DEXA 2022

The 33rd DEXA Conferences and Workshops
22-24 August 2022
Vienna, Austria

Mining Fluctuation Propagation Graph among Time Series with Active Learning

Mingjie Li, Minghua Ma, Xiaohui Nie, Kanglin Yin, Li Cao, Xidao Wen, Zhiyun Yuan, Duogang Wu, Guoying Li, Wei Liu, Xin Yang, Dan Pei

清華大學 Tsinghua University

必示 BIZSEER

中国建设银行 China Construction Bank

Background

Background

Online Service Systems

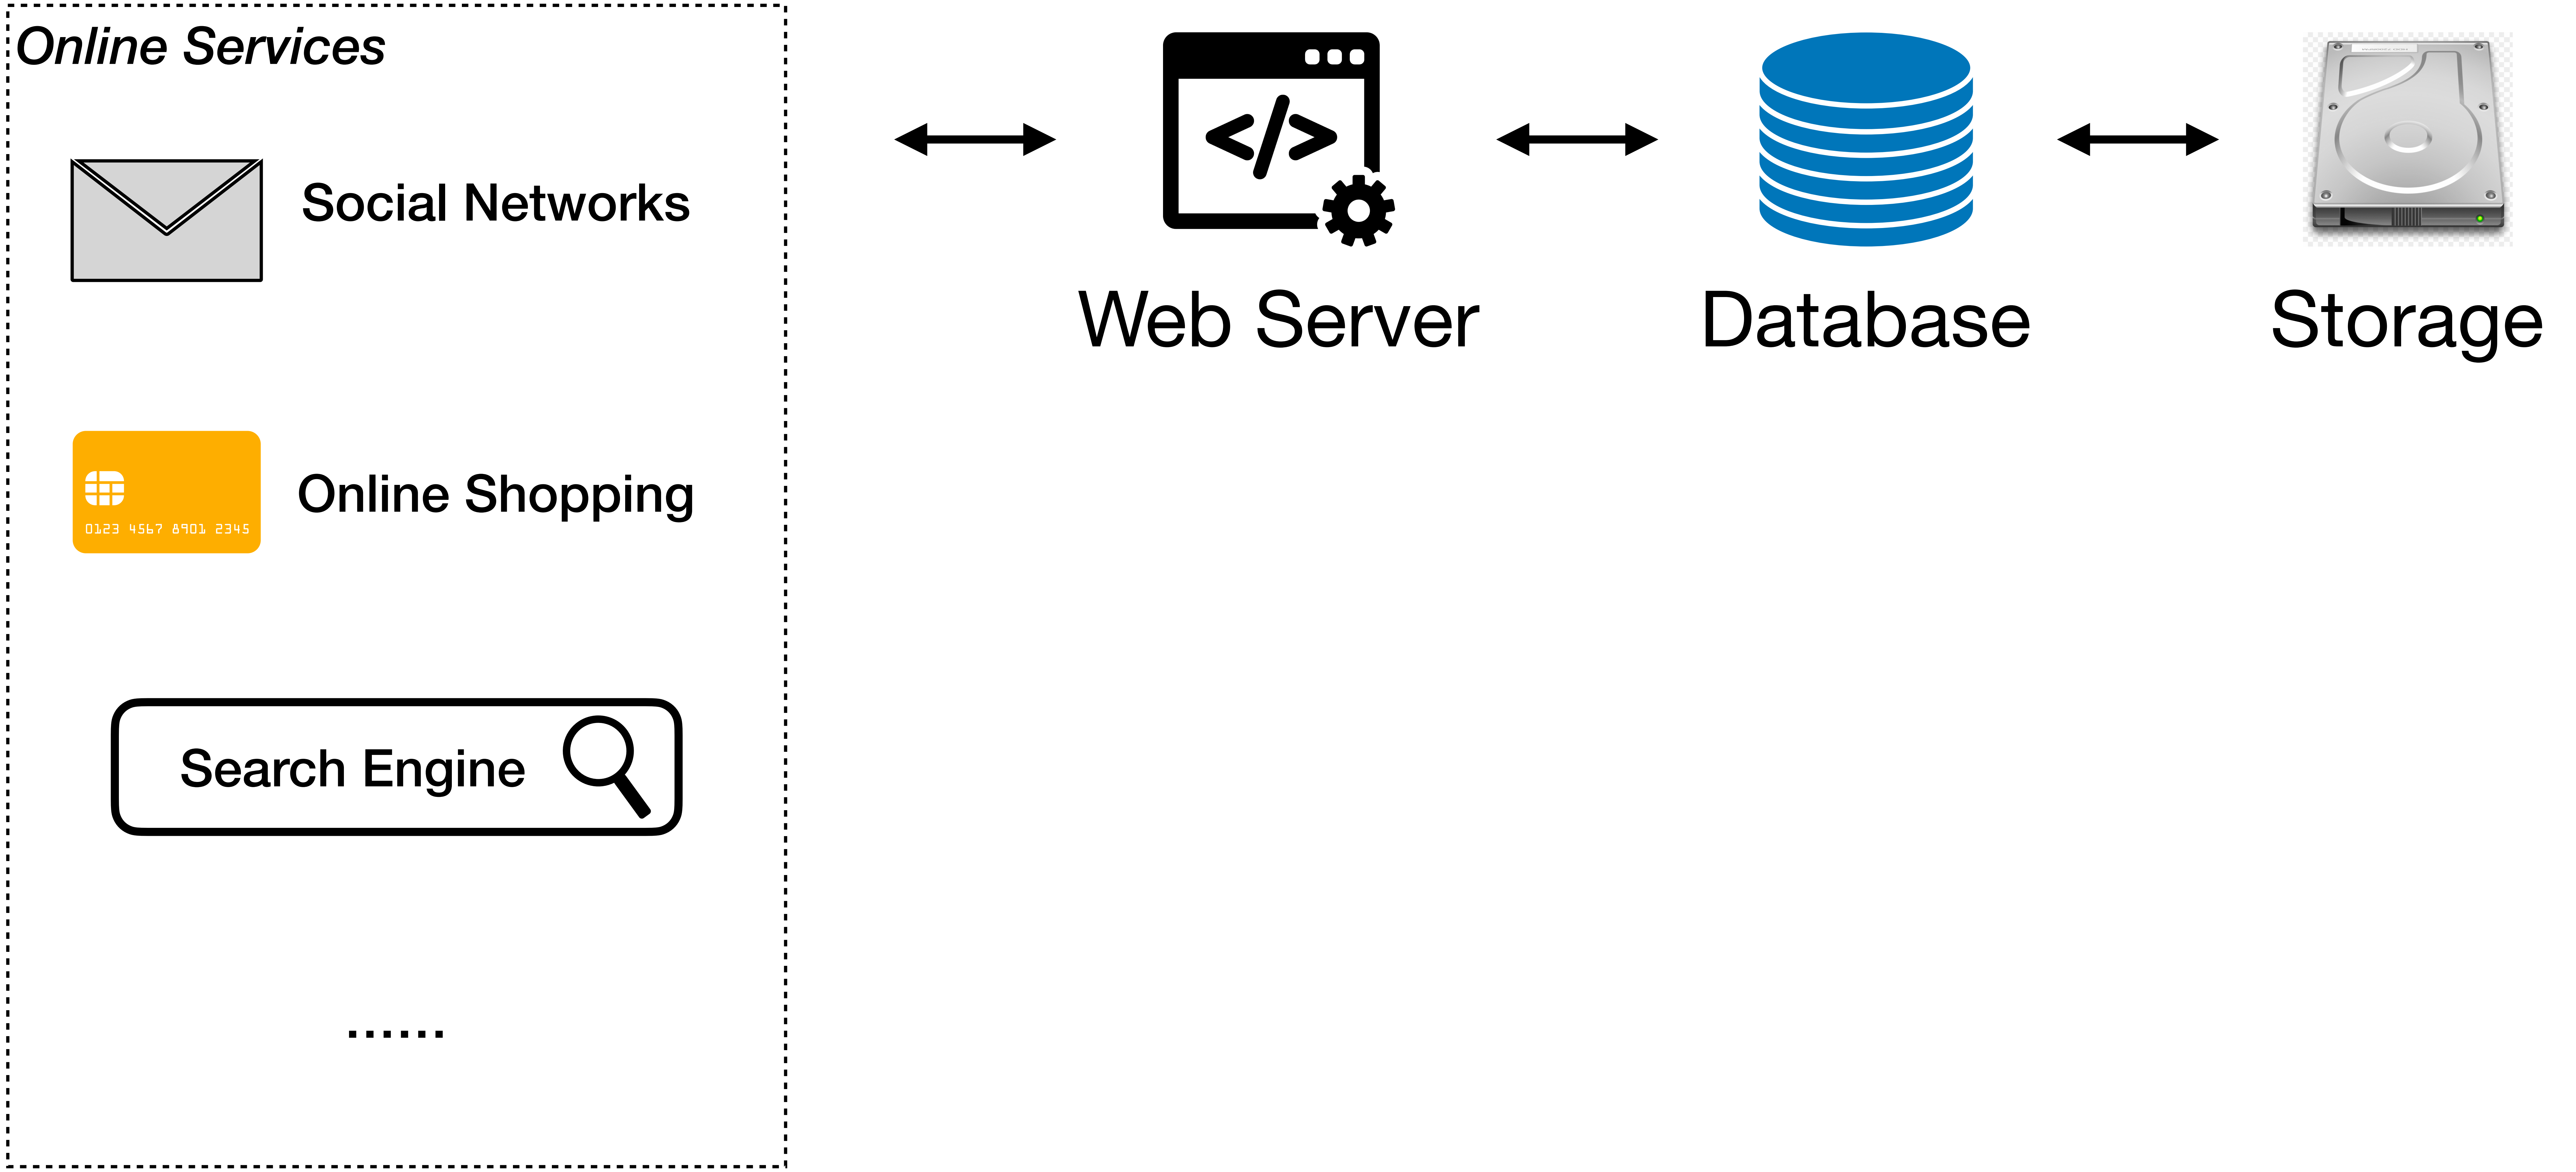

Background

Online Service Systems

Online Services

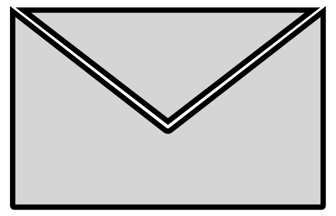
Social Networks

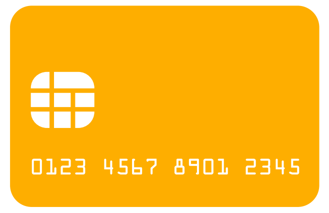

Online Shopping

Search Engine 

......

↔ 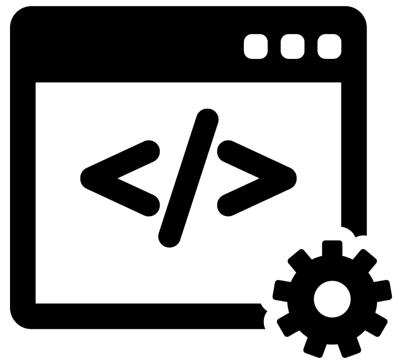 ↔  ↔ 

Web Server	Database	Storage
queries per second	average active session	disk utilization
response time	log file sync	IO per second
success rate	logic read per second	IO wait time
......	......	......

Background

Online Service Systems

Online Services

Social Networks

Online Shopping

Search Engine

......

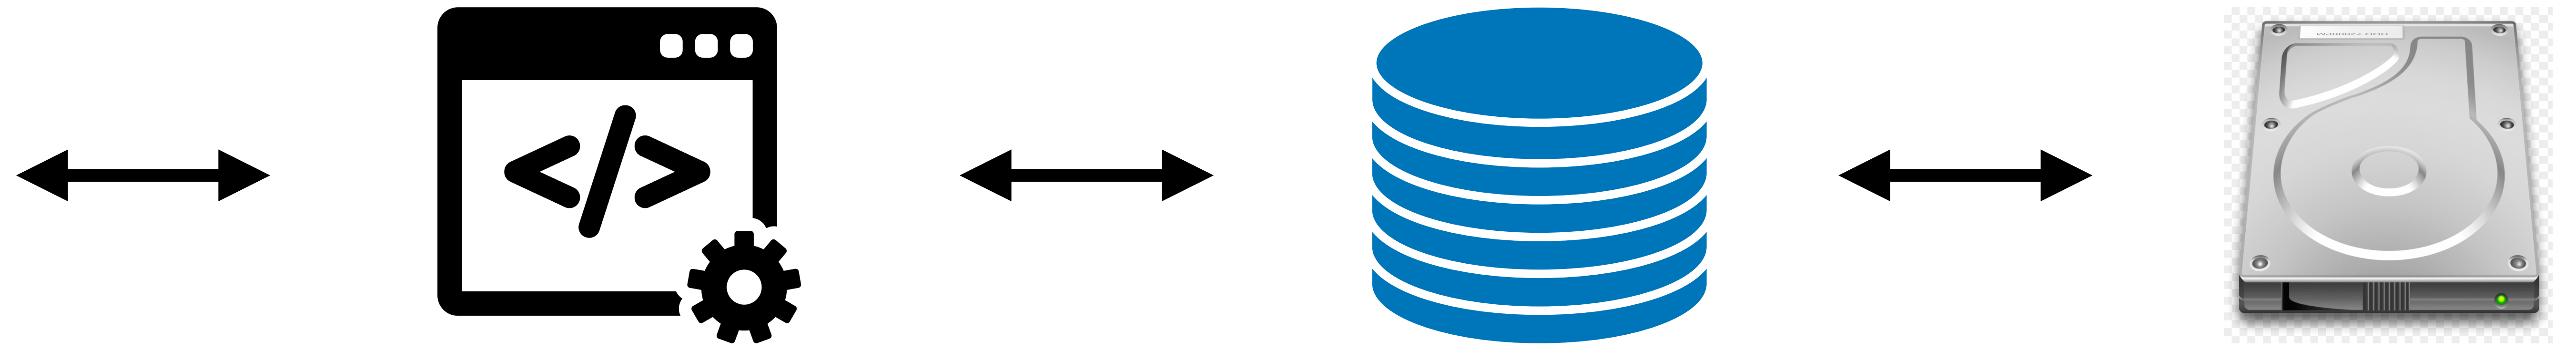

Web Server	Database	Storage
queries per second	average active session	disk utilization
response time	log file sync	IO per second
success rate	logic read per second	IO wait time
......	......	......

Metrics **are the most widely available data, often in the form of a time series**

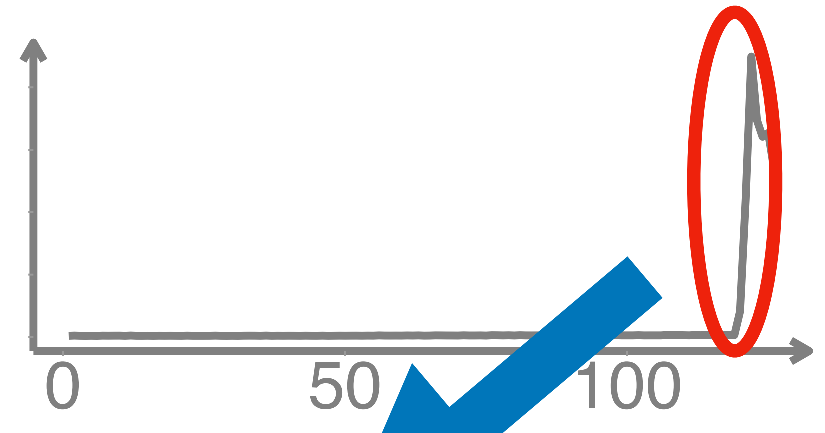

Time	09:23	09:24	09:25	09:26	09:27
Value	302	4095	22142	44936	34745

Background

Troubleshooting

Online Services

Social Networks

Online Shopping

Search Engine

……

Web Server

queries per second
response time
success rate
……

Database

average active session
log file sync
logic read per second
……

Storage

disk utilization
IO per second
IO wait time
……

Background

Troubleshooting

Online Services

Social Networks

Online Shopping

Search Engine

......

↔ Web Server ↔ Database ↔ Storage

Web Server

queries per second
response time
success rate
......

Database

average active session
log file sync
logic read per second
......

Storage

disk utilization
IO per second
IO wait time
......

A fault may propagate in the system

Background

Troubleshooting

- Recent automatic troubleshooting works model the propagation with a **graph**

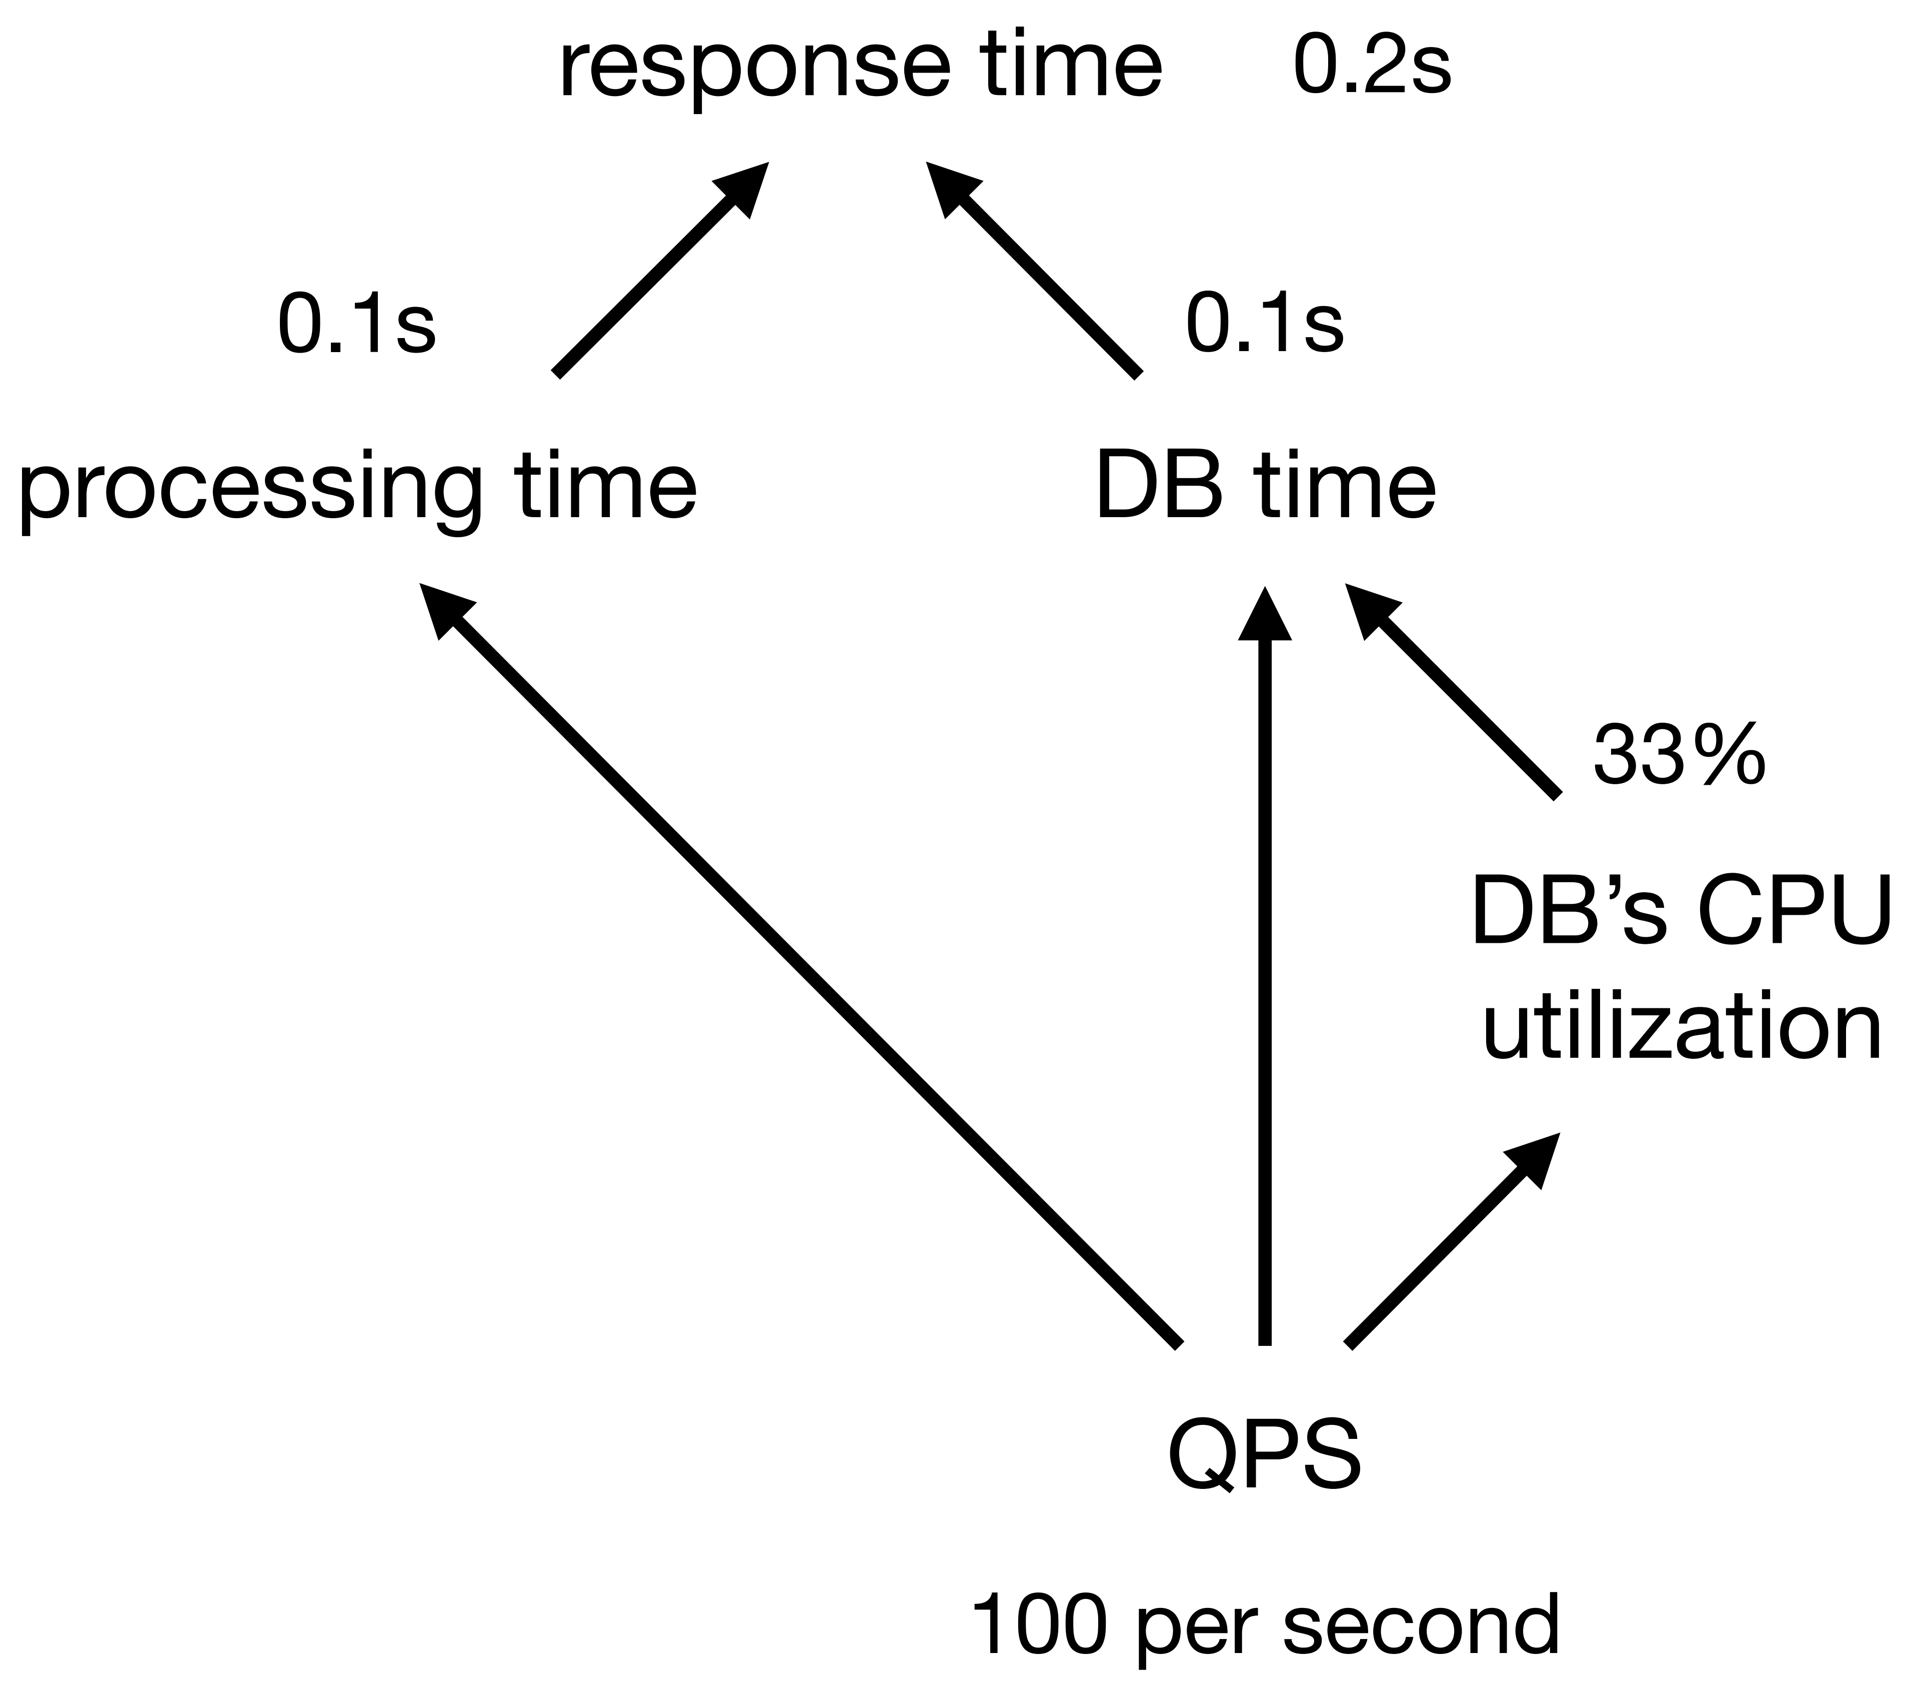

Background

Troubleshooting

- Recent automatic troubleshooting works model the propagation with a **graph**

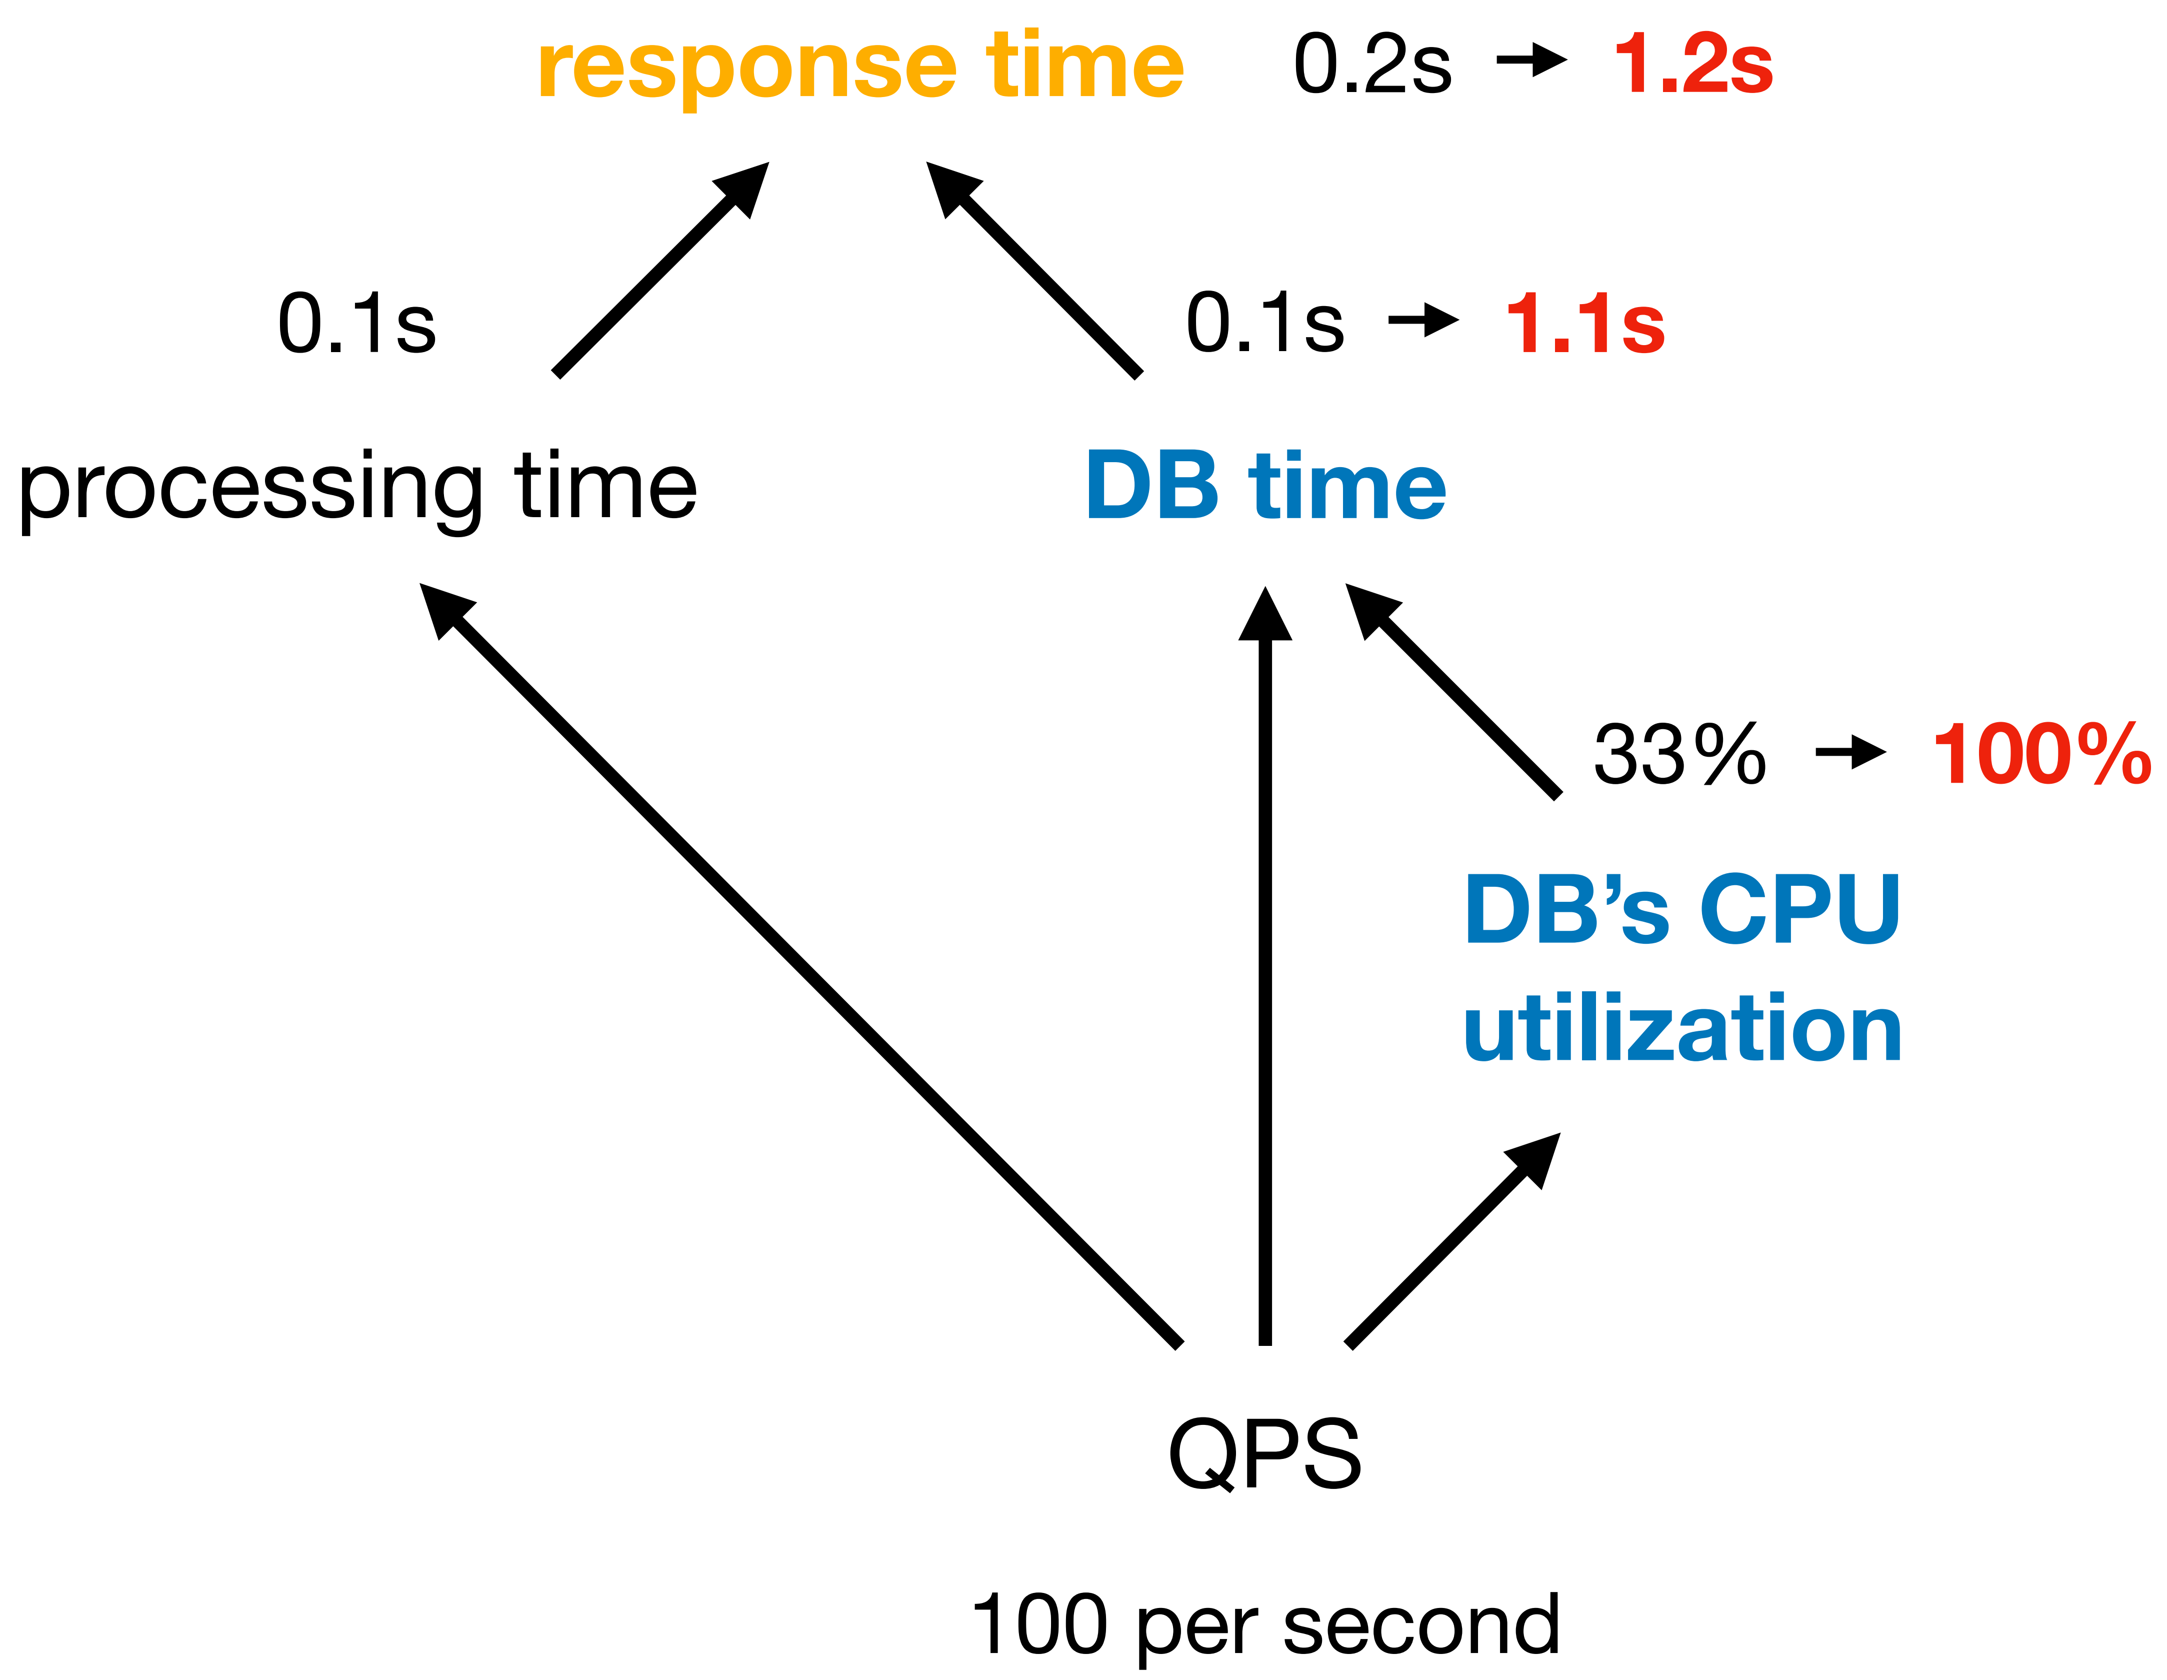

Background

Troubleshooting

- Recent automatic troubleshooting works model the propagation with a **graph**
- response time = processing time + DB time
  - This equation still holds after failure
  - response time is affected by DB time

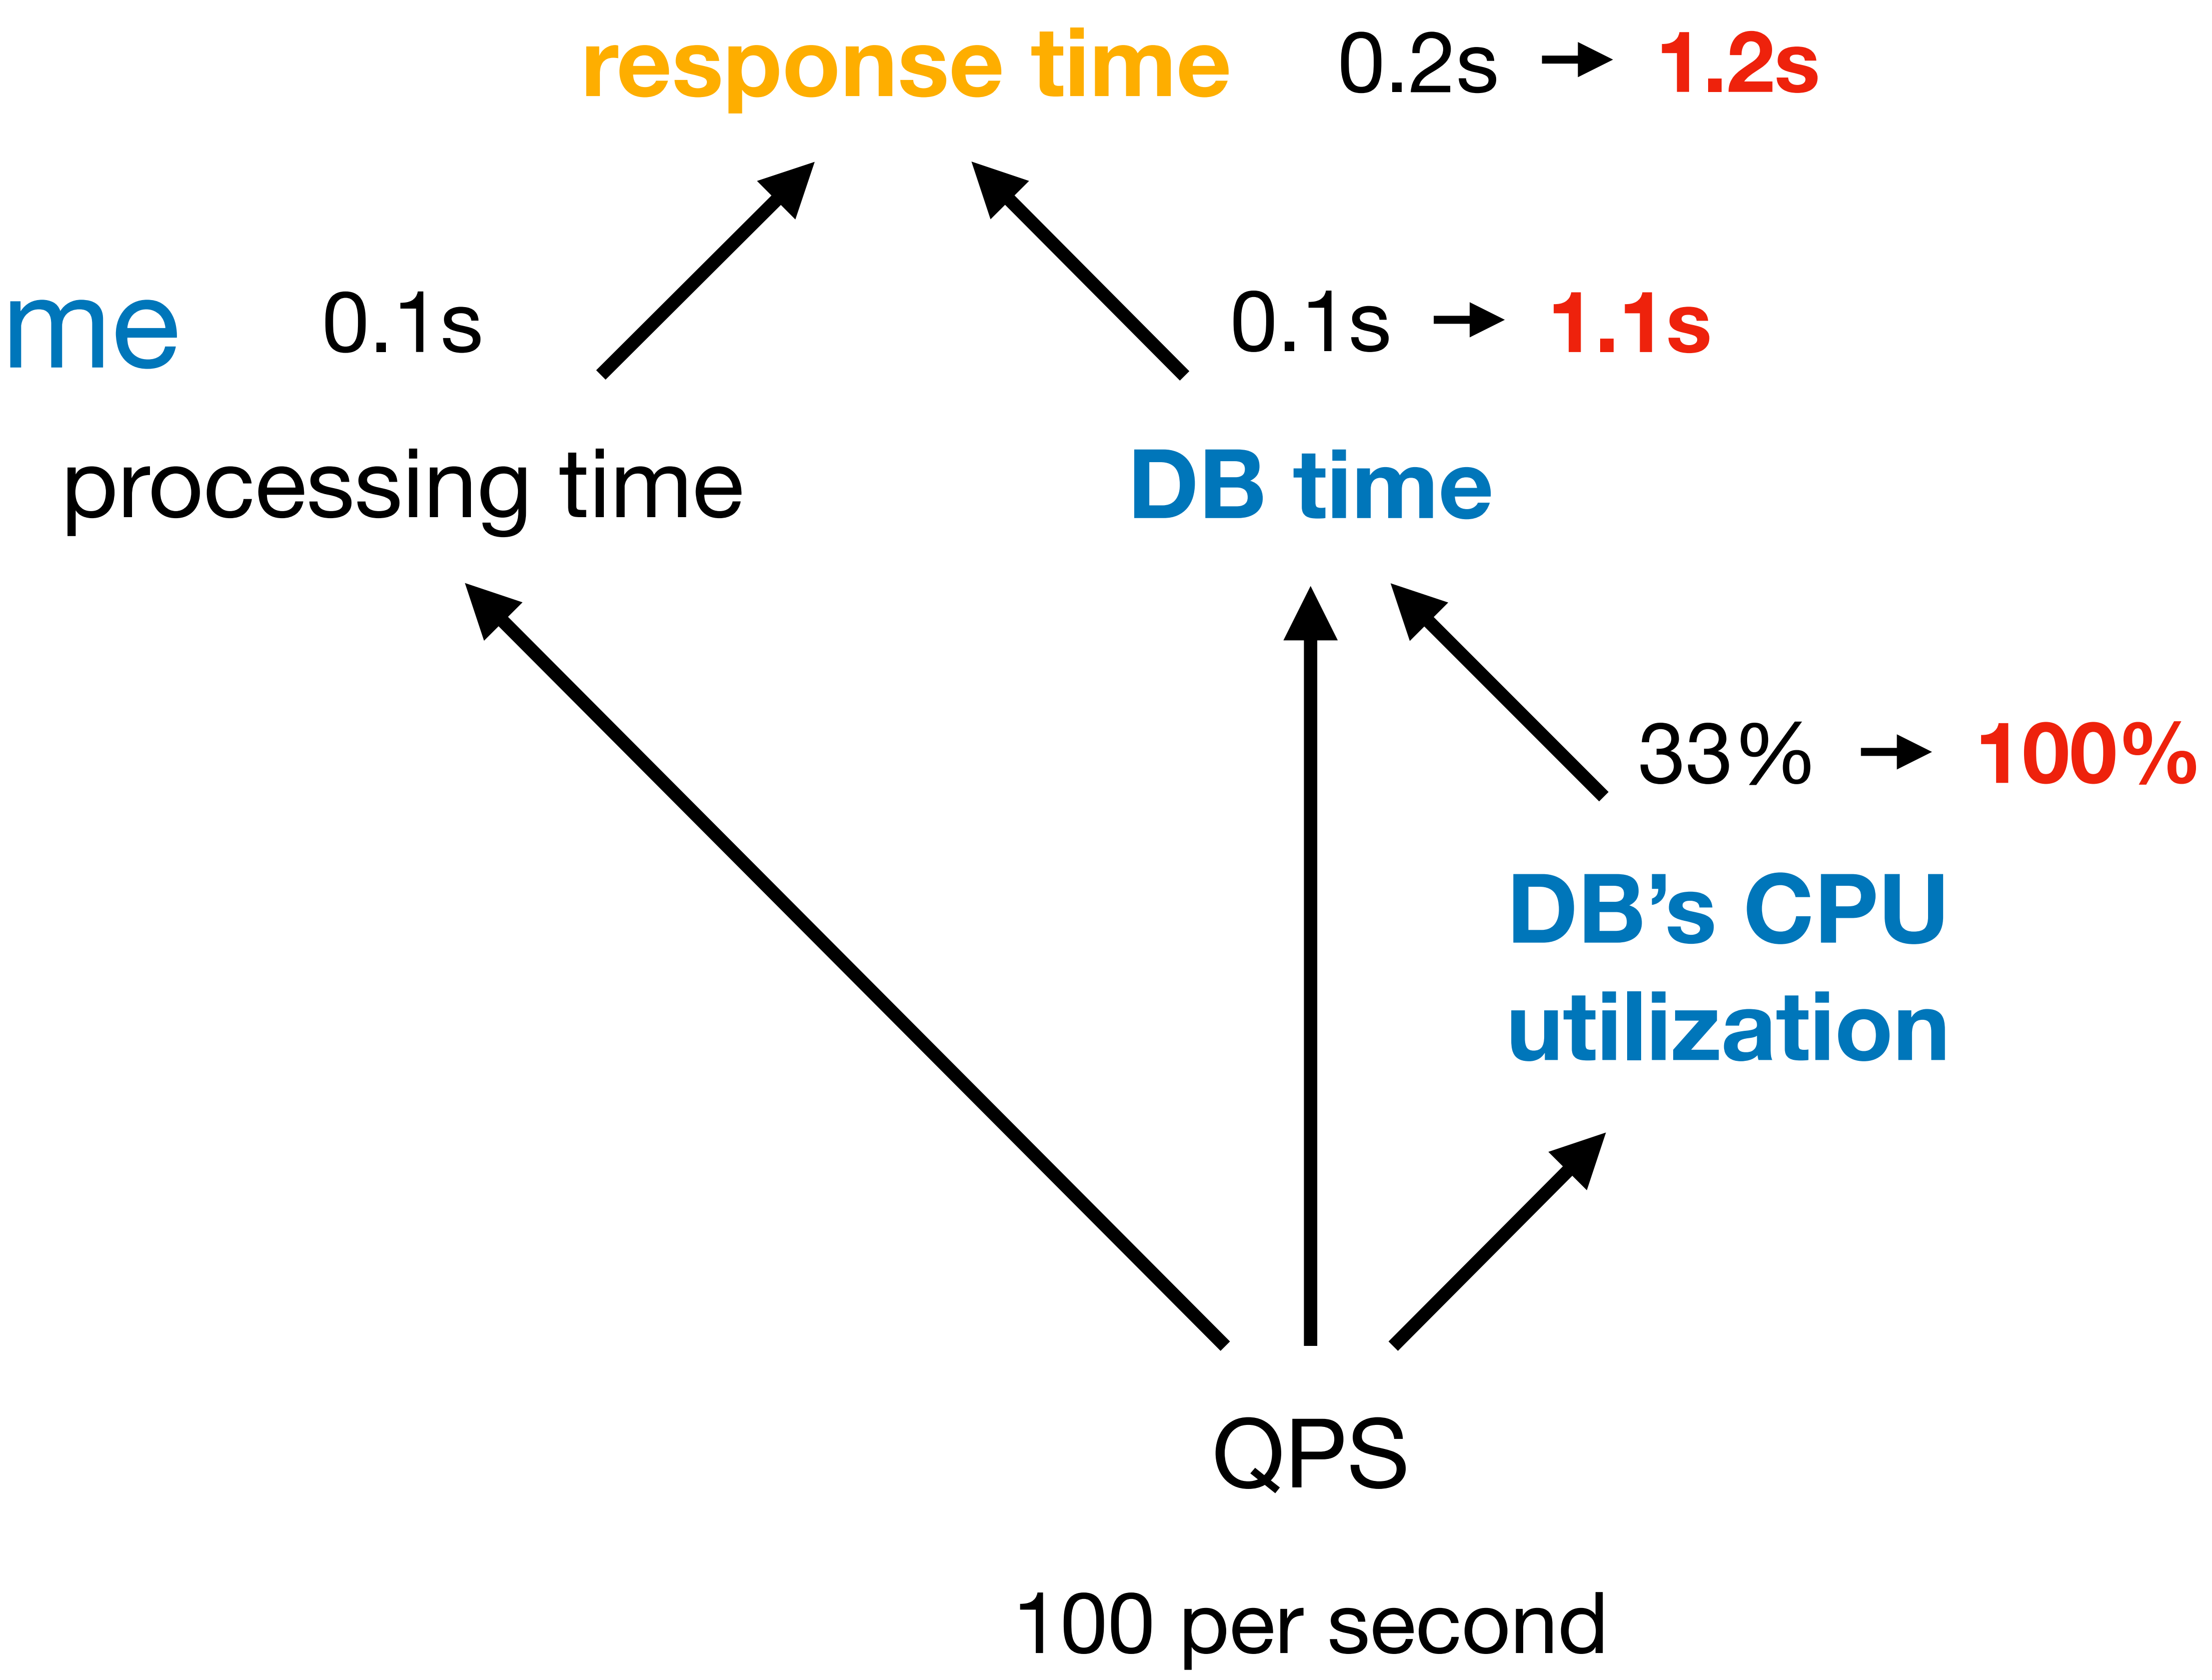

Background

Troubleshooting

- Recent automatic troubleshooting works model the propagation with a **graph**
- response time = processing time + DB time
  - This equation still holds after failure
  - response time is affected by DB time
- A search-based method will traverse the abnormal metrics in the graph

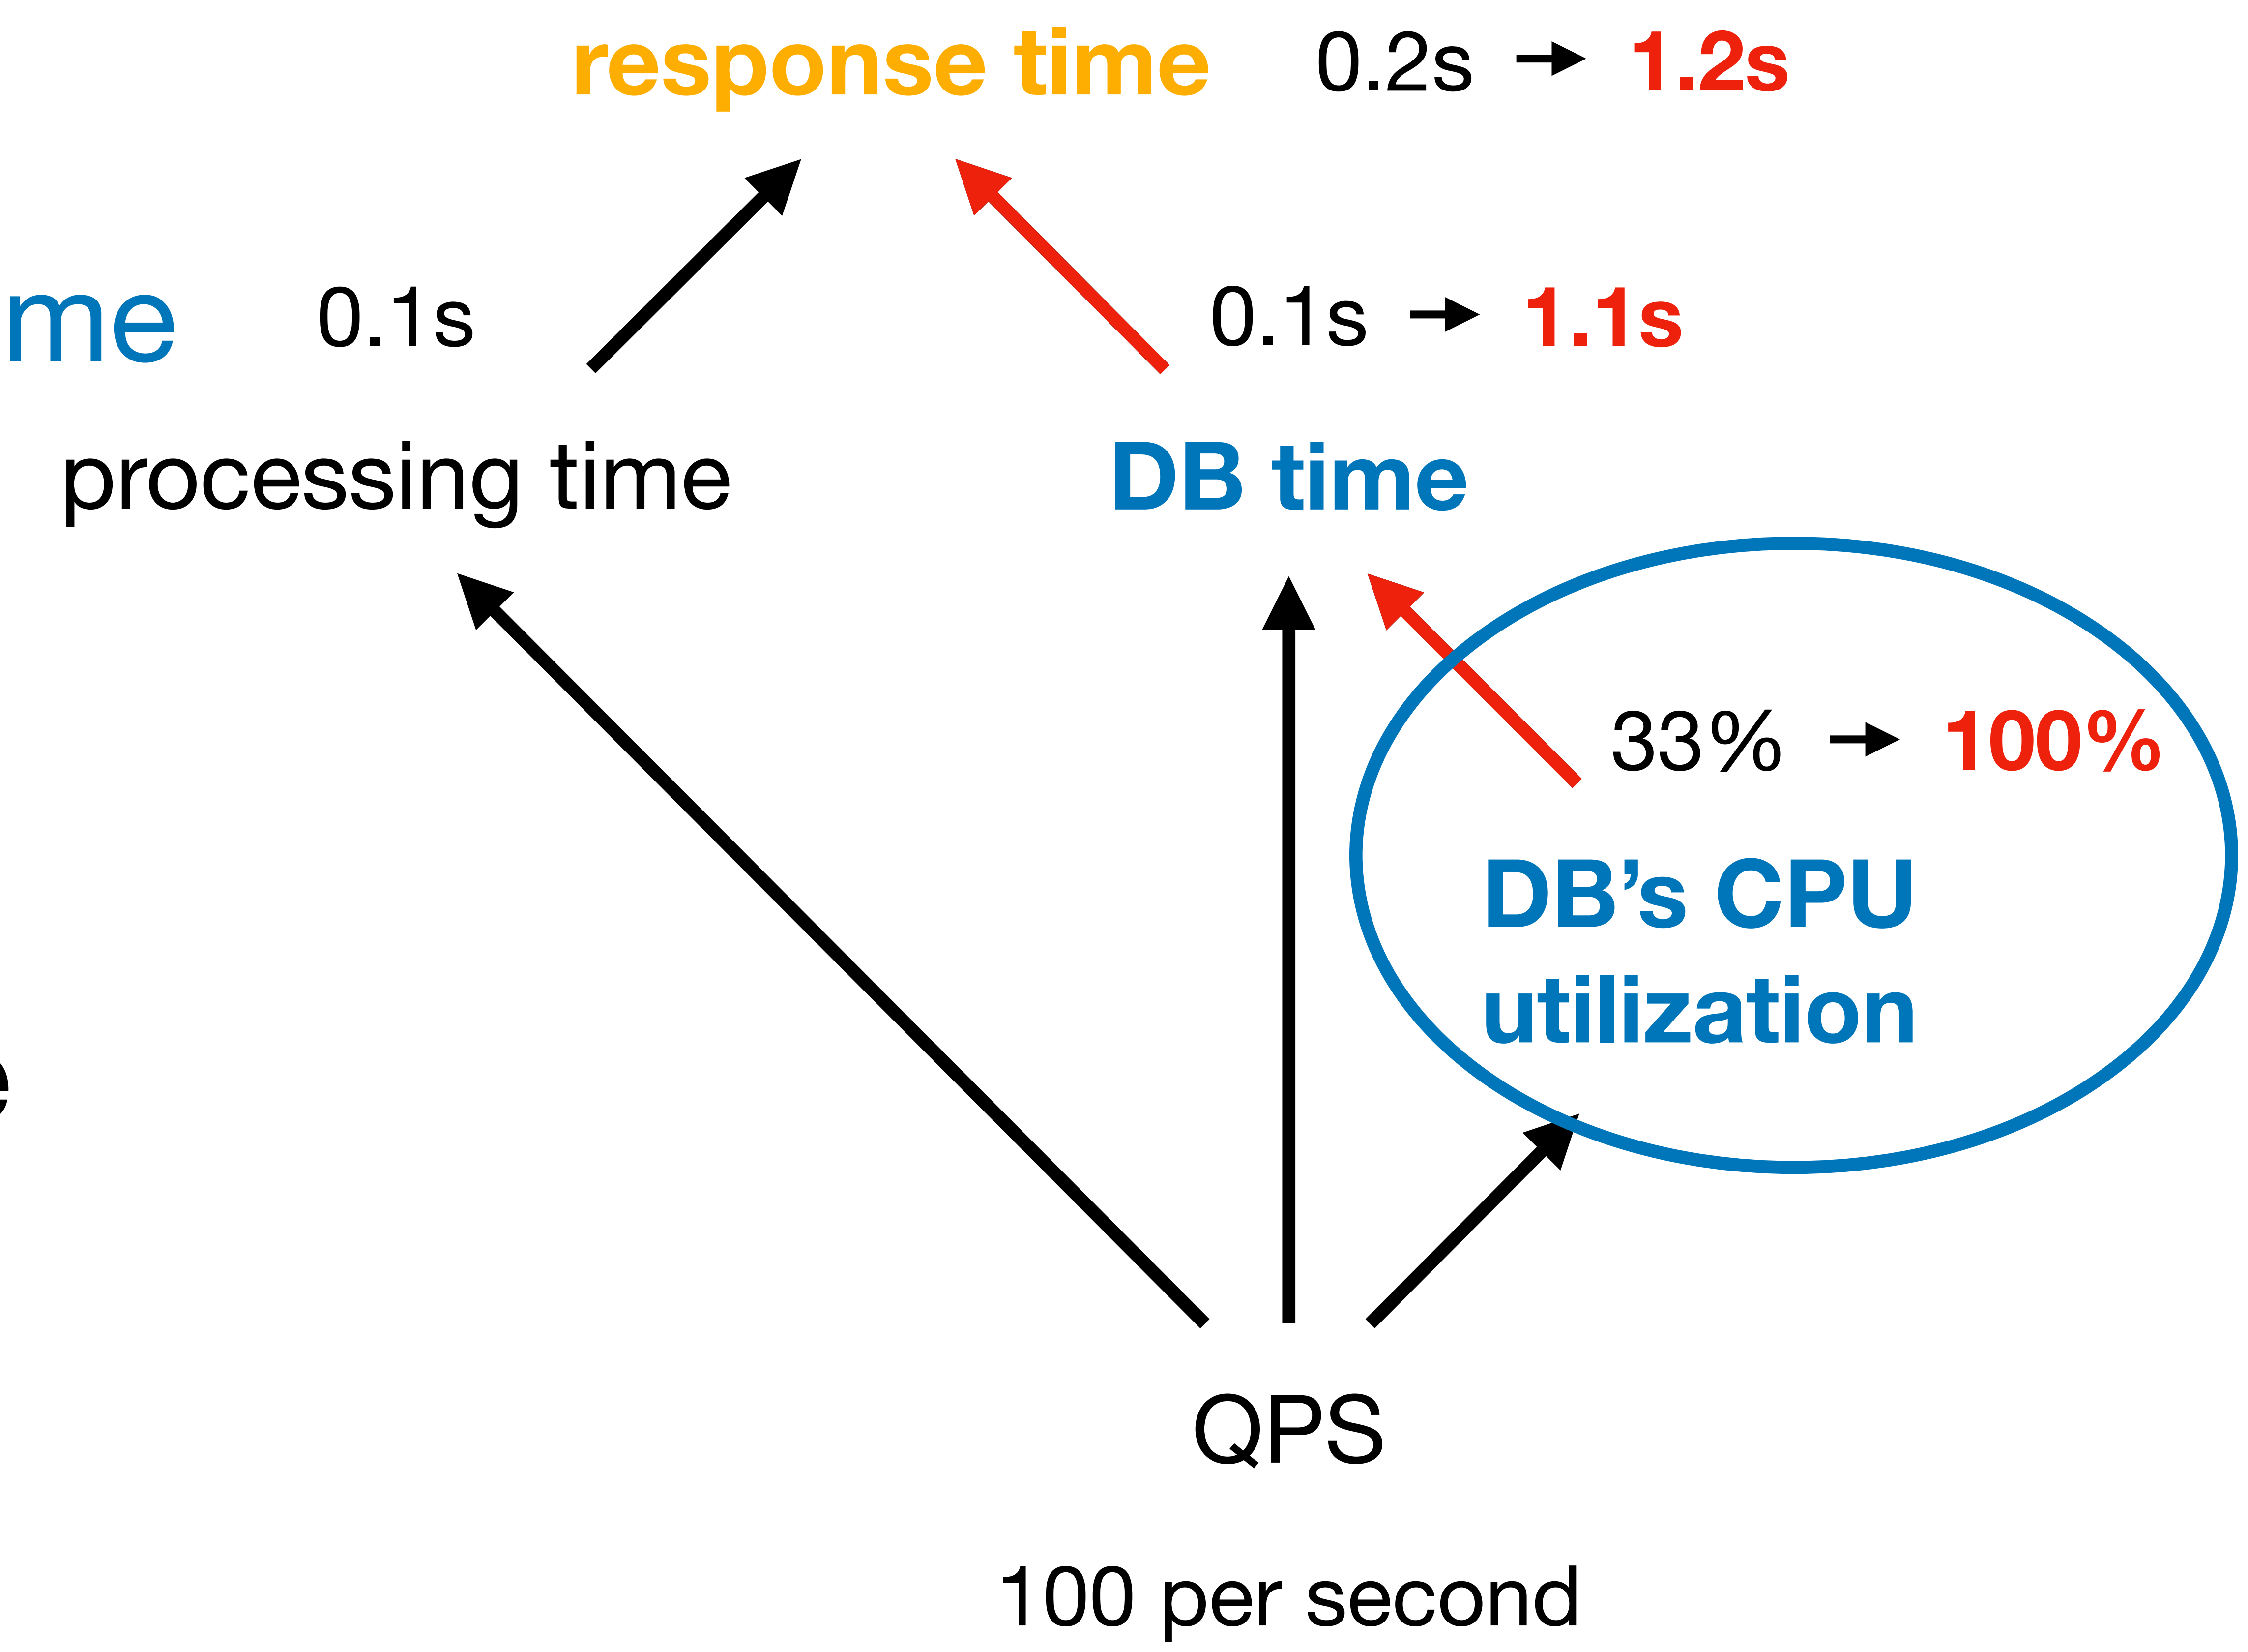

Background

Fluctuation Propagation Graph (FPG)

- In this work, we focus on the structure discovery of such a graph, named a *fluctuation propagation graph* (FPG)
  - An FPG describes how fluctuations, especially faults, propagate among metrics

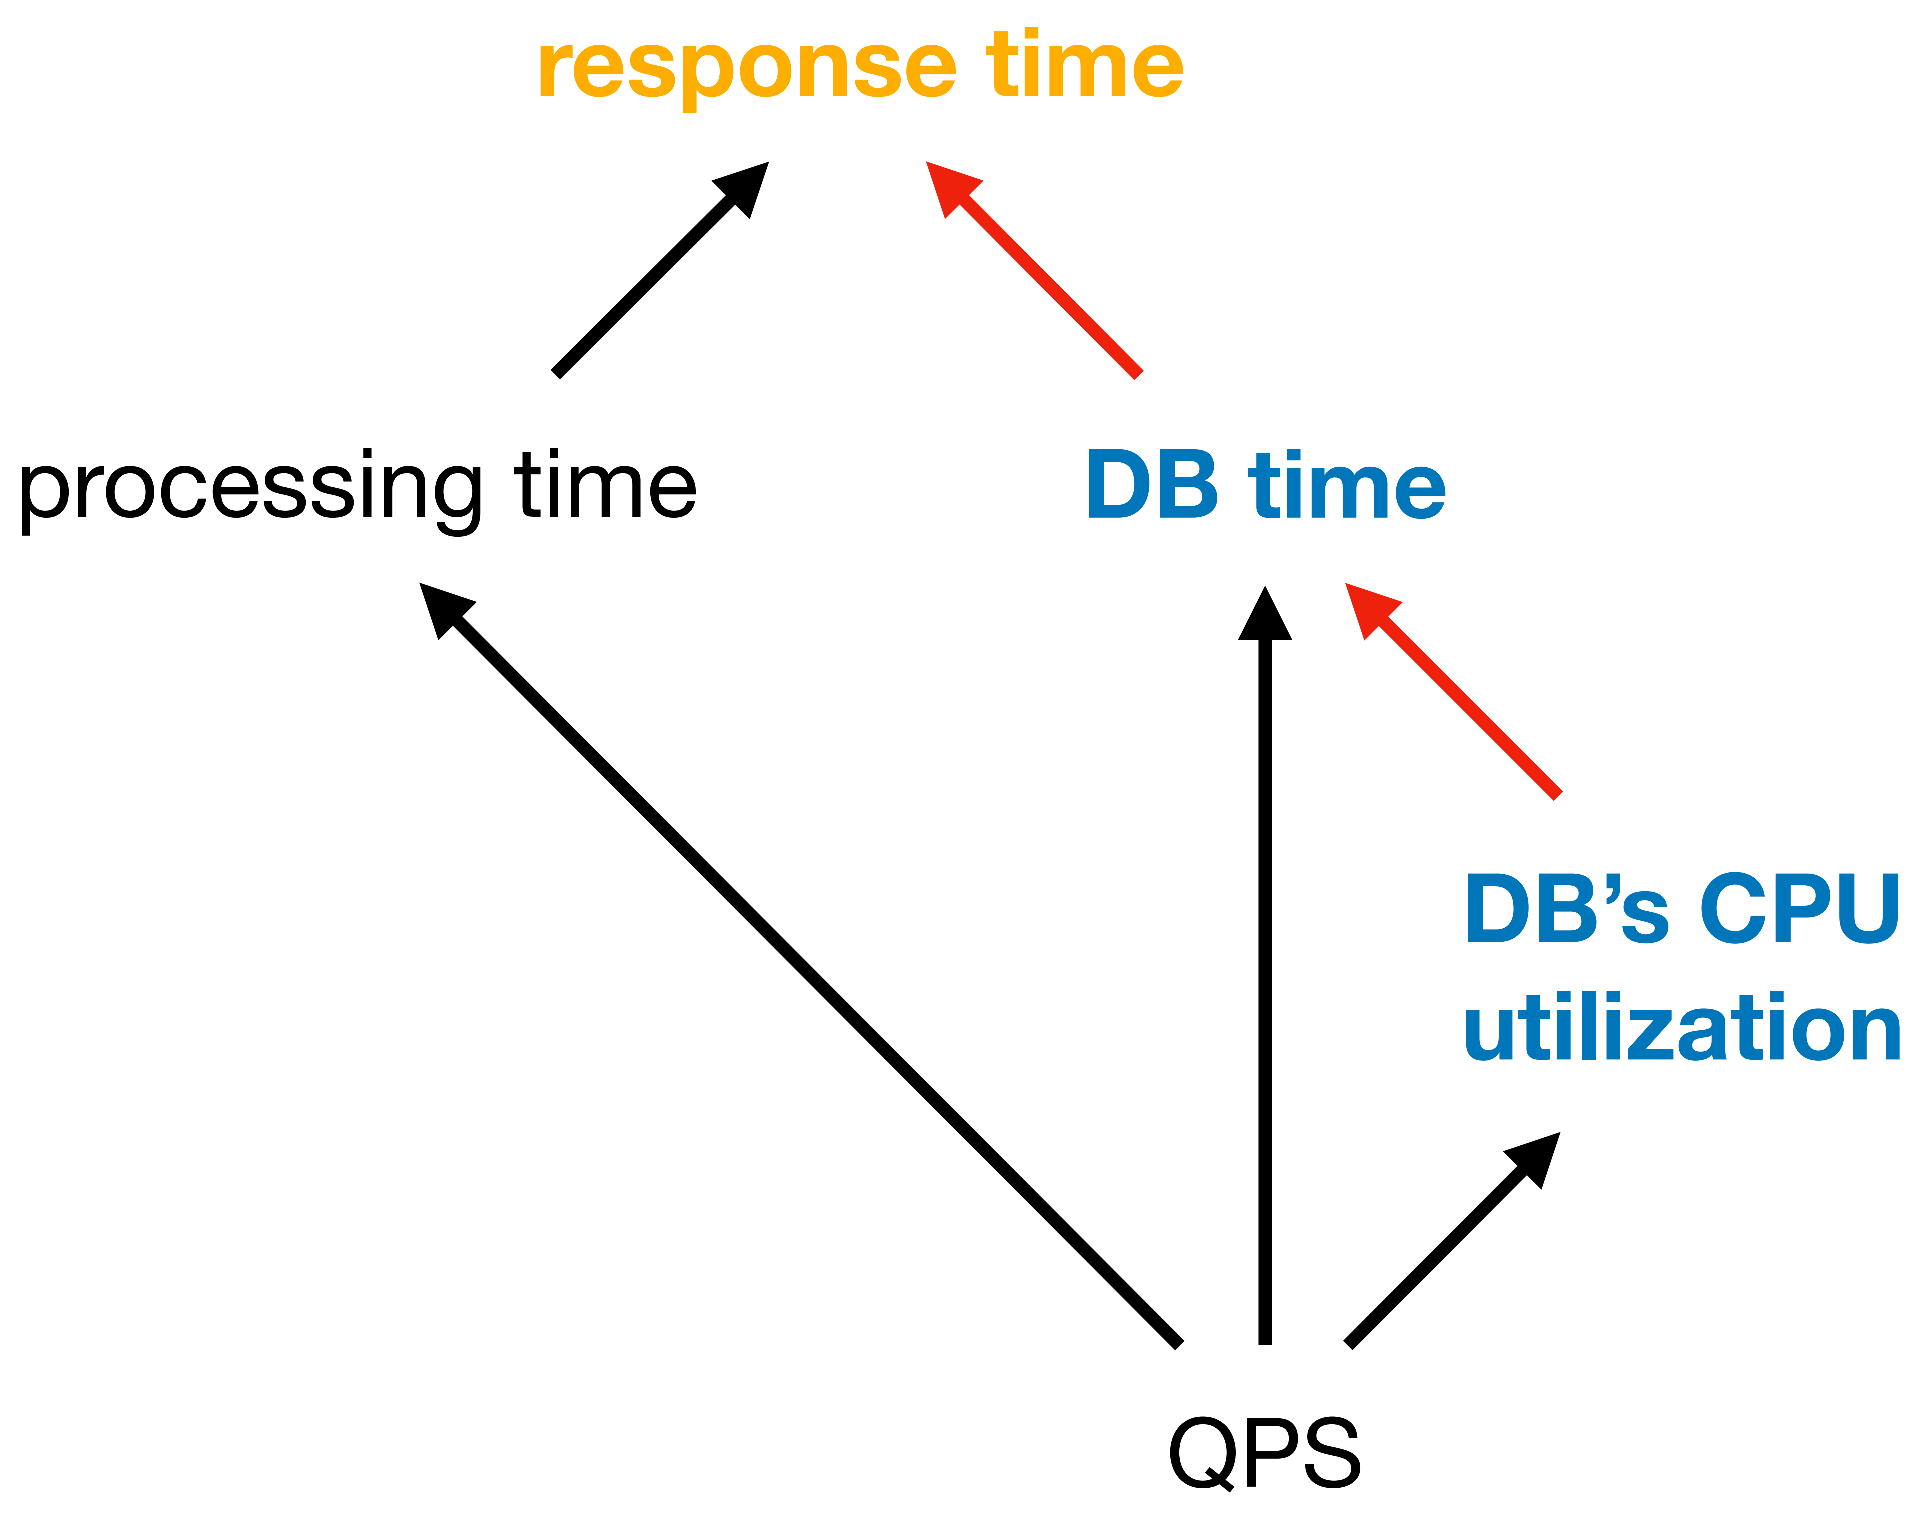

Background

Related Works

	Manual Construction	Mining from Data
Reference	TON’12, ASE’21	INFOCOM’14, KDD’16, CCGRID’18, ICSOC’20, IWQoS’20, NOMS’20, WWW’20
Limitation	Require extensive domain knowledge and significant efforts without a guideline	Neglect the correctness of the mined graph

Background

Challenges

FPG Construction $\xrightarrow{\text{FPG}}$ Graph-base Algorithms $\longrightarrow$ Application targets

Background

Challenges

FPG Construction $\xrightarrow{\text{FPG}}$ Graph-base Algorithms $\longrightarrow$ Application targets

Lack of effective tools for unsupervised mining

Existing mining methods fail to discover the ground truth graph, evaluated on two real-world datasets.

Background

Challenges

FPG Construction $\xrightarrow{\text{FPG}}$ Graph-base Algorithms $\longrightarrow$ Application targets

Lack of effective tools for unsupervised mining

Existing mining methods fail to discover the ground truth graph, evaluated on two real-world datasets.

Once a graph-based algorithm fails to achieve its goal, a trustworthy FPG can still provide basic situation awareness for operators, which requires operators' verification.

Background

Challenges

FPG Construction $\xrightarrow{\text{FPG}}$ Graph-base Algorithms $\longrightarrow$ Application targets

Lack of effective tools for unsupervised mining

Operators' verification is labor-intensive

Existing mining methods fail to discover the ground truth graph, evaluated on two real-world datasets.

Once a graph-based algorithm fails to achieve its goal, a trustworthy FPG can still provide basic situation awareness for operators, which requires operators' verification.

Empirical Study

- RQ1: How do existing mining methods perform among monitoring metrics?

Empirical Study

Formulation

- An FPG is a directed acyclic graph

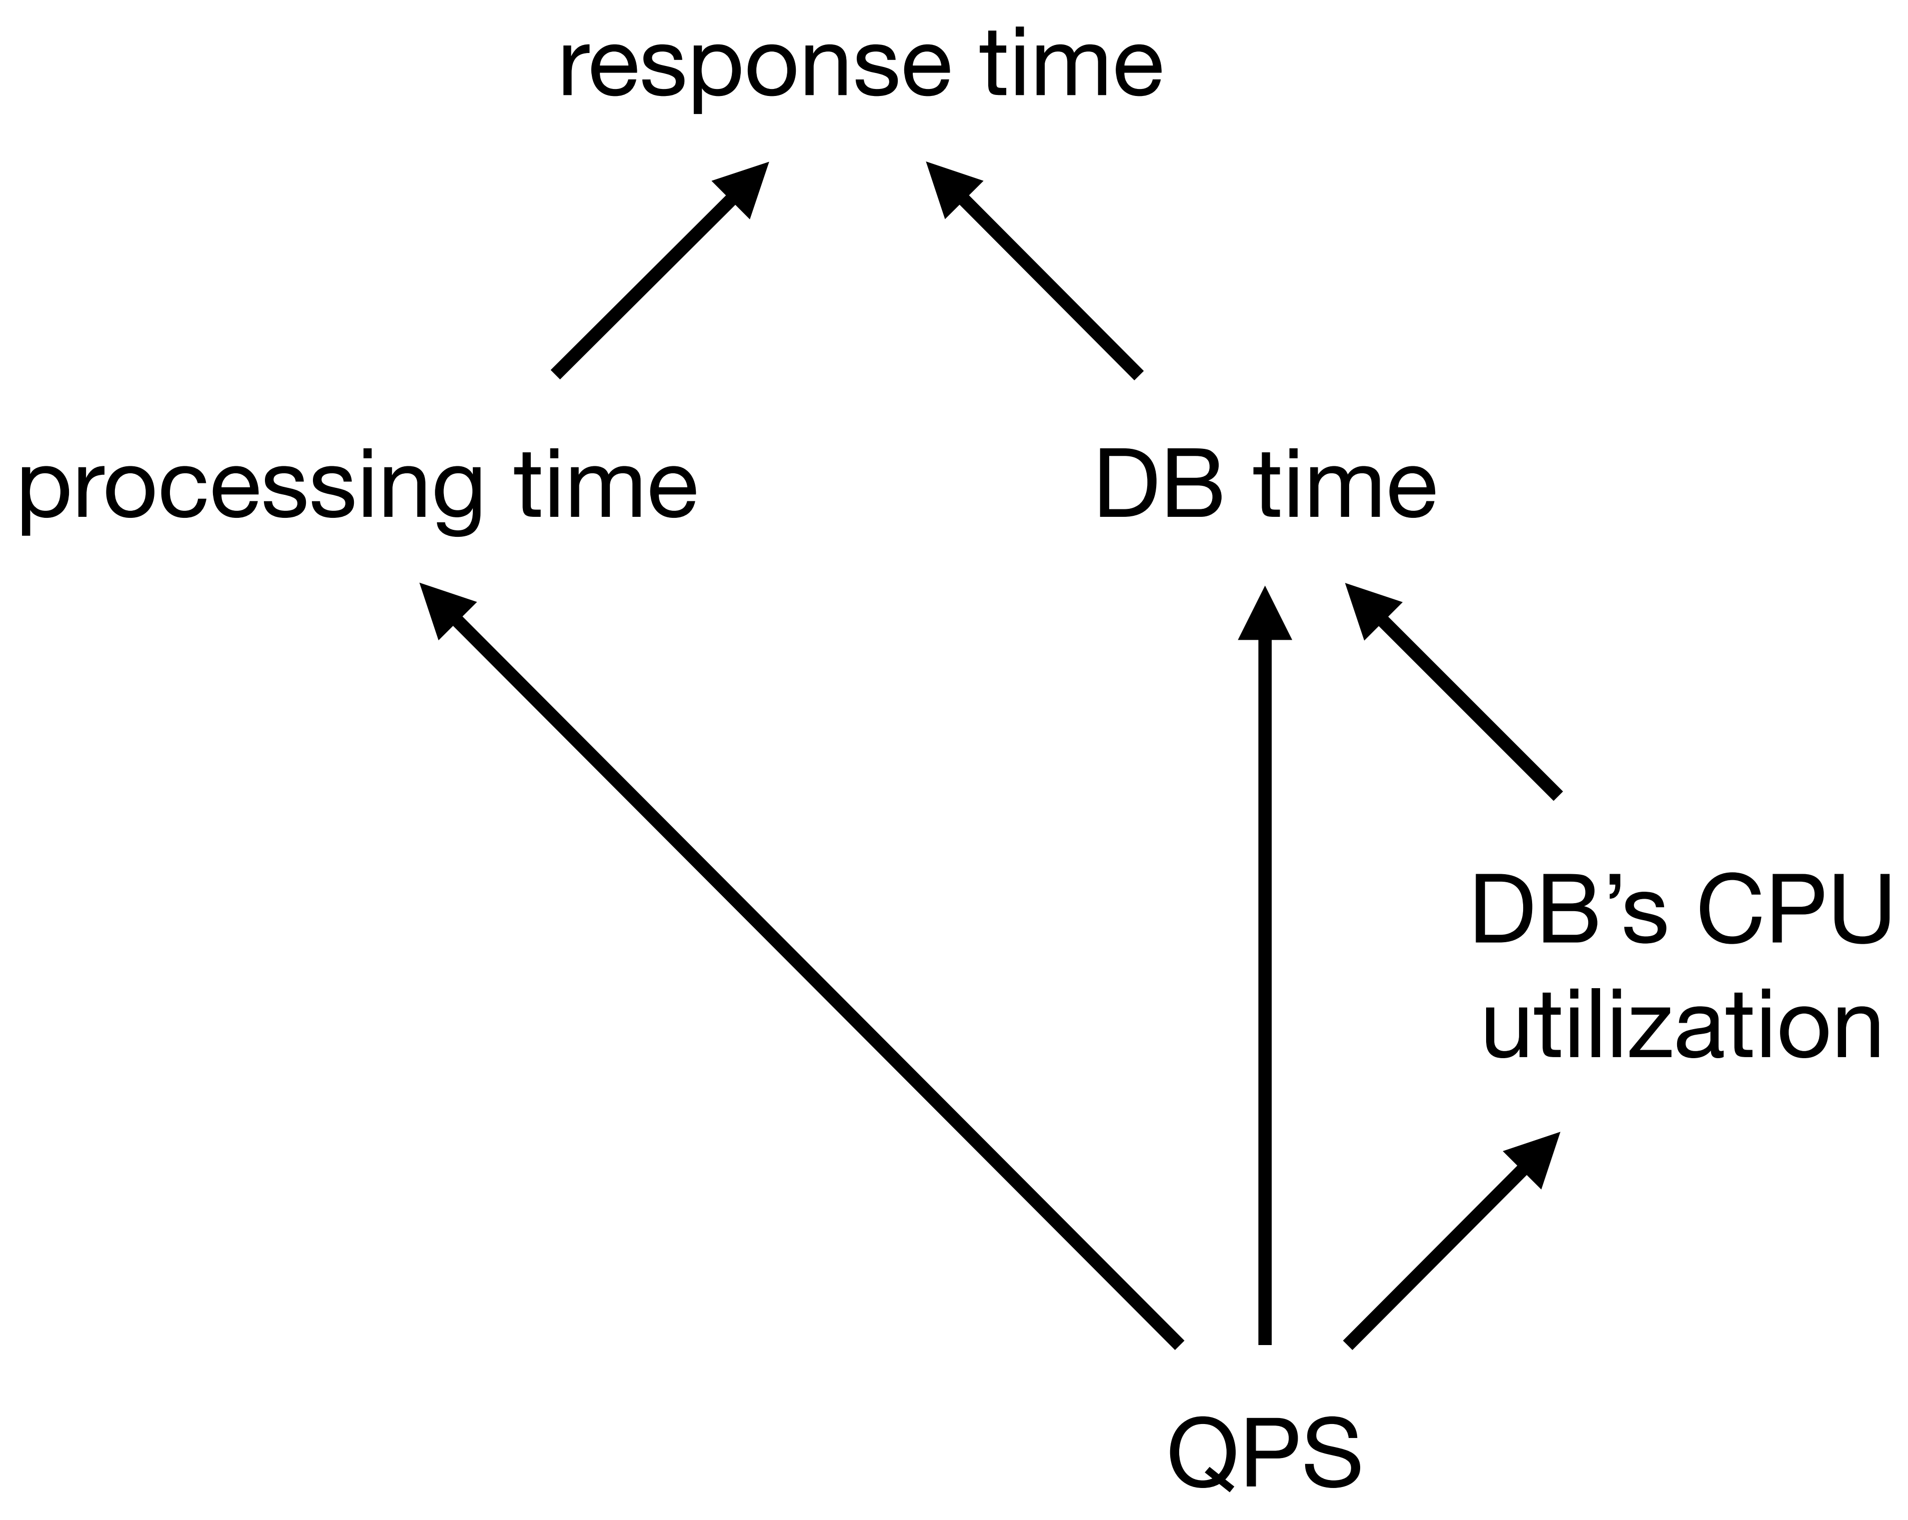

Empirical Study

Formulation

- An FPG is a directed acyclic graph
- Positive relations
  - Edges in the ground truth graph
  - e.g., "DB time" → "response time"

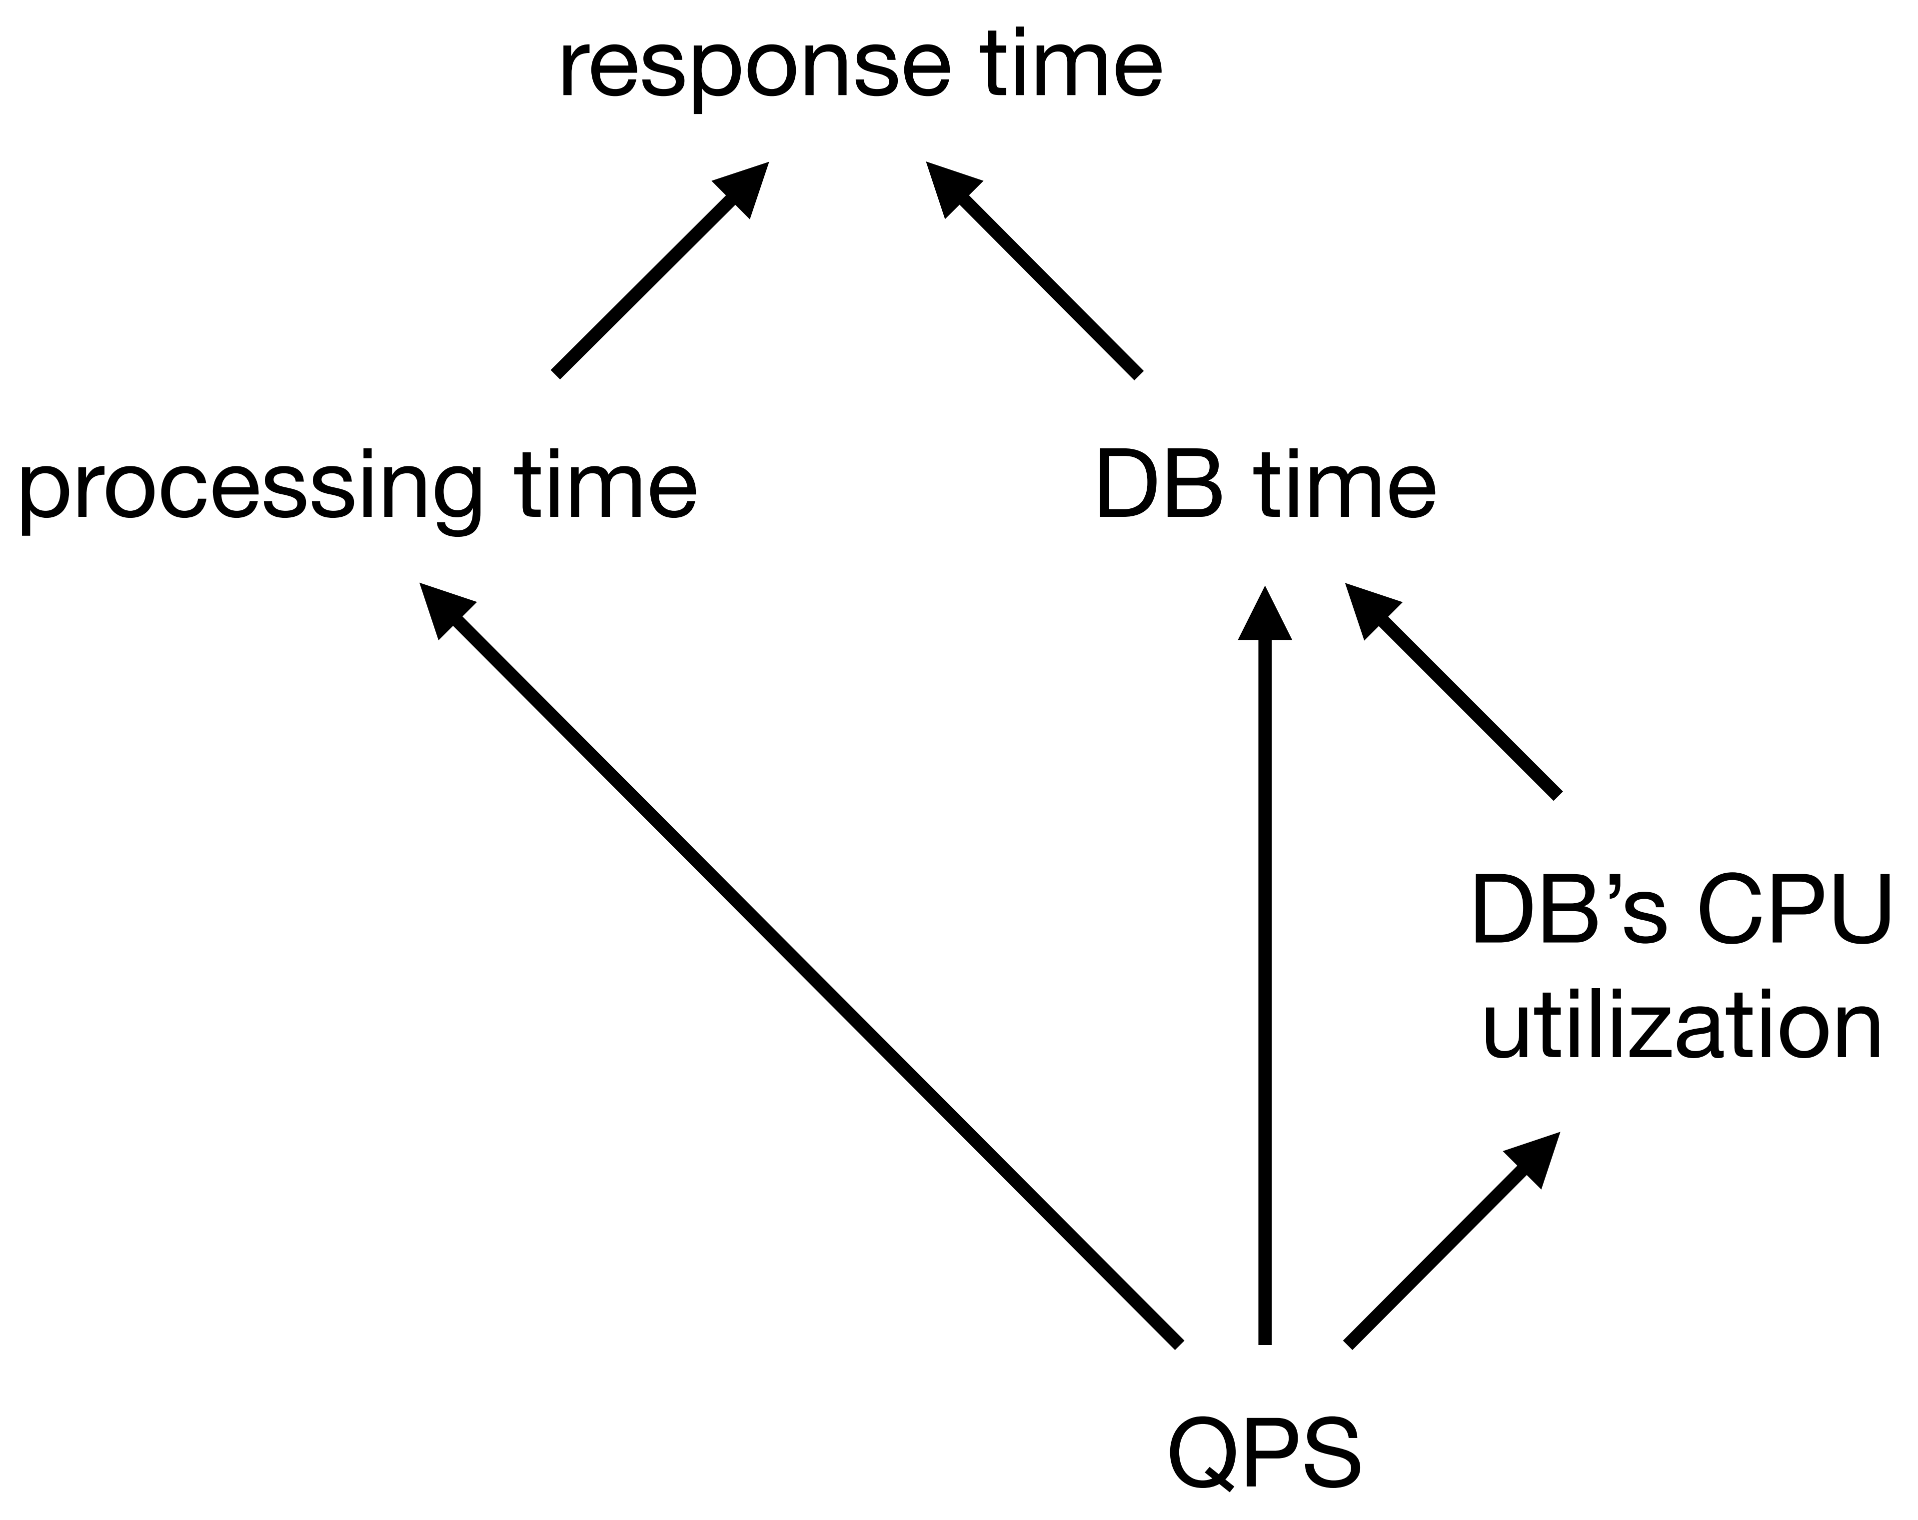

Empirical Study

Formulation

- An FPG is a directed acyclic graph
- Positive relations
  - Edges in the ground truth graph
  - e.g., "DB time" → "response time"
- Negative relations
  - Edges not in the ground truth graph
  - e.g., "response time" → "DB time"

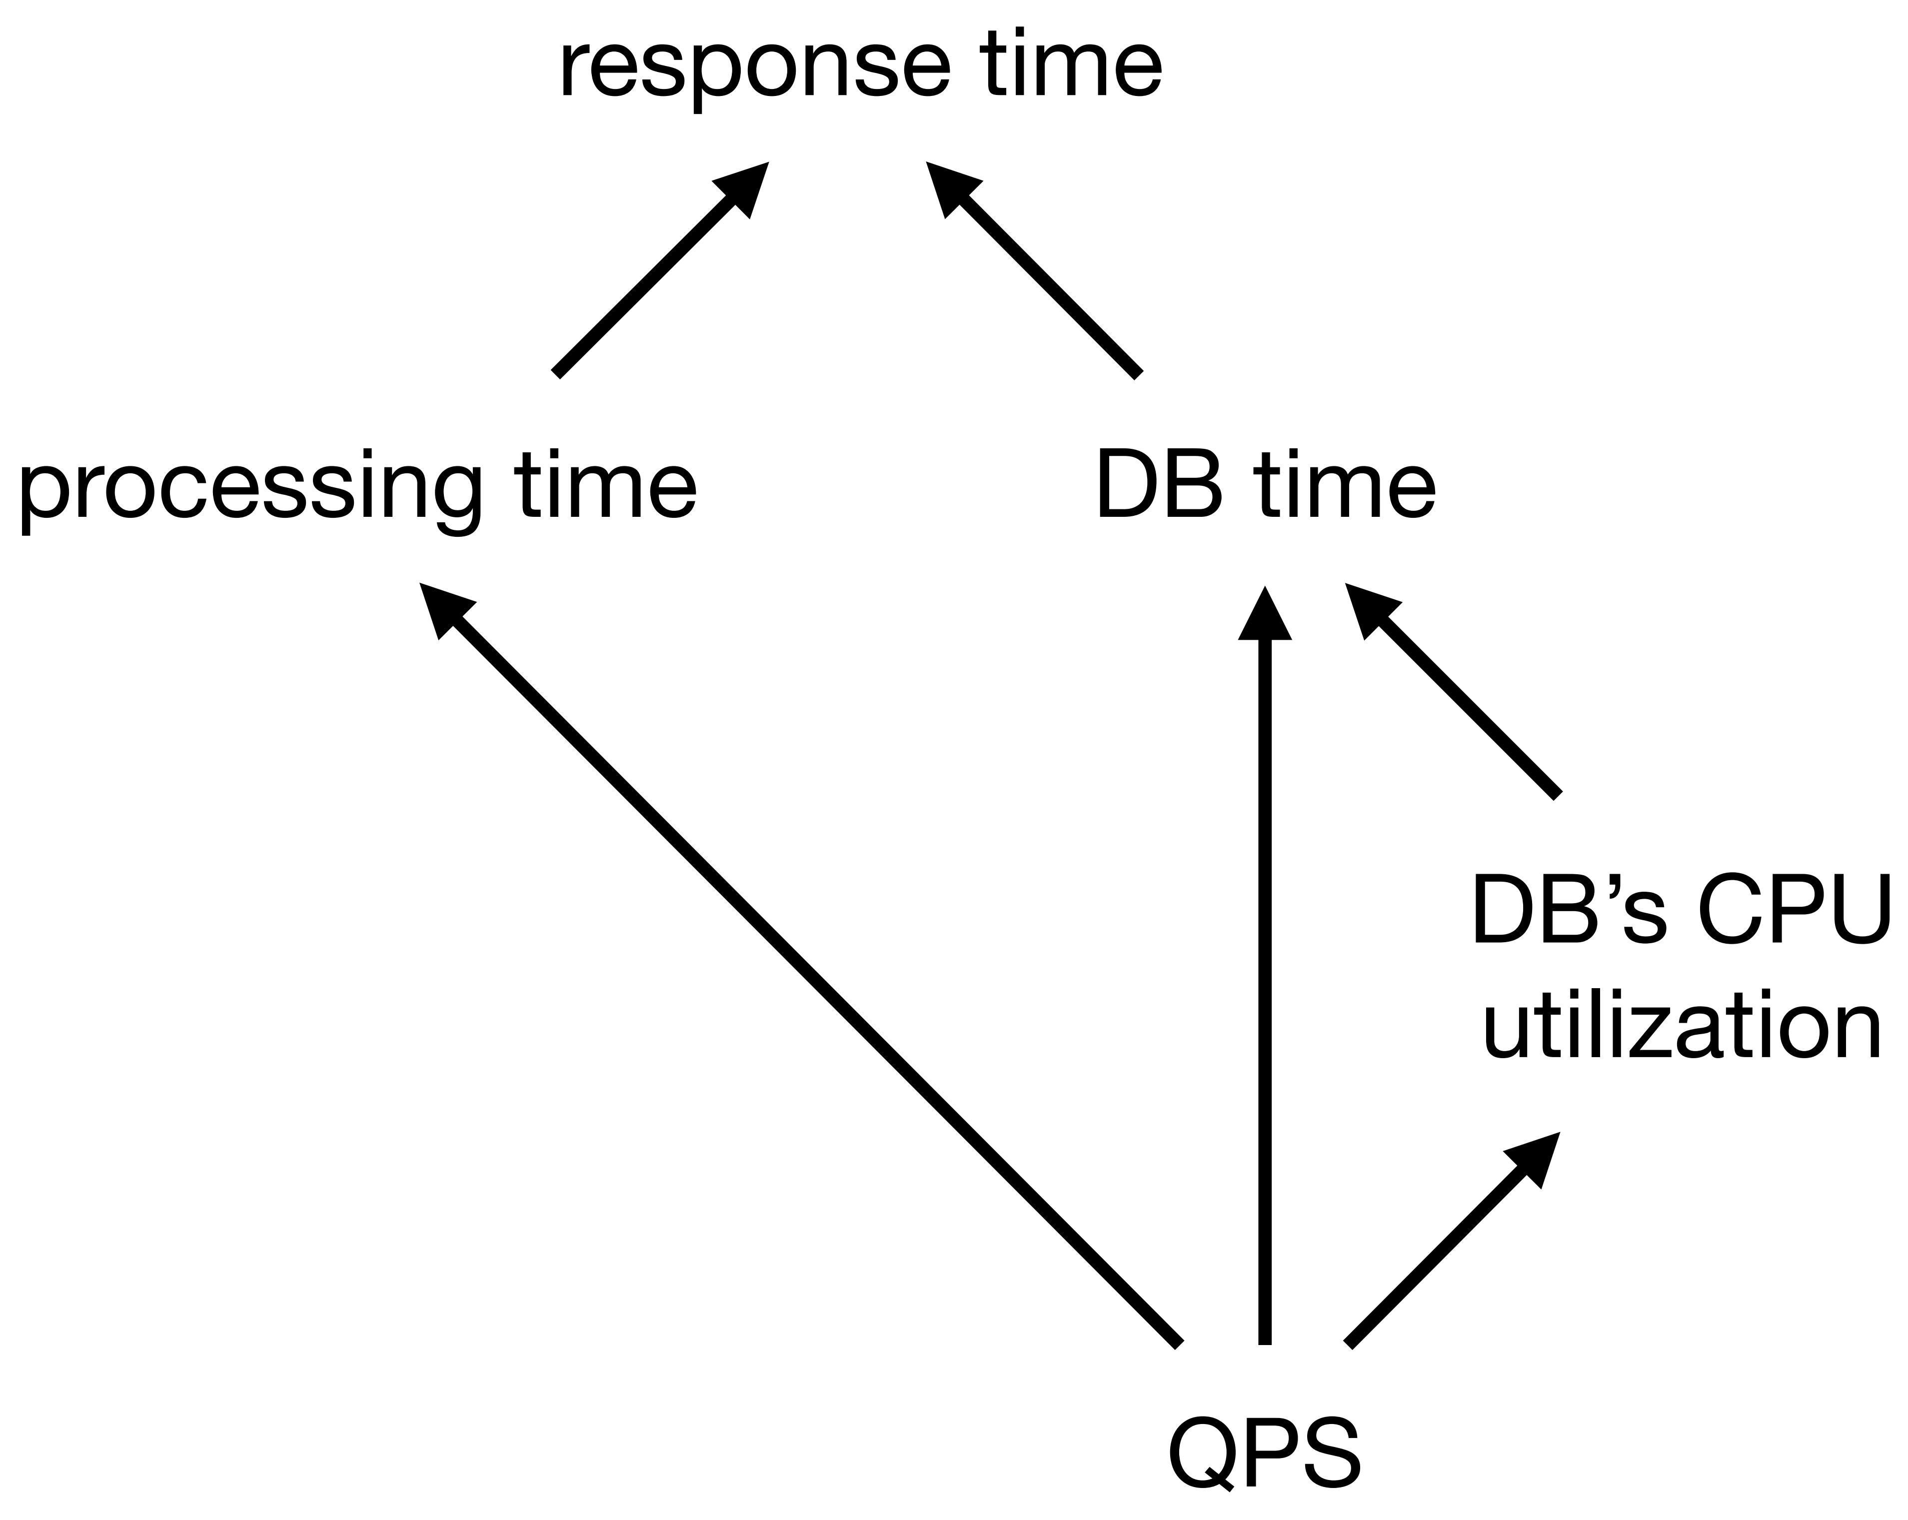

Empirical Study

Datasets

- $\mathscr{D}_{OD}$ is collected from an Oracle database with a real workload
- $\mathscr{D}_{TN}$ is a publicly available dataset from real telecommunication networks

Dataset	$\mathscr{D}_{OD}$	$\mathscr{D}_{TN}$
Scenario	**O**racle **D**atabase	**T**elecommunication **N**etwork
#Metric	51	55
#Length	1040	4032
Interval	6min	10min
#Label	490	1485
#Positive	210	563

Empirical Study

Mining Methods

Group		Mining Methods
Correlation		Pearson-r, Pearson-p, CC, CoFlux
Causality	**Constriant-based**	PC-gauss, PC-RCIT, PCTS-PCMCI, PCTS-PCMCI+
	Score-based	GES
	FCM-based	NOTEARS, NRI, TCDF

Empirical Study

Results

- Each method in the experiment suffers from either a low precision or low discovery ability on both datasets

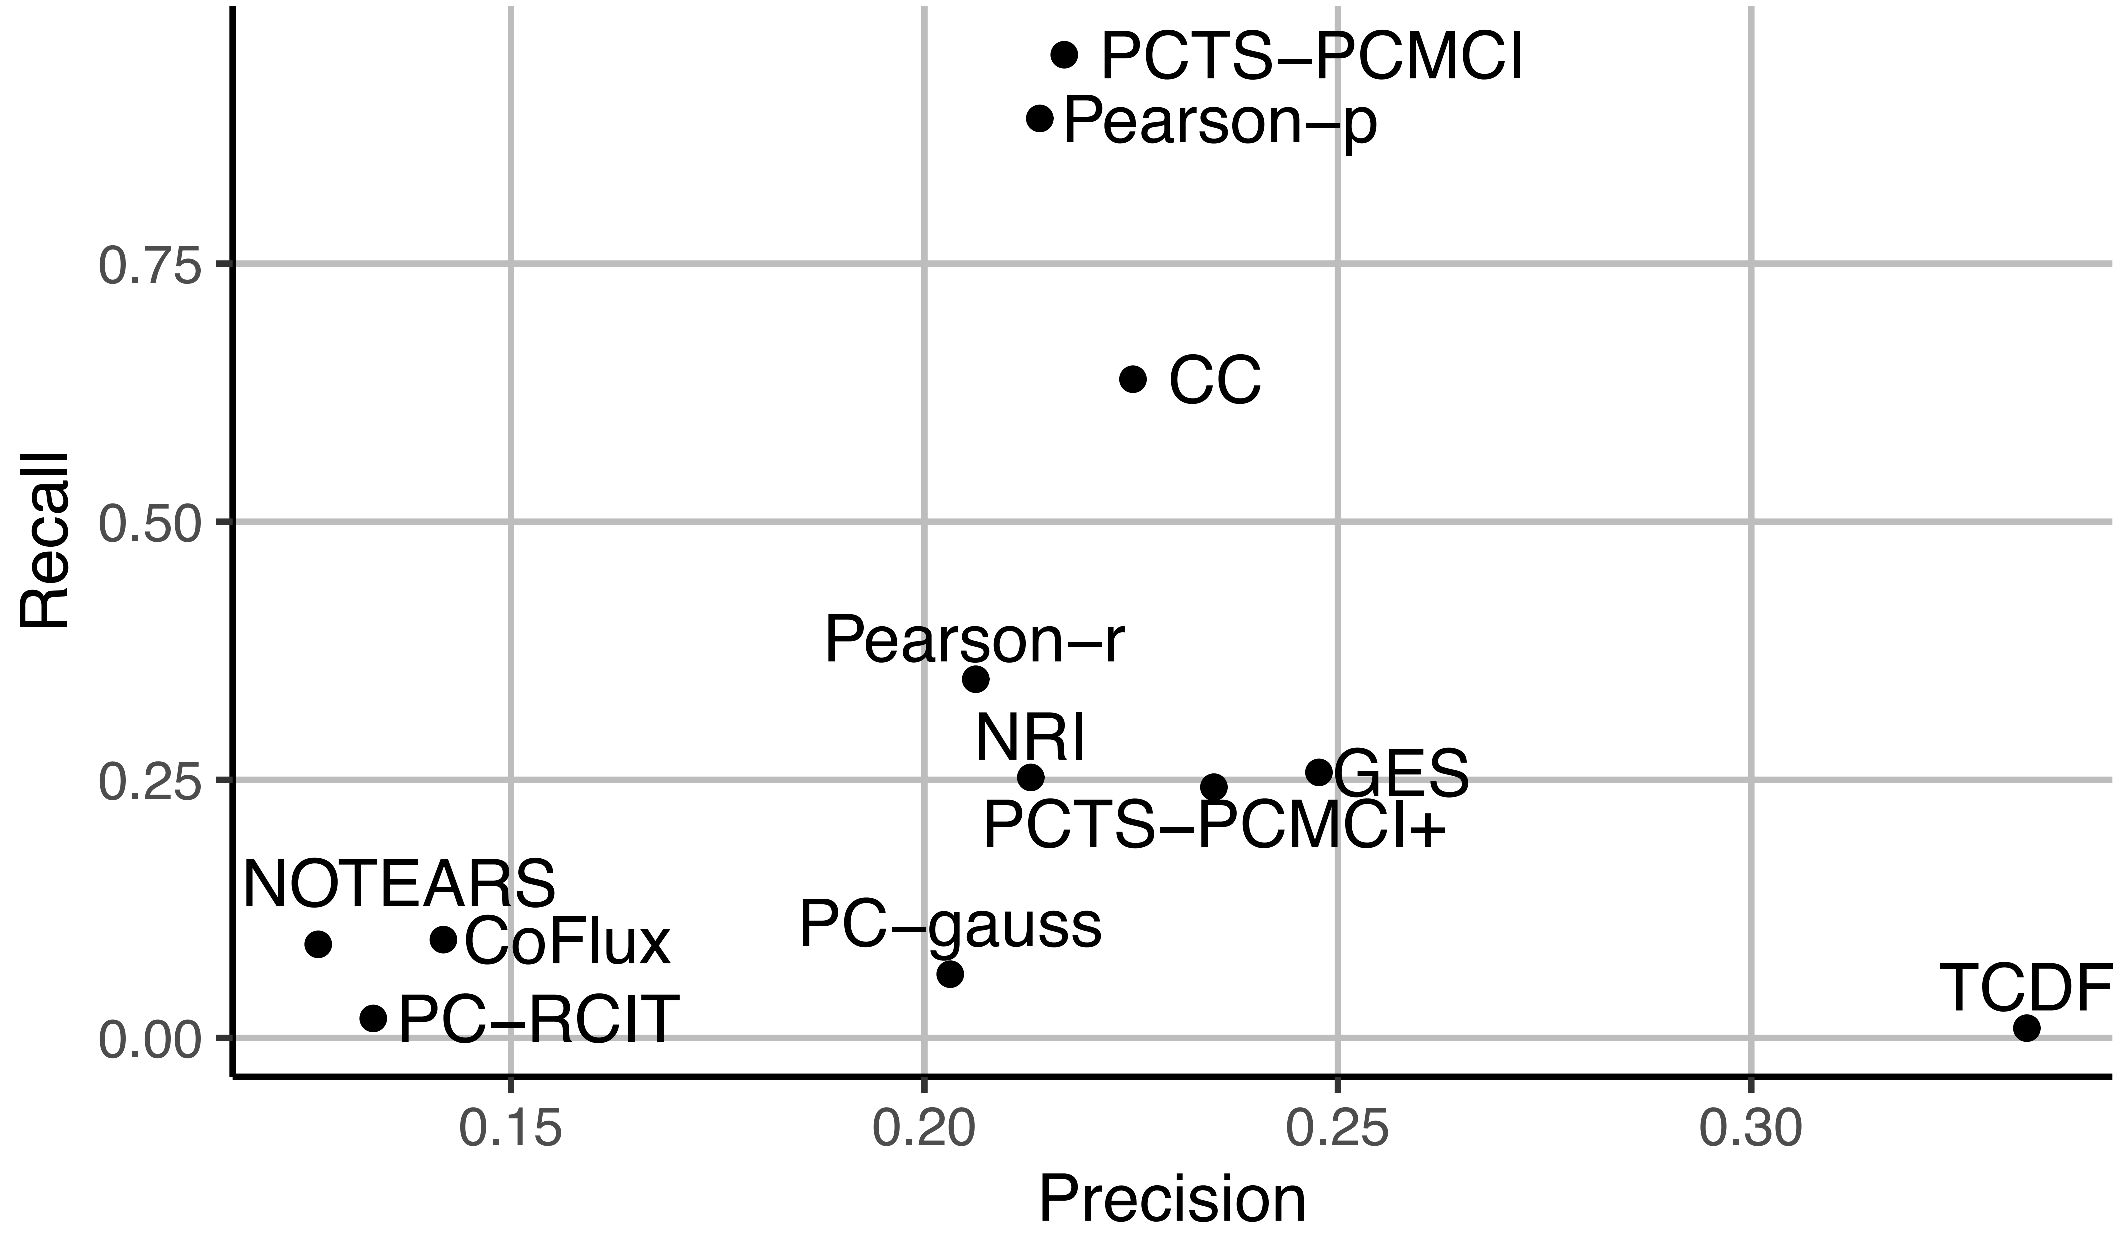

$\mathscr{D}_{OD}$

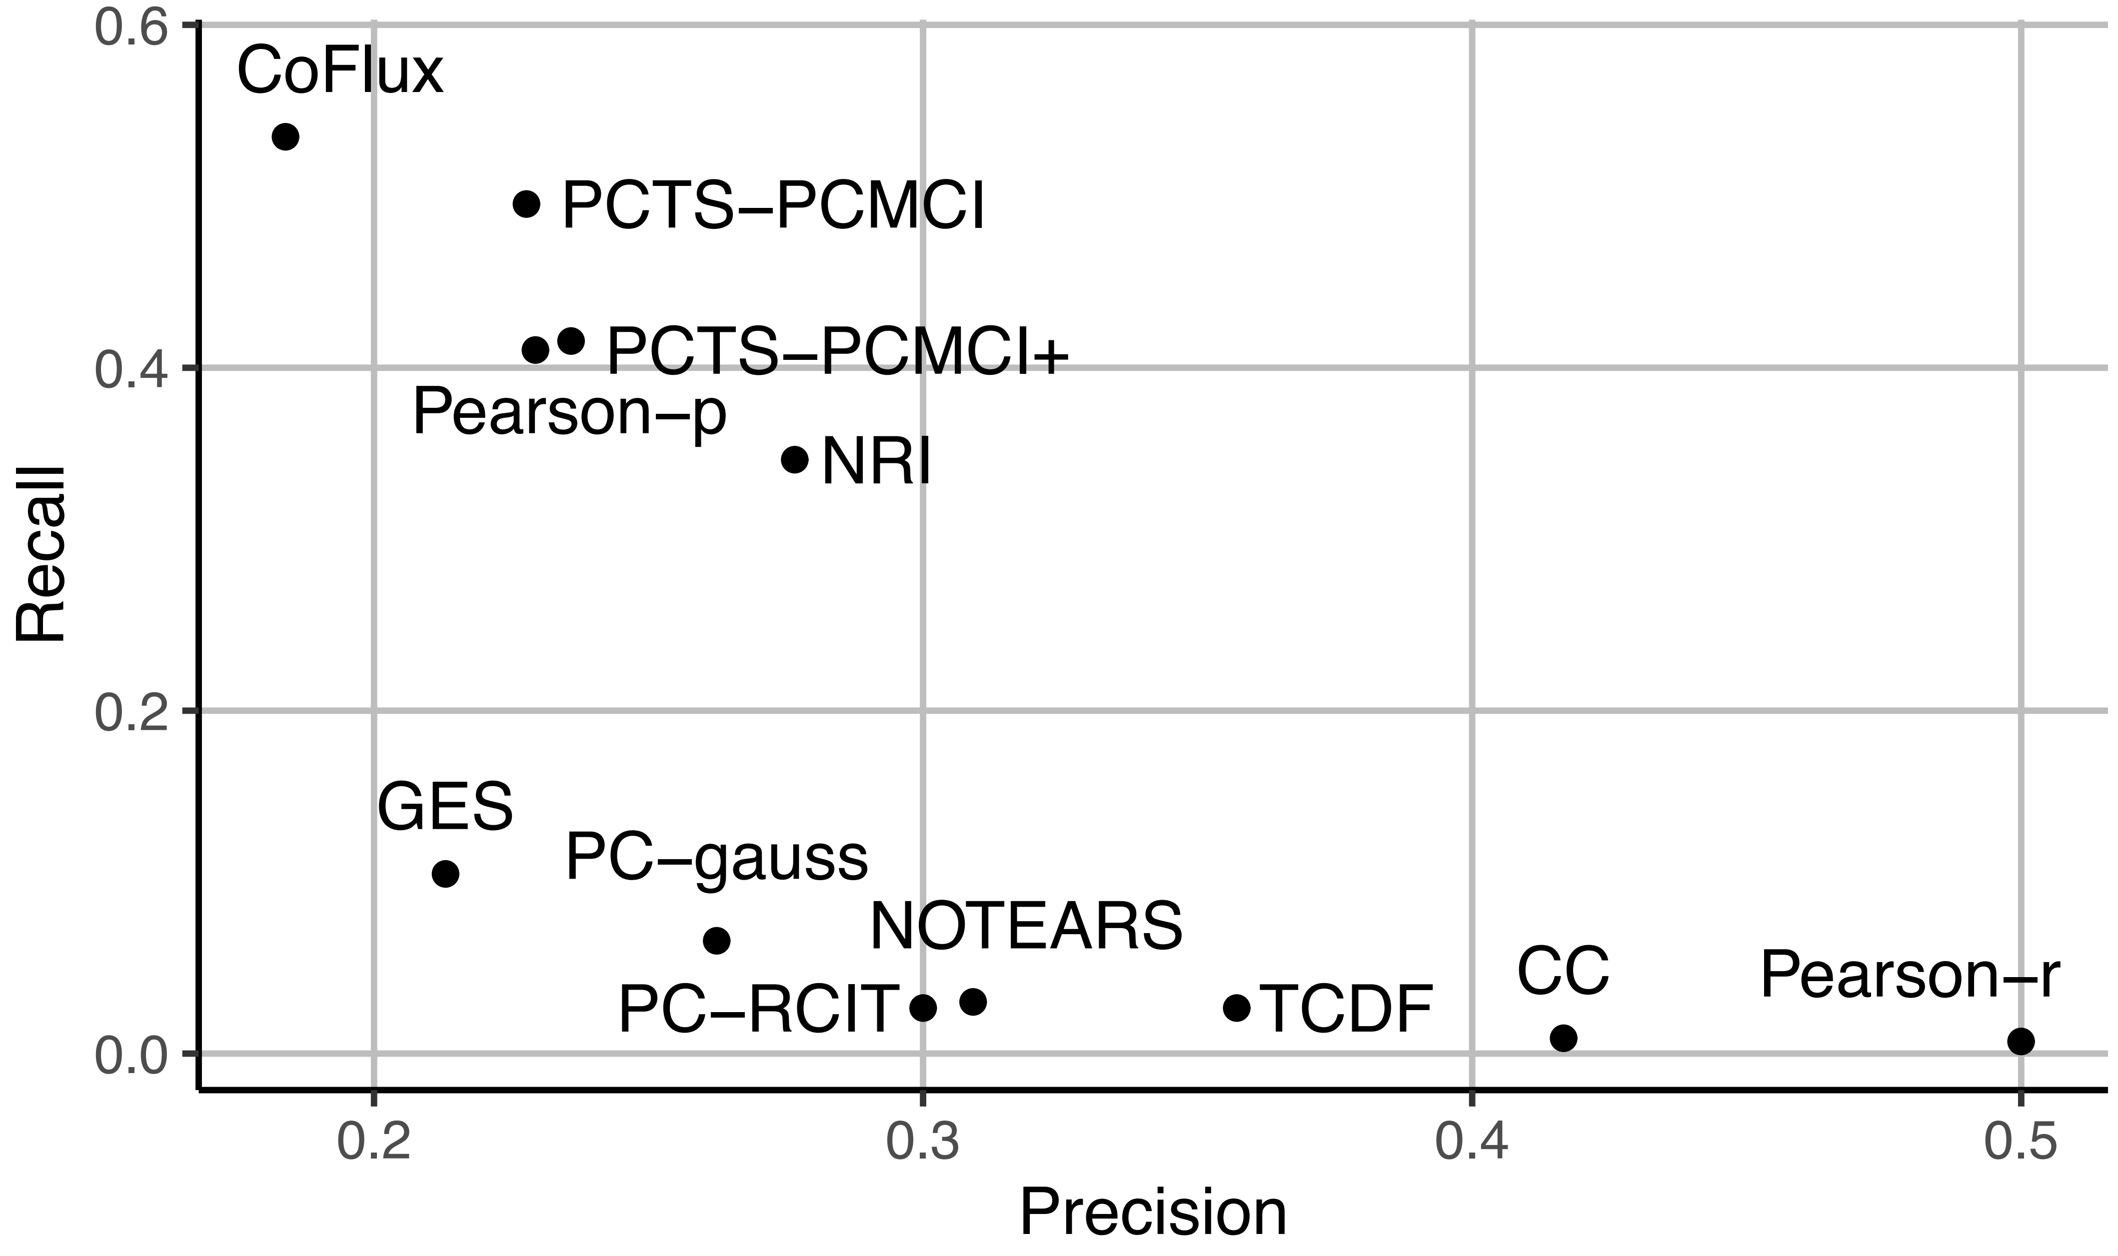

$\mathscr{D}_{TN}$

Empirical Study

Results

- Each method in the experiment suffers from either a low precision or low discovery ability on both datasets
- Existing mining methods lack domain knowledge
    - e.g., NOTEARS deals with linear relations, which generally does not suit the used datasets

Empirical Study

Results

- Each method in the experiment suffers from either a low precision or low discovery ability on both datasets
- Existing mining methods lack domain knowledge
  - e.g., NOTEARS deals with linear relations, which generally does not suit the used datasets

Bring operators' feedback (missing knowledge in the data) into the mining procedure

Methodology

Methodology

FPG-Miner: Recommendation Framework

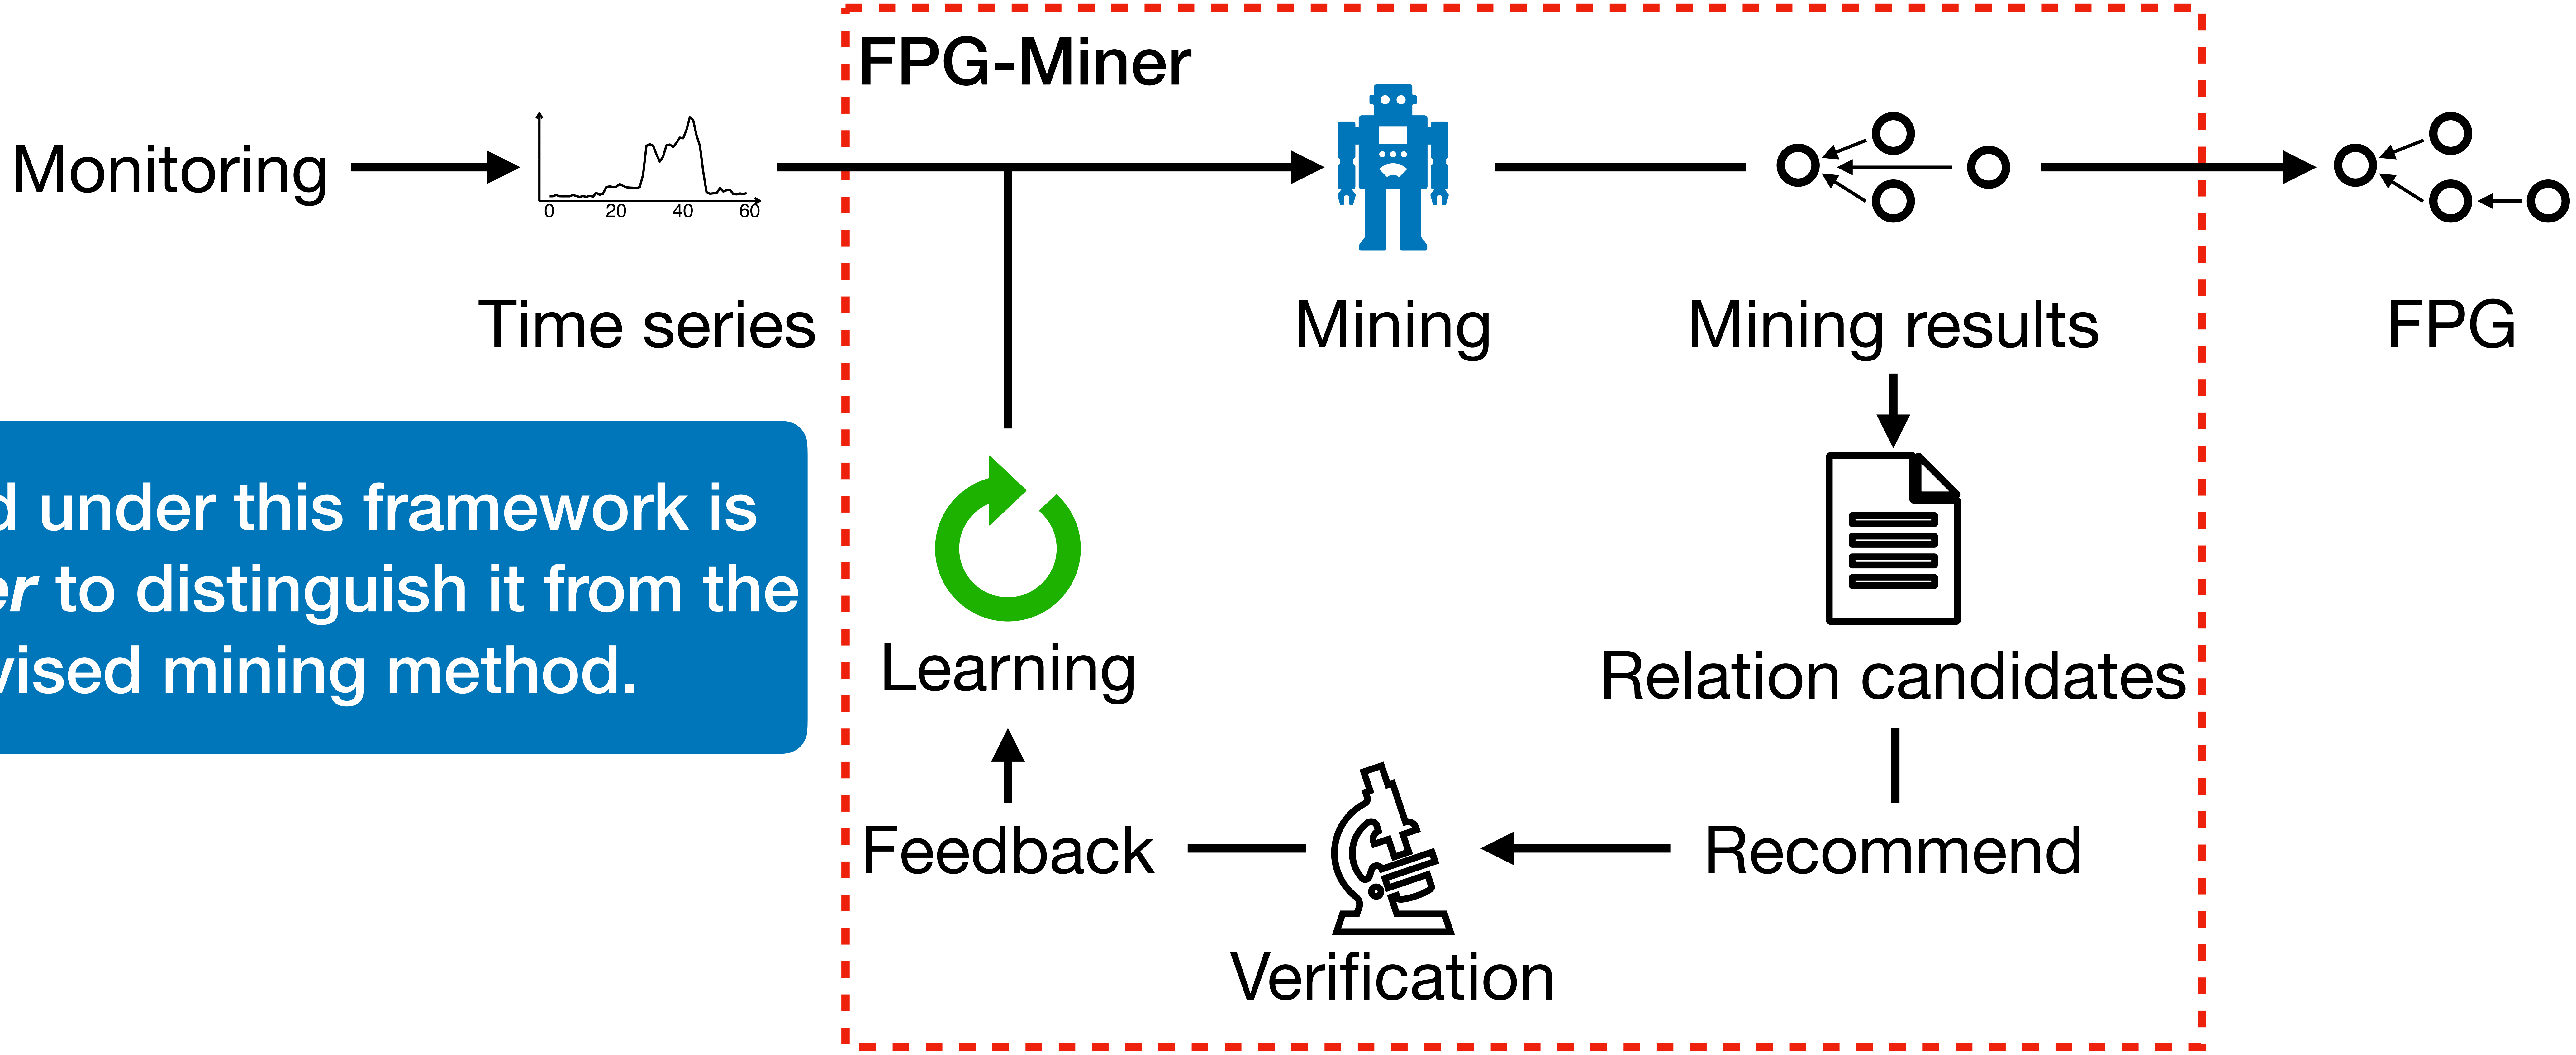

Each method under this framework is named a *miner* to distinguish it from the unsupervised mining method.

Methodology

FPG-Miner: Recommendation Framework

- Ideally, the recommendation flow contains correct relations (including reversed ones) in preference to incorrect ones
  - The process can stop after operators confront the first incorrect recommendation

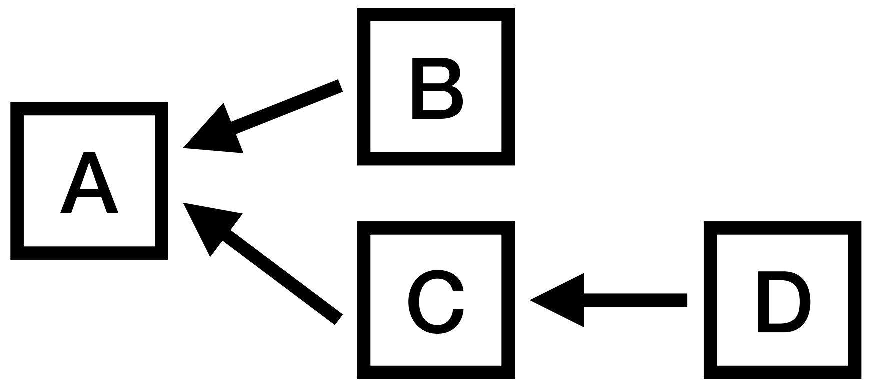

An FPG

	Recommendation Flow				
An ideal flow	B->A	C->A	D->C	D->A	Stop

Methodology

FPG-Miner: Recommendation Framework

- Each practical recommendation flow may mix correct and incorrect relations.
    - The process has to continue after the first incorrect recommendation arises

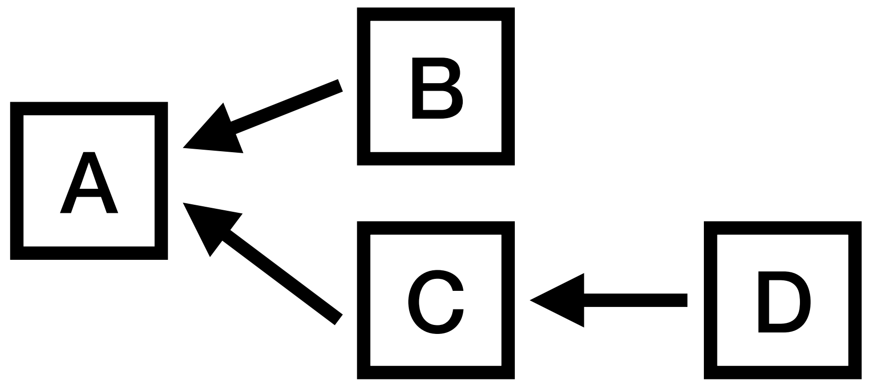

An FPG

	Recommendation Flow					
An ideal flow	B->A	C->A	D->C	D->A	Stop	
A practical flow	B->A	C->A	D->A	C->B	D->B	D->C

Methodology

FPG-Miner: Recommendation Framework

- Learn from mistakes to lessen incorrect recommendations
  - Make use of the feedback during the journey to obtain the whole graph

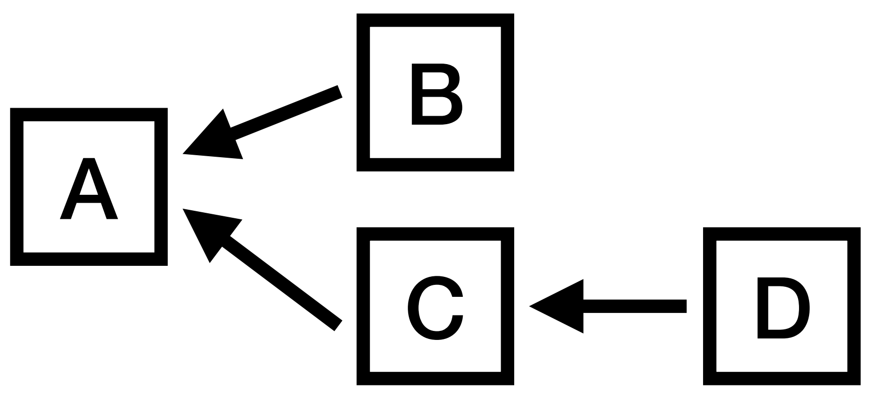

An FPG

	Recommendation Flow					
An ideal flow	B->A	C->A	D->C	D->A	Stop	
A practical flow	B->A	C->A	D->A	C->B	D->B	D->C
Learn from mistakes	B->A	C->A	D->A	D->B	D->C	C->B

Similar

Methodology

FPG-Miner: Recommendation Framework

- Inspired by active learning research, recommending uncertain relations may bring more information for long-term benefit

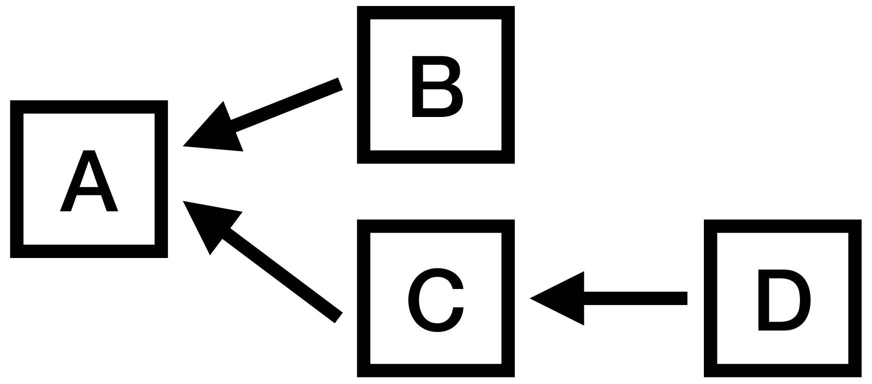

An FPG

Recommendation Flow						
An ideal flow	B->A	C->A	D->C	D->A	Stop	
A practical flow	B->A	C->A	D->A	C->B	D->B	D->C
Learn from mistakes	B->A	C->A	D->A	D->B	D->C	C->B
Make mistakes to learn	B->A	C->A	C->B	D->C	D->A	D->B

Similar Not Similar

Similar

Methodology

FPG-Miner: Recommendation Framework

- Inspired by active learning research, recommending uncertain relations may bring more information for long-term benefit
- Recommendation strategies
  - Confidence-first
  - Uncertainty-first
  - Mixed
  - Random

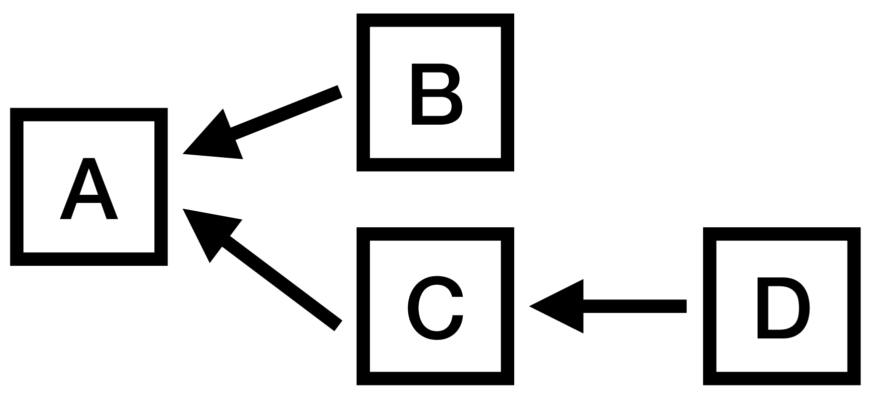

An FPG

	Recommendation Flow					
An ideal flow	B->A	C->A	D->C	D->A	Stop	
A practical flow	B->A	C->A	D->A	C->B	D->B	D->C
Learn from mistakes	B->A	C->A	D->A	D->B	D->C	C->B
Make mistakes to learn	B->A	C->A	C->B	D->C	D->A	D->B

Similar Not Similar

Similar

Methodology

FPG-Miner: Recommendation Framework

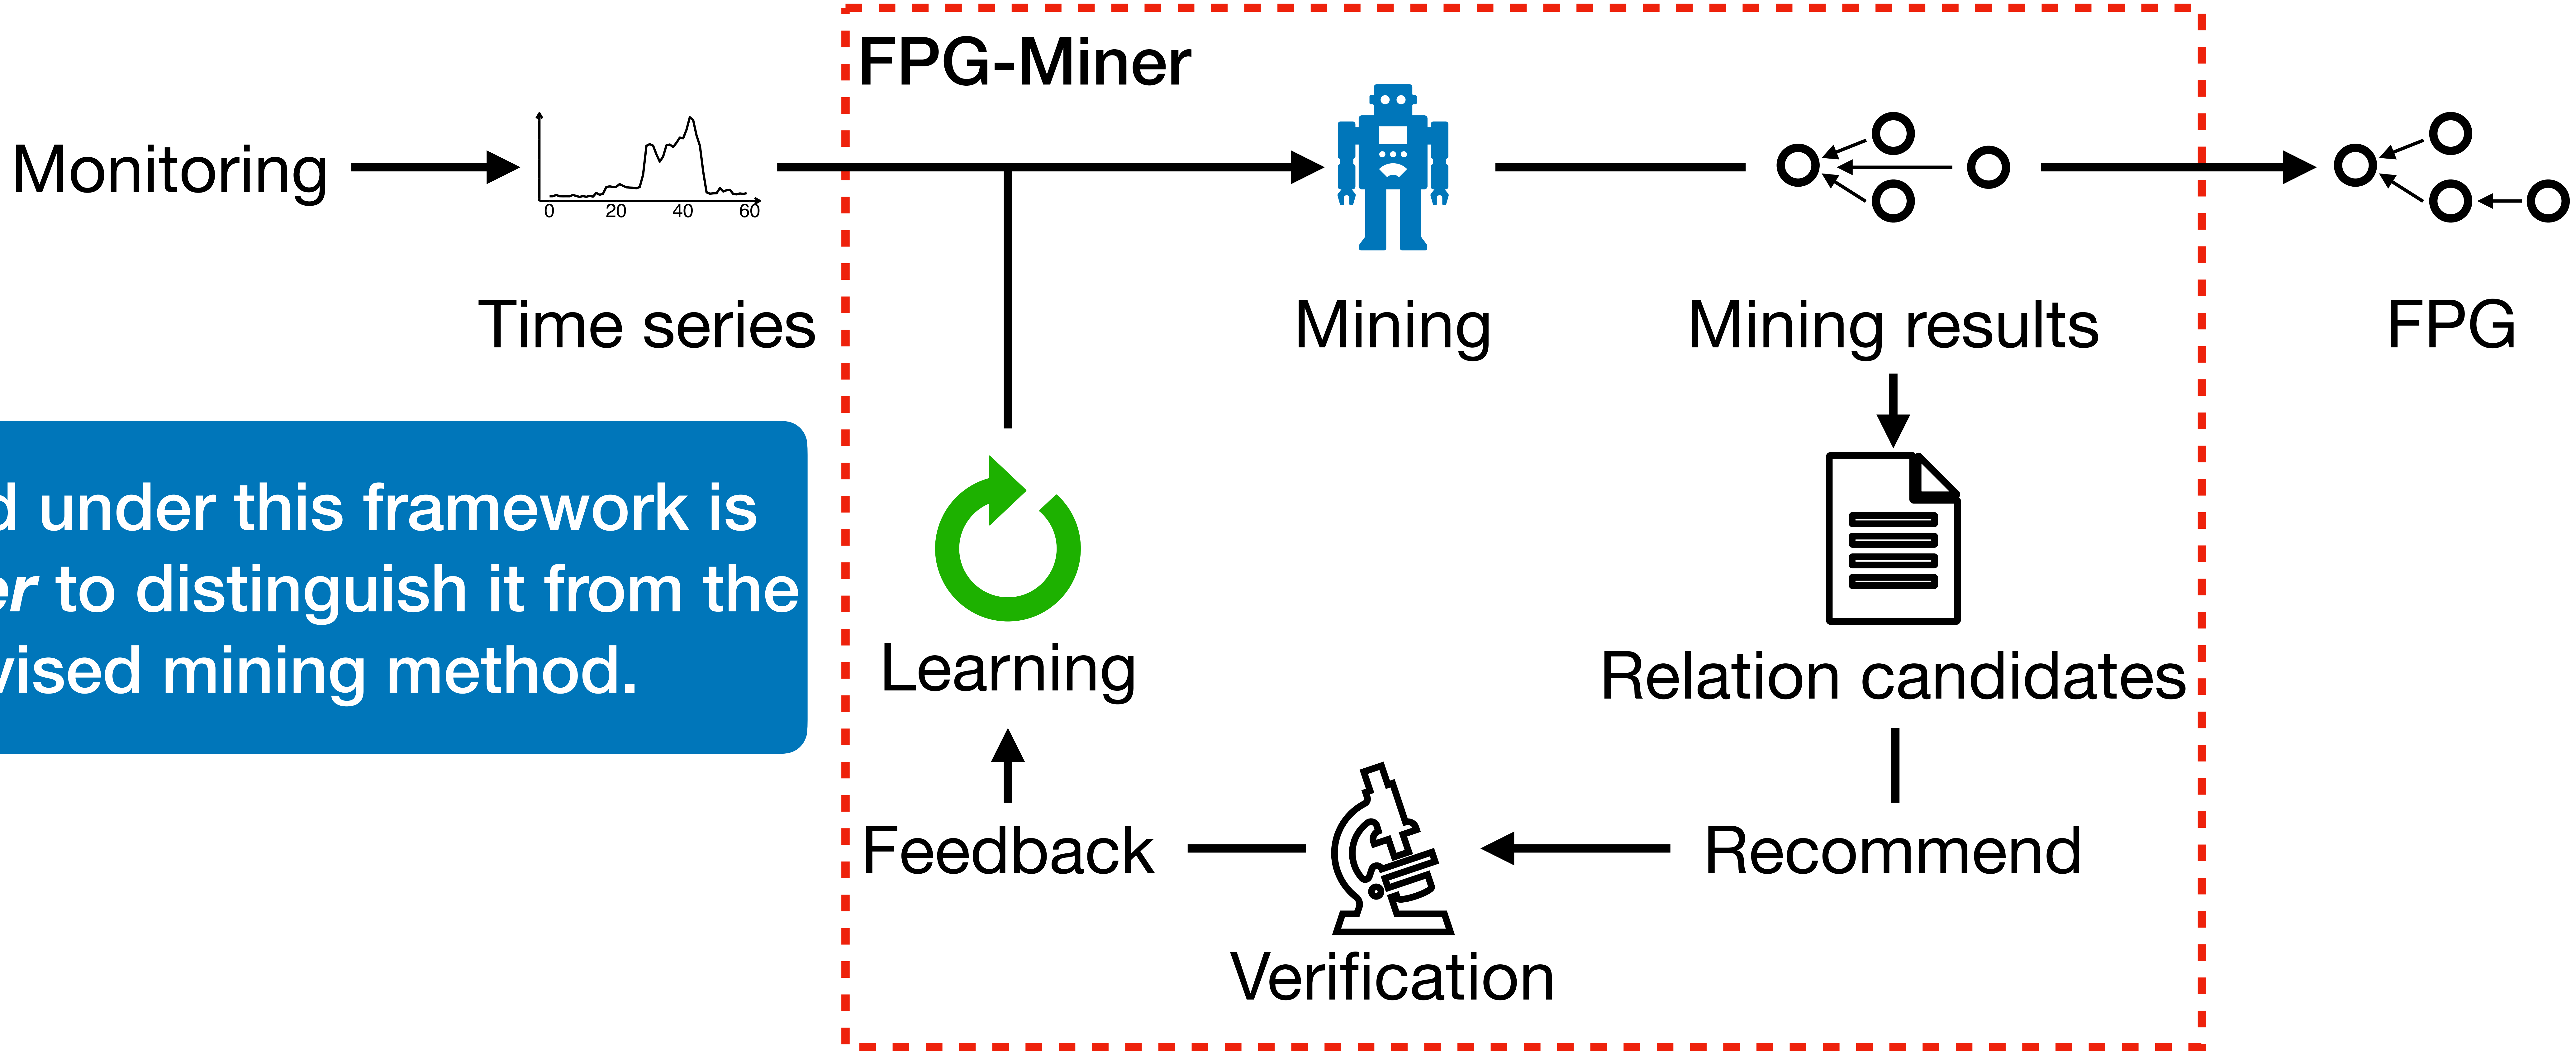

Each method under this framework is named a *miner* to distinguish it from the unsupervised mining method.

Methodology

FPG-Miner: Recommendation Framework

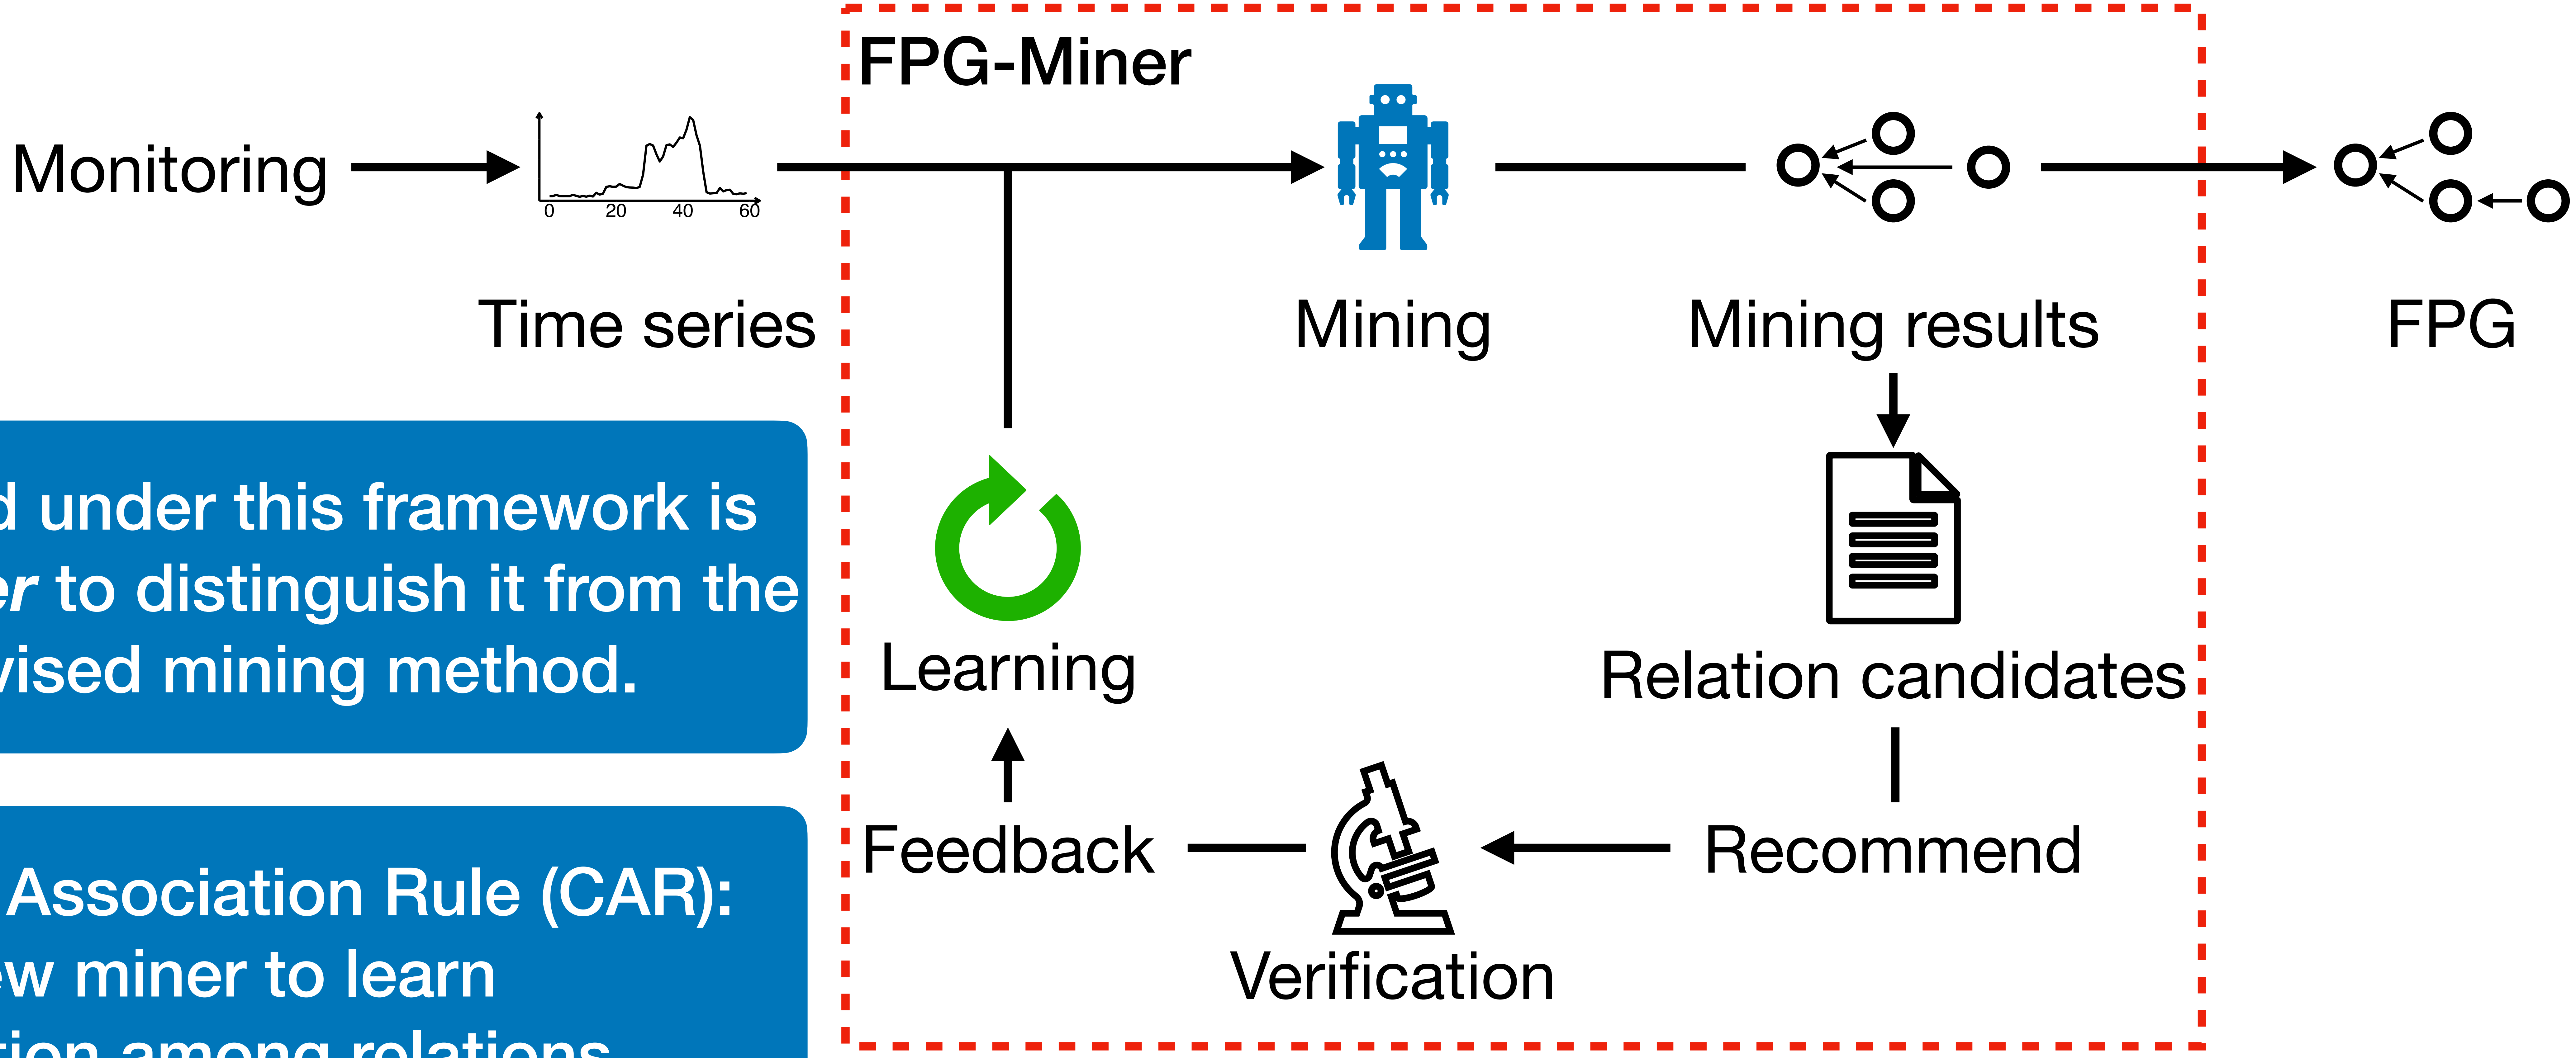

Each method under this framework is named a *miner* to distinguish it from the unsupervised mining method.

Continuous Association Rule (CAR): a new miner to learn the relation among relations

Methodology

Continuous Association Rule (CAR)

- Capture the co-fluctuations of two metrics when one changes large enough

    1. Partition each time series into sliding windows

    2. Focus on outliers and ignore natural fluctuations

    3. Count Support and Coverage based on co-fluctuations

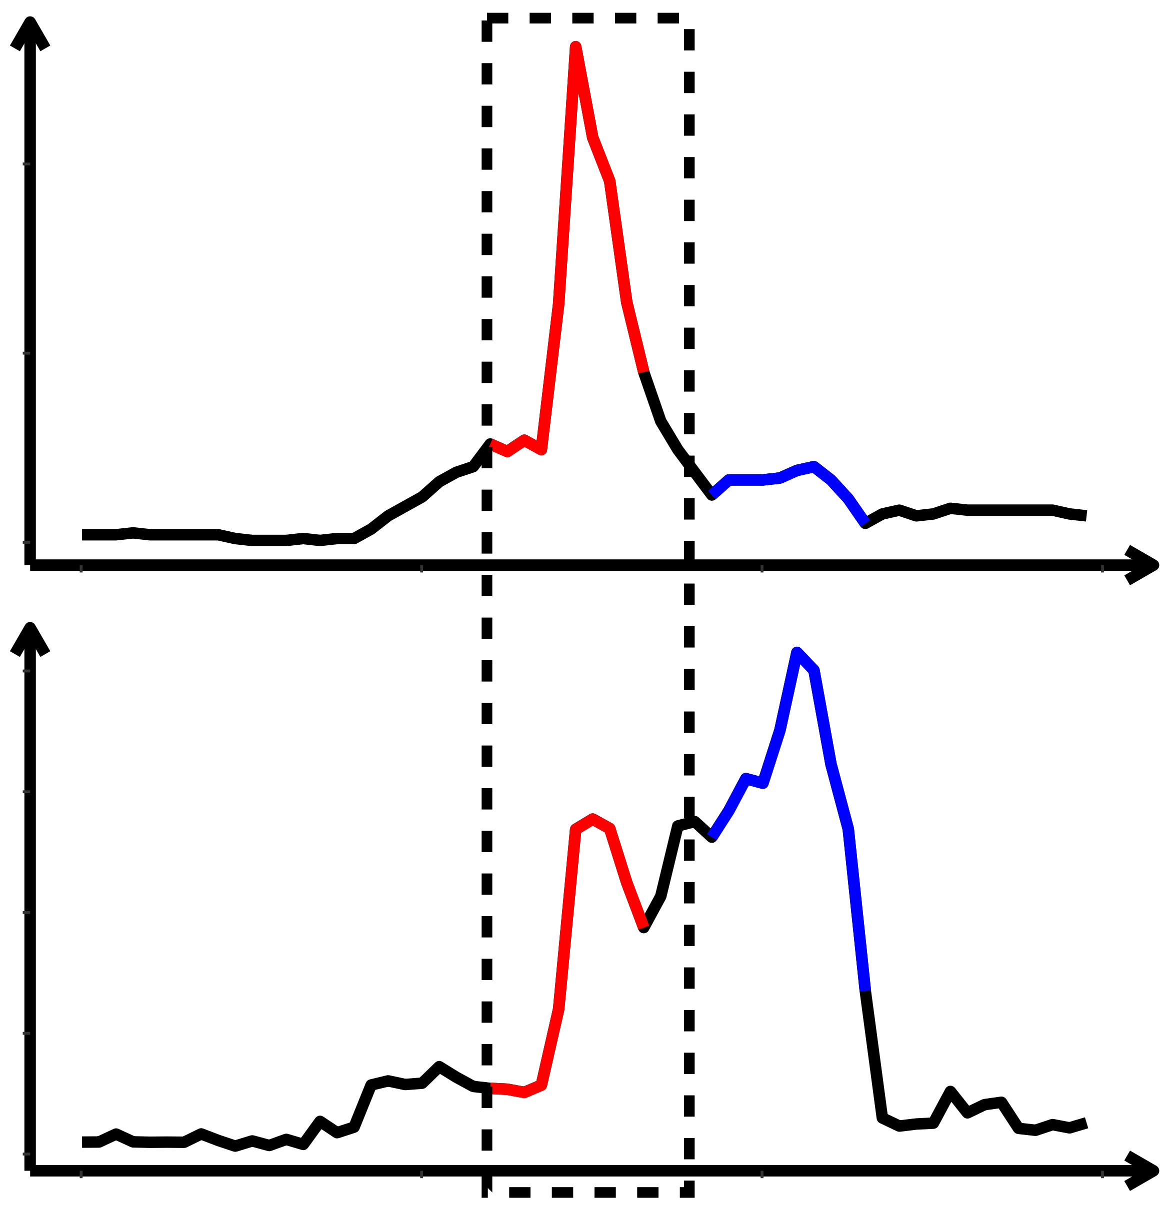

Methodology

Continuous Association Rule (CAR)

- Capture the co-fluctuations of two metrics when one changes large enough

- Calculate features based on Support and Coverage

Feature (A->B)	Definition
Coverage	P(A)
Support	P(AB)
Consequence coverage	P(B)
Confidence	P(B|A)
Reversed confidence	P(A|B)
Lift	P (AB)/ [P (A)P (B)]
KULC	[P (B|A) + P (A|B)] /2

Methodology

Continuous Association Rule (CAR)

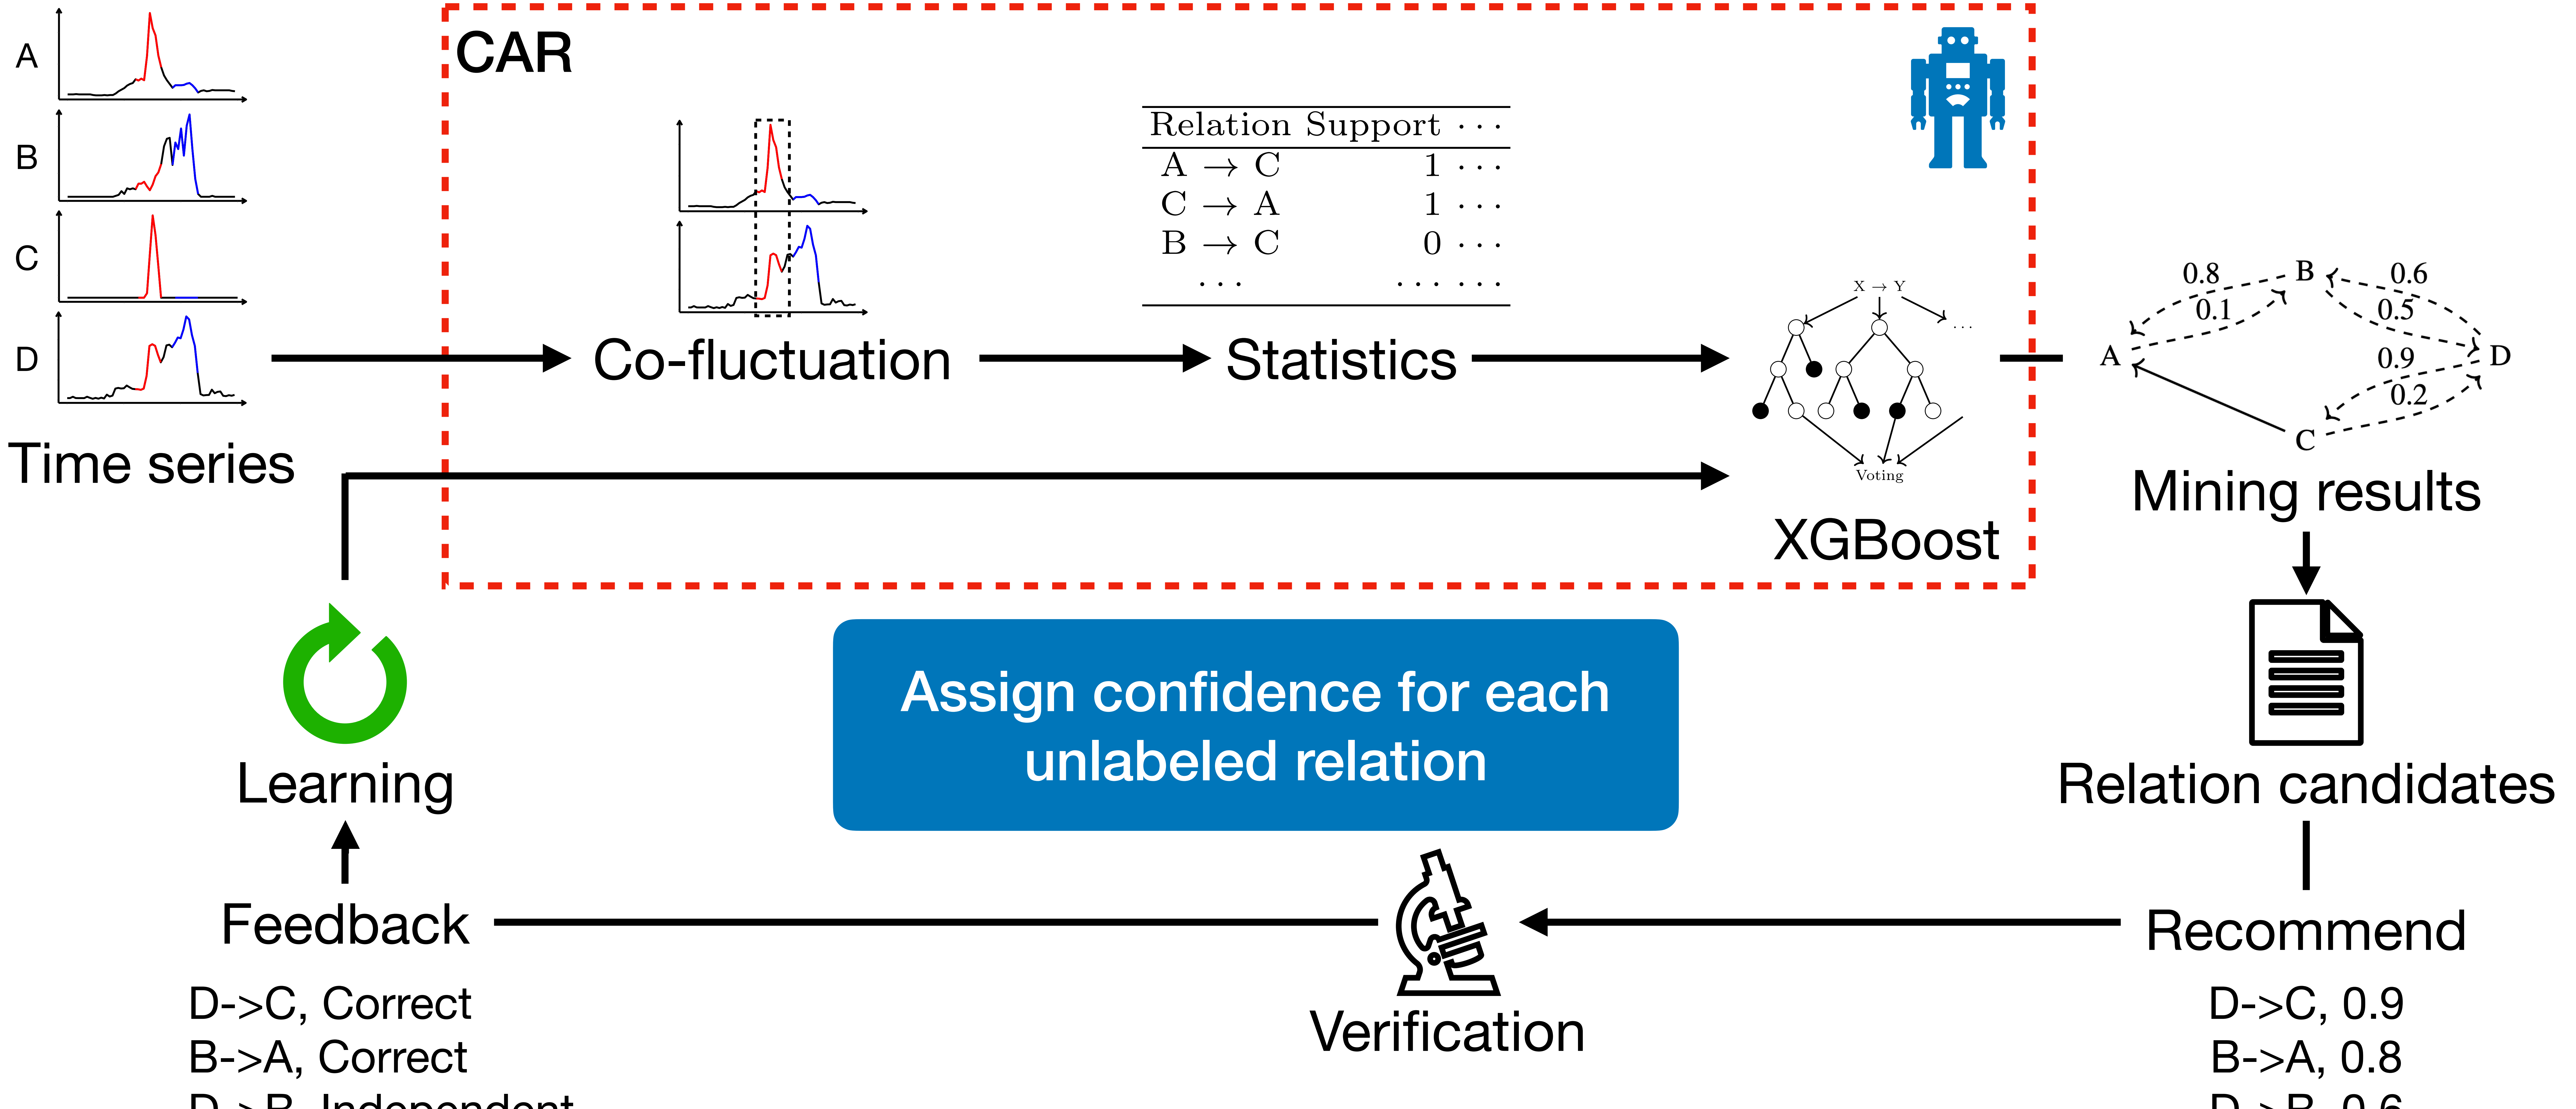

Experiment

- RQ2: Will a mining method perform better based on active learning than in an unsupervised manner?
- RQ3: How does CAR perform compared with other miners under the framework of FPG-Miner?
- RQ4: Are there some relations more important than other ones?

Experiment

Evaluation Metrics

- T@k
    - the number of times it takes a miner to recommend k correct relations
    - Precision@q = k / T@k, where q = T@k
- AUC (Area Under Curve)
    - calculate the k ~ T@k curve's area, compared with the ideal process

Experiment

RQ2: Improvement with Active Learning

Miner	Learning	$\mathscr{D}_{OD}$						$\mathscr{D}_{TN}$					
		AUC	T@k					AUC	T@k				
			10%	20%	30%	50%	100%		10%	20%	30%	50%	100%
PC-gauss	Without	0.589	47	**99**	161	291	490	**0.639**	**106**	**248**	**407**	**703**	**1483**
	With	**0.648**	**41**	102	**145**	**237**	**489**	0.619	112	259	428	746	1485
GES	Without	**0.690**	**28**	**76**	**125**	**214**	490	**0.651**	**90**	**223**	**370**	**684**	**1479**
	With	0.639	46	87	142	244	**488**	0.636	128	227	401	720	1483
NRI	Without	0.589	74	118	175	273	488	0.658	138	291	407	633	1485
	With	**0.741**	**53**	**83**	**113**	**192**	**478**	**0.731**	**85**	**177**	**285**	**575**	**1482**

Experiment

RQ2: Improvement with Active Learning

- Active learning can enhance some but not all relation mining methods
  - e.g., feedback may break the intrinsic mechanism of GES, as adding an extra relation ($A_1 \rightarrow A_{32}$) in $\mathscr{D}_{TN}$ will increase the score of GES

Miner	Learning	$\mathscr{D}_{OD}$						$\mathscr{D}_{TN}$					
		AUC	T@k					AUC	T@k				
			10%	20%	30%	50%	100%		10%	20%	30%	50%	100%
PC-gauss	Without	0.589	47	**99**	161	291	490	**0.639**	**106**	**248**	**407**	**703**	**1483**
	With	**0.648**	**41**	102	**145**	**237**	**489**	0.619	112	259	428	746	1485
GES	Without	**0.690**	**28**	**76**	**125**	**214**	490	**0.651**	**90**	**223**	**370**	**684**	**1479**
	With	0.639	46	87	142	244	**488**	0.636	128	227	401	720	1483
NRI	Without	0.589	74	118	175	273	488	0.658	138	291	407	633	1485
	With	**0.741**	**53**	**83**	**113**	**192**	**478**	**0.731**	**85**	**177**	**285**	**575**	**1482**

Experiment

RQ3: Performance of CAR

CAR recommends correct relations faster than baselines

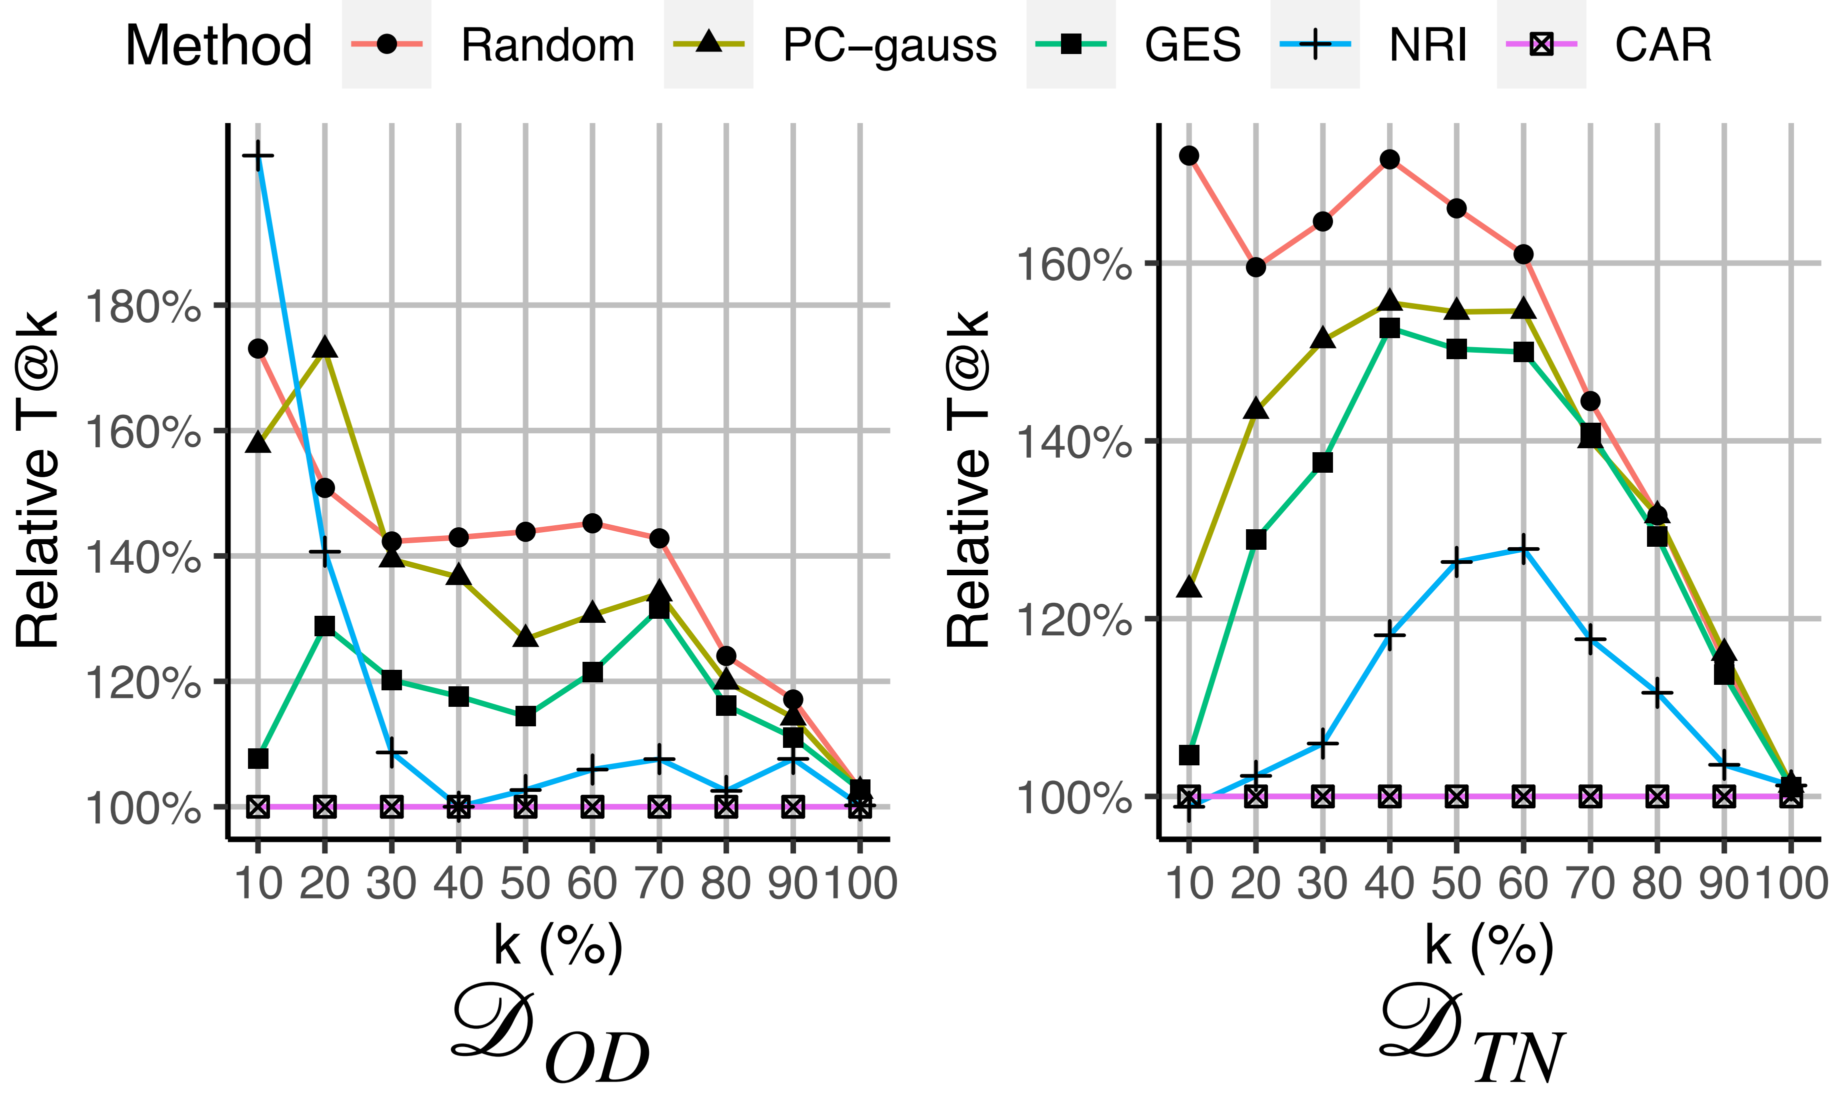

Miner	$\mathscr{D}_{OD}$						$\mathscr{D}_{TN}$					
	AUC	T@k					AUC	T@k				
		10%	20%	30%	50%	100%		10%	20%	30%	50%	100%
Random	0.617	45	89	148	269	490	0.617	148	276	443	756	1480
PC-gauss	0.648	41	102	145	237	489	0.639	106	248	407	703	1483
GES	0.690	28	76	125	214	490	0.651	90	223	370	684	1479
NRI	0.741	53	83	113	192	478	0.731	**85**	177	285	575	1482
CAR	**0.774**	**26**	**59**	**104**	**187**	**477**	**0.792**	86	**173**	**269**	**455**	**1464**

Experiment

RQ3: Performance of CAR

- AR, a variant of CAR
  1. Convert time series as events by outlier detection
  2. Count classic Support, Coverage, and the other features
- The proposed feature extraction shortens T@20% by 26% and 25% on $\mathscr{D}_{OD}$ and $\mathscr{D}_{TN}$, respectively

Miner	$\mathscr{D}_{OD}$						$\mathscr{D}_{TN}$					
	AUC	**T@k**					**AUC**	**T@k**				
		10%	**20%**	**30%**	**50%**	**100%**		**10%**	**20%**	**30%**	**50%**	**100%**
CAR	**0.774**	**26**	**59**	104	**187**	**477**	**0.792**	**86**	**173**	**269**	**455**	**1464**
AR	0.738	39	80	**103**	198	489	0.678	123	232	356	573	1484

Experiment

RQ4: Importance of Different Relations

- Take AC@5 as the quality indicator of the mined graph
  - AC@5 refers to the probability that the top 5 results given by MicroCause (IWQoS'20) include the root cause metrics
  - measured on 99 high AAS faults

Experiment

RQ4: Importance of Different Relations

- Take AC@5 as the quality indicator of the mined graph
  - AC@5 refers to the probability that the top 5 results given by MicroCause (IWQoS'20) include the root cause metrics
  - measured on 99 high AAS faults

- At least 33.8% of the relations seem neither helpful nor harmful to MicroCause in this case study

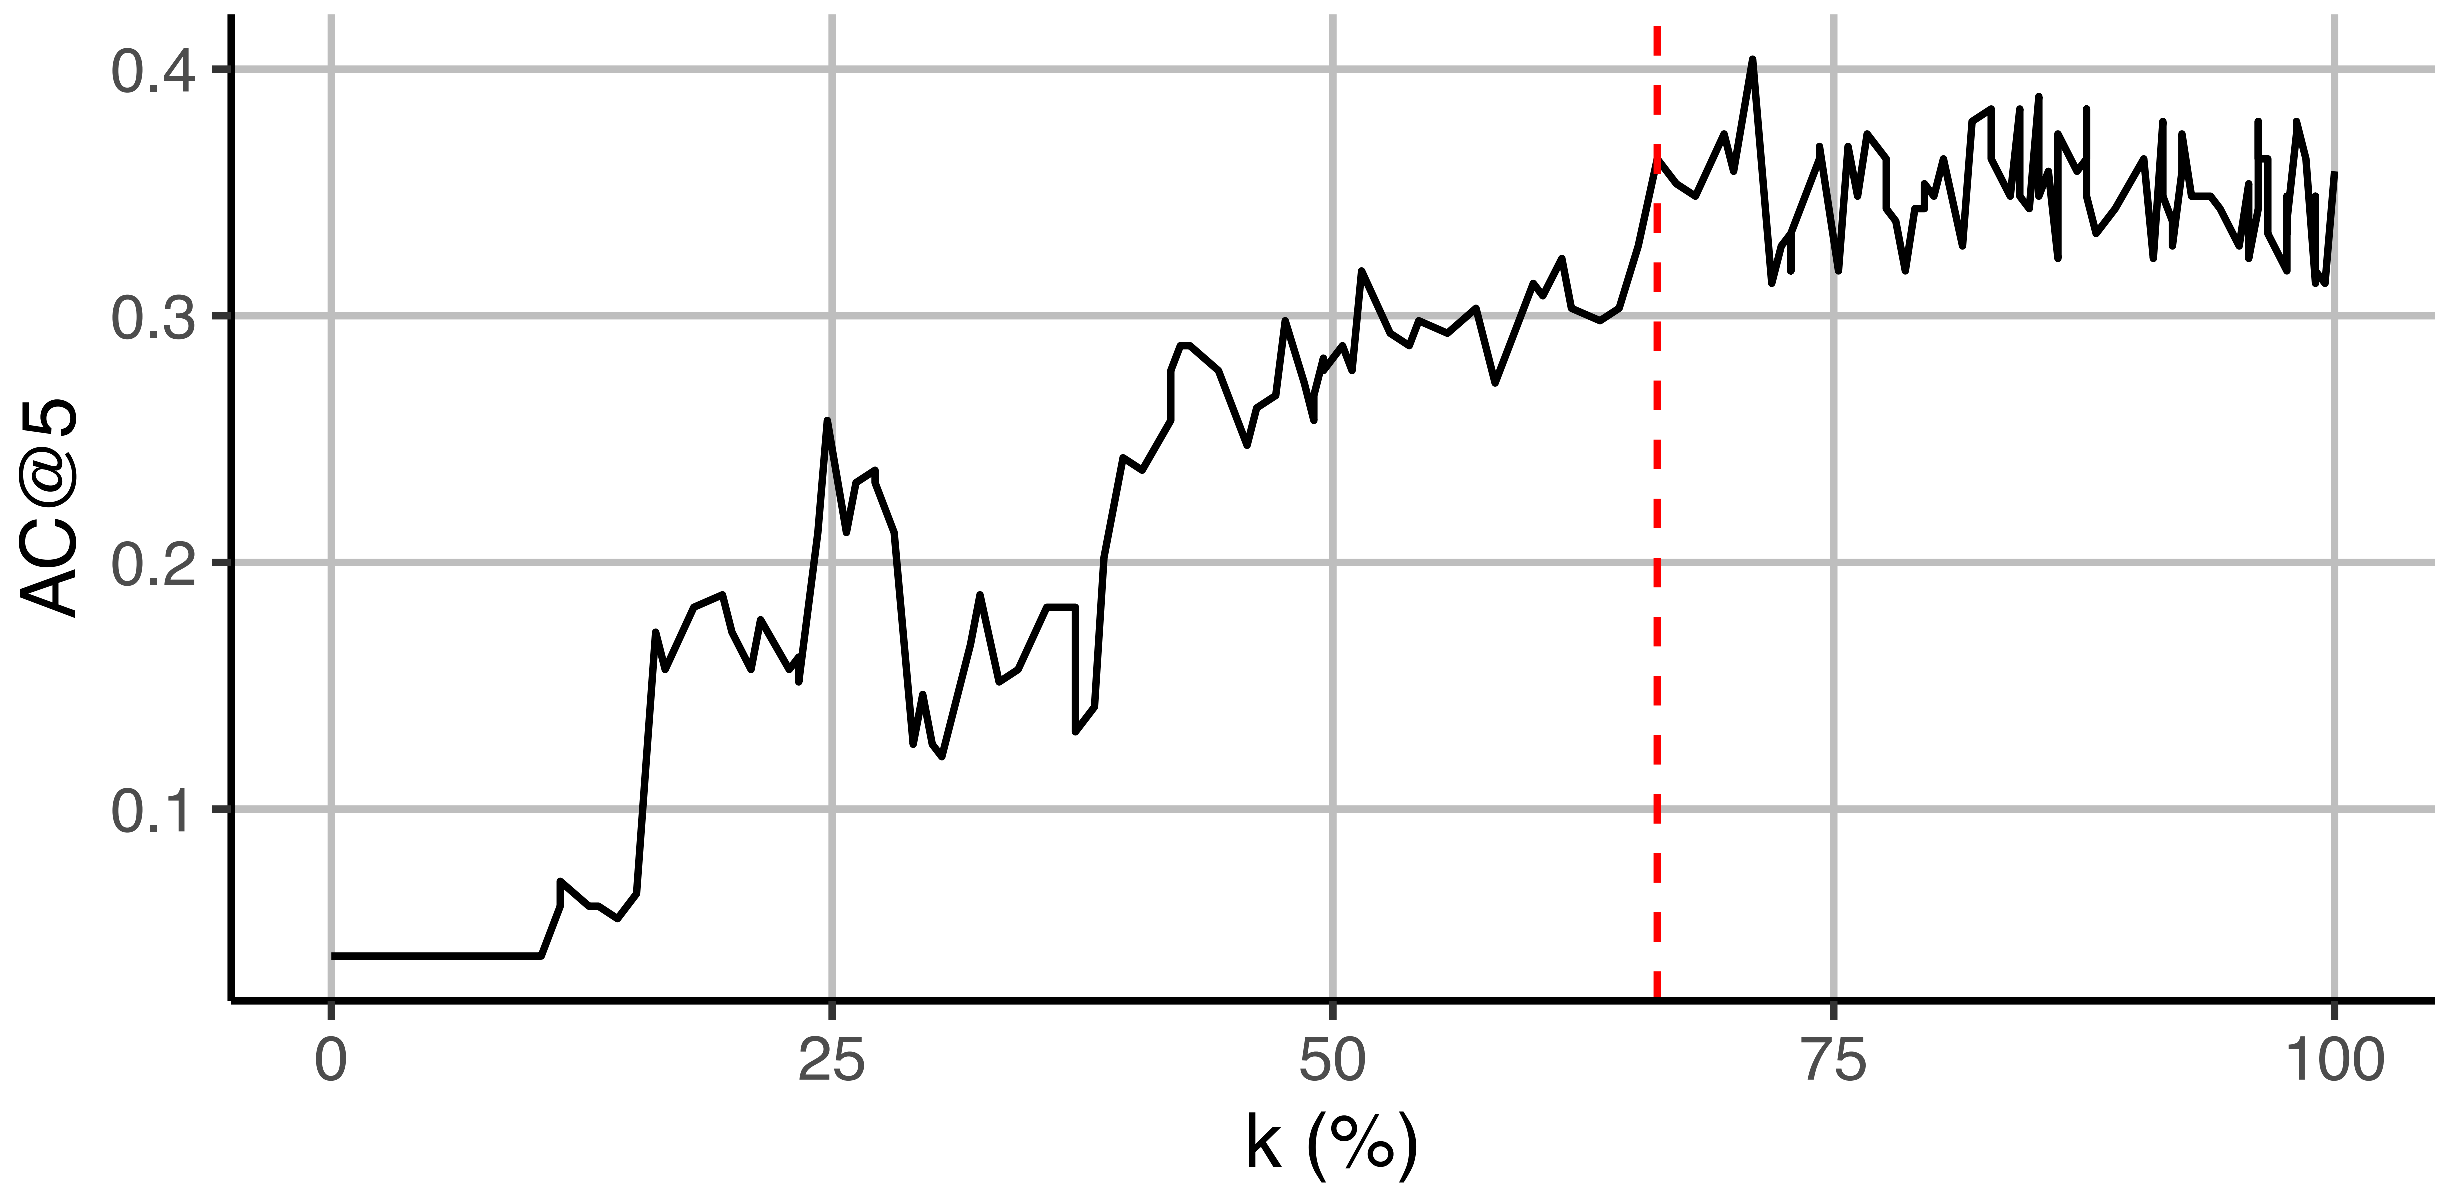

Experiment

RQ4: Importance of Different Relations

- Take AC@5 as the quality indicator of the mined graph
    - AC@5 refers to the probability that the top 5 results given by MicroCause (IWQoS'20) include the root cause metrics
    - measured on 99 high AAS faults

- **At least 33.8% of the relations seem neither helpful nor harmful to MicroCause in this case study**
- **It can be helpful to recommend correct and important relations faster**

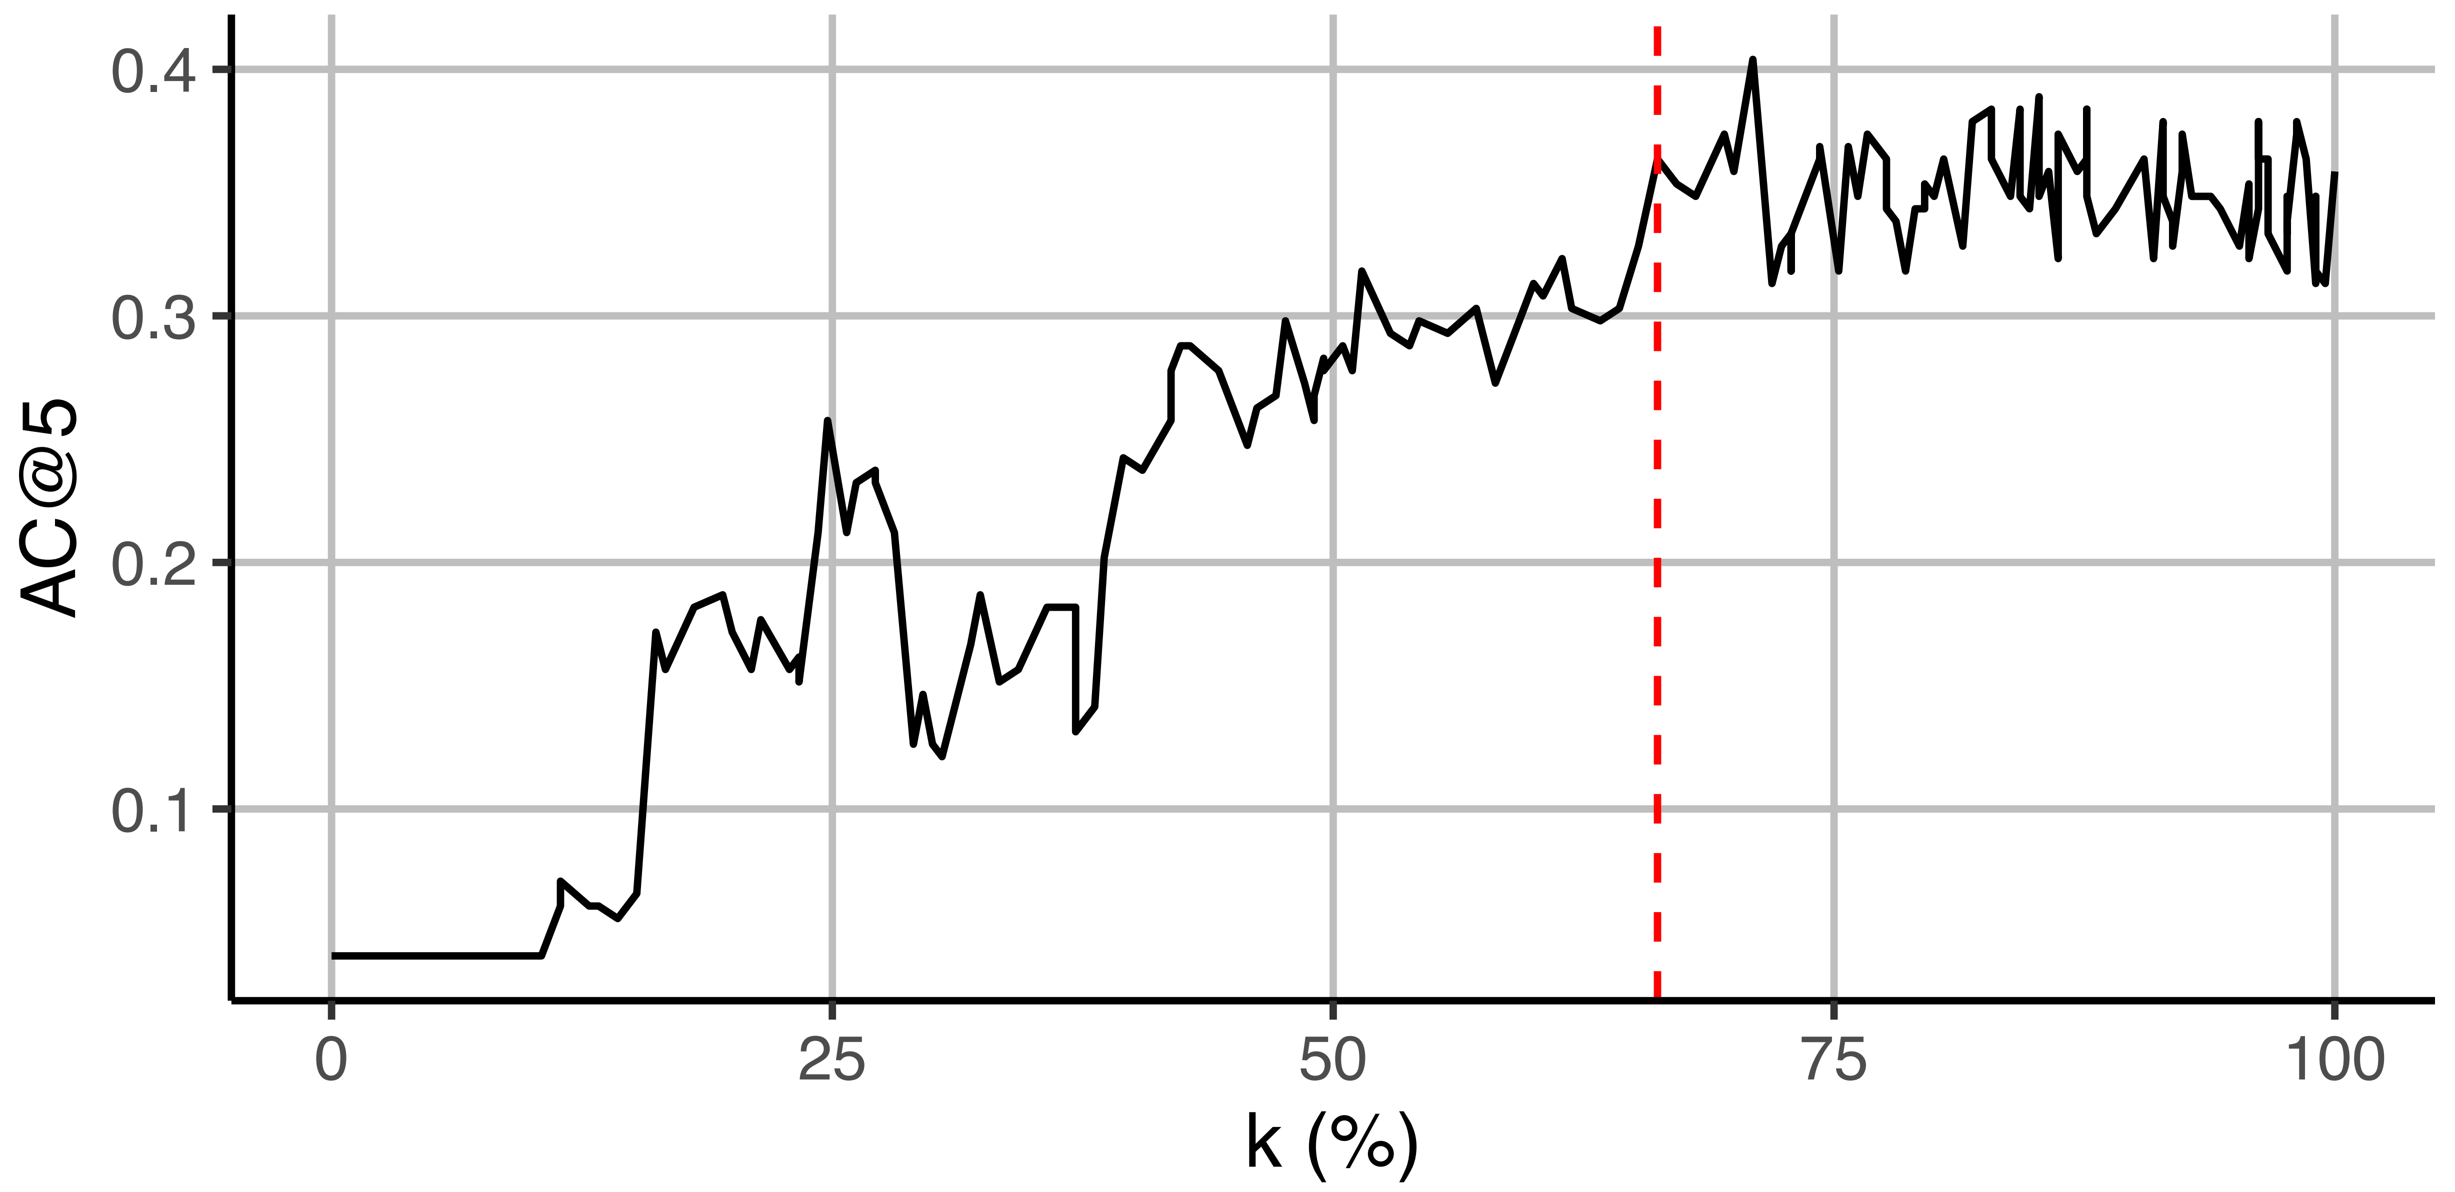

Conclusion

Background
Empirical Study
Methodology
Experiment
Conclusion

Conclusion

Mining Fluctuation Propagation Graph (FPG) with Active Learning

Empirical study with
two real-world datasets

- (RQ1) Each existing mining method in the experiment suffers from either a low precision or low discovery ability on both datasets

Conclusion

Mining Fluctuation Propagation Graph (FPG) with Active Learning

Empirical study with two real-world datasets

- (RQ1) Each existing mining method in the experiment suffers from either a low precision or low discovery ability on both datasets

FPG-Miner: an FPG construction framework with active learning

- (RQ2) Active learning can enhance some but not all relation mining methods
- (RQ4) It can be helpful to recommend correct and important relations faster

Conclusion

Mining Fluctuation Propagation Graph (FPG) with Active Learning

Empirical study with two real-world datasets

- (RQ1) Each existing mining method in the experiment suffers from either a low precision or low discovery ability on both datasets

FPG-Miner: an FPG construction framework with active learning

- (RQ2) Active learning can enhance some but not all relation mining methods
- (RQ4) It can be helpful to recommend correct and important relations faster

CAR: an implementation for FPG-Miner

- (RQ3) CAR recommends correct relations faster than baselines

Thanks for Listening

lmj18@mails.tsinghua.edu.cn