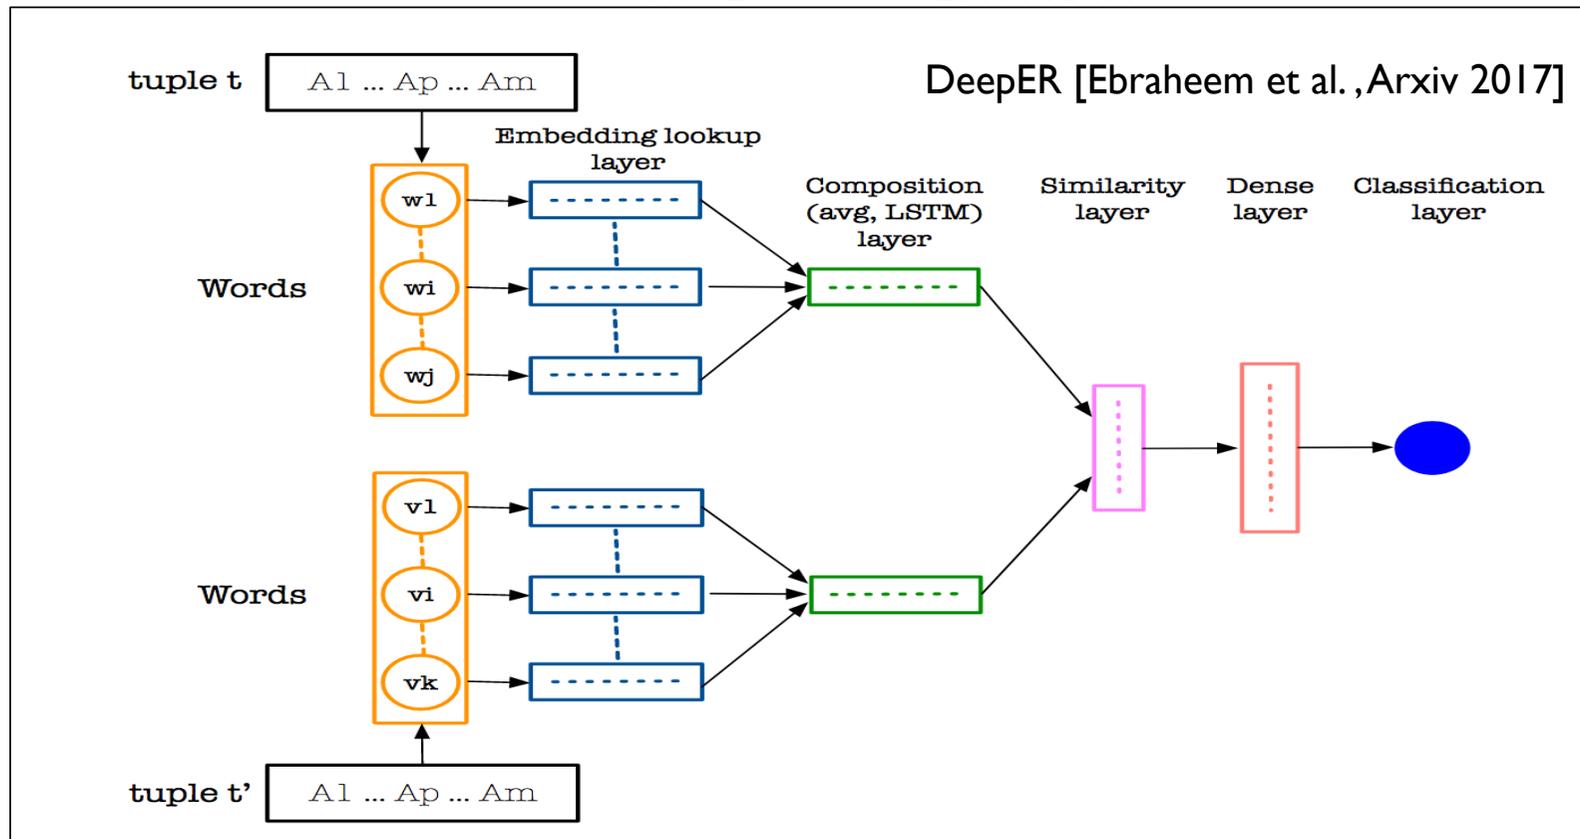
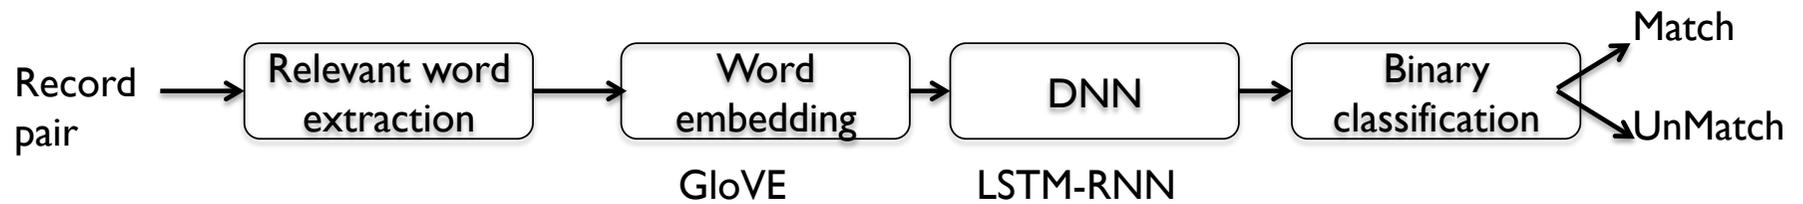
Human-In-The Loop for Entity Matching

[Doan et al., HILDA@SIGMOD'17]

Magellan project: Lessons learnt for How-to Guide for EM



Deep learning for Entity Resolution



Outline

1. **Detection of data quality problems:**

Profiling data quality

2. **Data cleaning**

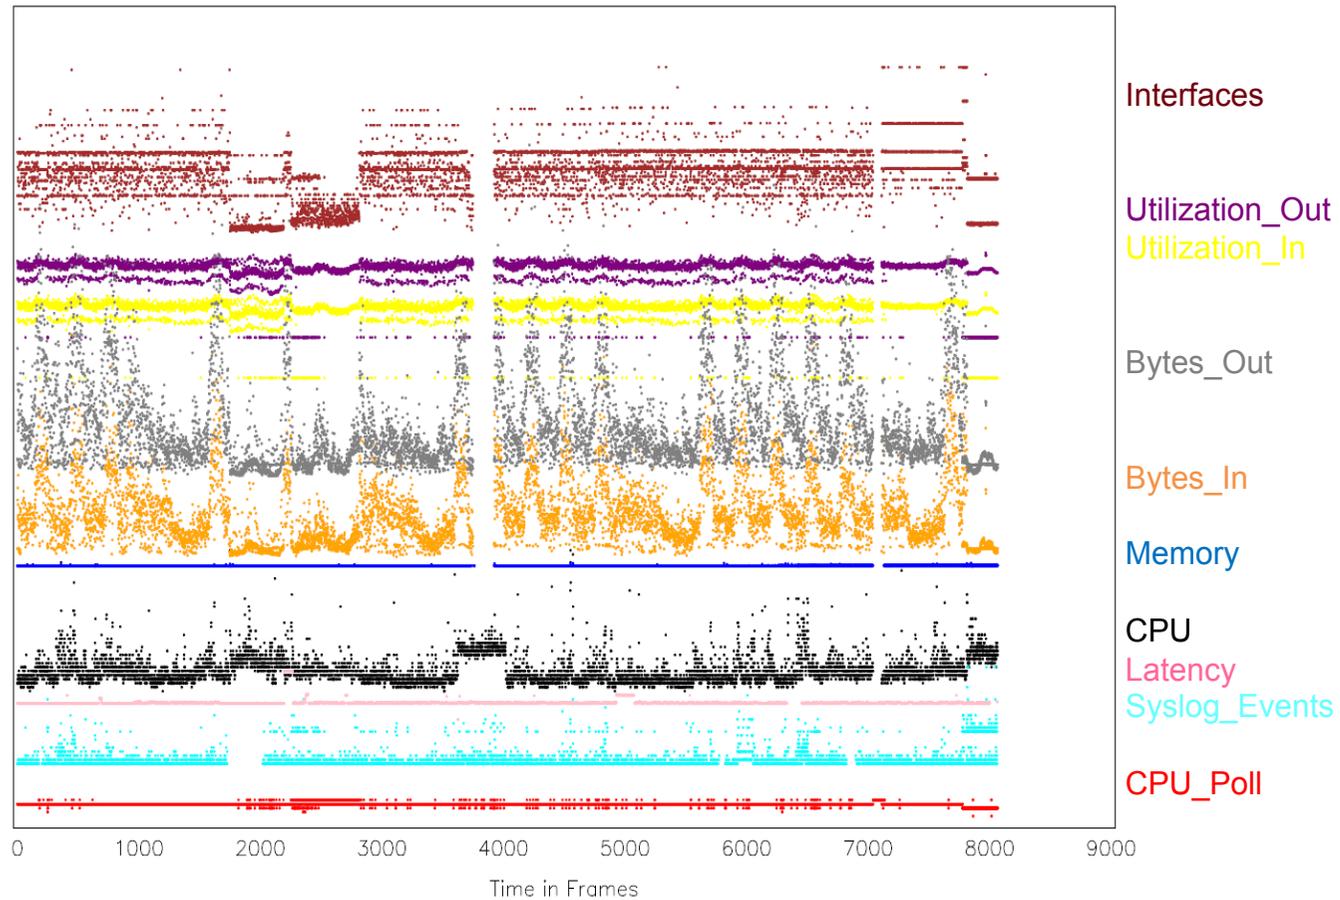
Leveraging the patterns of glitches

3. **Data preparation strategies:**

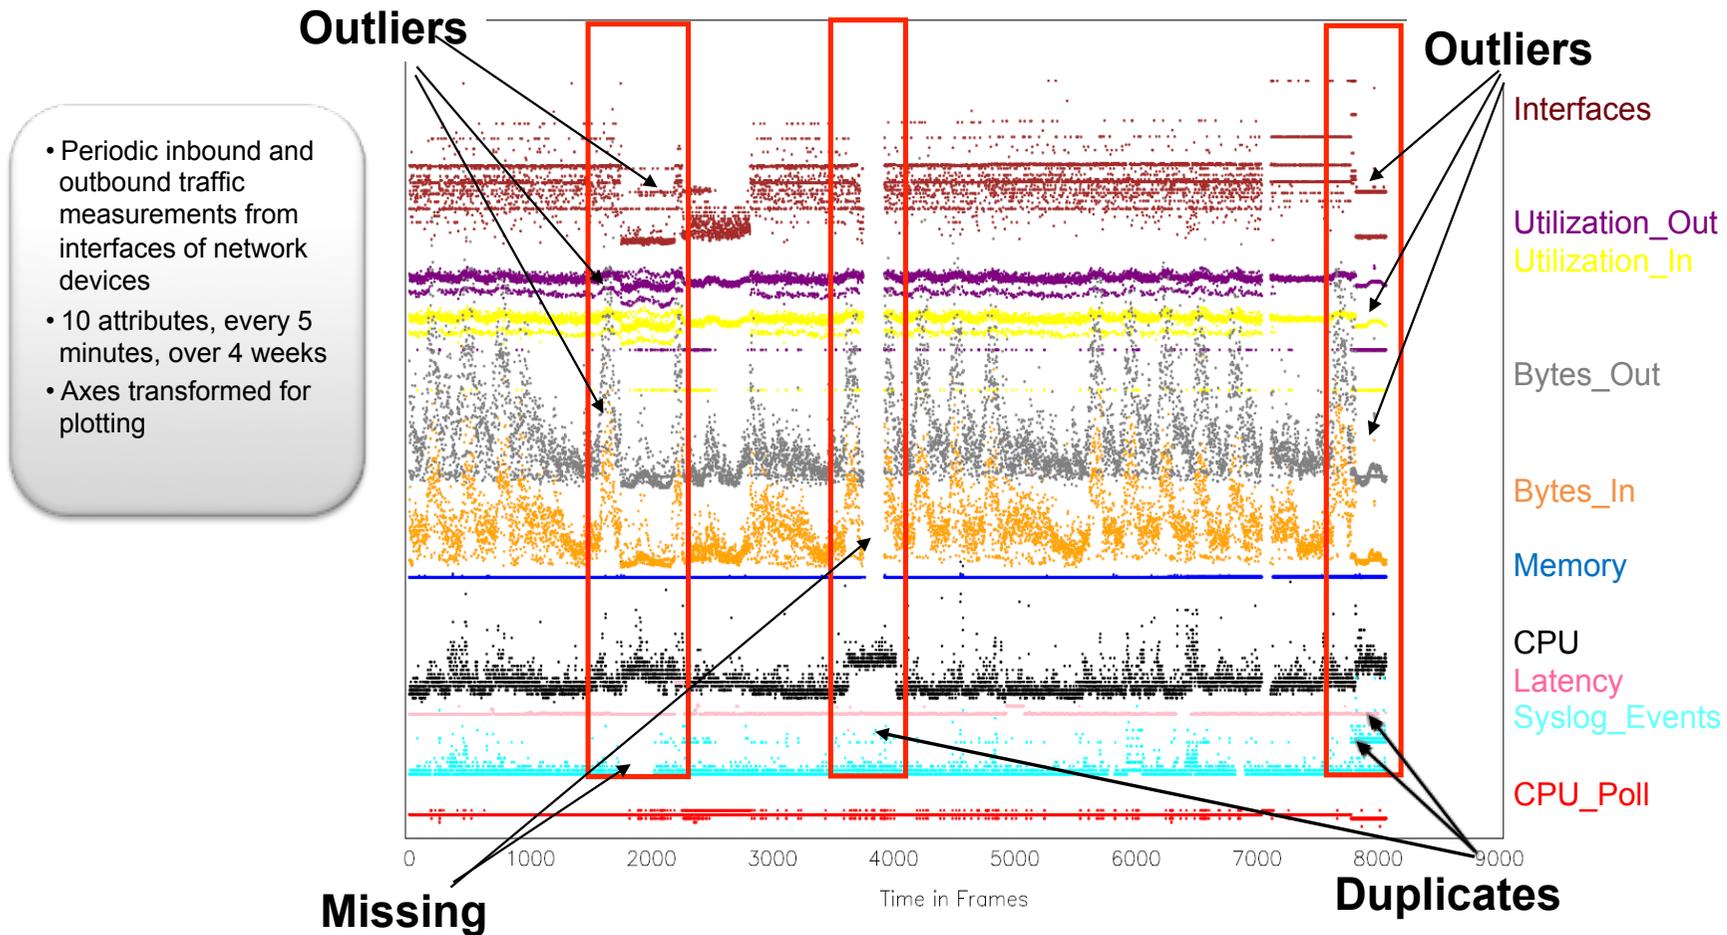
Learning to clean and prepare the data

SNMP Data Analysis

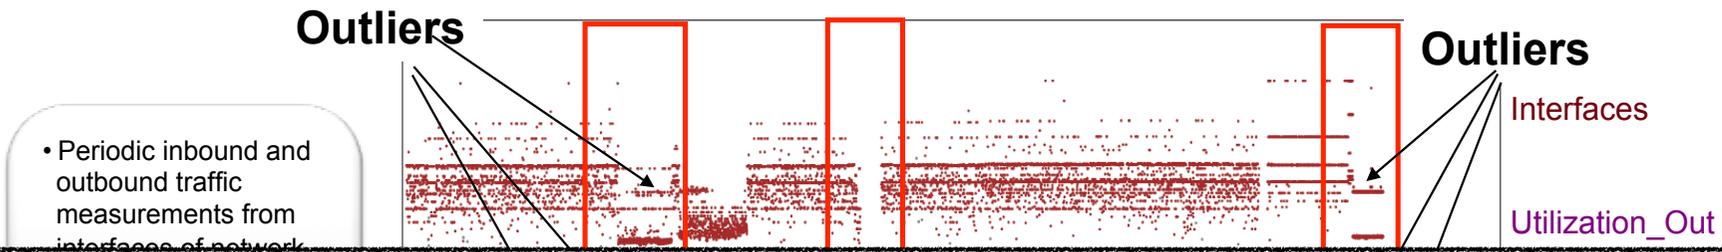
- Periodic inbound and outbound traffic measurements from interfaces of network devices
- 10 attributes, every 5 minutes, over 4 weeks
- Axes transformed for plotting



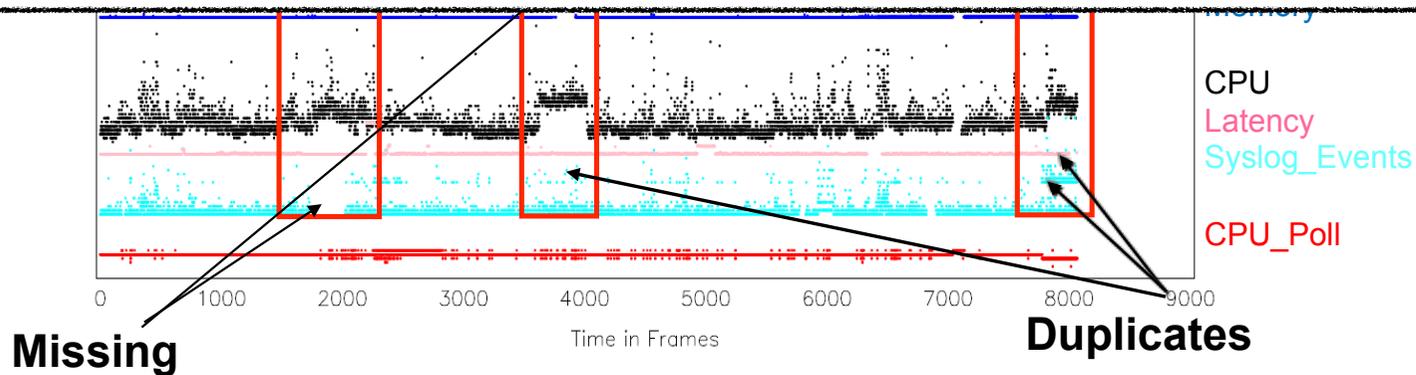
SNMP Data Analysis



SNMP Data Analysis



1. Detect patterns of multivariate, concomitant data anomalies
2. Use the anomaly patterns for consistent cleaning



Understanding Complex Glitch Patterns

Benefits

- A common root cause can generate correlated data errors
- In-depth anomaly analysis could help for:
 - Characterizing anomaly sources, processes, and propagation mechanisms
 - Systematizing data cleaning

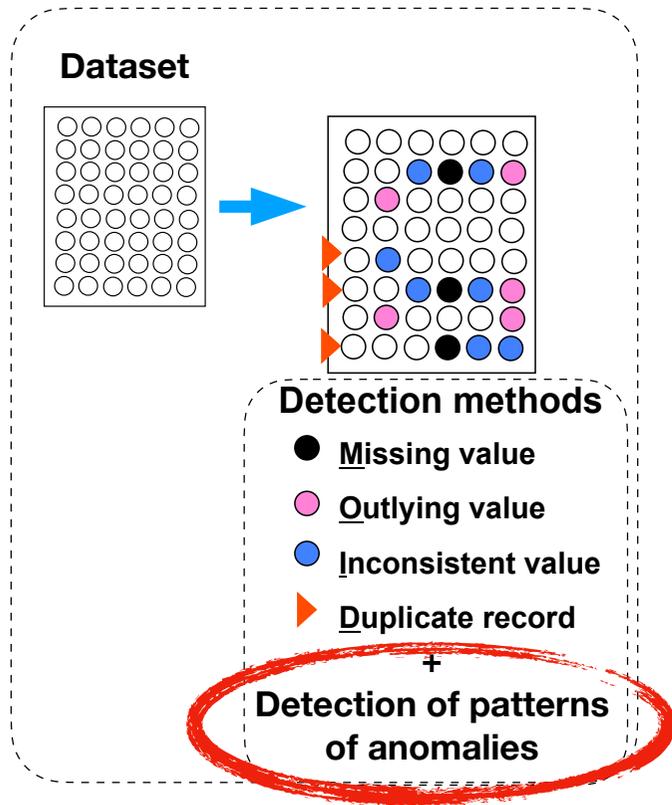
Current methods

- Make unrealistic assumptions (e.g., MAR)
- Treat glitches in isolation
- Are one-shot approaches (no reiteration between detection and cleaning)

Data cleaning and preprocessing may introduce new errors and distortions.

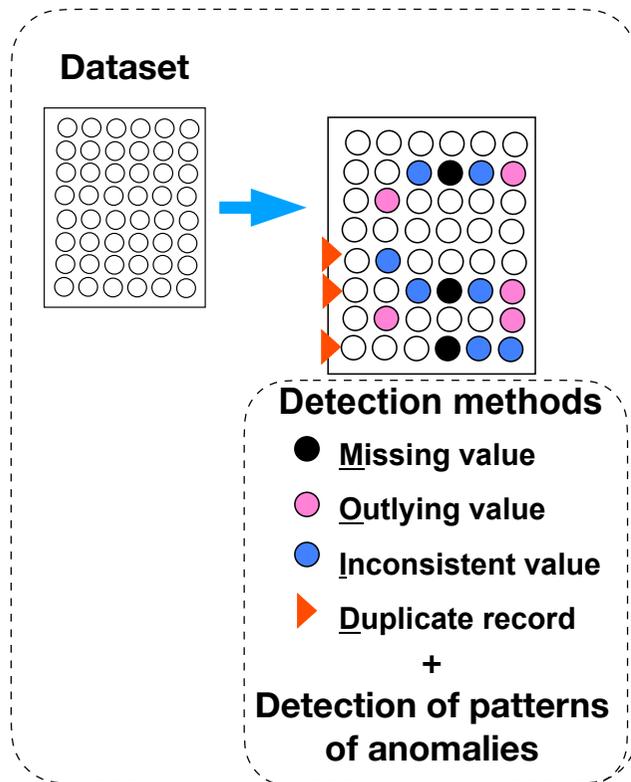
Detection-Exploration-Cleaning Framework

Input:

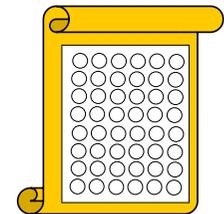


Detection-Exploration-Cleaning Framework

Input:



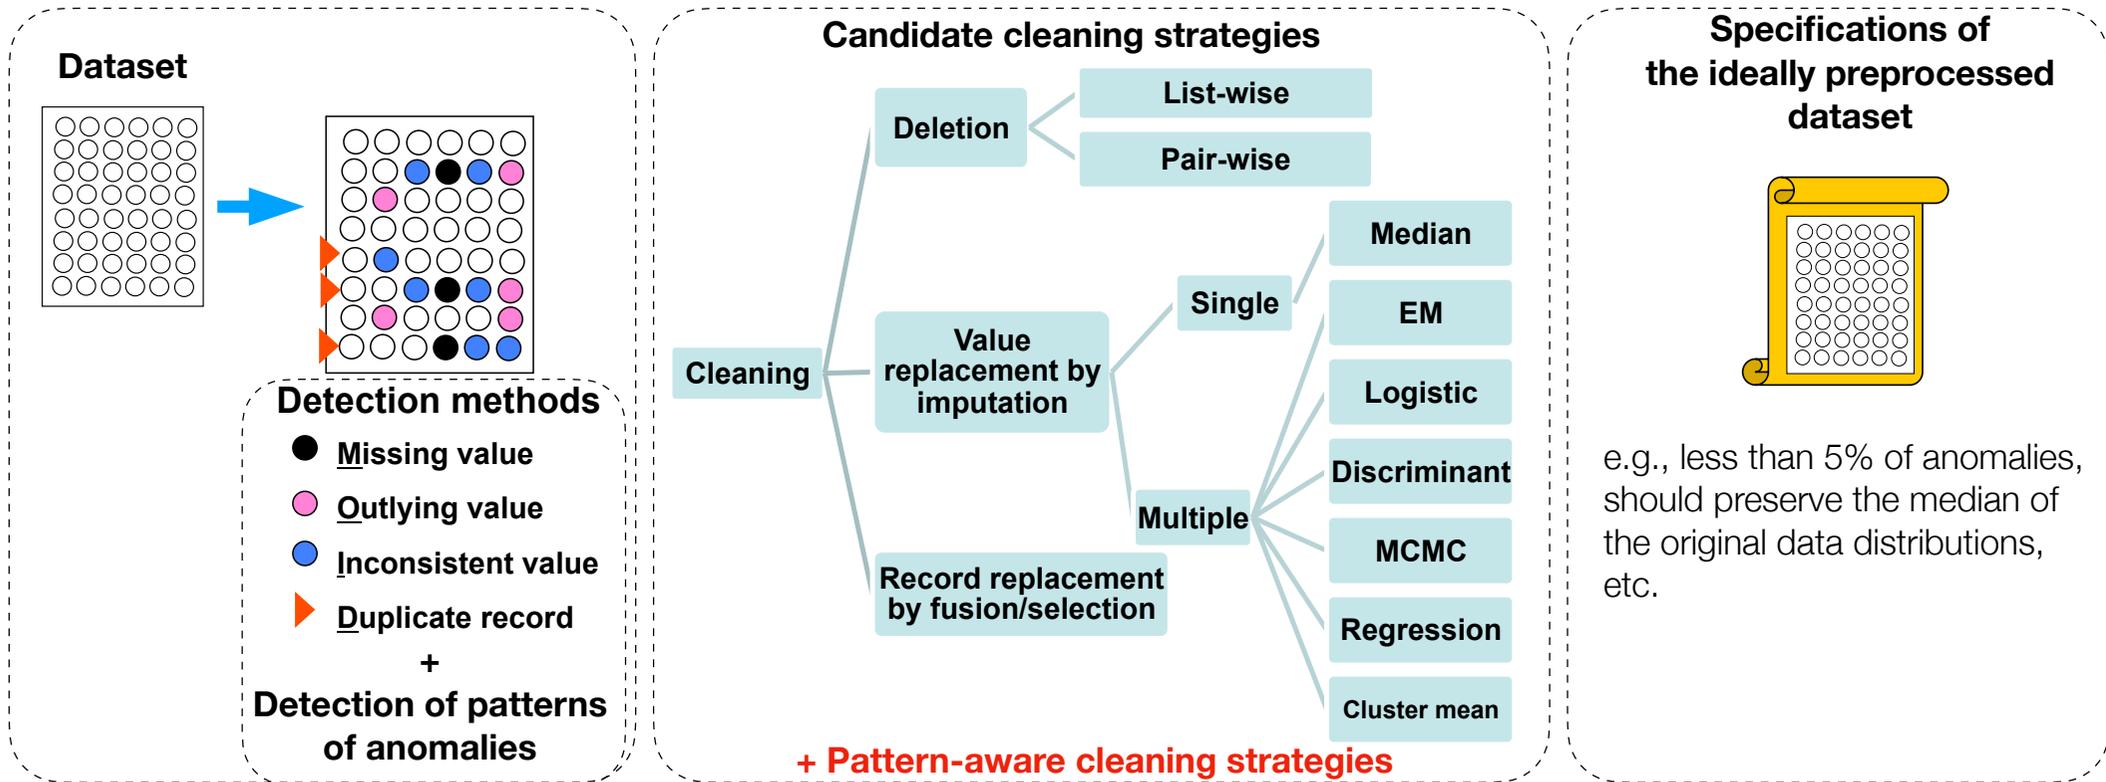
Specifications of the ideally preprocessed dataset



e.g., less than 5% of anomalies, should preserve the median of the original data distributions, etc.

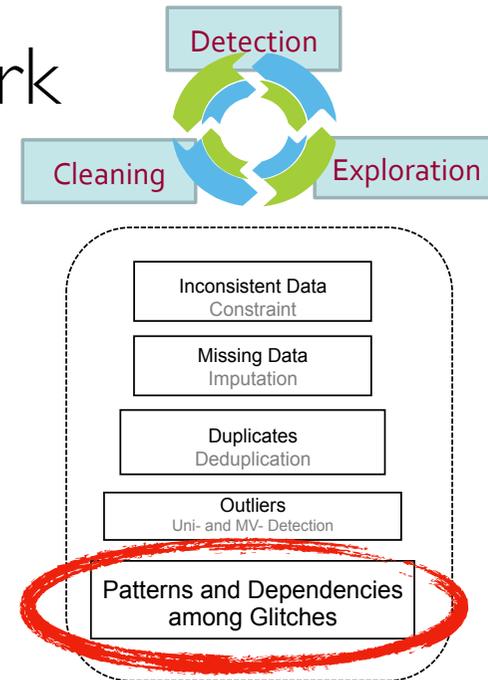
Detection-Exploration-Cleaning Framework

Input:



Detection-Exploration-Cleaning Framework

[Berti-Equille, Dasu, Srivastava, ICDE 2011]



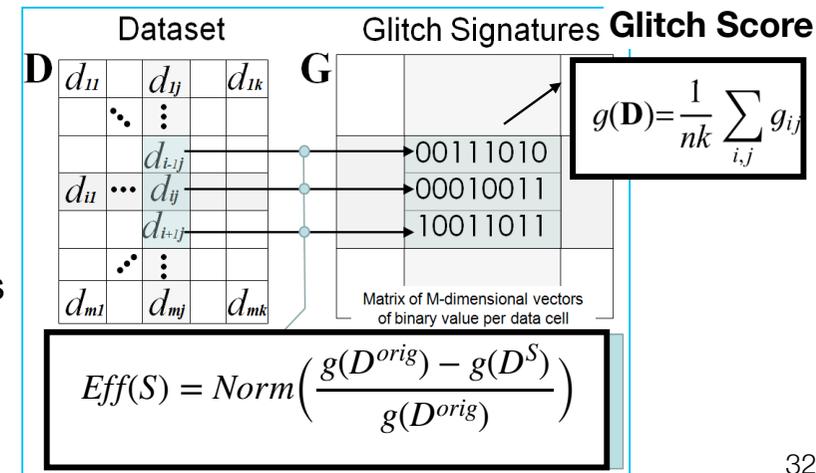
Problem Statement:

Find the quantitative cleaning strategy B composed of M methods among the candidate strategies S such that its resulting dataset D^B is the closest to an ideal dataset D^* specified from D as

$$D^B = \arg \min_{\{s \in S\}} (\text{dist}(D^s, D^*))$$

subject to $Cost(s) \leq U$ and $Eff(s) \geq \Gamma > 0$

- dist** is the Kullback-Leibler distance between two data distributions
- U** is a pre-defined upper bound for the cost of strategy s
- Γ** is the lower bound of $Eff(s)$, the effectiveness of strategy s



Experiments

Real-world and semi-synthetic data

- **EPO Dataset:** 754,075 records, 4 non-key attributes (string, categorical and numerical data)
- **Intel Berkeley Research lab Dataset:** 2,313,682 million readings, 8 attributes (timestamp, sensorID, temperature, light, voltage) collected every 31 seconds from 54 sensors deployed in the between February 28th and April 5th, 2
- **SNMP Dataset:** (8,632 tuples, 11 variables) collected every 5 minutes during one month (timestamps, categorical and numerical values)

Comparison of various cleaning strategies

- Cost-based
- Effectiveness-based
- Resource-driven to treat just p% of glitches (DEC-RD)
- Specification-driven to treat a particular glitch type (DEC-SD)
- Pattern-based (DEC-PD)

Outline

1. **Detection of data quality problems:**

Profiling data quality with MeSQuaL

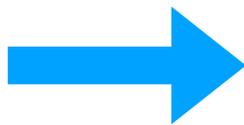
2. **Data cleaning**

Leveraging the patterns of glitches

3. **Data preparation strategies**

Learning to clean and prepare the data

Data preprocessing is challenging



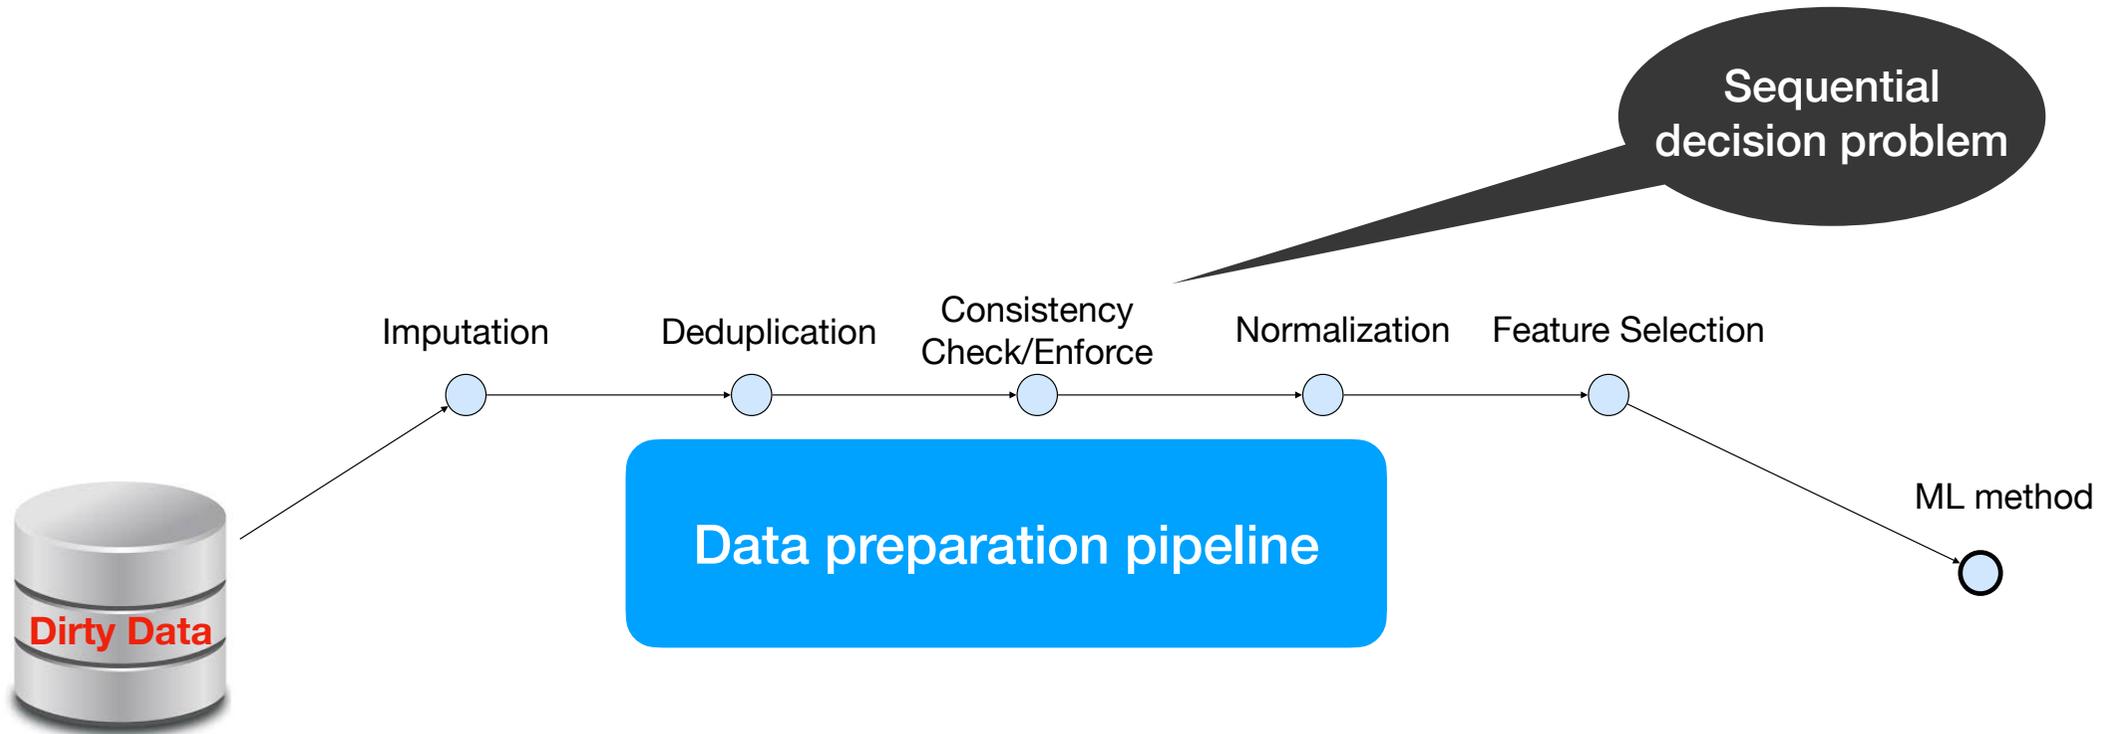
Data preparation pipeline



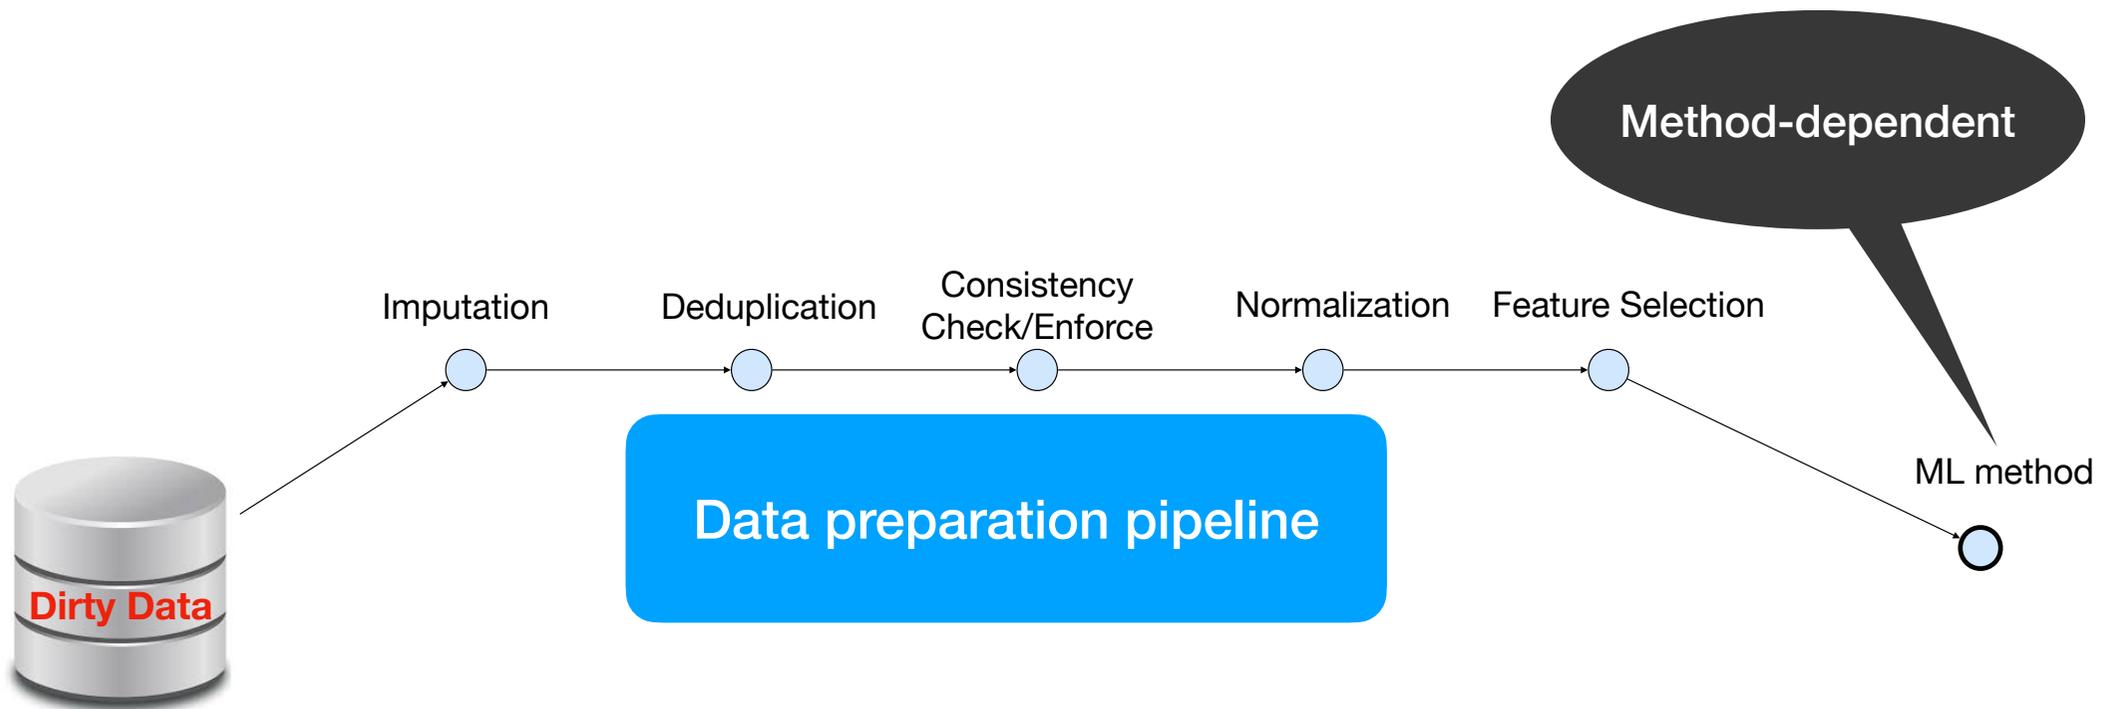
ML method



Data preprocessing is challenging

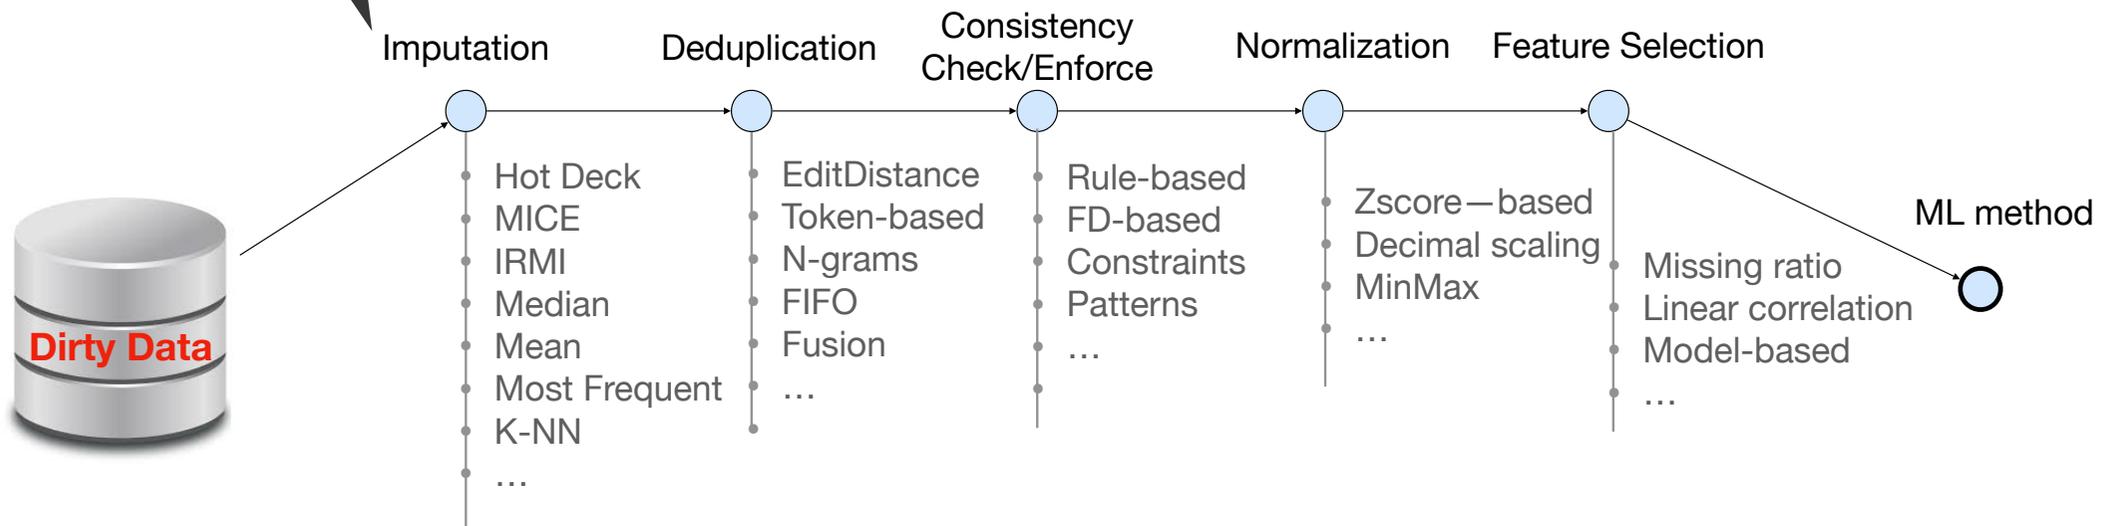


Data preprocessing is challenging

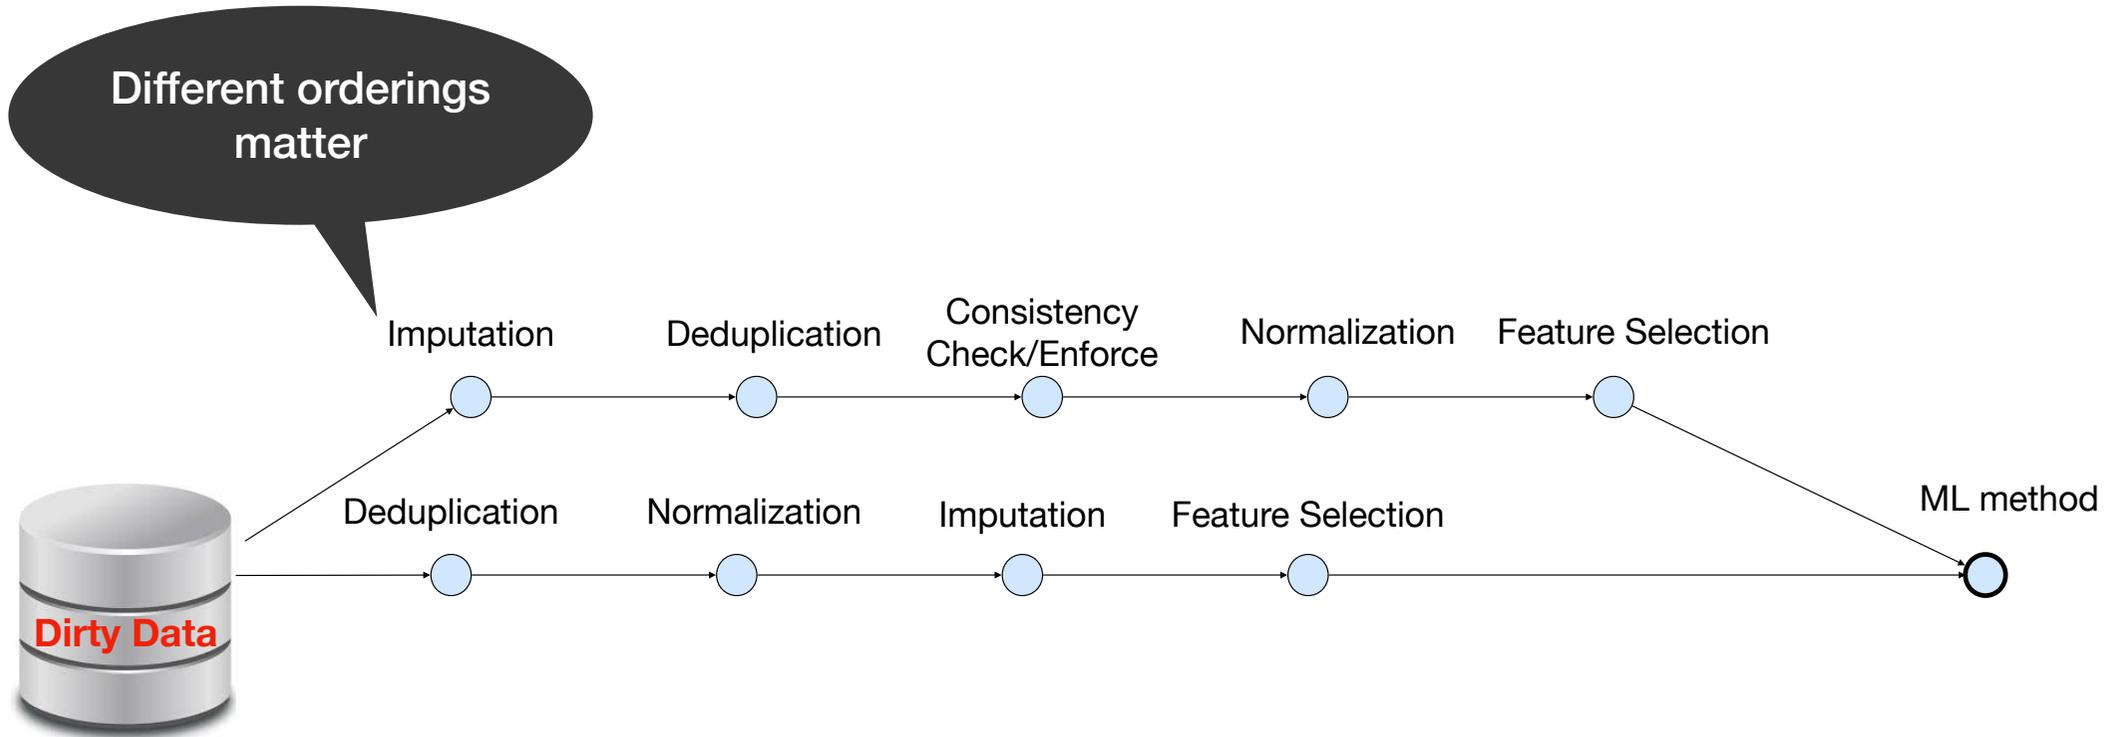


Data preprocessing is challenging

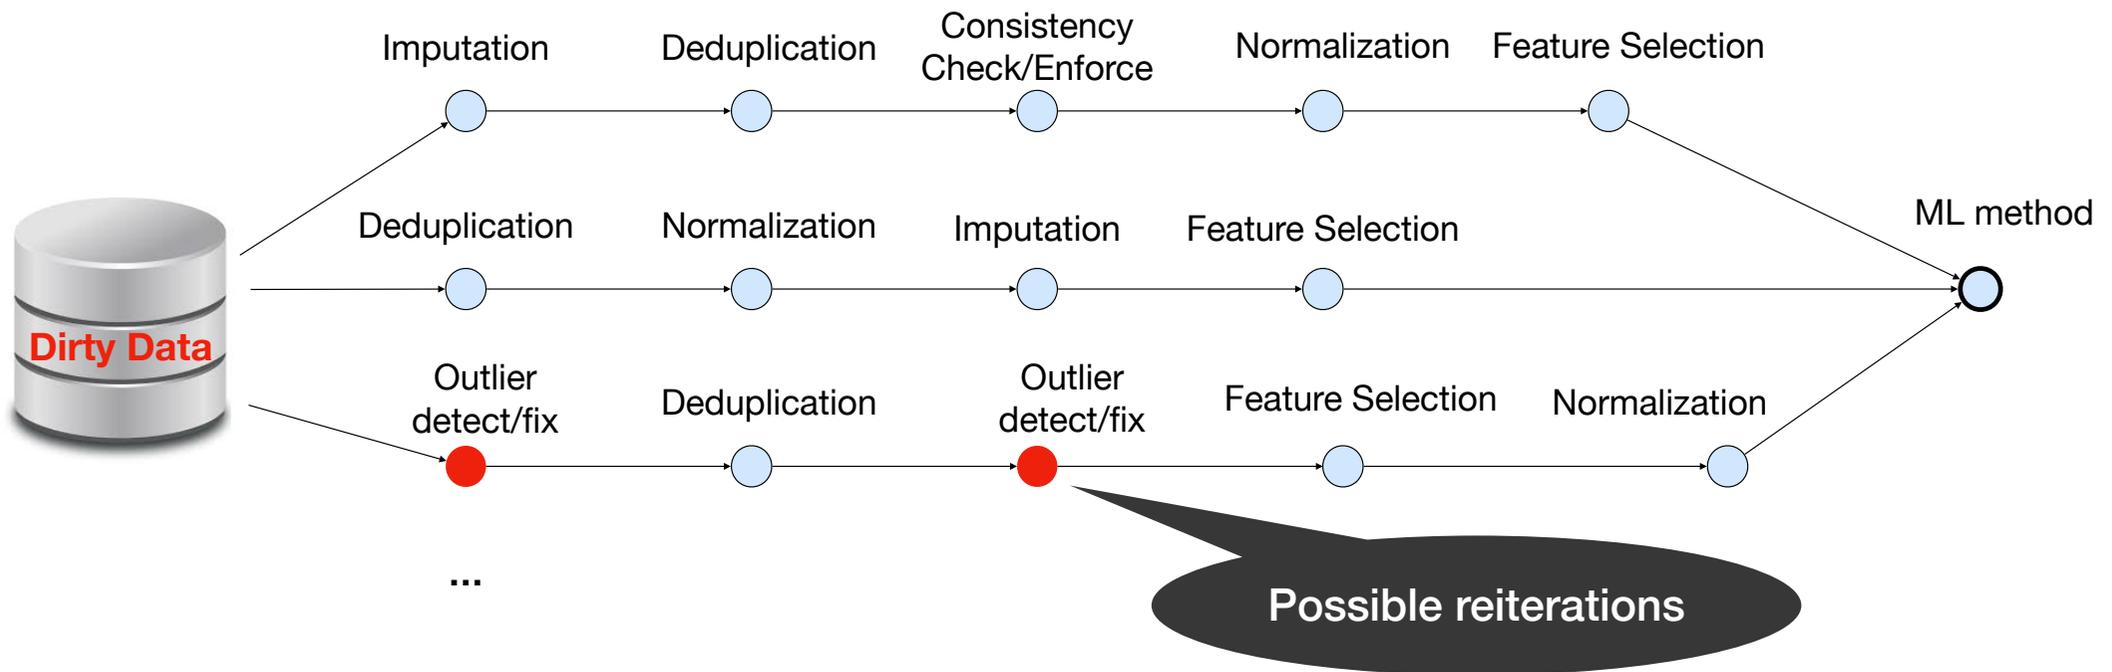
So many methods and parameter settings



Data preprocessing is challenging

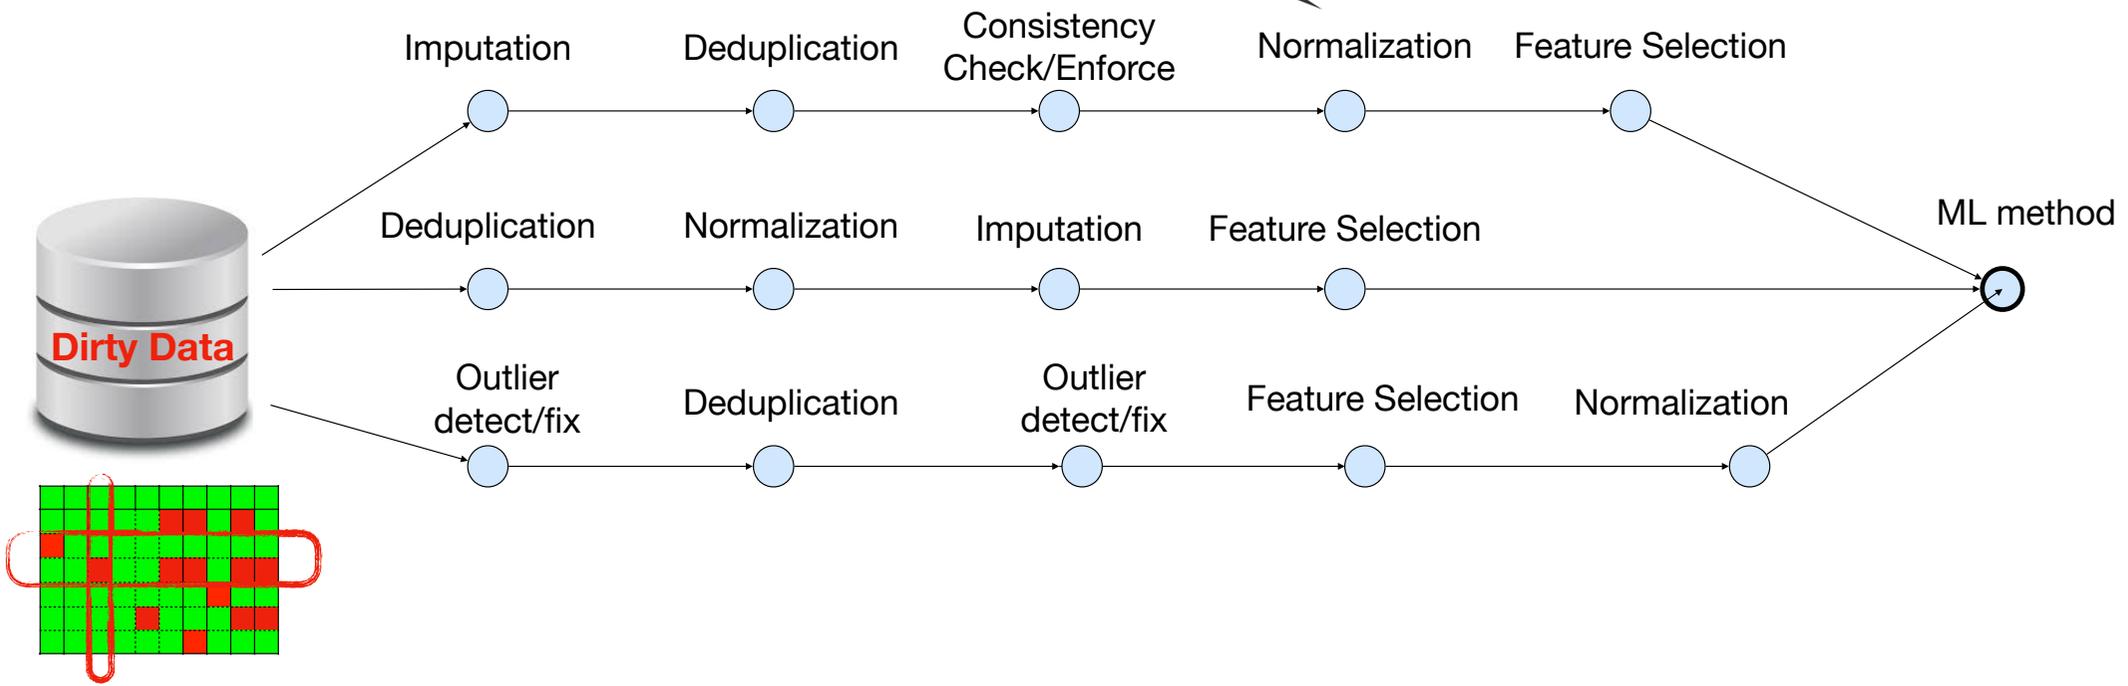


Data preprocessing is challenging



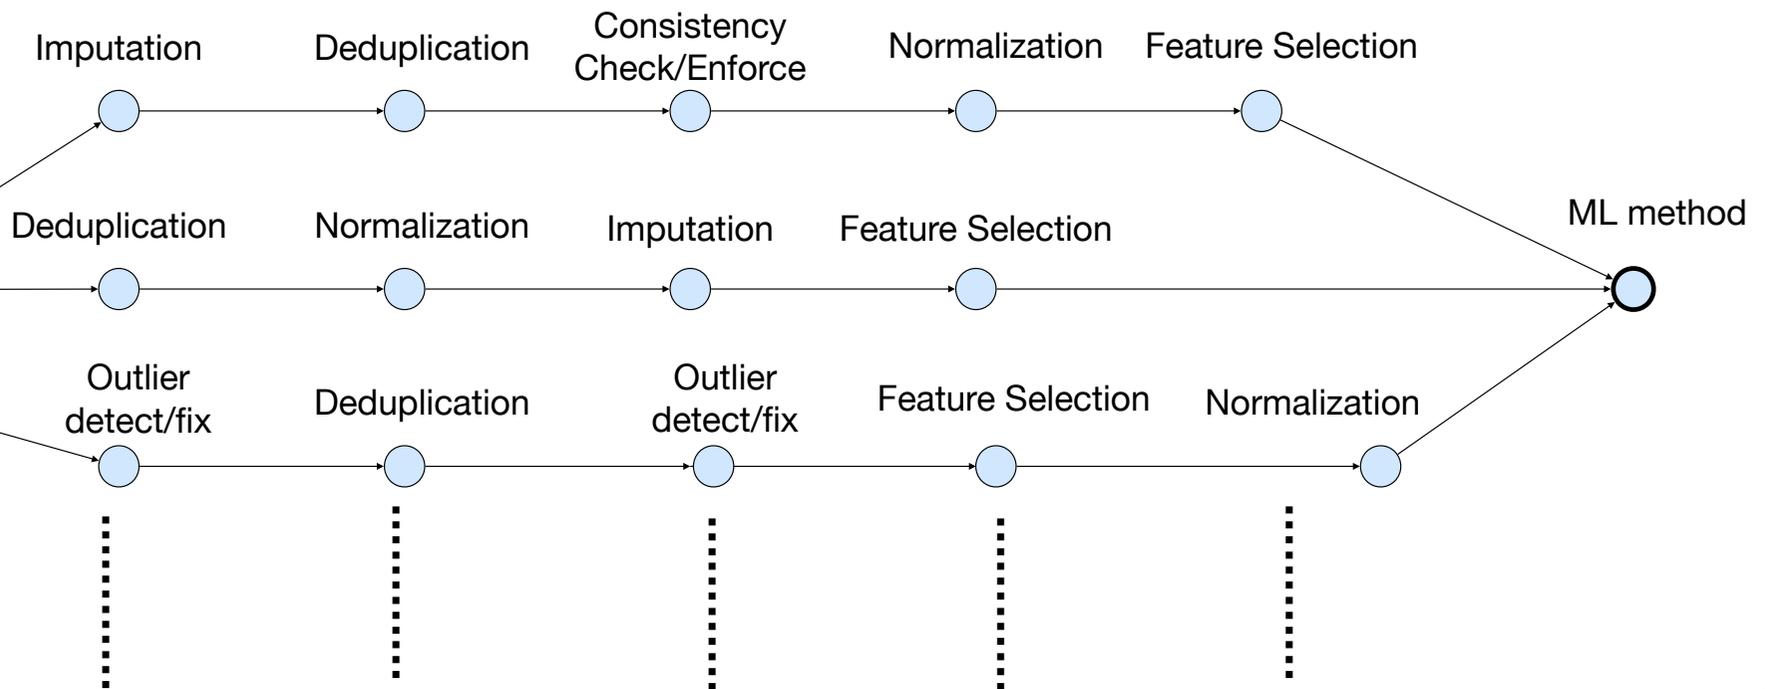
Data preprocessing is challenging

Selective processing of some parts of the dataset



Data preprocessing is challenging

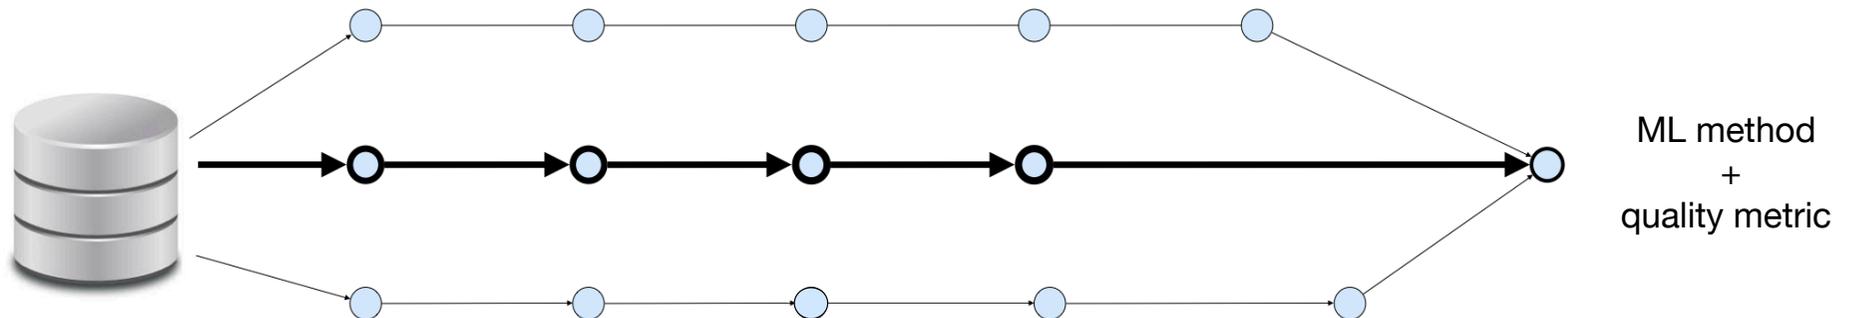
Infinite space of possible strategies



Optimization Problem



Can we help the user in composing the data preparation pipeline that maximizes the quality performance of the ML method ?



Optimization Problem



Can we help the user in composing the data preparation pipeline that maximizes the quality performance of the ML method ?

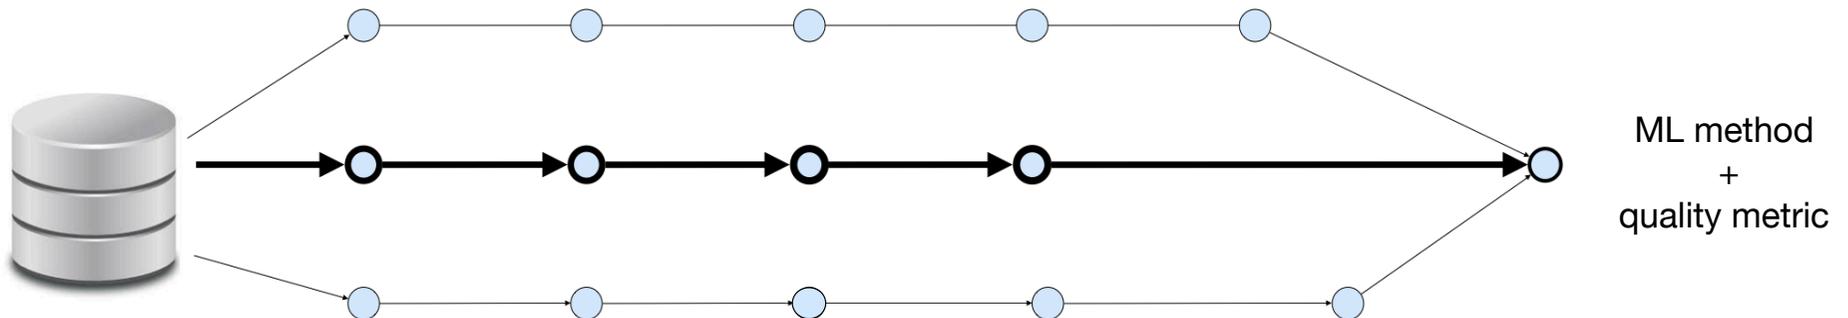
No training
example for "good"
data cleaning

AutoML-like approach

Metric-dependent

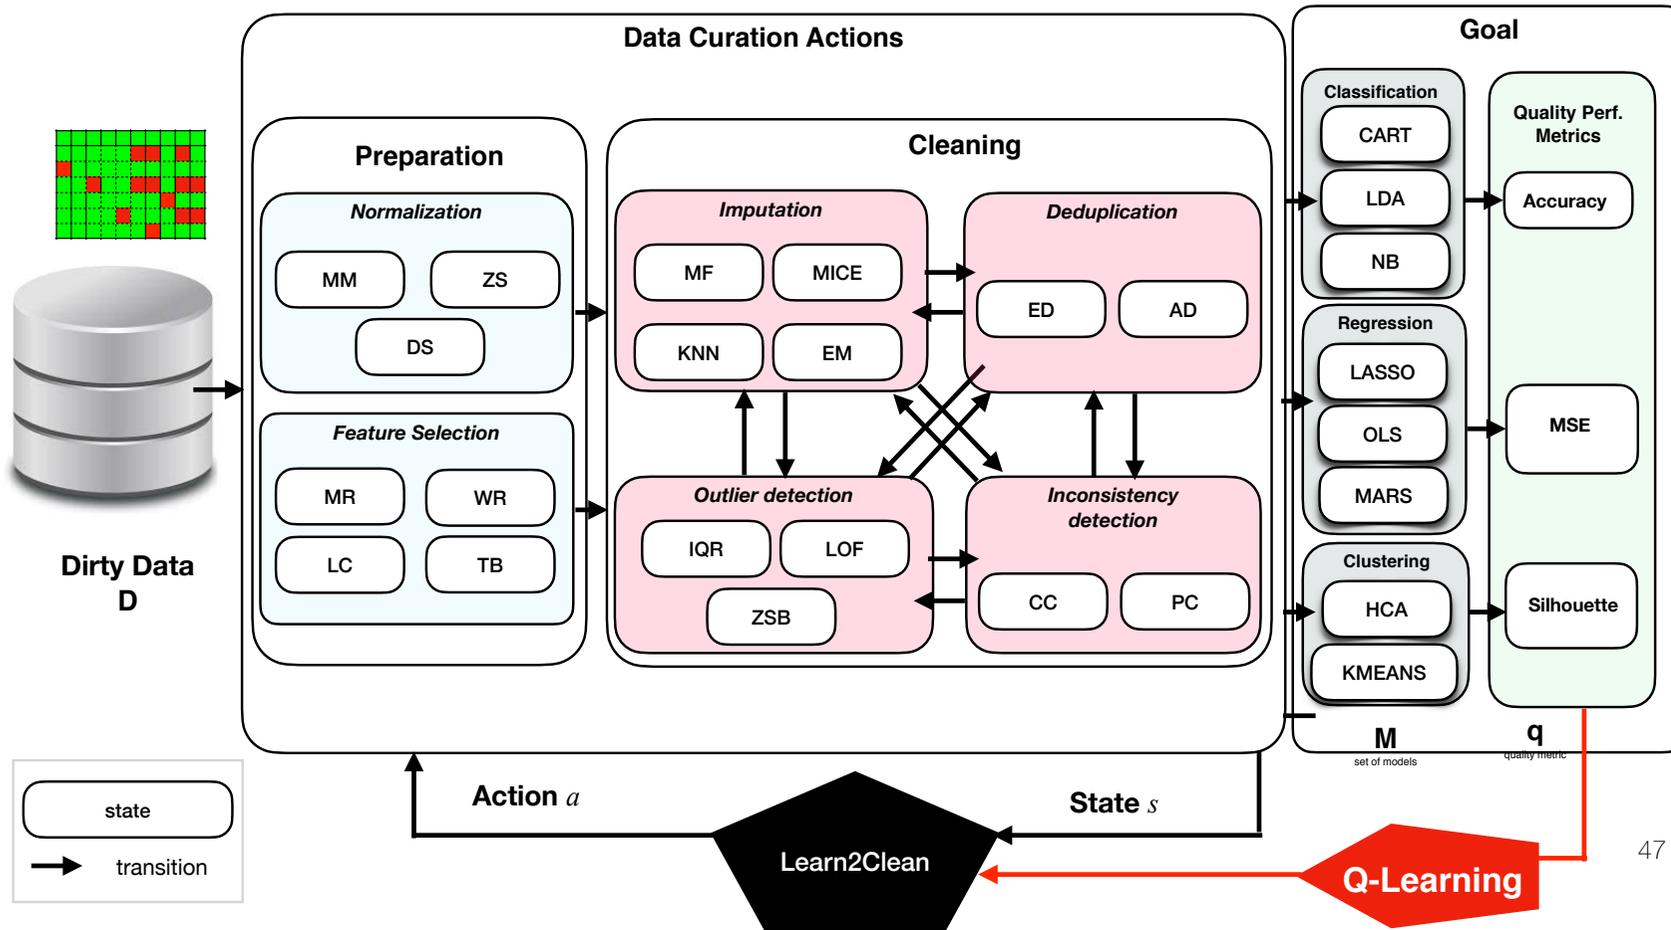
No model a priori

Human-In-The Loop



First Solution: Learn2Clean

<https://github.com/LaureBerti/Learn2Clean>



AutoML-like approach for data curation

Reinforcement Learning Framework

Markov Decision Process

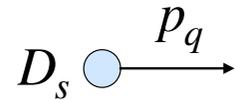
State

Action

Transition

Reward

Learn2Clean



Reinforcement Learning Framework

Markov Decision Process

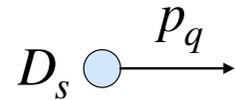
State

Action

Transition

Reward

Learn2Clean



Reinforcement Learning Framework

Markov Decision Process

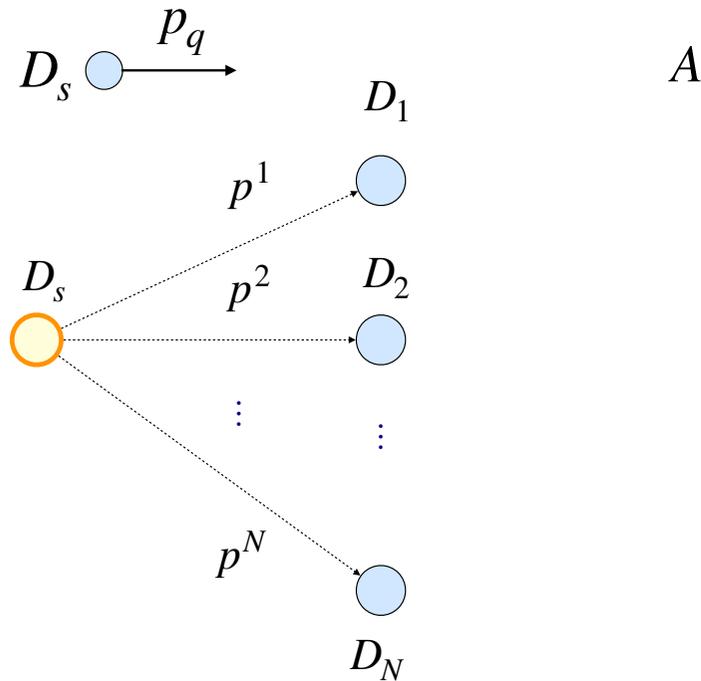
State

Action

Transition

Reward

Learn2Clean

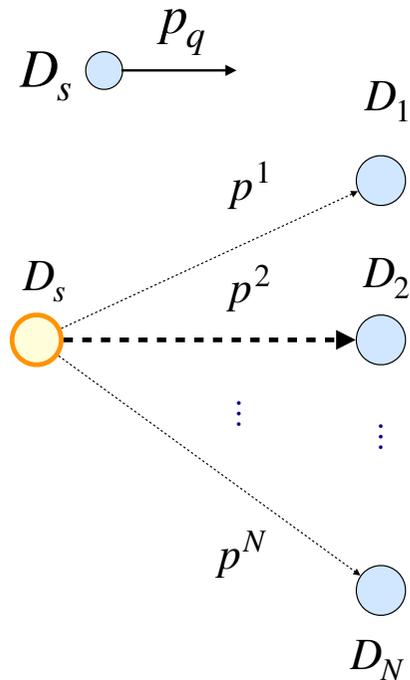


Reinforcement Learning Framework

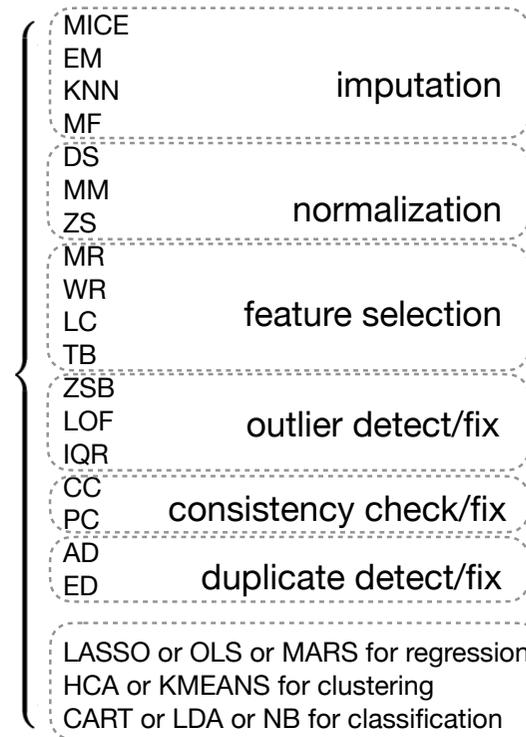
Markov Decision Process



Learn2Clean



A

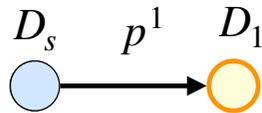
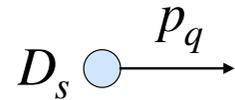


Reinforcement Learning Framework

Markov Decision Process



Learn2Clean



Reinforcement Learning Framework

Markov Decision Process

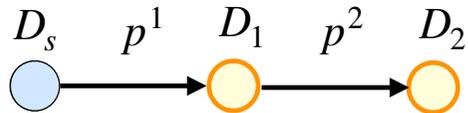
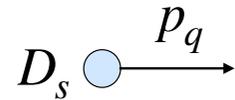
State

Action

Transition

Reward

Learn2Clean

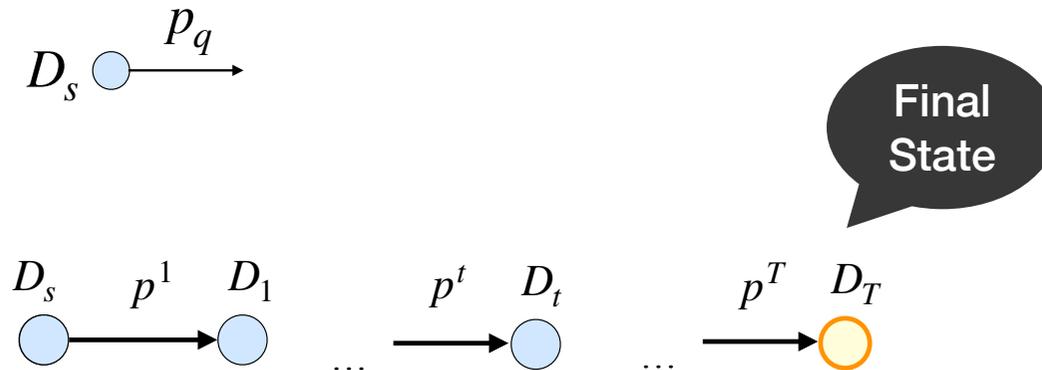


Reinforcement Learning Framework

Markov Decision Process



Learn2Clean



LASSO or OLS or MARS for regression
HCA or KMEANS for clustering
CART or LDA or NB for classification

Reinforcement Learning Framework

Markov Decision Process

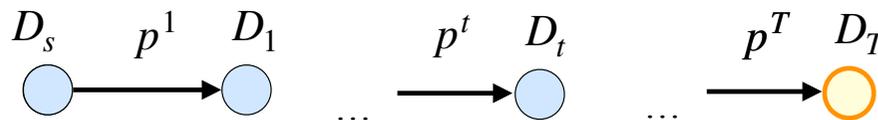
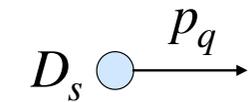
State

Action

Transition

Reward

Learn2Clean



deterministic

R =

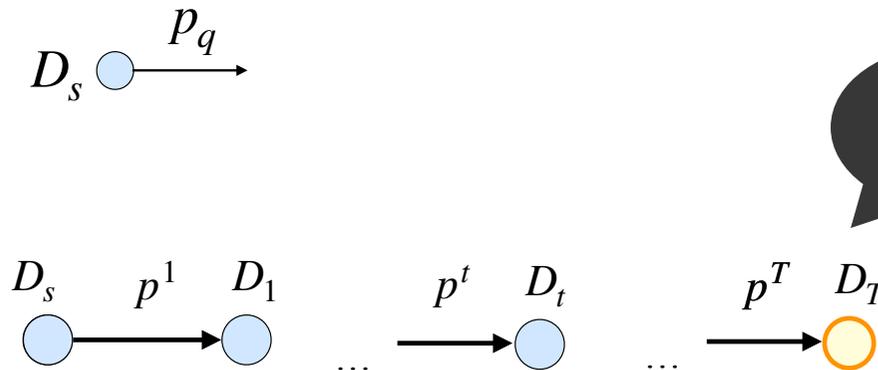
	MICE	EM	KNN	MF	DS	MM	ZS	MR	WR	LC	TB	ZSB	LOF	IQR	CC	PC	AD	ED	LASSO
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	0	0	0	0	-1	-1	0	0	-1
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	-1	-1	-1	0	0	0	0	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	-1
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100
0	0	0	0	-1	-1	-1	0	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	100
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-1	-1	100
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-1	100
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

Reinforcement Learning Framework

Markov Decision Process



Learn2Clean



deterministic

R =

	MICE	EM	KNN	MF	DS	MM	ZS	MR	WR	LC	TB	ZSB	LOF	IQR	CC	PC	AD	ED	LASSO
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	-1	-1	0	0	100	
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	-1	-1	0	0	100	
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	-1	-1	0	0	100	
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	0	0	0	-1	-1	0	0	-1	
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	0	0	0	-1	-1	0	0	-1	
0	0	0	0	-1	-1	-1	0	0	0	0	0	0	0	-1	-1	0	0	-1	
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	-1	-1	0	0	-1	
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	-1	-1	0	0	-1	
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	-1	-1	0	0	-1	
0	0	0	0	-1	-1	-1	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100	
-1	-1	-1	-1	-1	-1	-1	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100	
0	0	0	0	-1	-1	-1	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100	
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	100	
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	100	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-1	-1	100	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-1	-1	100	
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	

- > MSE
- > Silhouette
- > Accuracy

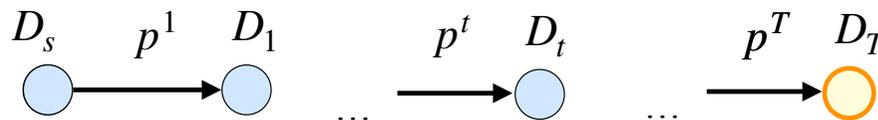
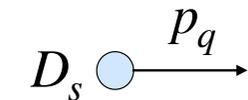
Quality metric

Reinforcement Learning Framework

Markov Decision Process



Learn2Clean



deterministic

R =

	MICE	EM	KNN	MF	DS	MM	ZS	MR	WR	LC	TB	ZSB	LOF	IQR	CC	PC	AD	ED	LASSO
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	0	0	0	-1	0	0	0	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	0	0	0	0	-1	-1	0	0	-1
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	-1	-1	-1	0	0	0	0	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	0	0	0	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	-1
0	0	0	0	-1	-1	-1	0	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100
-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100
0	0	0	0	-1	-1	-1	0	0	0	0	-1	-1	-1	-1	-1	-1	0	0	100
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	100
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	0	0	0	-1	-1	0	0	100
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-1	100
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-1	100
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

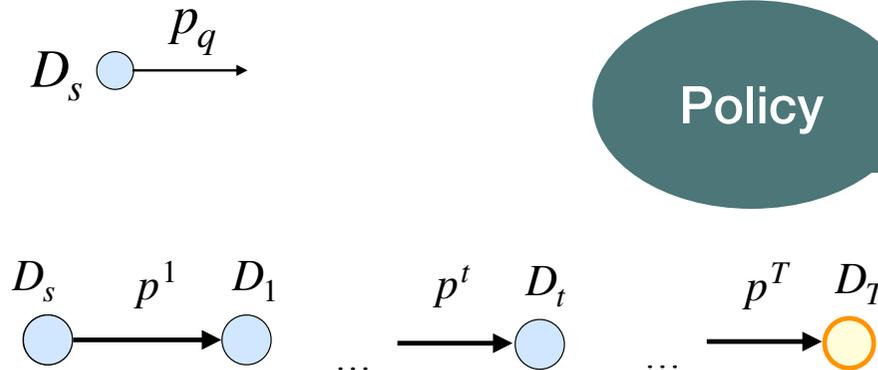
- > MSE
- > Silhouette
- > Accuracy

Quality metric

Learn2Clean selects the sequence of preprocessing actions that maximizes the quality metric (or minimizes the error)

Reinforcement Learning Framework

Markov Decision Process



Softmax action selection

$$\pi = P(a | s) = \frac{e^{Q(s,a)/k}}{\sum_j e^{Q(s,a_j)/k}}$$

Q-table

value iteration update

$$Q^\pi(s, a) \leftarrow (1 - \alpha) \cdot Q(s, a) + \alpha \cdot \left(R(s, a) + \gamma \cdot \max_{a'} Q(s', a') \right)$$

↑ new value ↗ learning rate ↑ old value ↑ reward ↑ discount factor ↘ optimal future value

learned value

Experiments

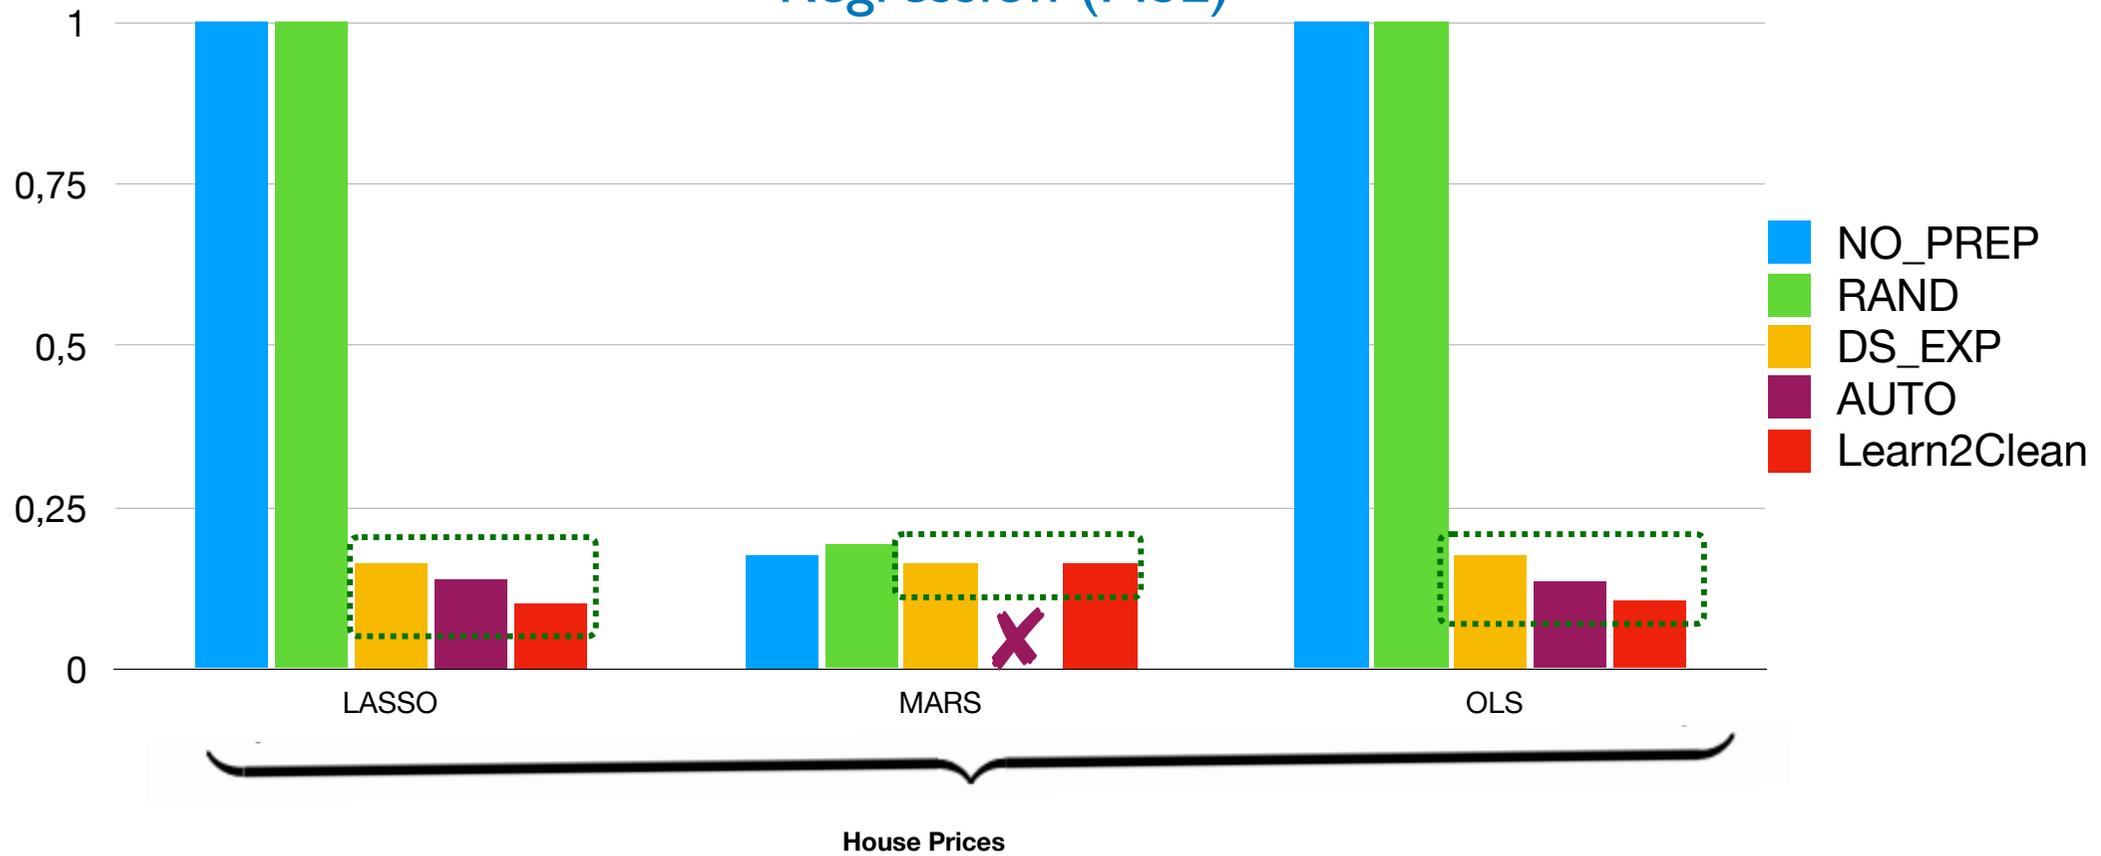
Datasets

Name	# Att.	# Rows	Clustering	Regression	Classification
House Prices	81	1.46k	✓	✓	✓
Google Playstore Users	5	64.3k	✓		
Google Playstore Apps	13	10.8k	✓		✓

Evaluation : Silhouette for Clustering
MSE for Regression
Accuracy for Classification

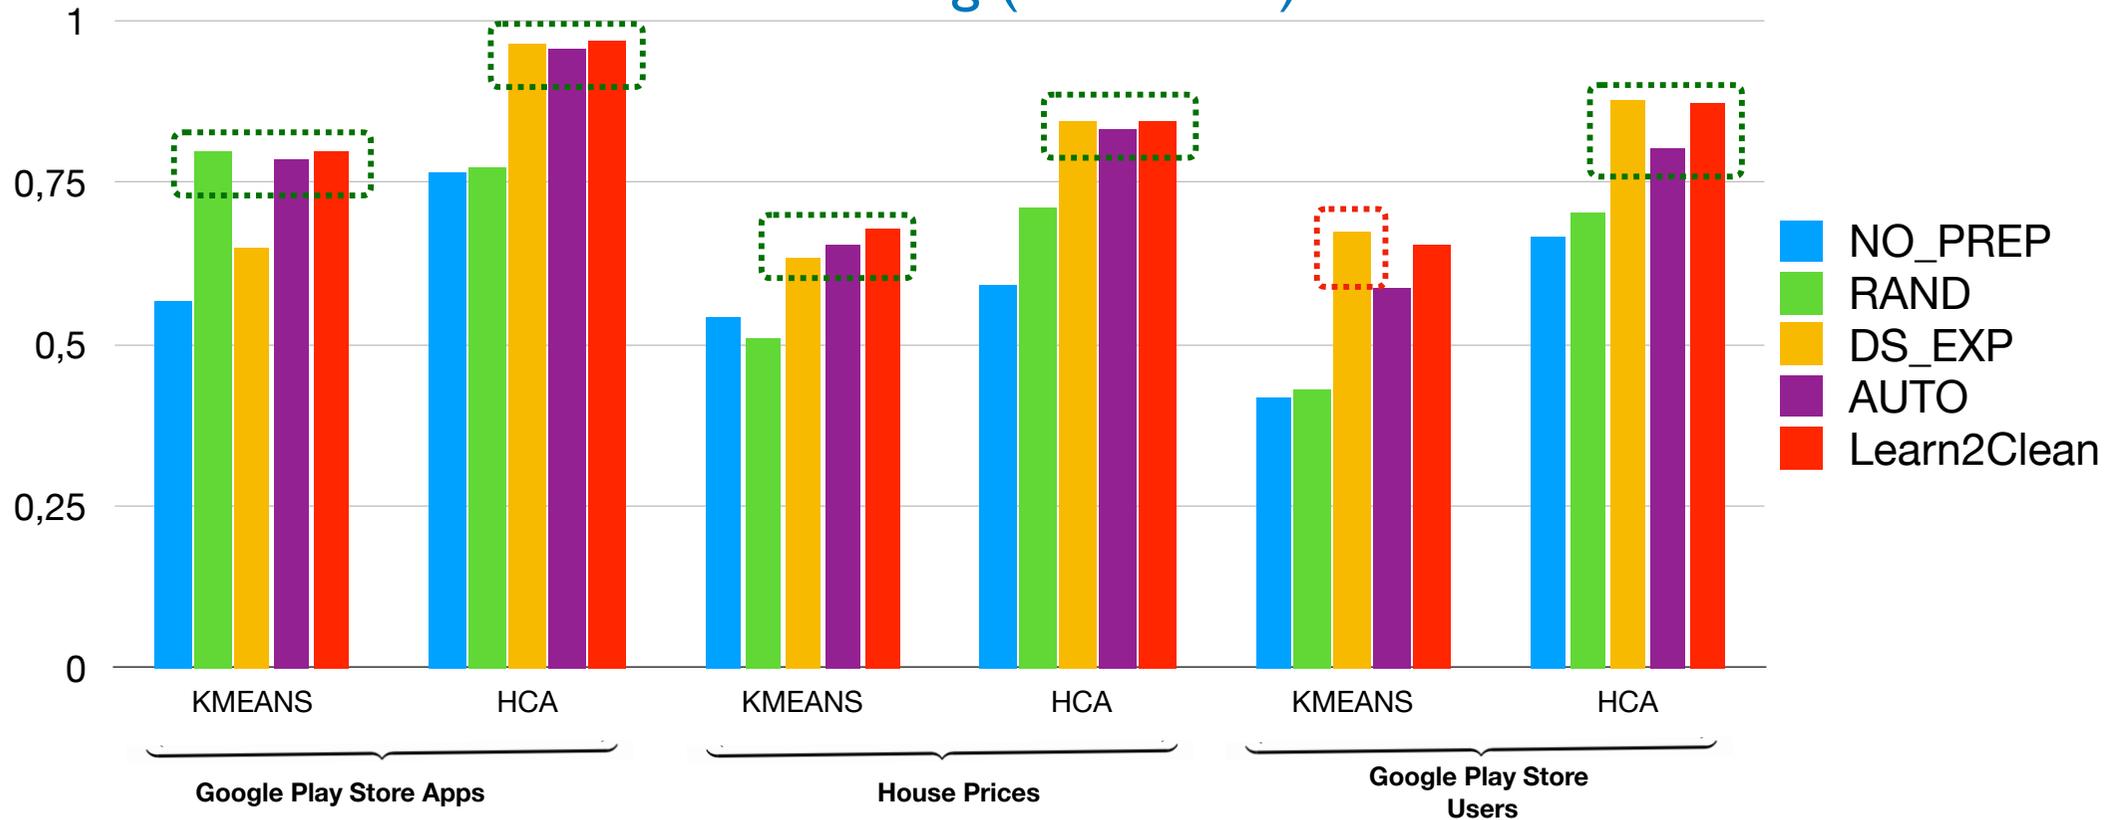
Experimental Results

Regression (MSE)



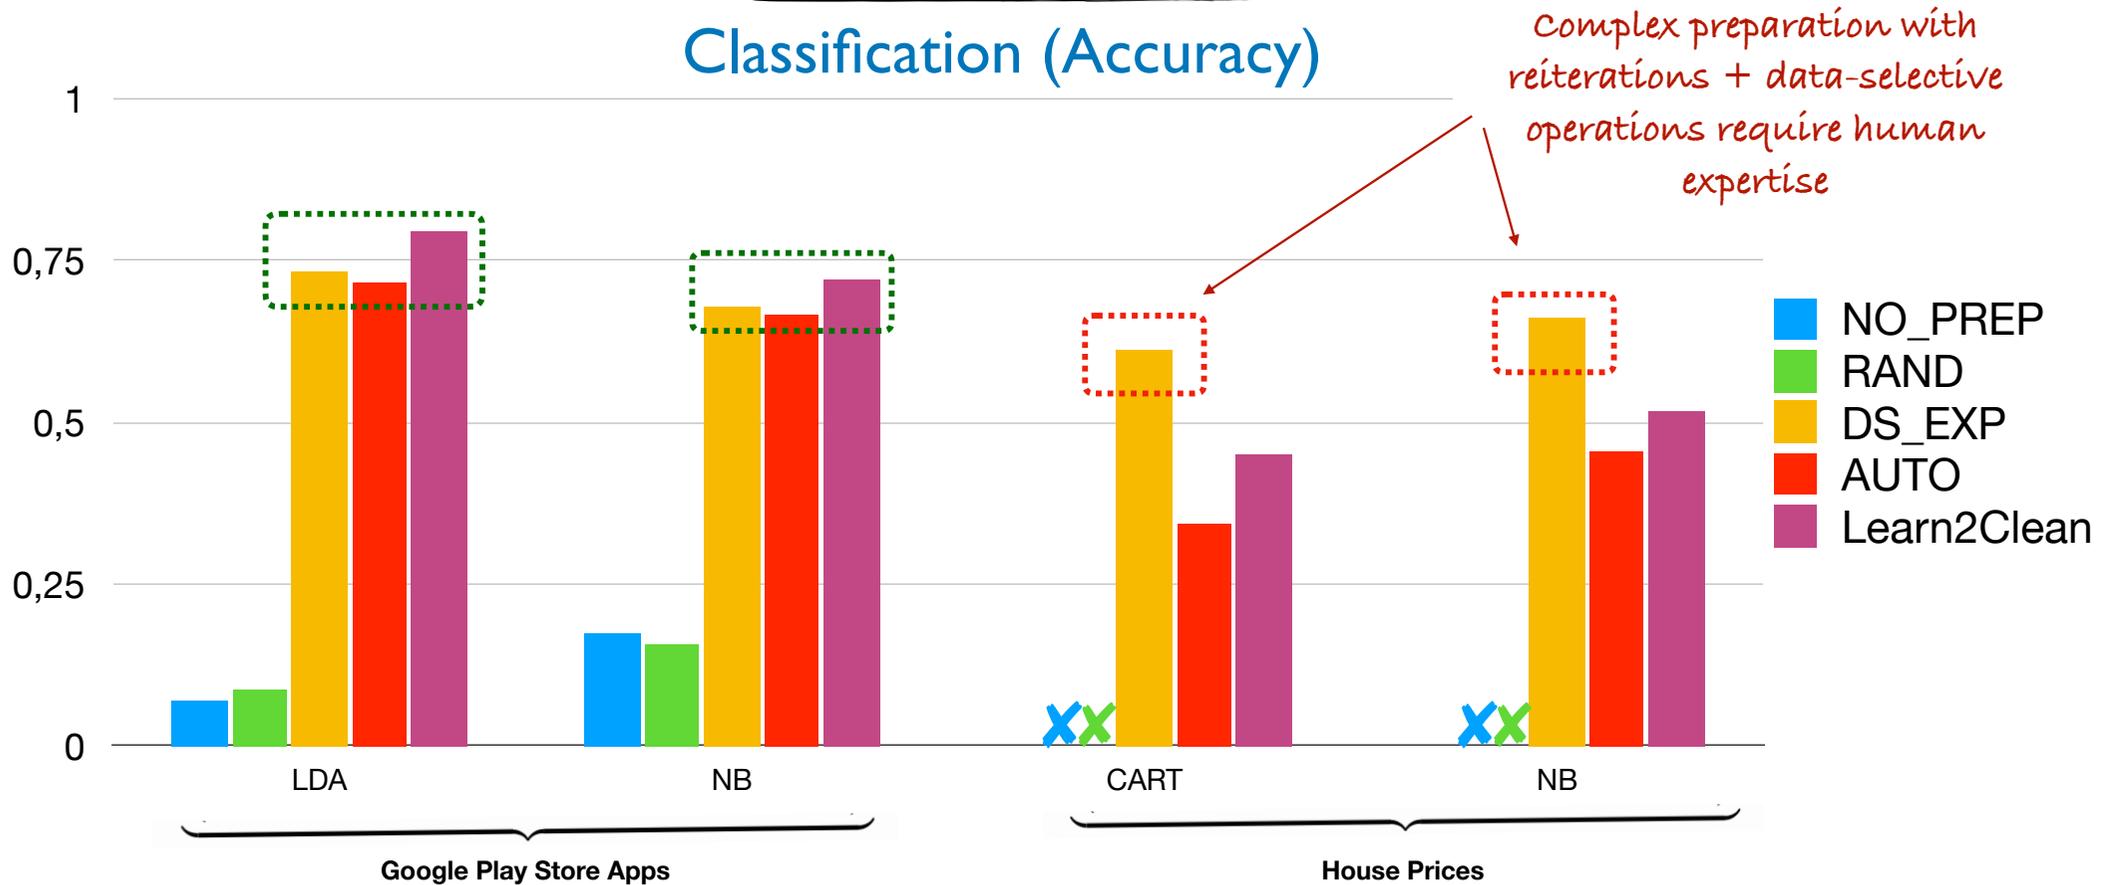
Experimental Results

Clustering (Silhouette)

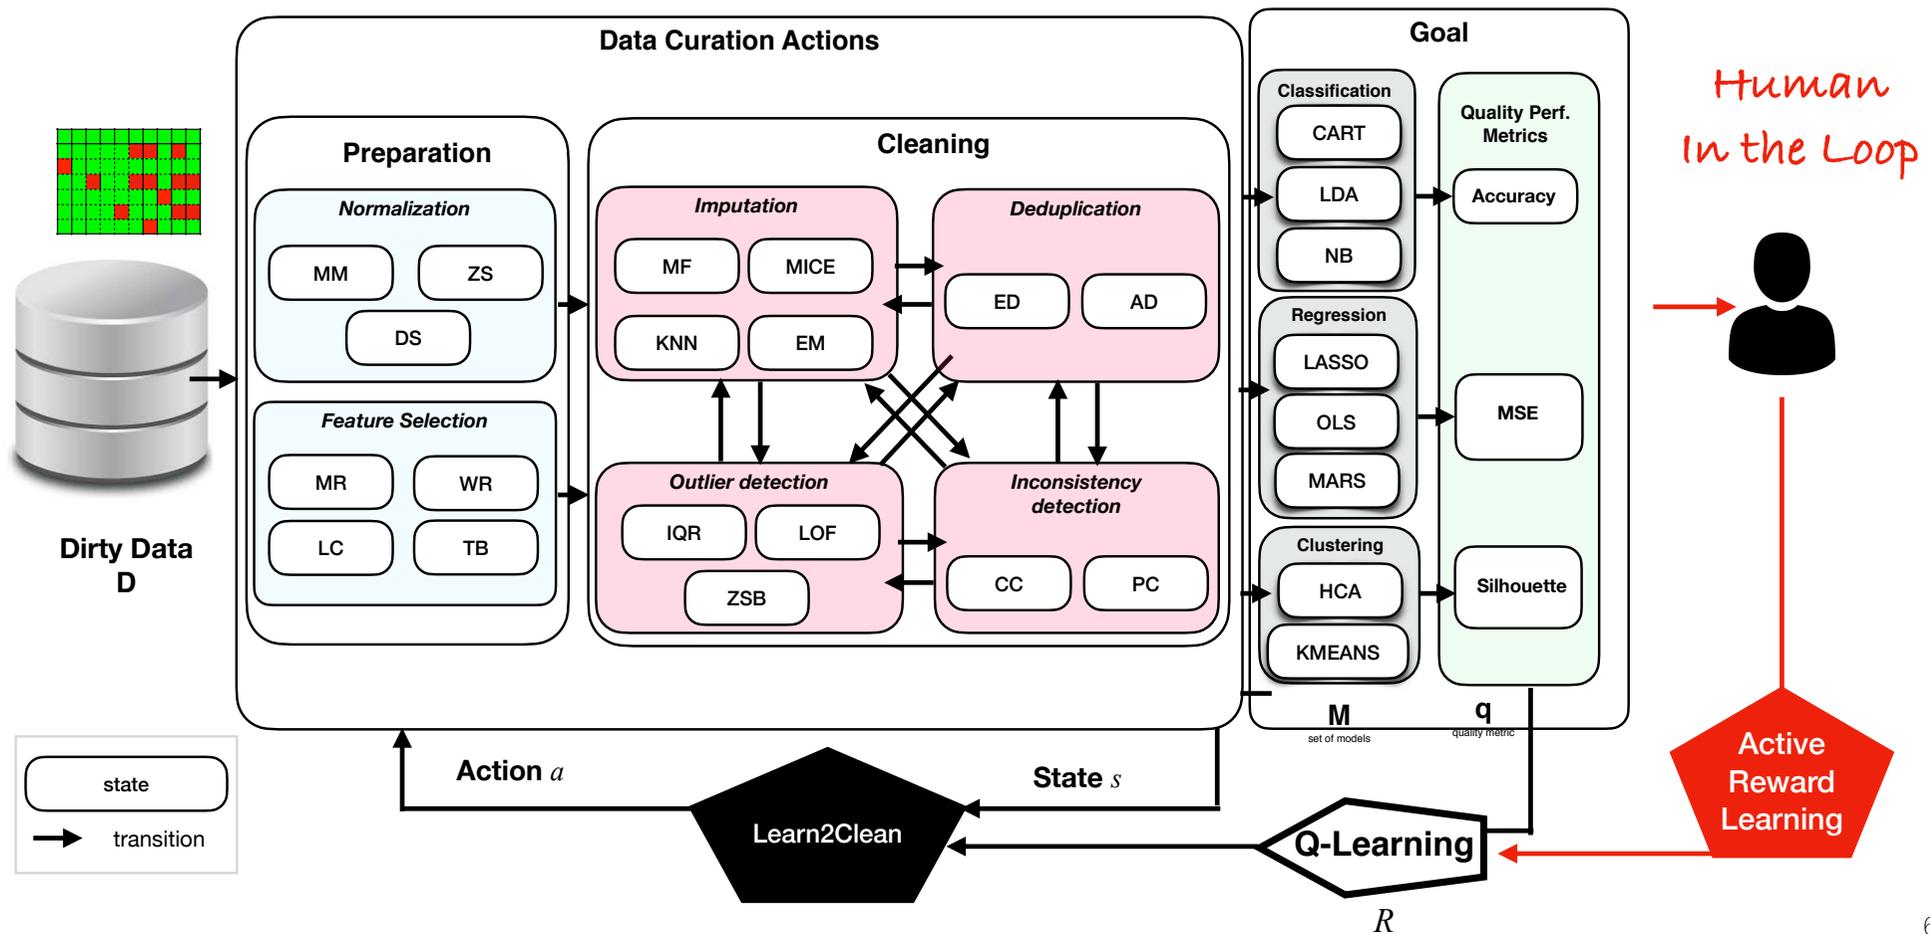


Experimental Results

Classification (Accuracy)

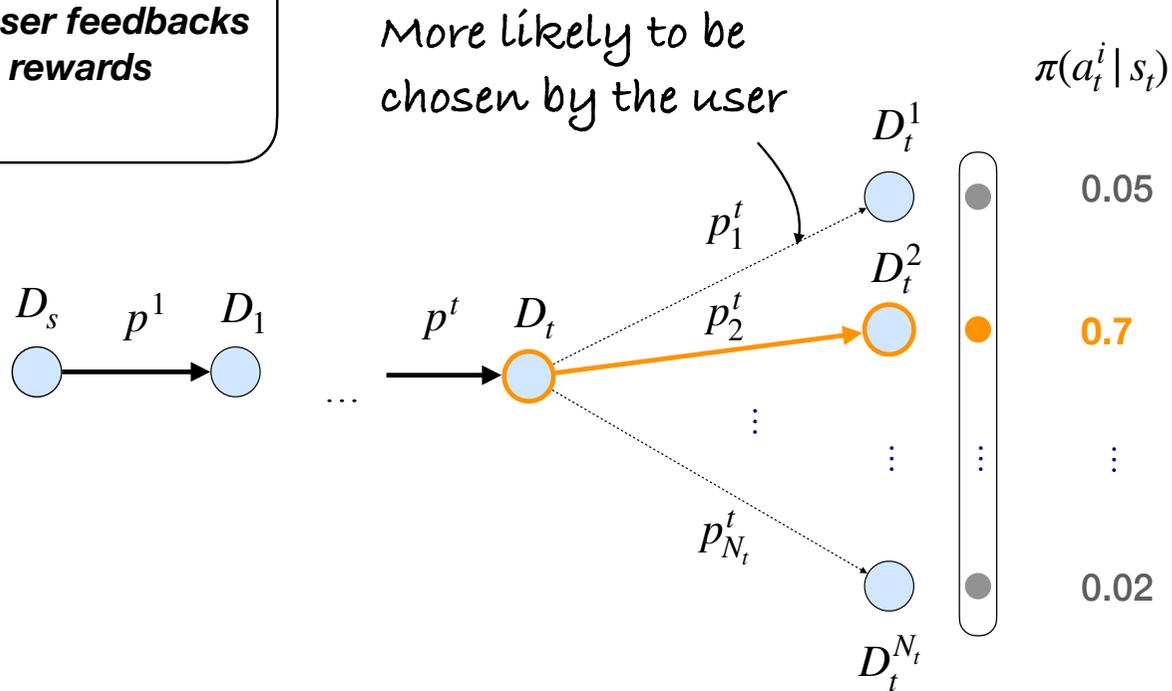


HIL with Active Reward Learning



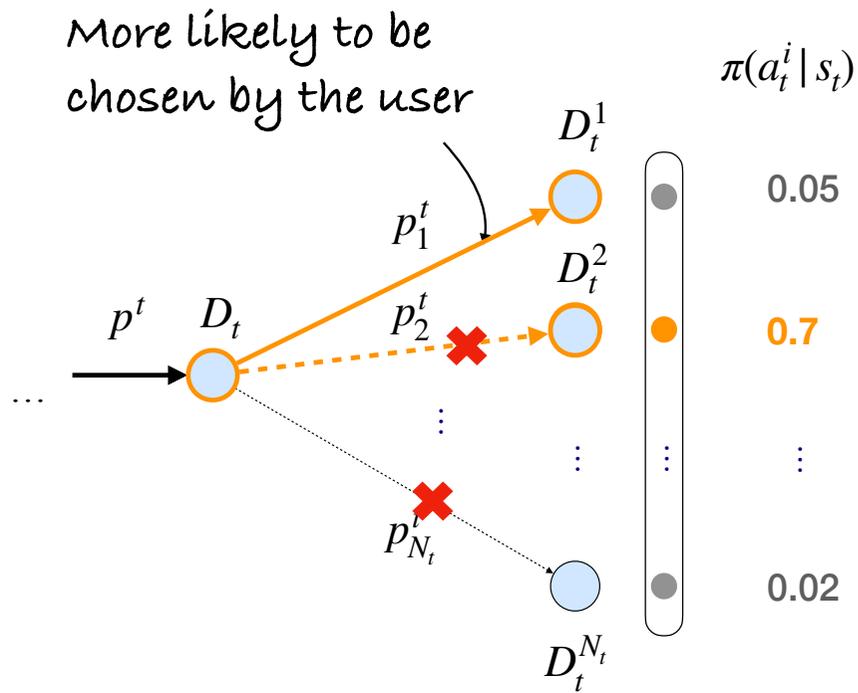
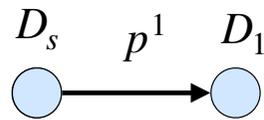
Active Reward Learning

Goal: *learn from user feedbacks to adapt the rewards*



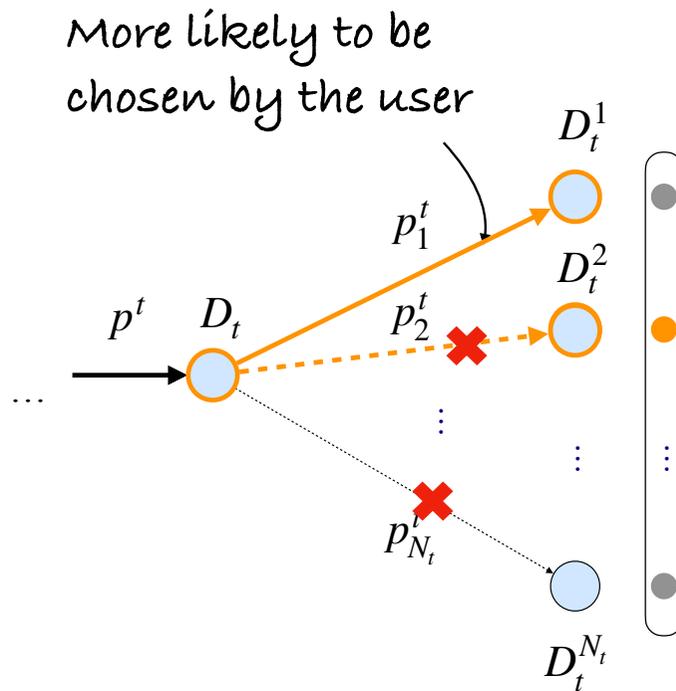
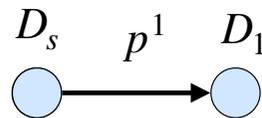
Active Reward Learning

Goal: *learn from user feedbacks to adapt the rewards*



Active Reward Learning

Goal: *learn from user feedbacks to adapt the rewards*

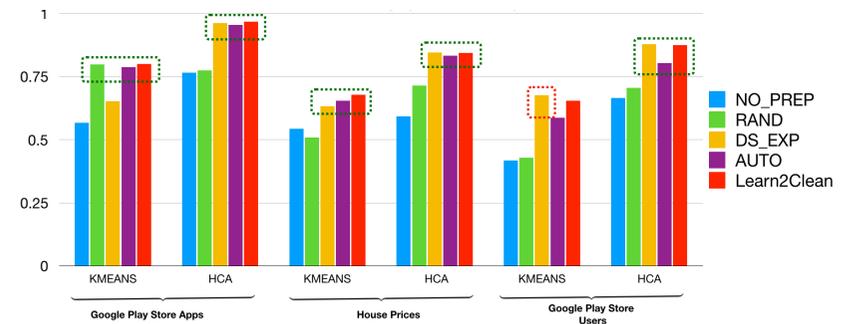
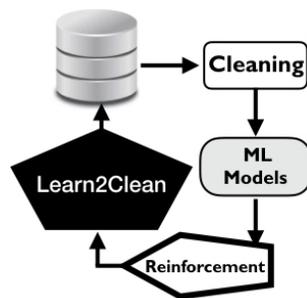


$\tilde{\pi}(a_t^i s_t)$	$\pi(a_t^i s_t)$
0.9	0.05
0	0.7
⋮	⋮
0	0.02

A green arrow points from the text "Force exploration" to the value 0.7 in the second row of the table.

Ongoing work

- New version of Learn2Clean with deep RL agents
- Combine AutoML, AutoCuration, and HIL
- Learn better reward functions
- Extend the library of ML and data preparation methods
- Extend experiments with more intricate data glitches and various glitch distributions



Code: <https://github.com/LaureBerti/Learn2Clean>

Concluding Remarks

- ML crucially needs principled data curation and preparation, adequate tooling, and user assistance
- The impact of data preprocessing variability is largely underestimated in ML
- Many data preprocessing tasks require seamless integration of Human-in-the-Loop and automated ML-based solutions
- Perfect timing for many R&D opportunities:
 - Manage and orchestrate human/machine resources
 - Challenge and transfer research ideas to operational and very large-scale contexts



Thank you!