




Maverick: Discovering Exceptional Facts from Knowledge Graphs

Exceptional Facts

	Denzel Washington followed Sidney Poitier as only the second black to win the Best Actor award .	Entity of interest	Denzel Washington	Given an entity x	find	such that
	This was Brazil's first own goal in World Cup history.	Context	Best Actor award winners	A context		the context has many entities, including x
	Hillary Clinton becomes first female presidential nominee.	Attributes	Ethnicity	A set of attributes (subspace)		
		Peculiar value	African American (only two satisfy)			x bears a peculiar value w.r.t. the subspace (few in the context have the value)

Applications

- Computational Journalism**
- Fact-finding
 - Fact-checking
- E.g., The first female presidential nominee was Victoria Woodhull, not Hillary Clinton (snopes.com)
- Data Cleaning**
- Recommendation Systems**
- Friends, news, and product promotion

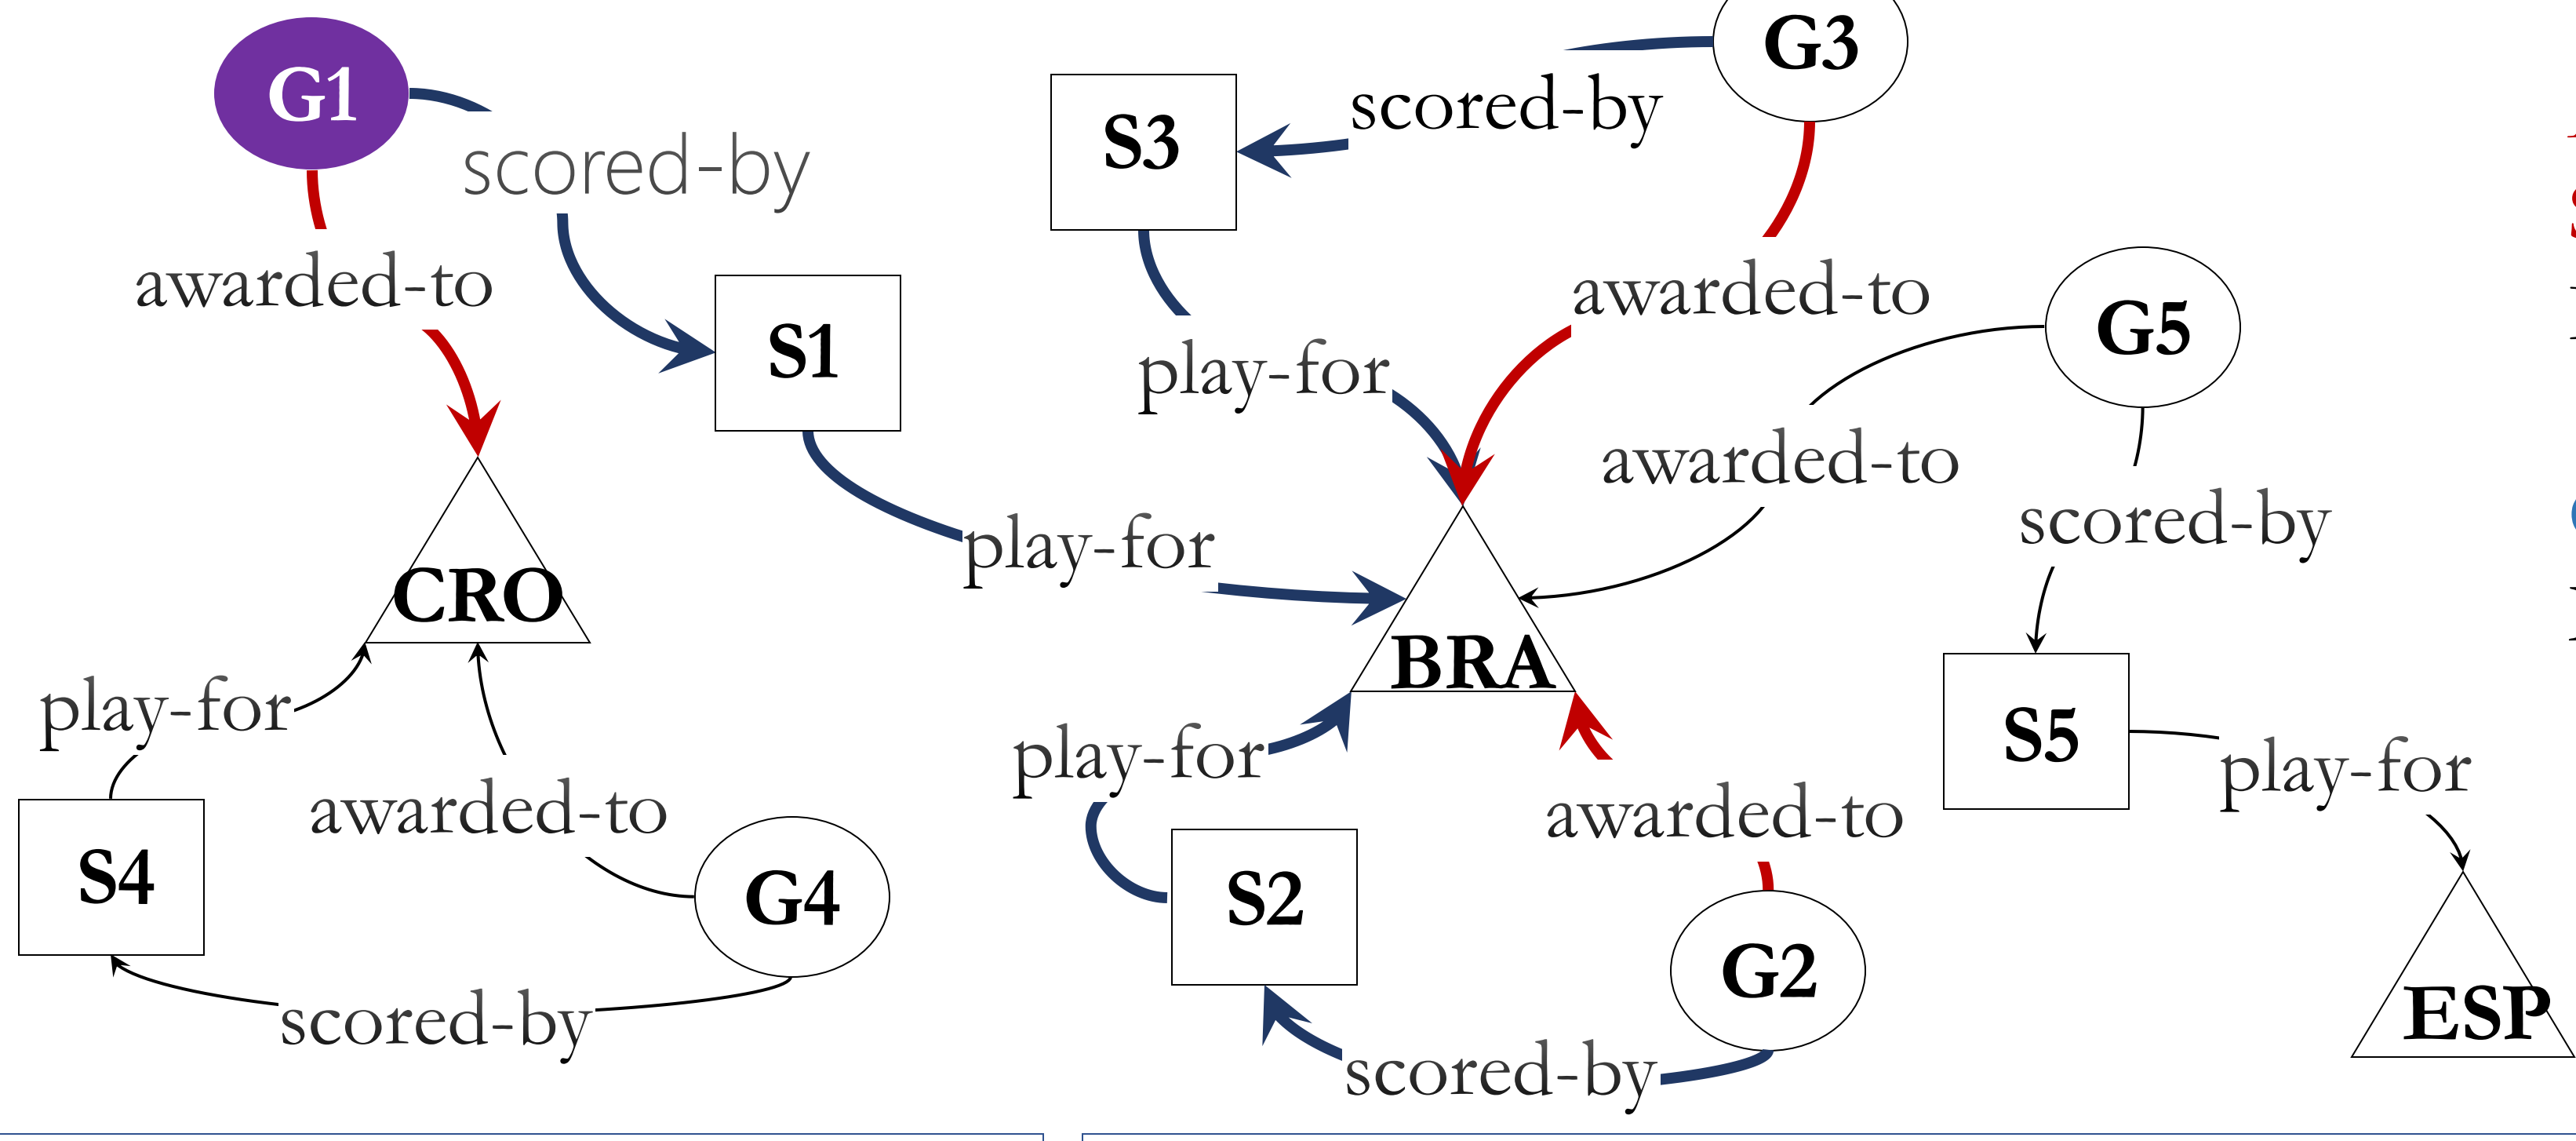


Related Work

- Outlier detection**
- Maverick finds conditions that make an object stand out, although the object may not necessarily be an outlier.
- Outlying aspect mining**
- Challenges in adopting existing algorithms:
- Many assume a single-table model: a graph can be an extremely large and sparse table
 - Conjunctive queries on a single table \neq pattern queries
 - Multiple tables: unclear how to handle joins
 - Unclear how to handle set values

Exceptional Facts from Knowledge Graphs

What is exceptional about G1? Among all the goals scored by BRA players, G1 is the only own goal.



Attributes: labels of incoming/outgoing edges
Subspace: a subset of attributes
 E.g., G1. awarded-to = CRO

Context: entities sharing some common characteristics. Defined by a pattern-variable pair

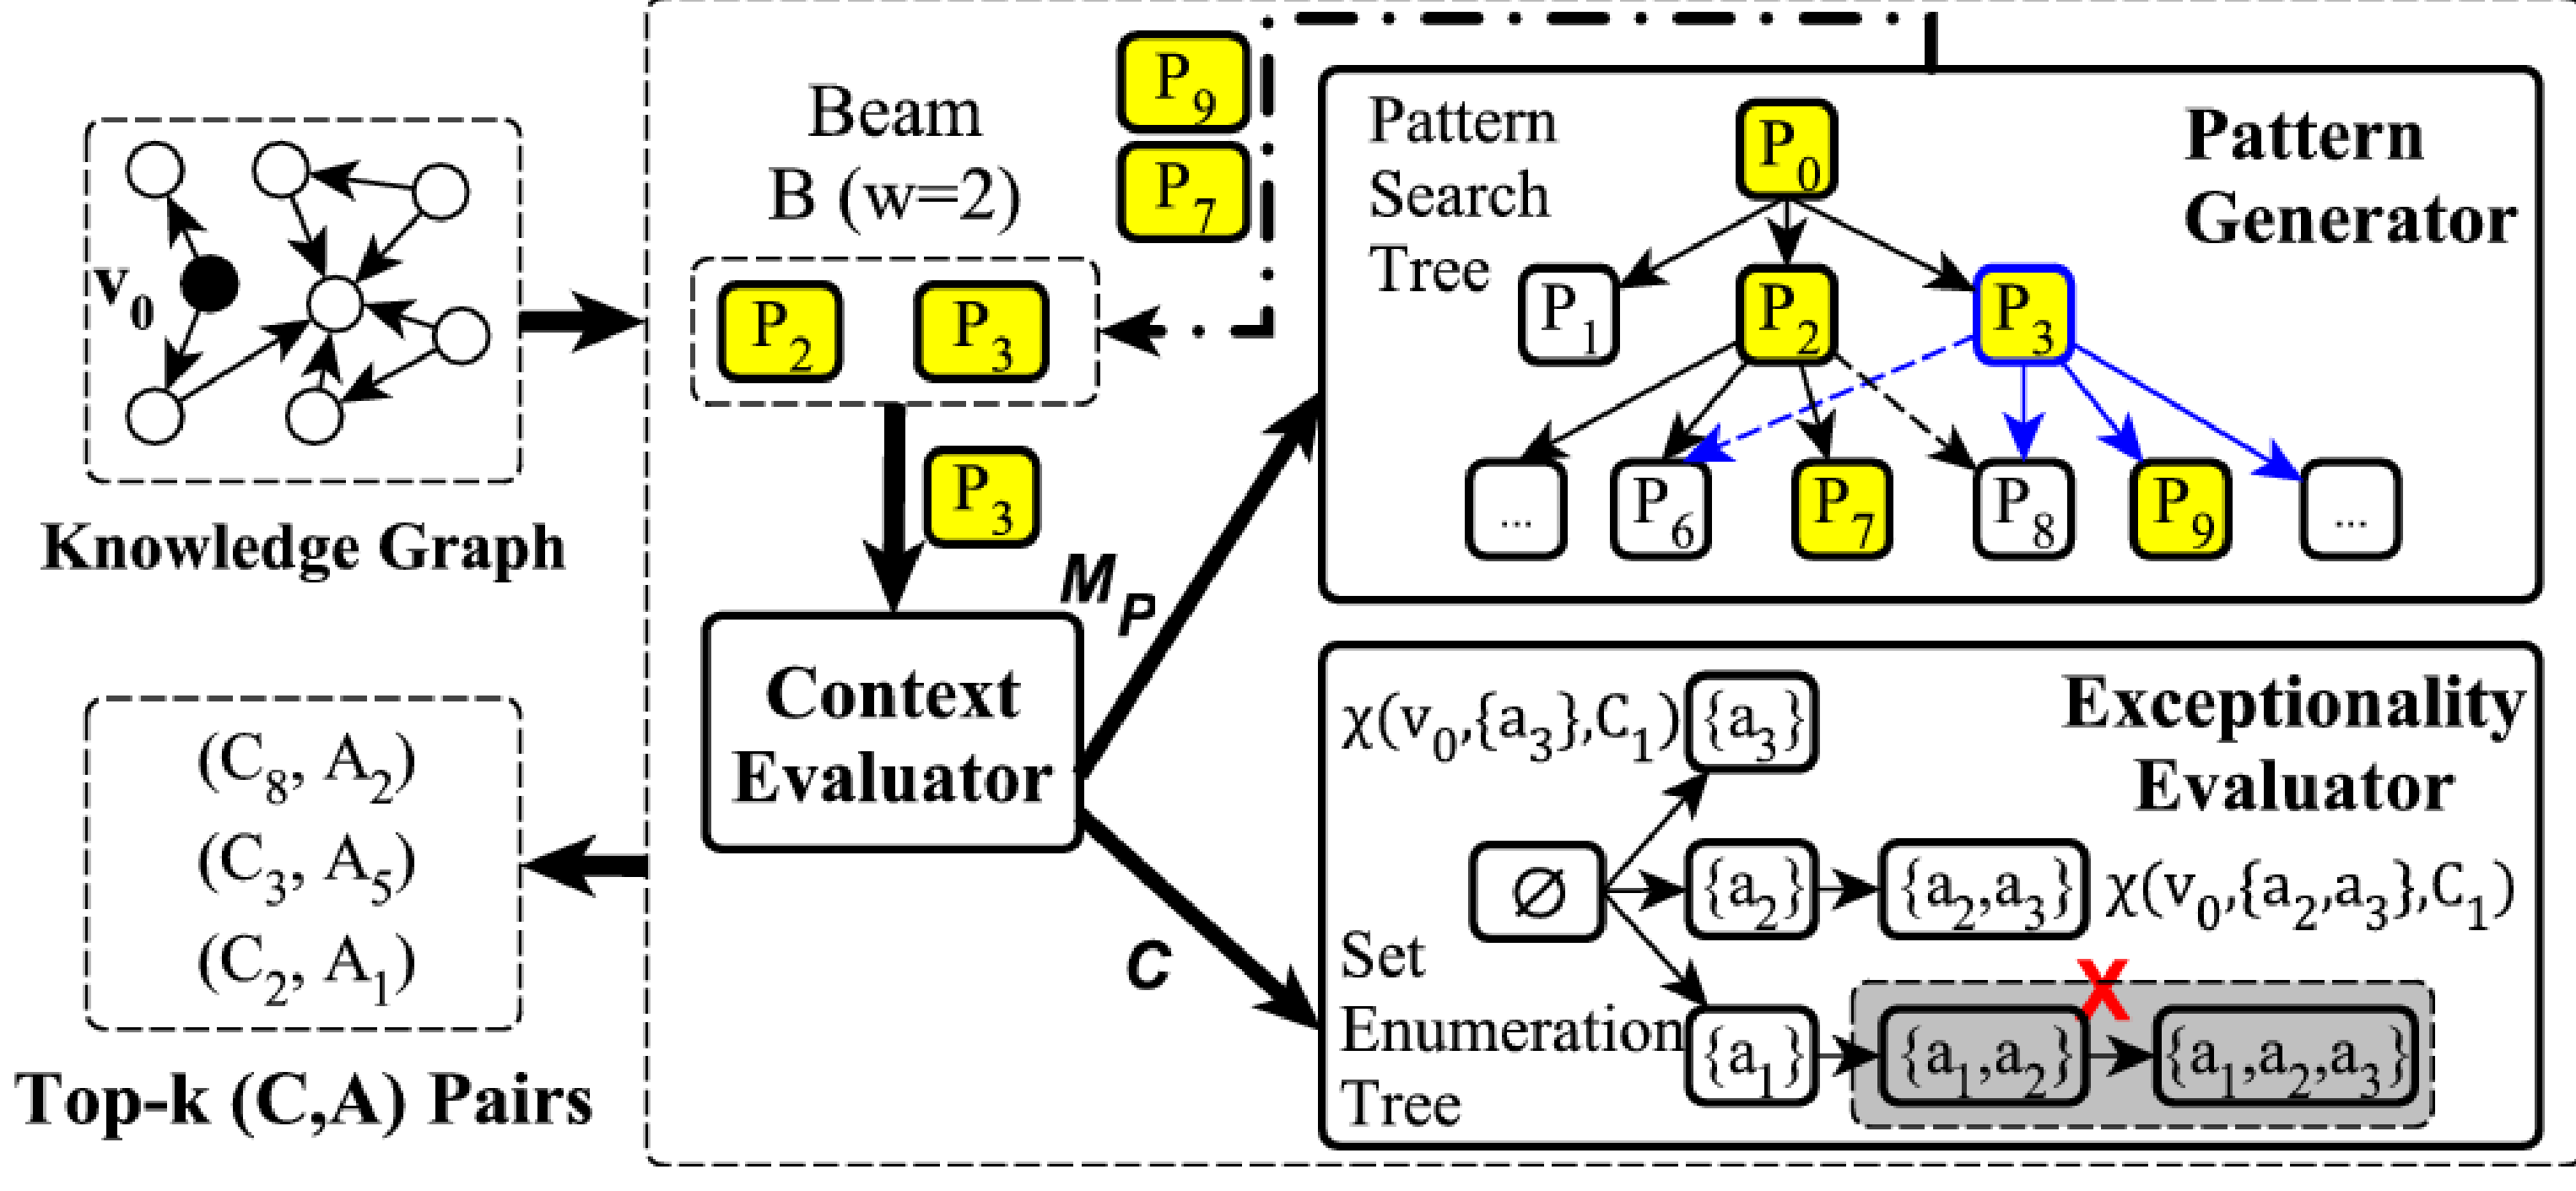
$?g \rightarrow \text{awarded-to} \rightarrow ?s \rightarrow \text{play-for} \rightarrow \text{BRA}$

E.g., Goals scored by Brazilian players

Problem Formulation

- Input**
- Entity of interest v_0
 - Exceptionality function χ
 - Result size k
- Output**
- Top- k (context, subspace) pairs with regard to χ , in which v_0 stands out

Maverick

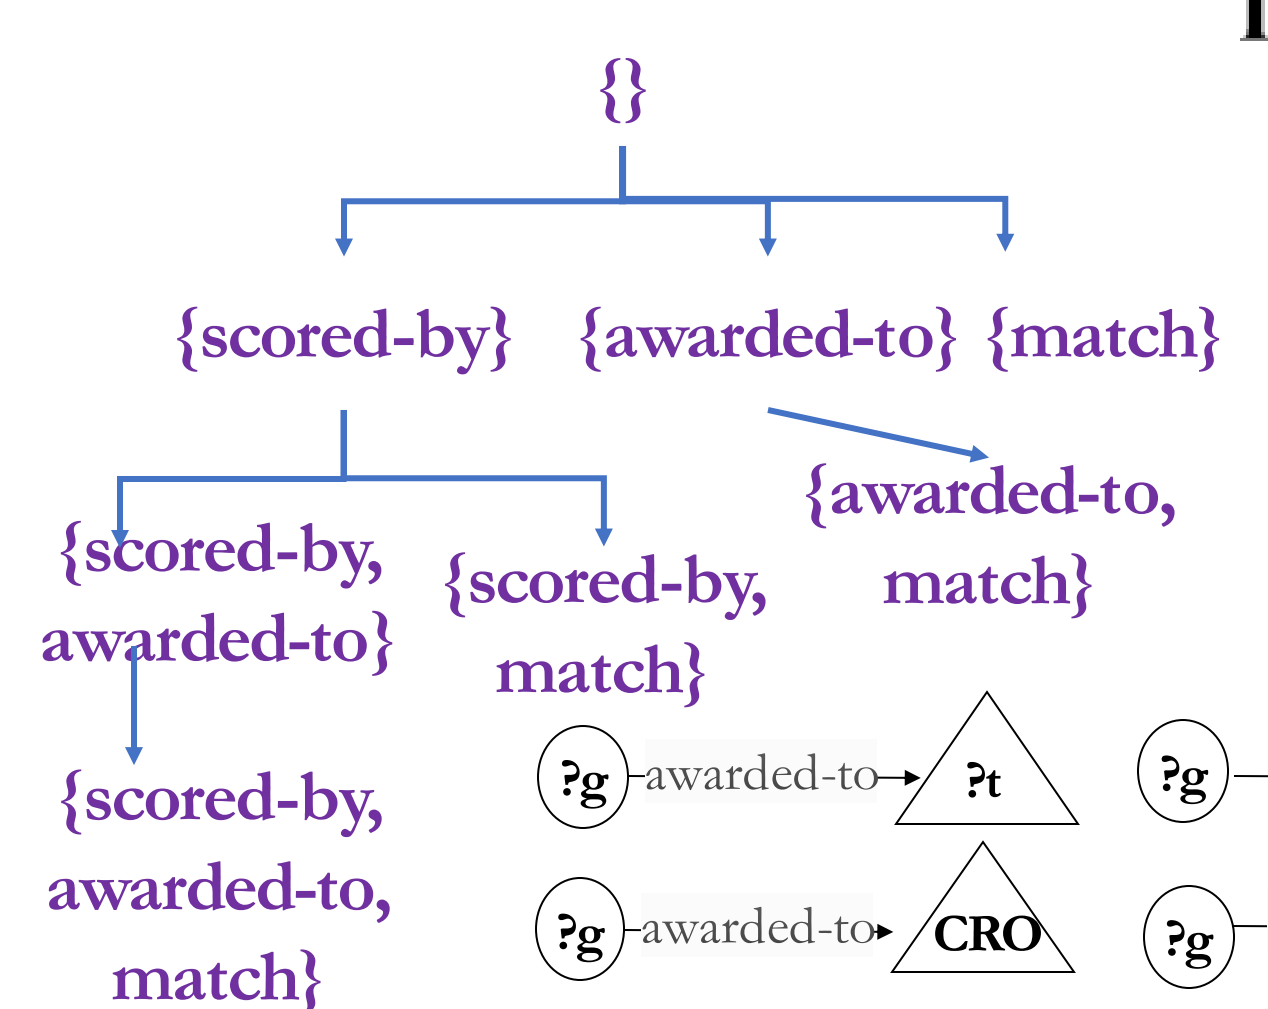


Experiments


WCGoals
 Created based on FIFA.com
 11 node types, 13 edge types
 49,078 nodes, 158,114 edges

Film-Award
 A subgraph of Freebase
 95 node types, 117 edge types
 5,437,628 nodes, 10,879,448 edges

Challenges



Number of attribute subspaces: $O(2^{|A_{v_0}|})$
 Number of patterns (contexts): $\Omega(2^{|V_{G_1}|})$



See you in Rio!
 VLDB2018 demo

