

WatDiv: How to Tune-Up Your RDF Data Management System

Güneş Aluç

Olaf Hartig

M. Tamer Özsu

Khuzaima Daudjee



This presentation is sponsored in part by the *Linked Data Benchmark Council* (LDBC).

Questions

- Which of the existing SPARQL benchmarks, if any, should I use to diagnose (and fix) potential problems with the physical design of my system?
- How can I use the Waterloo SPARQL Diversity Test Suite (WatDiv) where existing benchmarks fall short?

Contributions

Waterloo SPARQL Diversity Test Suite
(WatDiv)



<http://db.uwaterloo.ca/watdiv/>

Contributions

Waterloo SPARQL Diversity Test Suite
(WatDiv)



<http://db.uwaterloo.ca/watdiv/>


Measures to Evaluate Diversity
in SPARQL Workloads

Structural

Data-driven

Contributions

Waterloo SPARQL Diversity Test Suite
(WatDiv)



<http://db.uwaterloo.ca/watdiv/>

Measures to Evaluate Diversity
in SPARQL Workloads

Structural

Data-driven

Analysis of WatDiv and Popular
SPARQL Benchmarks

Contributions

Waterloo SPARQL Diversity Test Suite
(WatDiv)



<http://db.uwaterloo.ca/watdiv/>

Measures to Evaluate Diversity
in SPARQL Workloads


Structural

Data-driven


Analysis of WatDiv and Popular
SPARQL Benchmarks

Debugging with WatDiv


Structural Features



Structural Features




Structural Features




Structural Features

[Triple Pattern Count]




Structural Features

[Join Vertex Count]




Structural Features

[Join Vertex Degree]




Structural Features

[Join Vertex Degree]




Structural Features

[Join Vertex Degree]




Structural Features






Join Vertex Count	Mean Join Vertex Degree
8	2.0

Structural Features






Join Vertex Count	Mean Join Vertex Degree
8	2.0
5	2.6

Structural Features



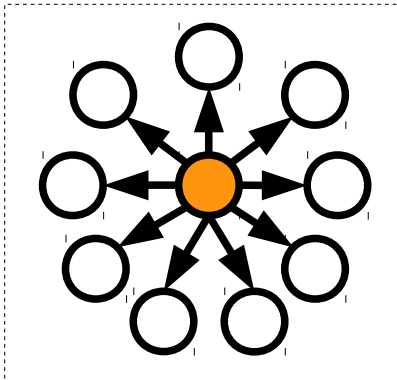



Join Vertex Count	Mean Join Vertex Degree
8	2.0
5	2.6
3	~3.7

Structural Features



Join Vertex Count	Mean Join Vertex Degree
8	2.0
5	2.6
3	~3.7
2	5.0


Structural Features



Join Vertex Count	Mean Join Vertex Degree
8	2.0
5	2.6
3	~3.7
2	5.0
1	9.0

Structural Features


[Join Vertex Type]



SS⁺ Type

Structural Features


[Join Vertex Type]



OO⁺ Type

Structural Features


[Join Vertex Type]



SO⁺ Type


How Diverse are SPARQL Benchmarks?

[Triple Pattern Count]




How Diverse are SPARQL Benchmarks?

[Triple Pattern Count]




How Diverse are SPARQL Benchmarks?

[Join Vertex Count]



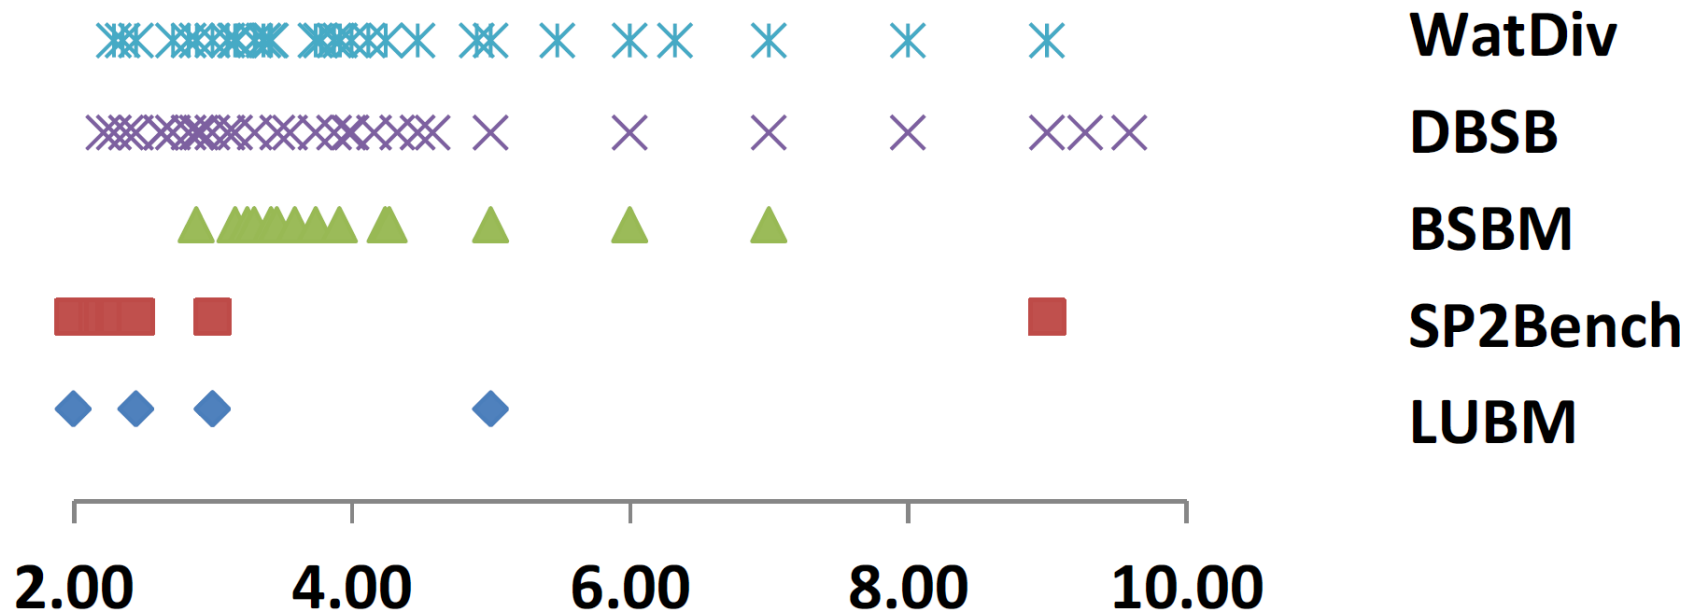
How Diverse are SPARQL Benchmarks?

[Join Vertex Count]




How Diverse are SPARQL Benchmarks?

[Join Vertex Degree – mean]



How Diverse are SPARQL Benchmarks?

[Join Vertex Degree – mean]



How Diverse are SPARQL Benchmarks?

[Join Vertex Type – % Queries w/in Workload]

	SS ⁺	OO ⁺	SO ⁺
LUBM	78.6 %	0.0 %	42.9 %
SP ² Bench	81.0 %	33.3 %	57.1 %
BSBM	84.8 %	5.6 %	52.8 %
DBSB	41.1 %	4.4 %	5.4 %
WatDiv	61.3 %	18.0 %	61.3 %

How Diverse are SPARQL Benchmarks?

[Join Vertex Type – % Queries w/in Workload]


	SS ⁺	OO ⁺	SO ⁺
LUBM	78.6 %	0.0 %	42.9 %
SP ² Bench	81.0 %	33.3 %	57.1 %
BSBM	84.8 %	5.6 %	52.8 %
DBSB	41.1 %	4.4 %	5.4 %
WatDiv	61.3 %	18.0 %	61.3 %

Data-Driven Features

- Why are data-driven query features important?
 - *Why are structural features not sufficient?*
 - *Why is analysis based purely on the data not sufficient?*

Data-Driven Features


[Result Cardinality]



	?a	...	?g
1			
2			
.			
.			
.			
k			

Data-Driven Features

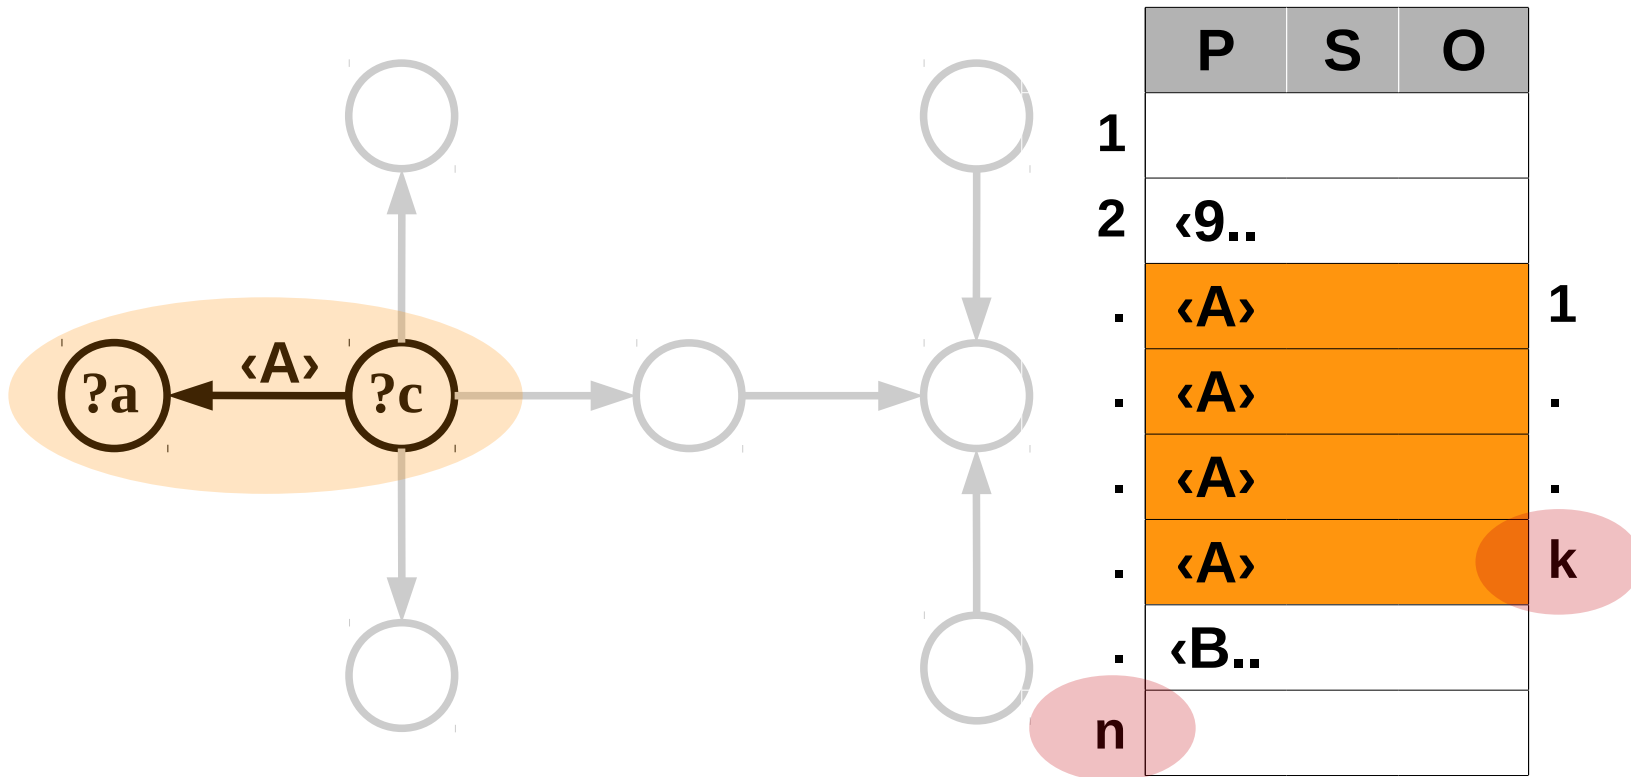
[Filtered Triple Pattern (f-TP) Selectivity]



P	S	O
<9..		
<A>		
<A>		
<A>		
<A>		
<B..		

Data-Driven Features


[Filtered Triple Pattern (f-TP) Selectivity]



$$\text{f-TP Selectivity} = k / n$$

Data-Driven Features


[f-TP Selectivity, **BGP-Restricted**]



P	S	O
		<9..
		<A>
		<A>
		<A>
		<A>
		<B..

Data-Driven Features

[f-TP Selectivity, **BGP-Restricted**]




P	S	O
	<9..	
	<A>	
	<A>	
	<A>	
	<A>	
	<B..	

?a	...	?g

BGP-Restricted
f-TP Selectivity = |blue| / |orange|

Data-Driven Features

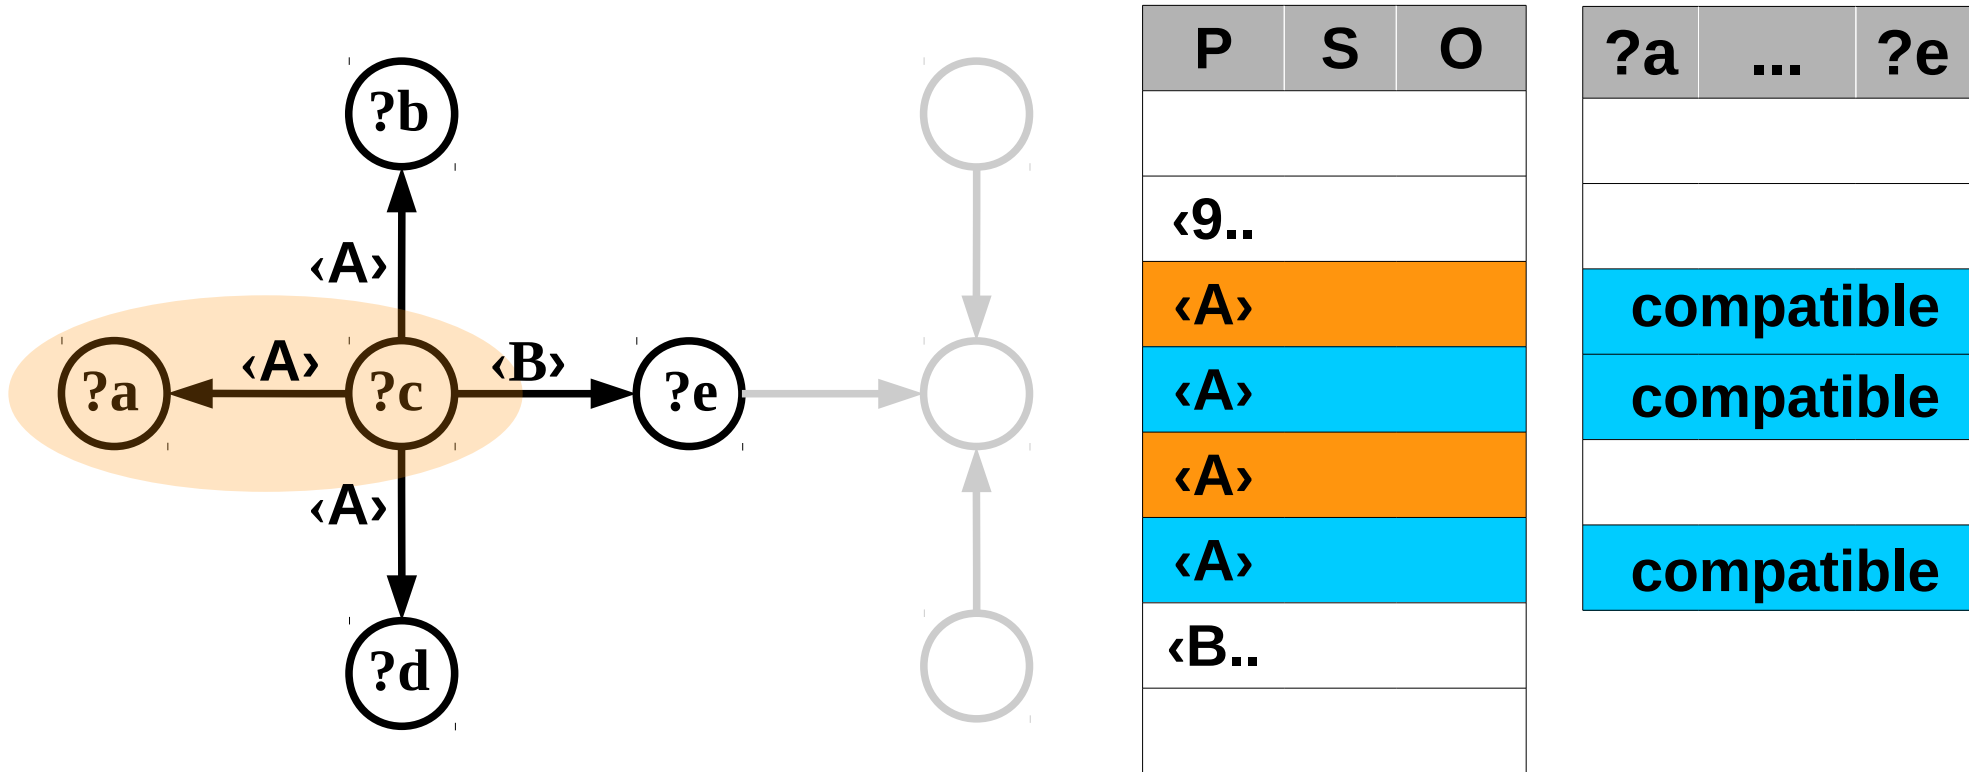
[f-TP Selectivity, **Join-Restricted**]



P	S	O
		<9..
		<A>
		<A>
		<A>
		<A>
		<B..

Data-Driven Features

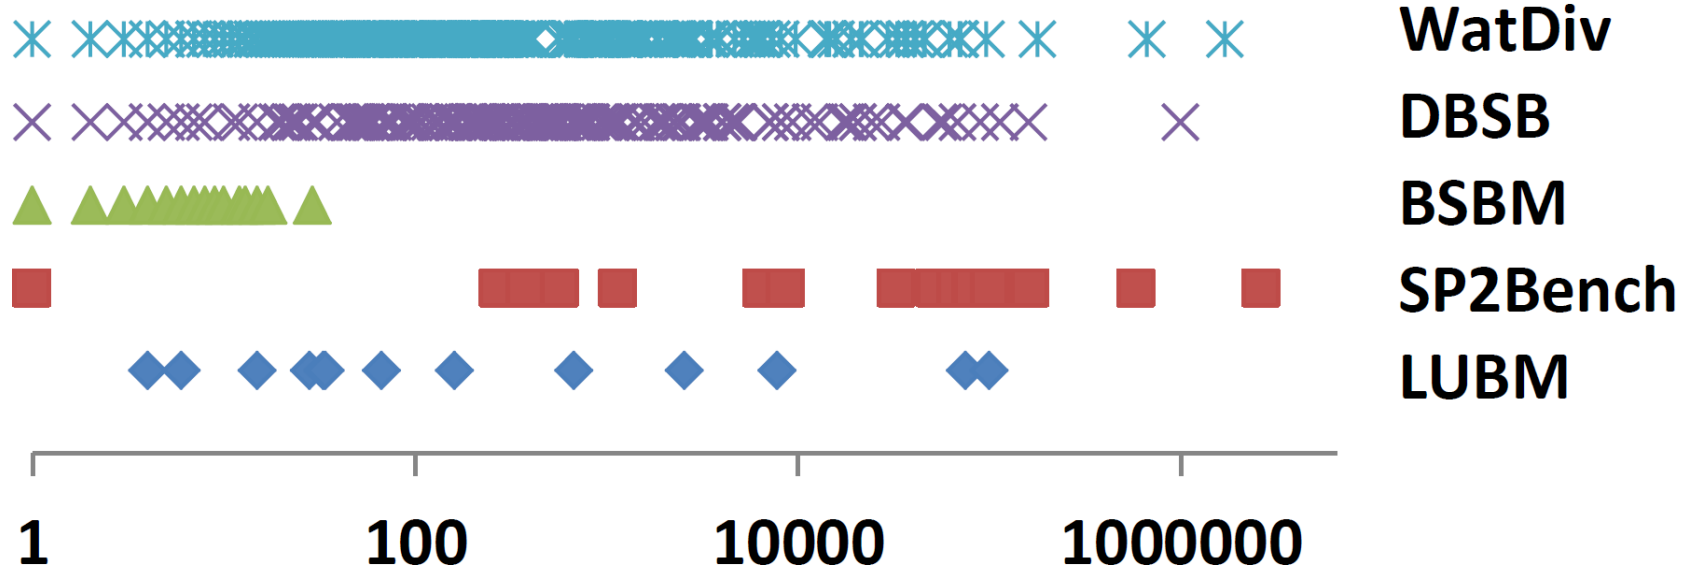
[f-TP Selectivity, **Join-Restricted**]



Join-Restricted
f-TP Selectivity = |blue| / |orange|


How Diverse are SPARQL Benchmarks?

[Result Cardinality]




How Diverse are SPARQL Benchmarks?

[Result Cardinality]




How Diverse are SPARQL Benchmarks?

[f-TP Selectivity – mean]




How Diverse are SPARQL Benchmarks?

[f-TP Selectivity – mean]




How Diverse are SPARQL Benchmarks?

[f-TP Selectivity – **stdev**]




How Diverse are SPARQL Benchmarks?

[f-TP Selectivity – **stdev**]




How Diverse are SPARQL Benchmarks?

[BGP-Restricted f-TP Selectivity – mean]




How Diverse are SPARQL Benchmarks?

[BGP-Restricted f-TP Selectivity – mean]




How Diverse are SPARQL Benchmarks?

[BGP-Restricted f-TP Selectivity – stdev]




How Diverse are SPARQL Benchmarks?

[BGP-Restricted f-TP Selectivity – stdev]




How Diverse are SPARQL Benchmarks?

[Join-Restricted f-TP Selectivity – mean]



How Diverse are SPARQL Benchmarks?

[Join-Restricted f-TP Selectivity – mean]



How Diverse are SPARQL Benchmarks?

- WatDiv
- DBSB
- BSBM
- SP²Bench
- LUBM
- Other SPARQL benchmarks
- Production workload
- ...

WatDiv Tools

- Data Generator
 - Customizable data description model
- Query Template Generator
- Query Instantiator

WatDiv Dataset

Entities generated according to the *default* data description model

Entity	No. of Instances (per scale factor)
wsdbm:Purchase	1500
wsdbm:User	1000
wsdbm:Offer	900
wsdbm:Product	250
wsdbm:Website	50
wsdbm:Retailer	12

Entity	No. of Instances
wsdbm:Topic	250
wsdbm:City	240
wsdbm:SubGenre	145
wsdbm:Language	25
wsdbm:Country	25
wsdbm:Genre	21
wsdbm:ProductCategory	15
wsdbm:AgeGroup	9
wsdbm:Role	3
wsdbm:Gender	2

The entities above do not scale.

WatDiv Dataset

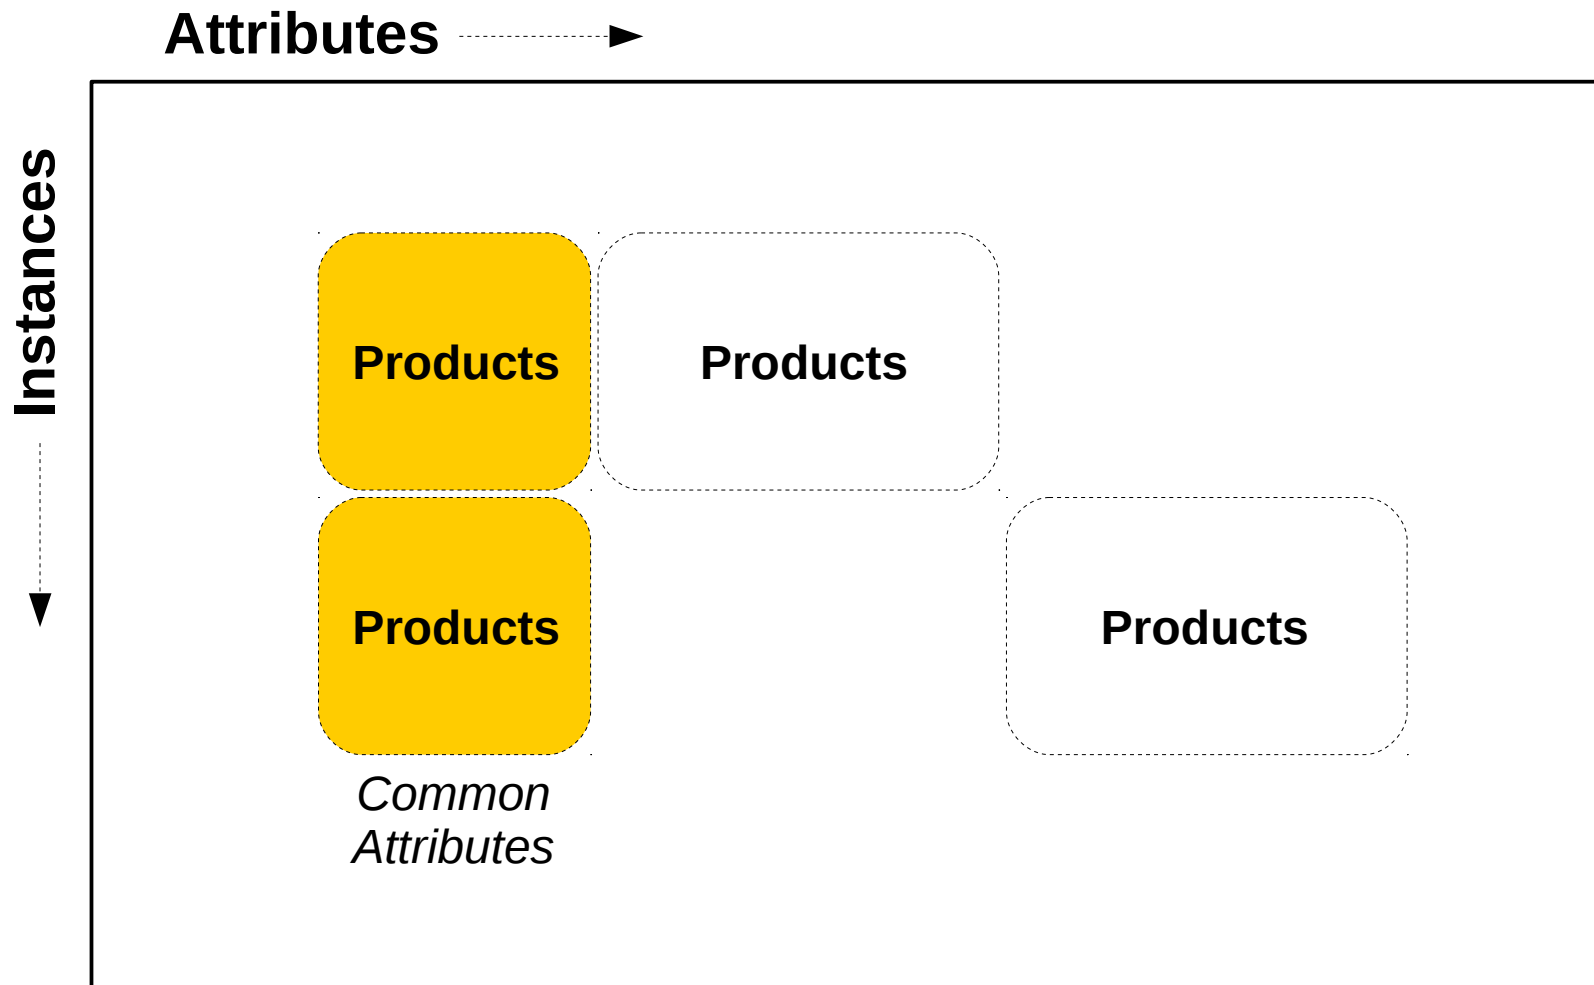
Characteristics of the dataset at
scale-factor=1 (default model)

Triples	105257
Distinct subjects	5597
Distinct predicates	85
Distinct objects	13258
URIs	5947
Literals	14286
Distinct literals	8018

Numbers are approximate and may vary slightly in each dataset generation

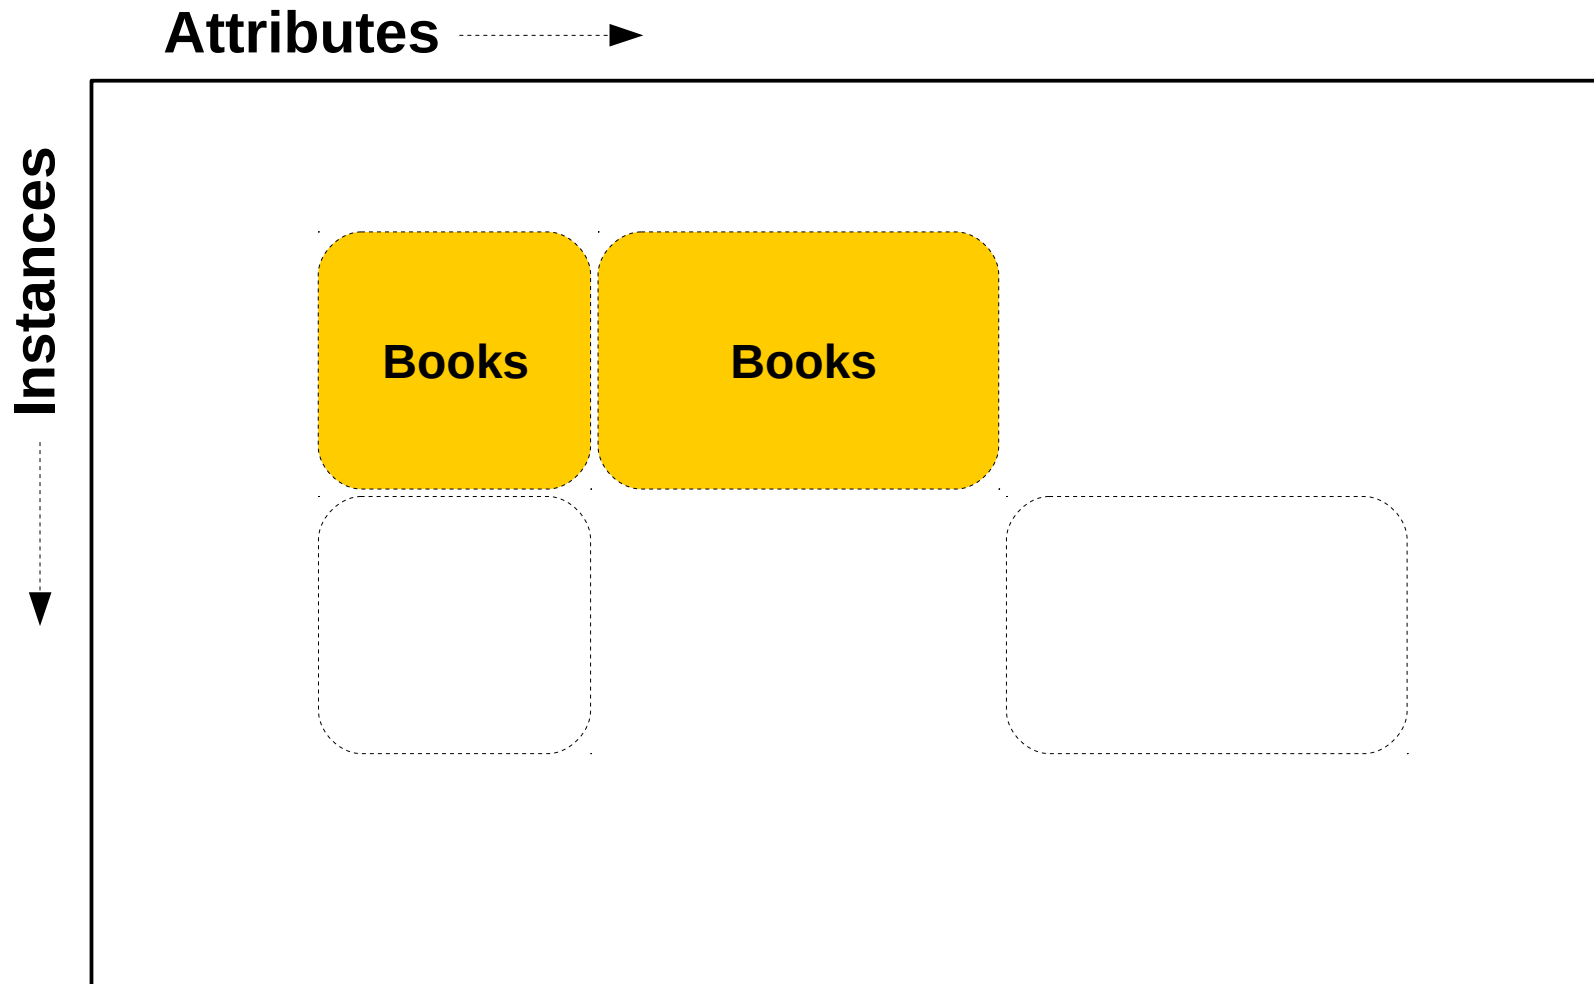
WatDiv Dataset

[What Sets It Apart]



(1) Heavily Relies on Optional Attributes

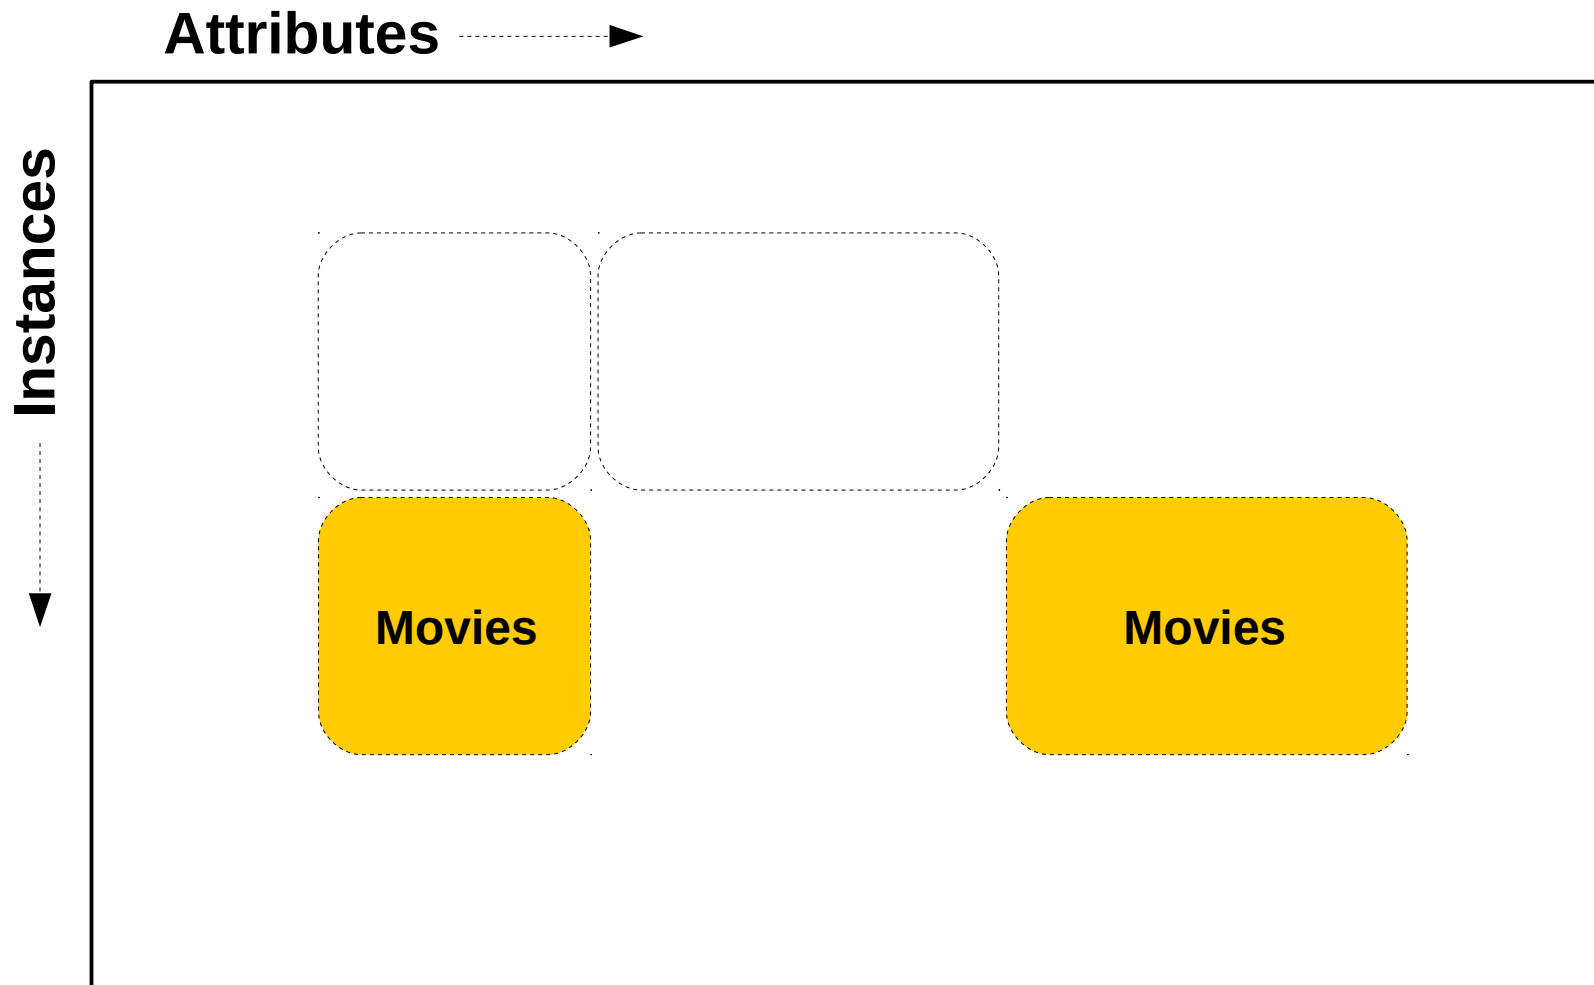
WatDiv Dataset [What Sets It Apart]



(1) Heavily Relies on Optional Attributes

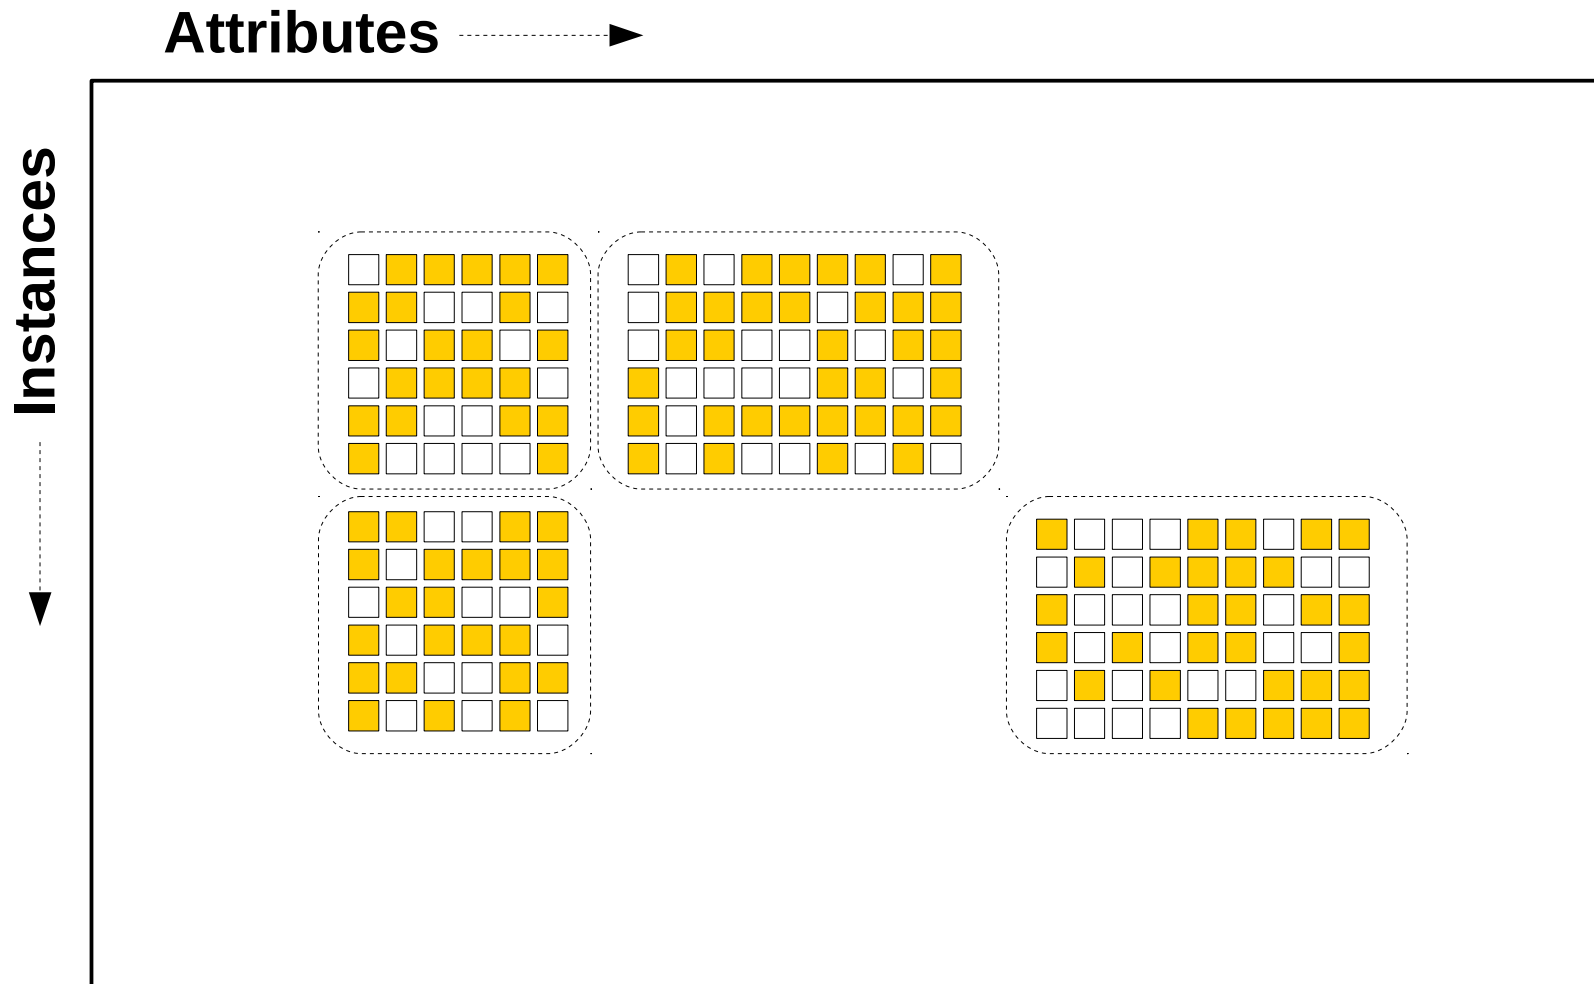
WatDiv Dataset

[What Sets It Apart]



(1) Heavily Relies on Optional Attributes

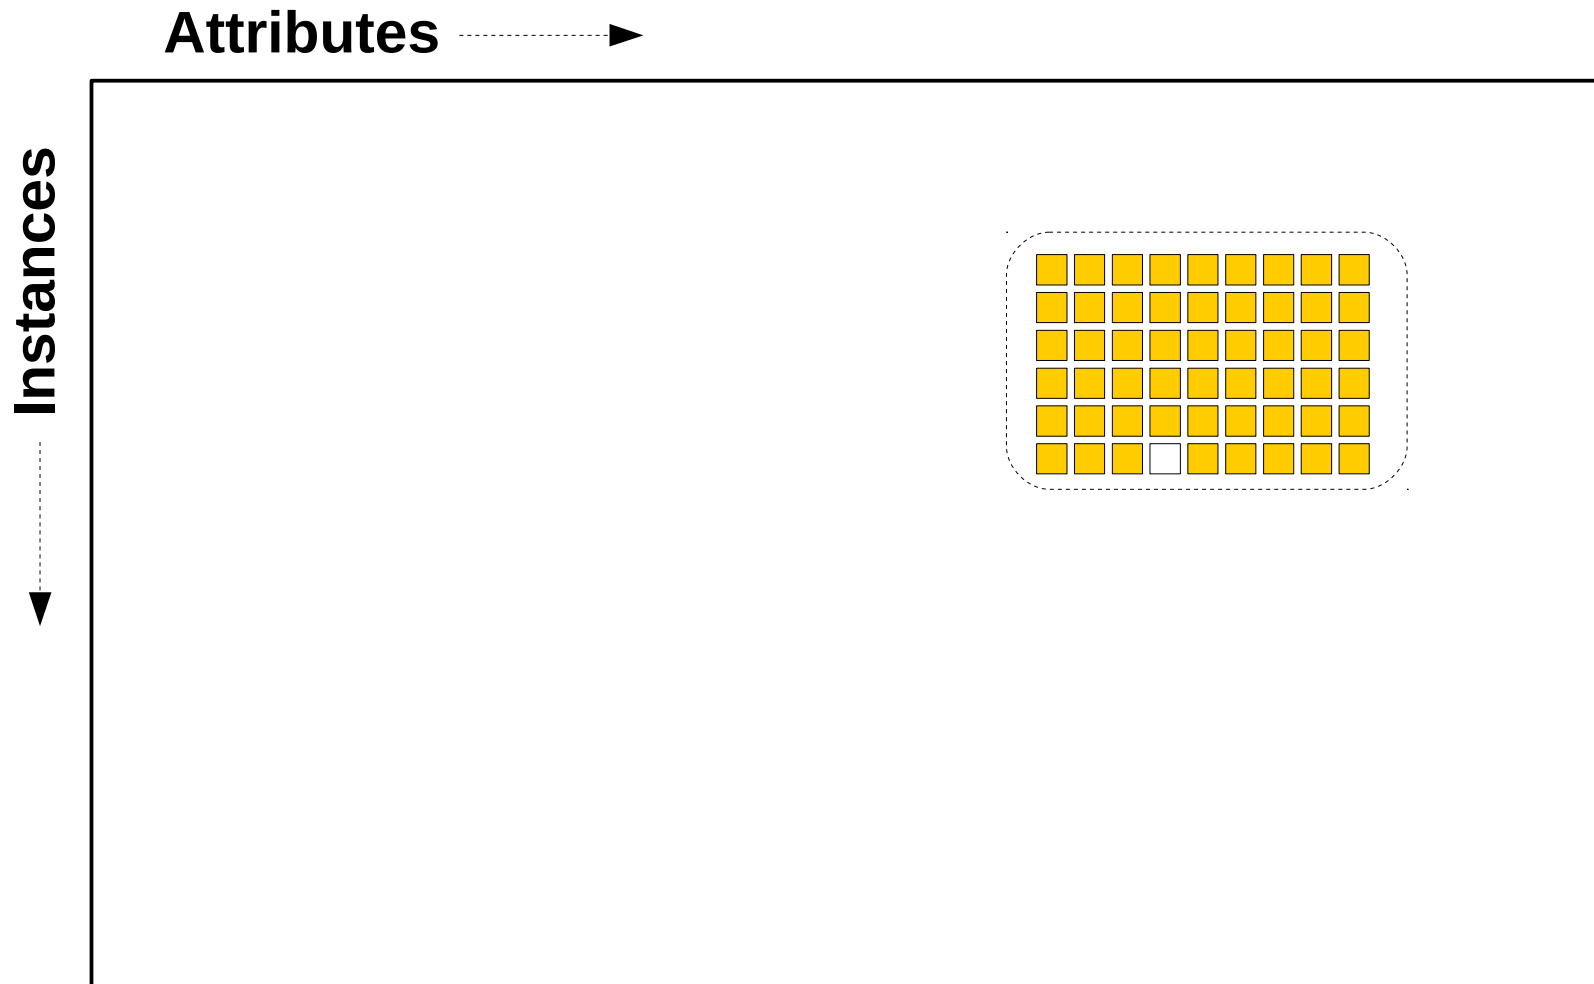
WatDiv Dataset [What Sets It Apart]



(1) Heavily Relies on Optional Attributes

WatDiv Dataset

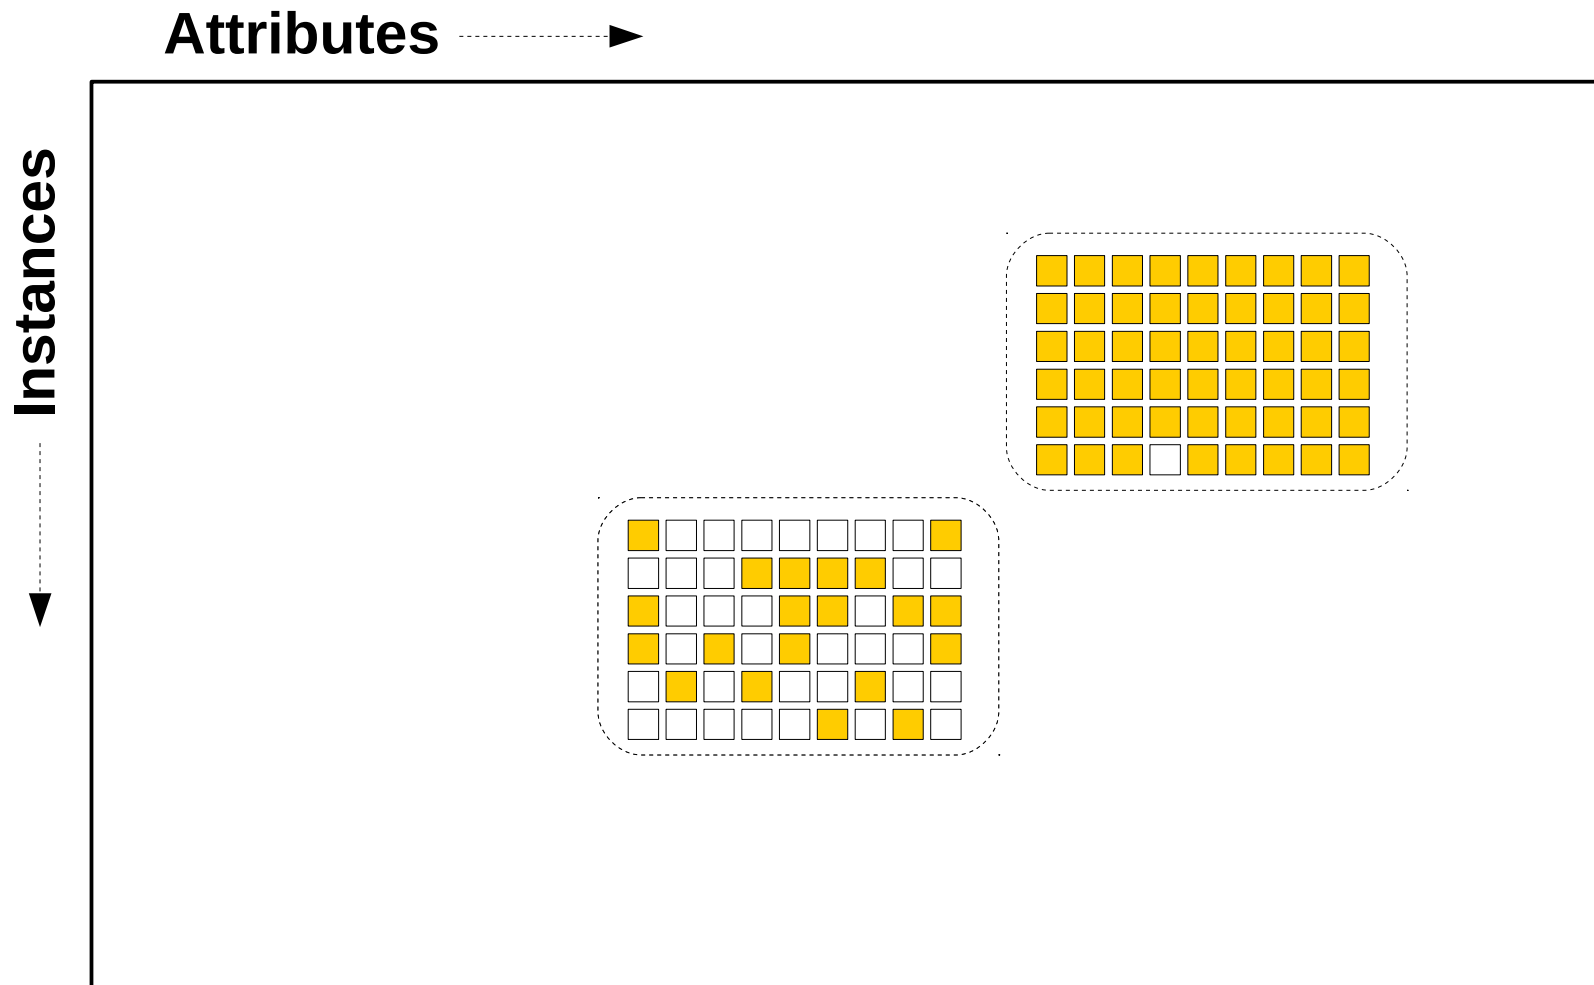
[What Sets It Apart]



(2) Parts of the database are well-structured

WatDiv Dataset

[What Sets It Apart]



(2) Parts of the database are well-structured while remaining parts are less well-structured

WatDiv Dataset

[What Sets It Apart]

```
// Attributes for wsdbm:ProductCategory4 (i.e., NewsArticle)////////////////////////////////////
```

```
...
```

```
<pgroup> 0.8 @wsdbm:ProductCategory4
```

```
  #predicate sorg:publisher string
```

```
</pgroup>
```

```
<pgroup> 0.7 @wsdbm:ProductCategory4
```

```
  #predicate sorg:datePublished date
```

```
</pgroup>
```

```
<pgroup> 0.2 @wsdbm:ProductCategory4
```

```
  #predicate sorg:printPage integer 1 999
```

```
  #predicate sorg:printSection integer 1 9
```

```
</pgroup>
```

WatDiv Dataset


[What Sets It Apart]

- // Associations for wsdbm:ProductCategory2 (i.e., Movie)////////////////////////////////////
- #association
wsdbm:Product sorg:actor wsdbm:User
2 25[normal] 0.8 UNIFORM
@wsdbm:ProductCategory2 @wsdbm:Role2
- #association
wsdbm:Product sorg:director wsdbm:User
2 1 0.8 ZIPFIAN
@wsdbm:ProductCategory2 @wsdbm:Role2
- #association
wsdbm:Product sorg:trailer wsdbm:Website
2 3[uniform] 0.1 UNIFORM
@wsdbm:ProductCategory2 @null

How Robust are Systems across WatDiv Workloads?

[BGP-restricted f-TP selectivity – mean]


■ [0.000, 0.003) ▨ [0.030, 1.000)



WatDiv 100M triples, queries w/ single join vertex, result cardinality ≤ 2000

How Robust are Systems across WatDiv Workloads?

GOOD NEWS



WatDiv 10M triples

linear = { mean join vertex degree ≤ 3.0 , join vertex count ≥ 3 }
star/snowflake = { mean join vertex degree ≥ 5.0 , join vertex count ≤ 2 }

Conclusions

- Which of the existing SPARQL benchmarks, if any, should I use to diagnose (and fix) potential problems with the physical design of my system?
 - ***Analyze your production workload and find the best-matching benchmarks***
- How can I use the Waterloo SPARQL Diversity Test Suite (WatDiv) where existing benchmarks fall short?
 - ***Drill down into different classes of queries until you hit problematic spots***

Questions

Waterloo SPARQL Diversity Test Suite
(WatDiv)



<http://db.uwaterloo.ca/watdiv/>