Bench-Ranking:

Towards prescriptive analysis of big graph processing: the case of SparkSQL

The 14th LDBC TUC Meeting

Mohamed Ragab — University of Tartu— mohamed.ragab@ut.ee

Under Sup.: Riccardo Tommasini

RDF data is exploding…

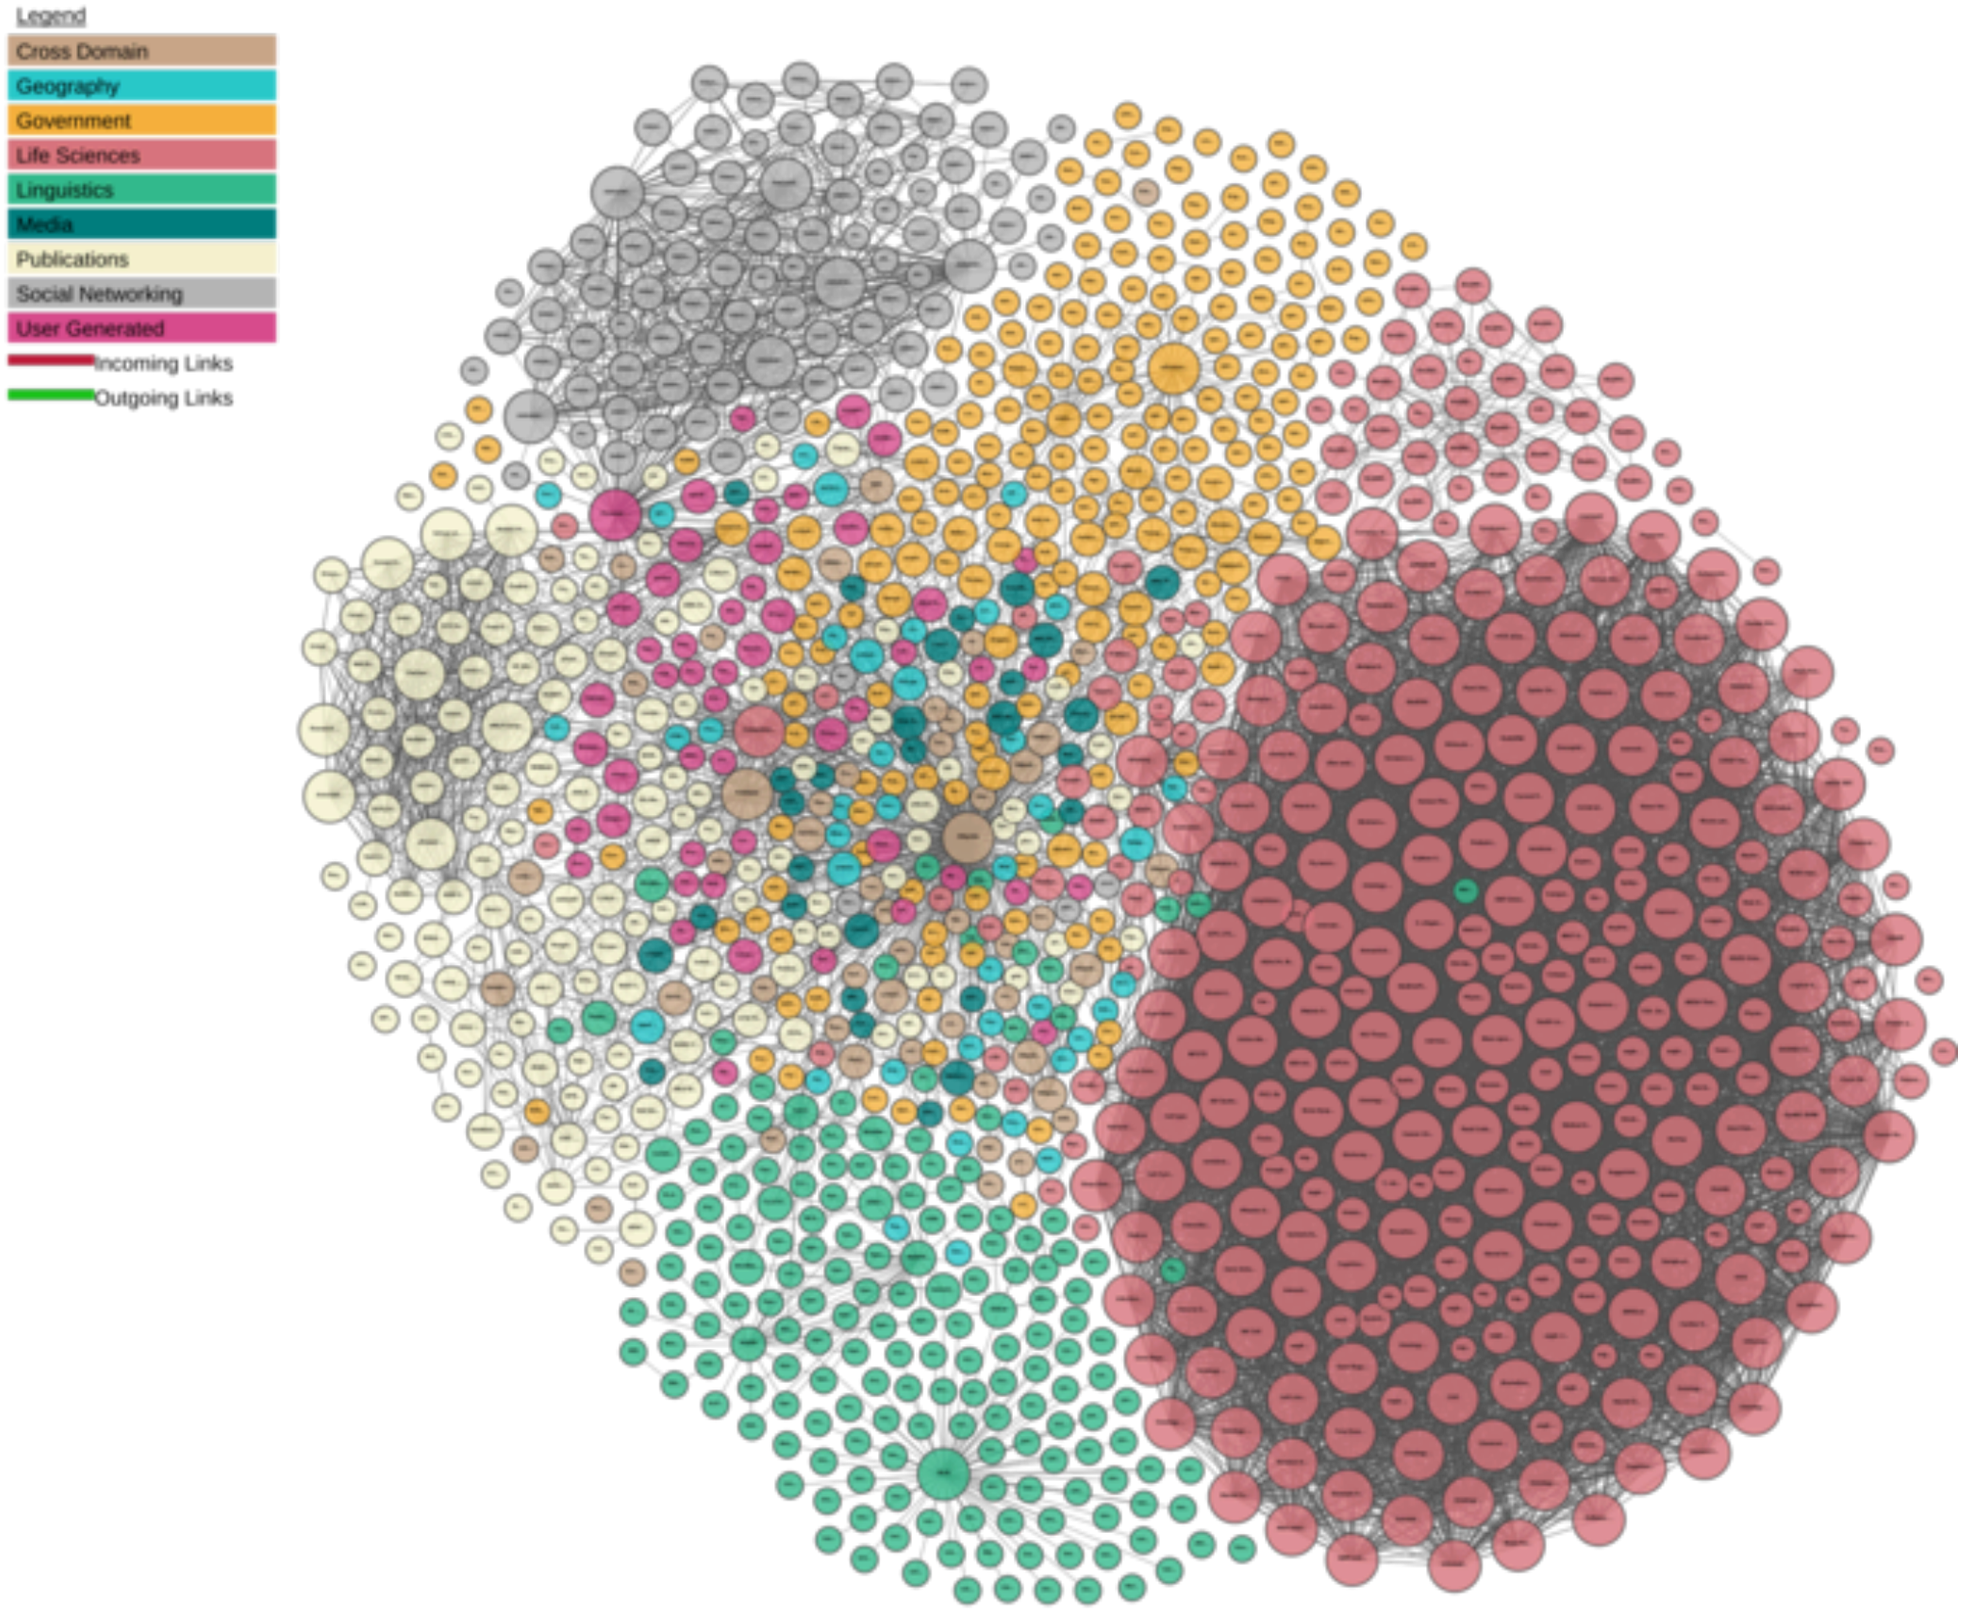

The Linked Open Data Cloud: https://lod-cloud.net/


DO YOU THINK THAT'S A
Graph
YOU ARE QUERYING?

Native Vs. Non-native (Relational) RDF Proc.

Native Triple Stores

- Centralized
- E.g, Jena, RDF3X,..

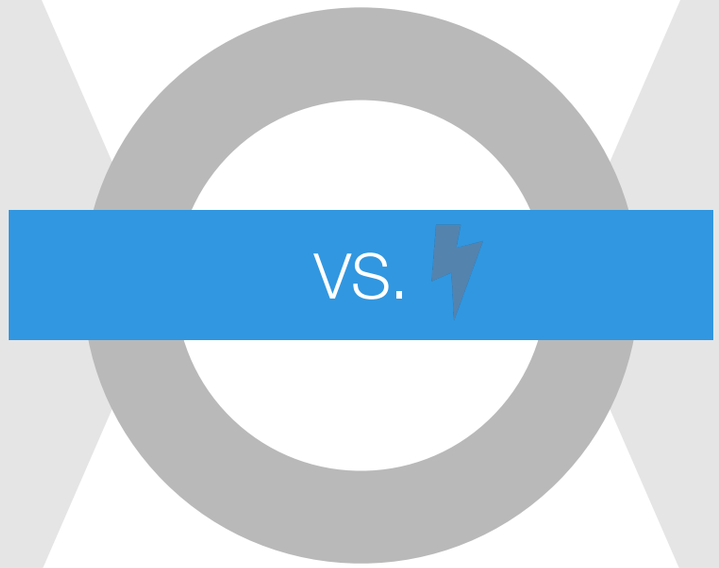

Big Relational Systems

- Not tailored for RDF Processing.
- E.g, Apache Spark, Hive, Impala.

The Datafloq Open Source Landscape 2.0

Data Analysis & Platforms

HPCC Systems
Storm
Distributed and fault-tolerant realtime computation
Hadoop
Dremel
hadoop MapReduce
APACHE Spark
APACHE DRILL
SAMOA
HD
IKANOW | BRILLIANT DECISIONS
Hortonworks

Databases / Data warehousing

bigdata
INFOBRIGHT
Cassandra
4store
H2
GlobalsDB
InfiniDB
riak
Infinispan
HYPERTABLE
MariaDB
Drizzle
SQLite
RethinkDB
terrastore
HIVE
Firebird
ORACLE BERKELEY DB
HyperSQL
monetdb

In-Memory Computing

GridGain IN-MEMORY COMPUTING
hazelcast
TERRACOTTA by Software AG
GORA
NMemory

Business Intelligence

ERP BI Solutions
Open Source Business Intelligence Solutions for ERP
openi.org Open Intelligence
talend *open data solutions
Jaspersoft
Palo Open Source Business Intelligence
spagobi
pentaho
jedox.
BIRT

Data Mining

orange
rapidminer
KNIME Data Analytics Made Easy
mahout
WEKA The University of Waikato
KEEL
togaware
SPMF

Big Data search

lucene
Apache Solr
elasticsearch.

Multivalue database

Rocket
U2 REVELATION SOFTWARE
Open QM Ladybridge Systems
northgate Reality
jBASE INTERNATIONAL FREEDOM. OPENNESS. POWER.
ScarletDME

Programming

R
julia

Data aggregation

sqoop
FLUME
chukwa

KeyValue

AEROSPIKE
leveldb
redis
Chordless Beta
Tokyo Cabinet 8192PIB
MEMCACHED
SCALIEN
Project Voldemort A distributed database.
hamsterdb
RAPTORDB
FairCom
STSDB DATABASE & VIRTUAL FILE SYSTEM
HyperDex
OpenLDAP
IQLECT
ioremap.net STORAGE AND BEYOND
Scalaris

Document Store

mongoDB
CouchBase
Raven DB
CLUSTERPOINT
Tokutek
RaptorDB
EJDB
djondb
JasDB
SchemafreeDB
sisodb
Apache CouchDB relax

Object databases

db4objects
ZOPE
mobject precision data management
Magma
Picolisp
NEOPPOD Distributed Transactional NoSQL for the Cloud
siaqodb
RAMER D
PERSEVERE
EyeDB
STARCOUNTER
Sterling
NDatabase C# Lightweight Object Database

Graph databases

Gephi makes graphs handy
Gremlin
GraphBuilder
FRANZ INC.
*Sparksee
INFOGRID
InfiniteGraph Powered by Objectivity
HYPERGRAPHDB
Neo4j the graph database
GraphBase
BrightstarDB
FlockDB

Operational

VOLTDB SMART DATA FAST.

Social

Apache Kafka A high-throughput distributed messaging system.
ThinkUp
Corona

Multidimensional

FIS
SciDB
rasdaman raster data manager

Multimodel

ArangoDB
alchemydatabase A Hybrid Relational-Database/NOSQL-Datastore

XML Databses

existdb
BASEX
Qizx
sedna
LIQUIBASE

Grid Solutions

GIGASPACES
Galaxy

Created by: www.Datafloq.com

Let's take (Apache Spark-SQL) as example

01

Schema

- Single Statement Table (ST)
- Vertically Partitioned Tables (VT)
- Property Tables (PT)

02

Partitioning

- Horizontal Partitioning (HP)
- Subject-based Partitioning (SBP)
- Predicate-based Partitioning(PBP)

03

Storage Formats

- Row-oriented (Avro, CSV).
- Columnar-oriented (ORC, Parquet).

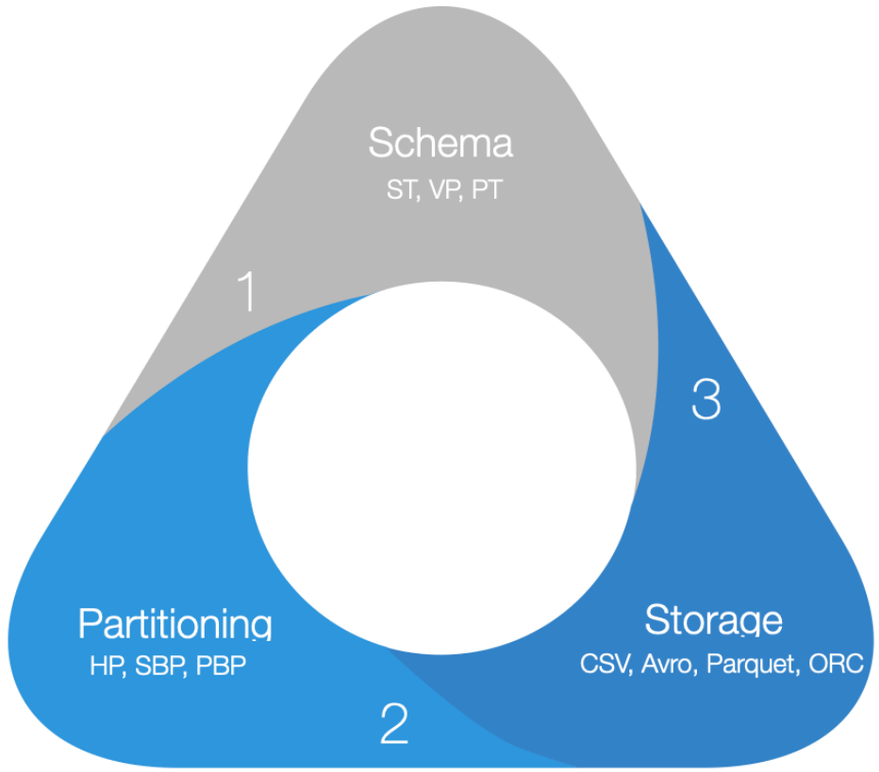

Experimental Solution Space

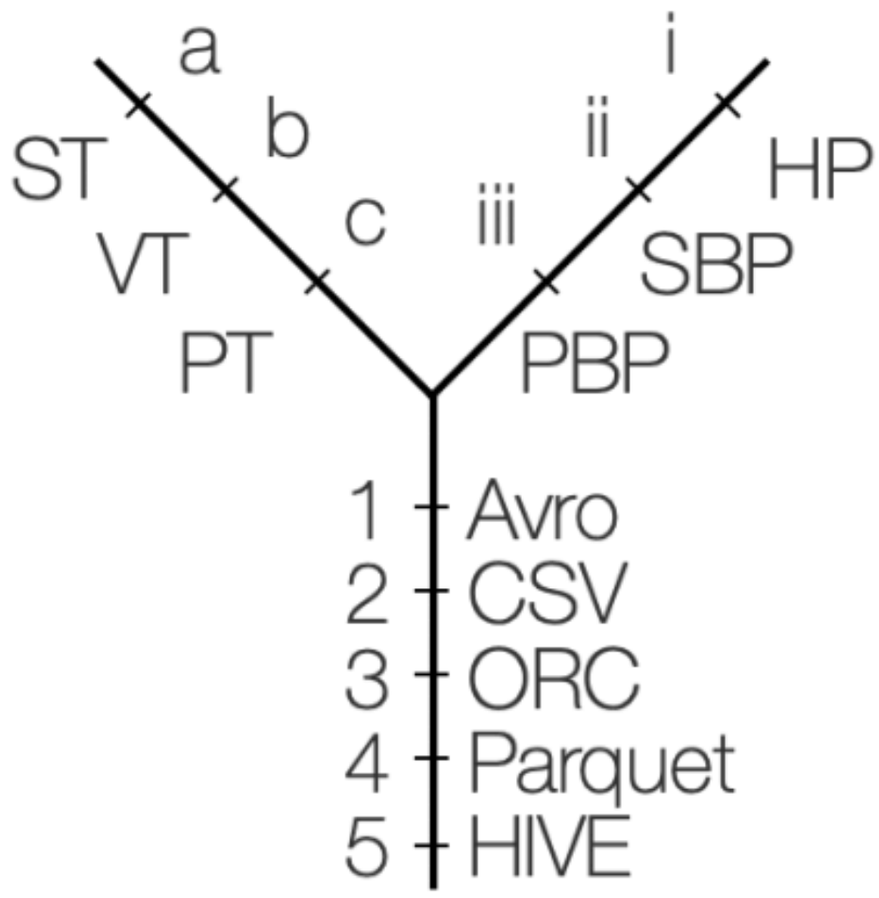

a.i.1 == ST.HP.Avro

Experimental Solution Space

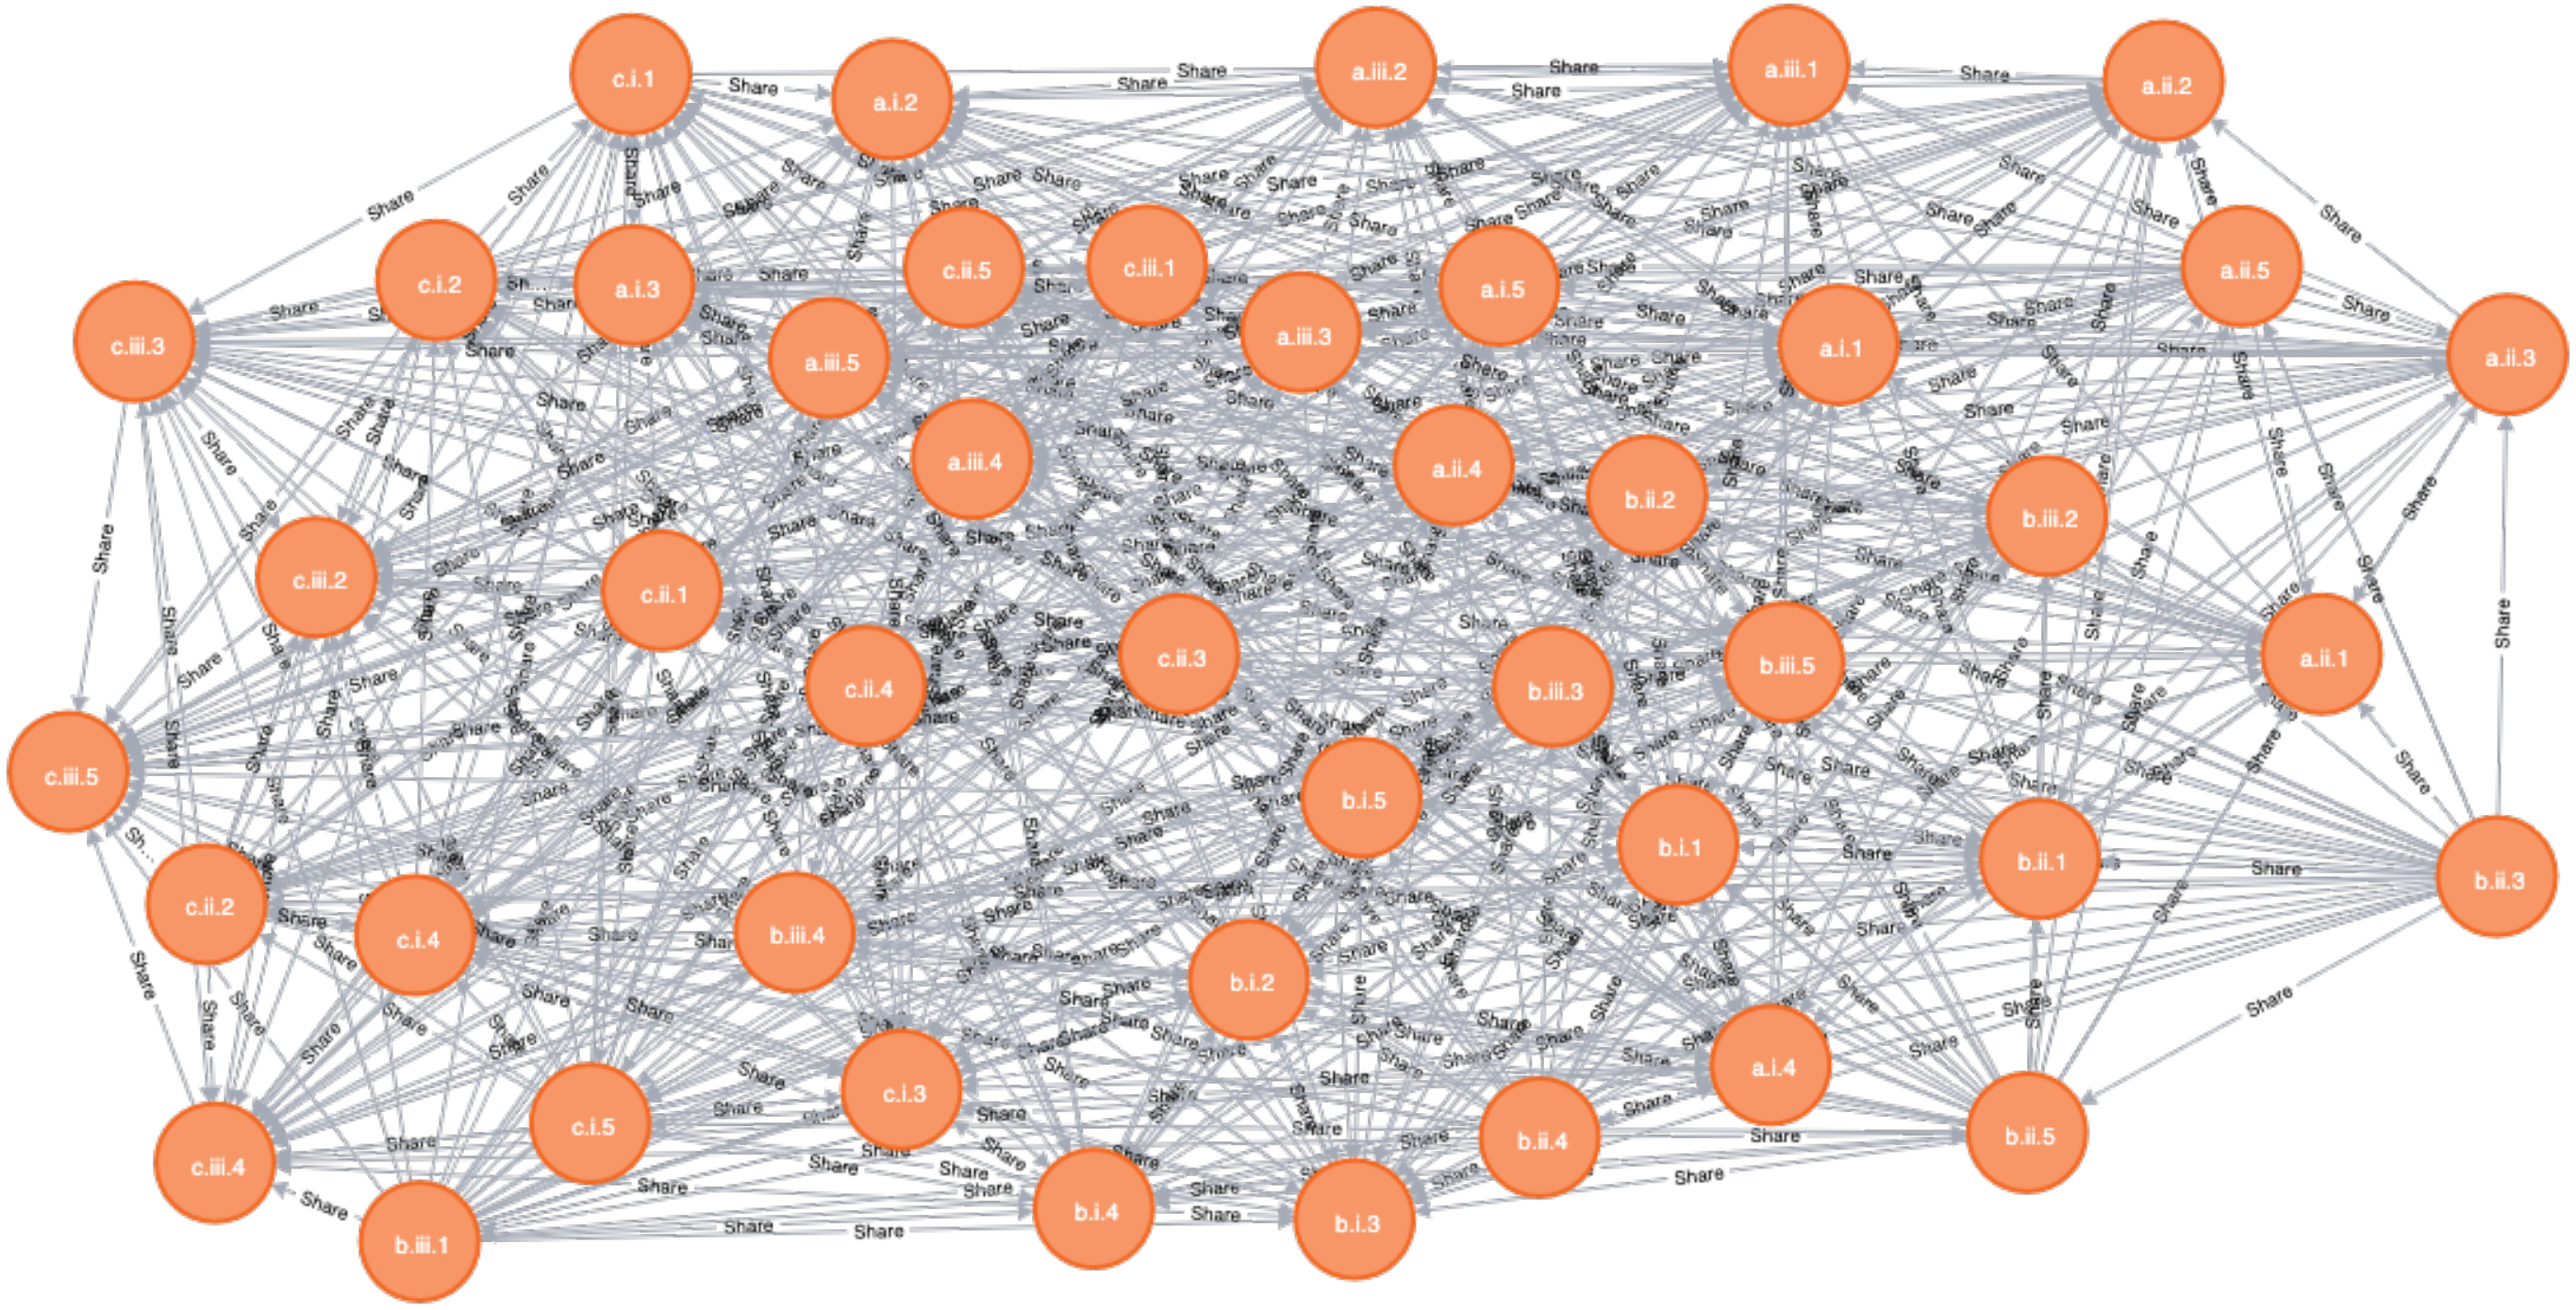

Which configuration combination the best to choose ?!

The 4 Levels of Analysis

01 Descriptive Analysis

- - Describe results (Which dimension was better).
- By how much.

02 Diagnostic Analysis

- Try to describe why it happened
- using the Domain knowledge.

03 Predictive Analysis

- Predict what will happen?!
- Applying ML and other stat.

04 Prescriptive Analysis

- What should be done?
  - E.g. what is the best conf

The 4 levels of Data Analytics: https://koopingshung.com/blog/four-levels-of-analytics-data-science-descriptive-diagnostic/

Bench-Ranking

01

Looking at the descriptive analysis is not enough!

- looking at a lot of performance data might be overwhelming.
- Sometimes, **Contradicting results**.

02

Dimensions Trade-offs

There are always clear trade-offs between these dimensions.

03

Selecting Best Combination

Selecting the best configuration combination out of this complex solution space is not an easy task.

04

Bench-Ranking: Simple Yet Accurate

Bench-ranking criteria provide an accurate yet simple way that supports the practitioners in this task even in the existence of dimensions' trade-offs.

05

Firsts Steps of BD Ranking: Saleem et.al

Saleem et al. Proposed a ranking criteria for ranking 7 approaches for Partitioning RDF graphs.

An Empirical Evaluation of RDF Graph Partitioning Techniques: https://svn.aksw.org/papers/2018/EKAW_Graph_Partitioning/public.pdf

Individual Ranking Criteria

- For each dimension, we rank how the alternatives of this dimension are ranked.
- Example of Rank Scores.

$$R = \sum_{r=1}^{d} \frac{O_{dim}(r) * (d - r)}{Q(d-1)}, 0 < R \leq 1$$

Schema Rank Scores (for HP and Avro)

HP/ Avro	1st	2nd	3rd	*Rank(R)*
ST	1	3	7	0.23
VT	6	4	1	0.73
PT	4	4	3	0.55

Partiton. Rank Scores (for ST and Avro)

ST/ Avro	1st	2nd	3rd	*Rank(R)*
HP	2	6	3	0.45
SBP	8	3	0	0.86
PBP	1	2	8	0.18

Storage Rank Scores (for ST and HP)

ST/ HP	1st	2nd	3rd	4th	5th	Rank (R)
Avro	0	0	1	10	0	0.27
CSV	0	0	0	0	11	0.00
ORC	4	6	1	0	0	0.82
Parquet	2	3	5	1	0	0.64
Hive	5	2	4	0	0	0.77

An Empirical Evaluation of RDF Graph Partitioning Techniques: https://svn.aksw.org/papers/2018/EKAW_Graph_Partitioning/public.pdf

Ranking towards each dimension (Rf, Rp, Rs)

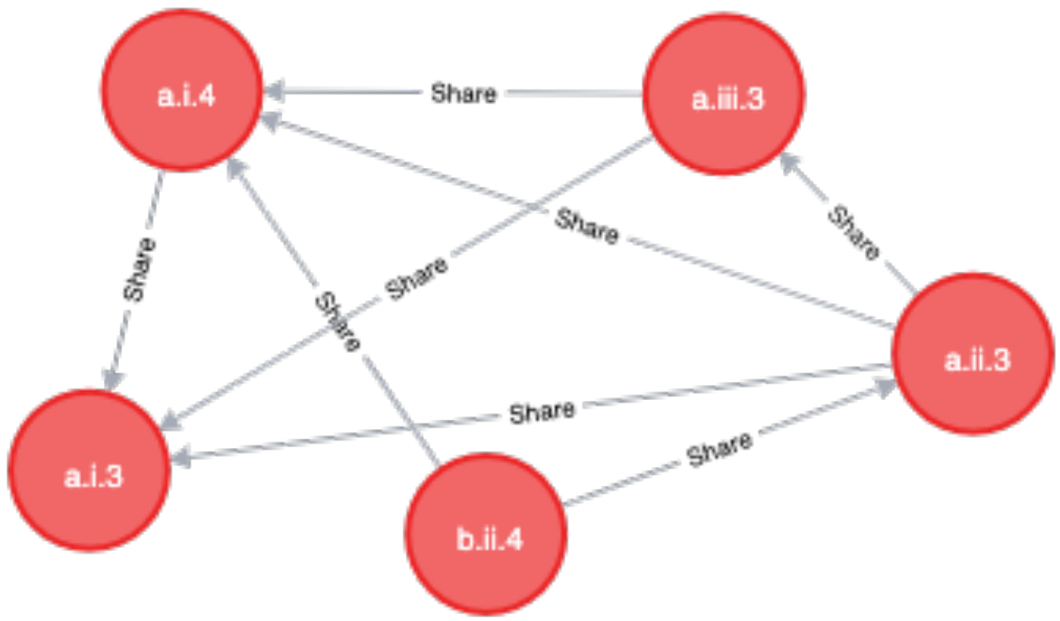

Rf

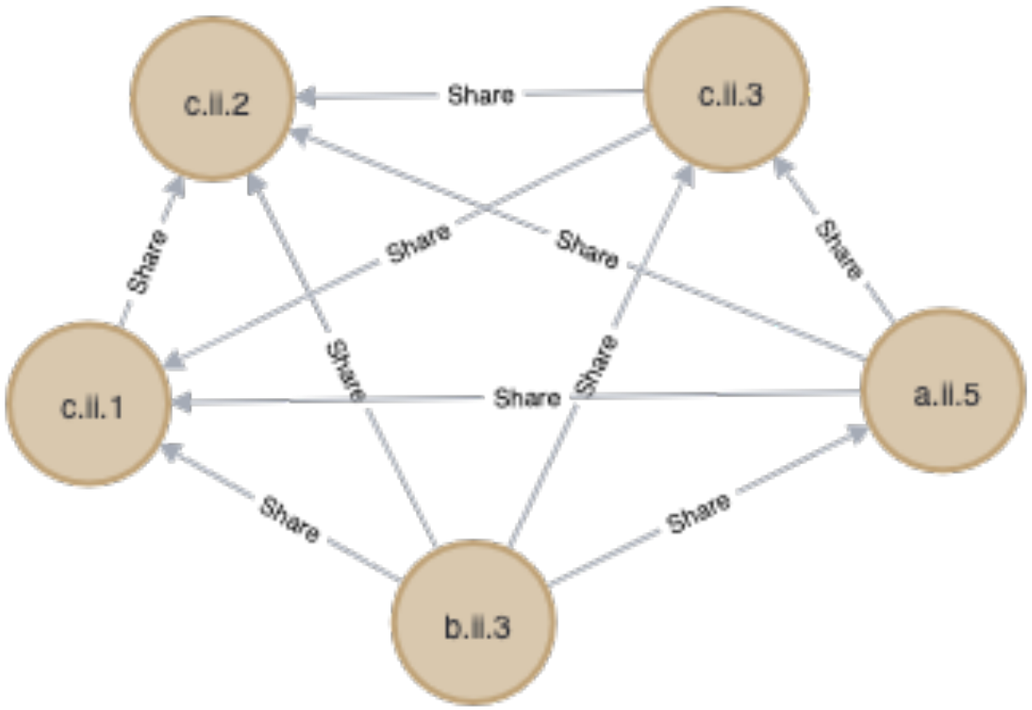

Rp

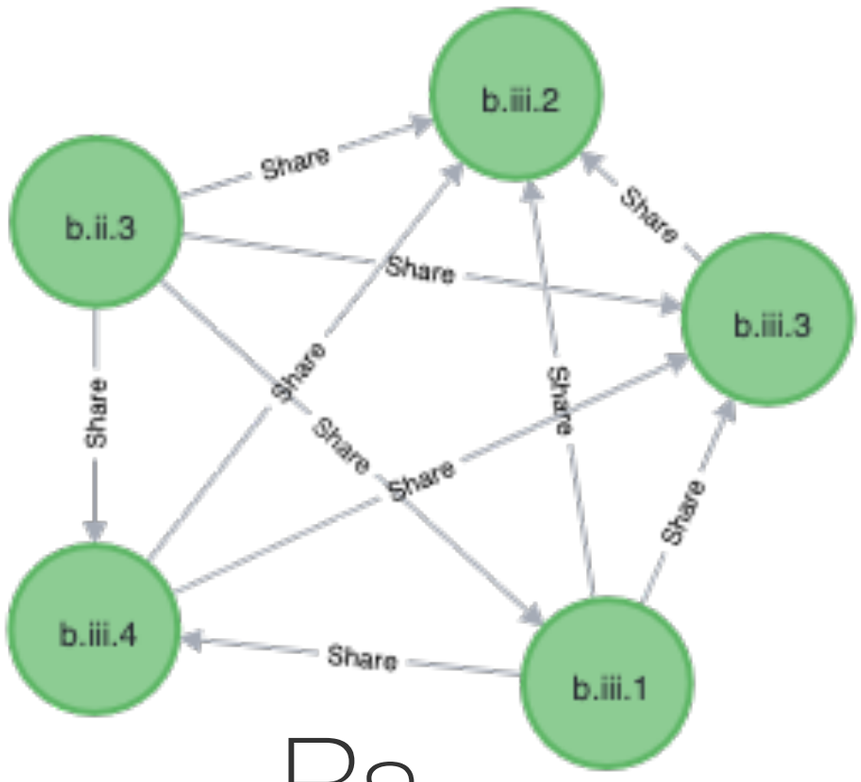

Rs

Individual ranking Limitations

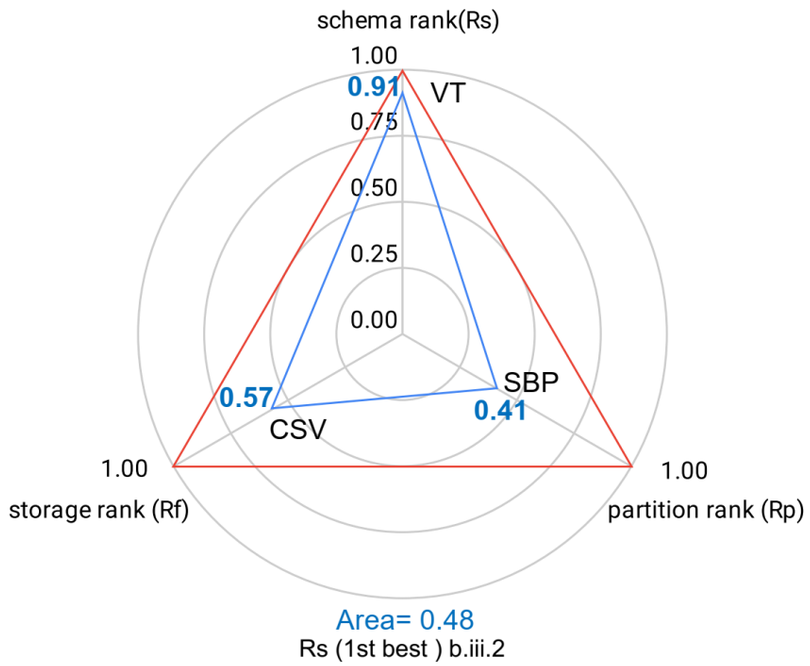

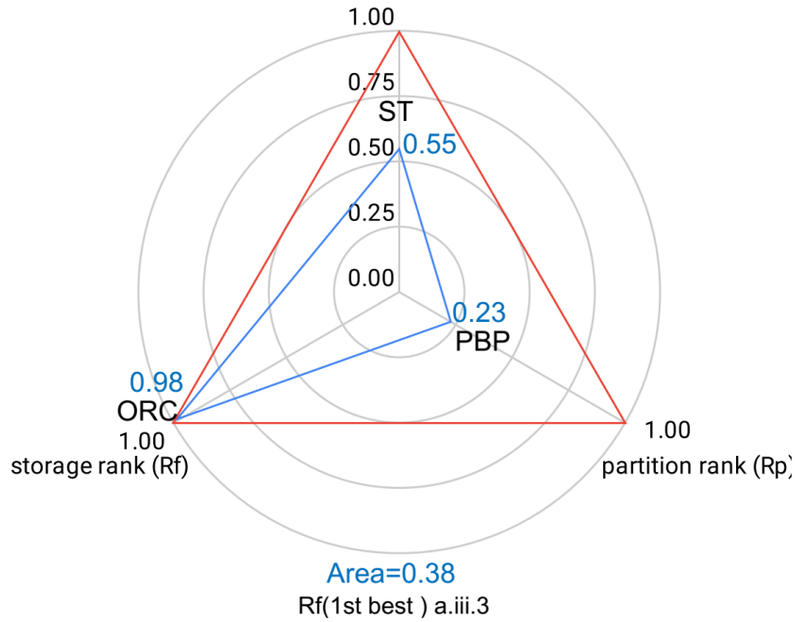

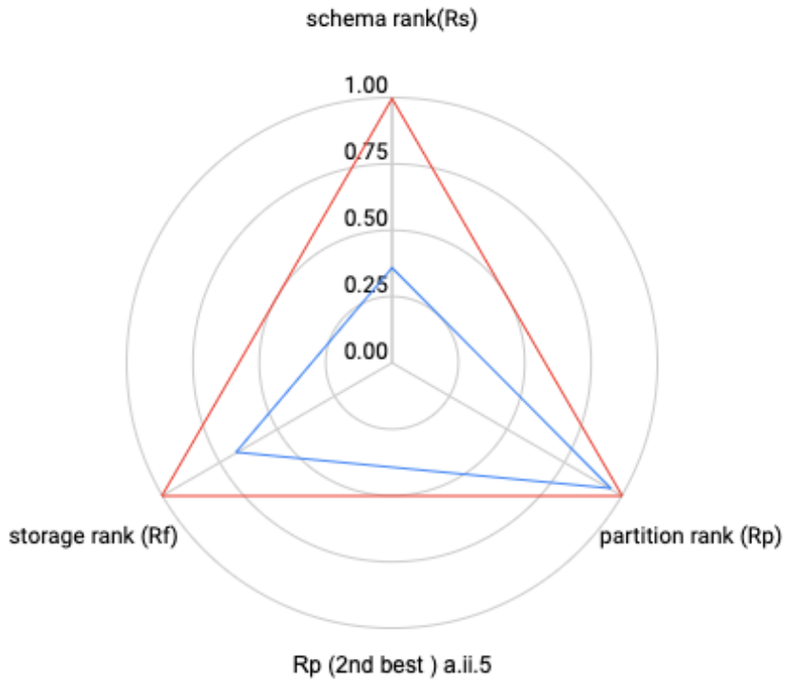

Rank_Schema (Rs)

Rank_Format (Rf)

Rank_Partition. (Rp)

Ranking Criteria Goodness Metrics

Conformance

Conformance with the actual query rankings

1

2

Coherence

How Coherent the ranking across different dataset scales.

1- Ranking Goodness

- A ranking criterion “**good**” if it does not suggest a low-performing configuration.
- We are interested to be the best at any particular query as long as we are never the worst.
- The ranking criterion is confident if it’s top ranked configurations are not actually performing bad.

$$A(\mathcal{R}^k) = 1 - \sum_{i=0}^{Q}\sum_{j=0}^{k}\frac{\bar{A}(i,j)}{Q*k}, \quad \bar{A}(i,j) = \begin{cases} 1 & \mathcal{R}^k[j] \in \mathcal{Q}i_h^i \\ 0 & otherwise \end{cases}$$

Confs.

500M	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11
a.i.1	35	40	37	43	41	34	19	36	22	25	32
a.i.2	45	45	45	45	45	45	30	45	30	36	40
..	..	..	..	..	..	..	..	..	..	..	..
c.iii.5	41	42	39	37	38	28	31	2	31	43	43

Ranks

Confs. Query Ranking

Goodness example

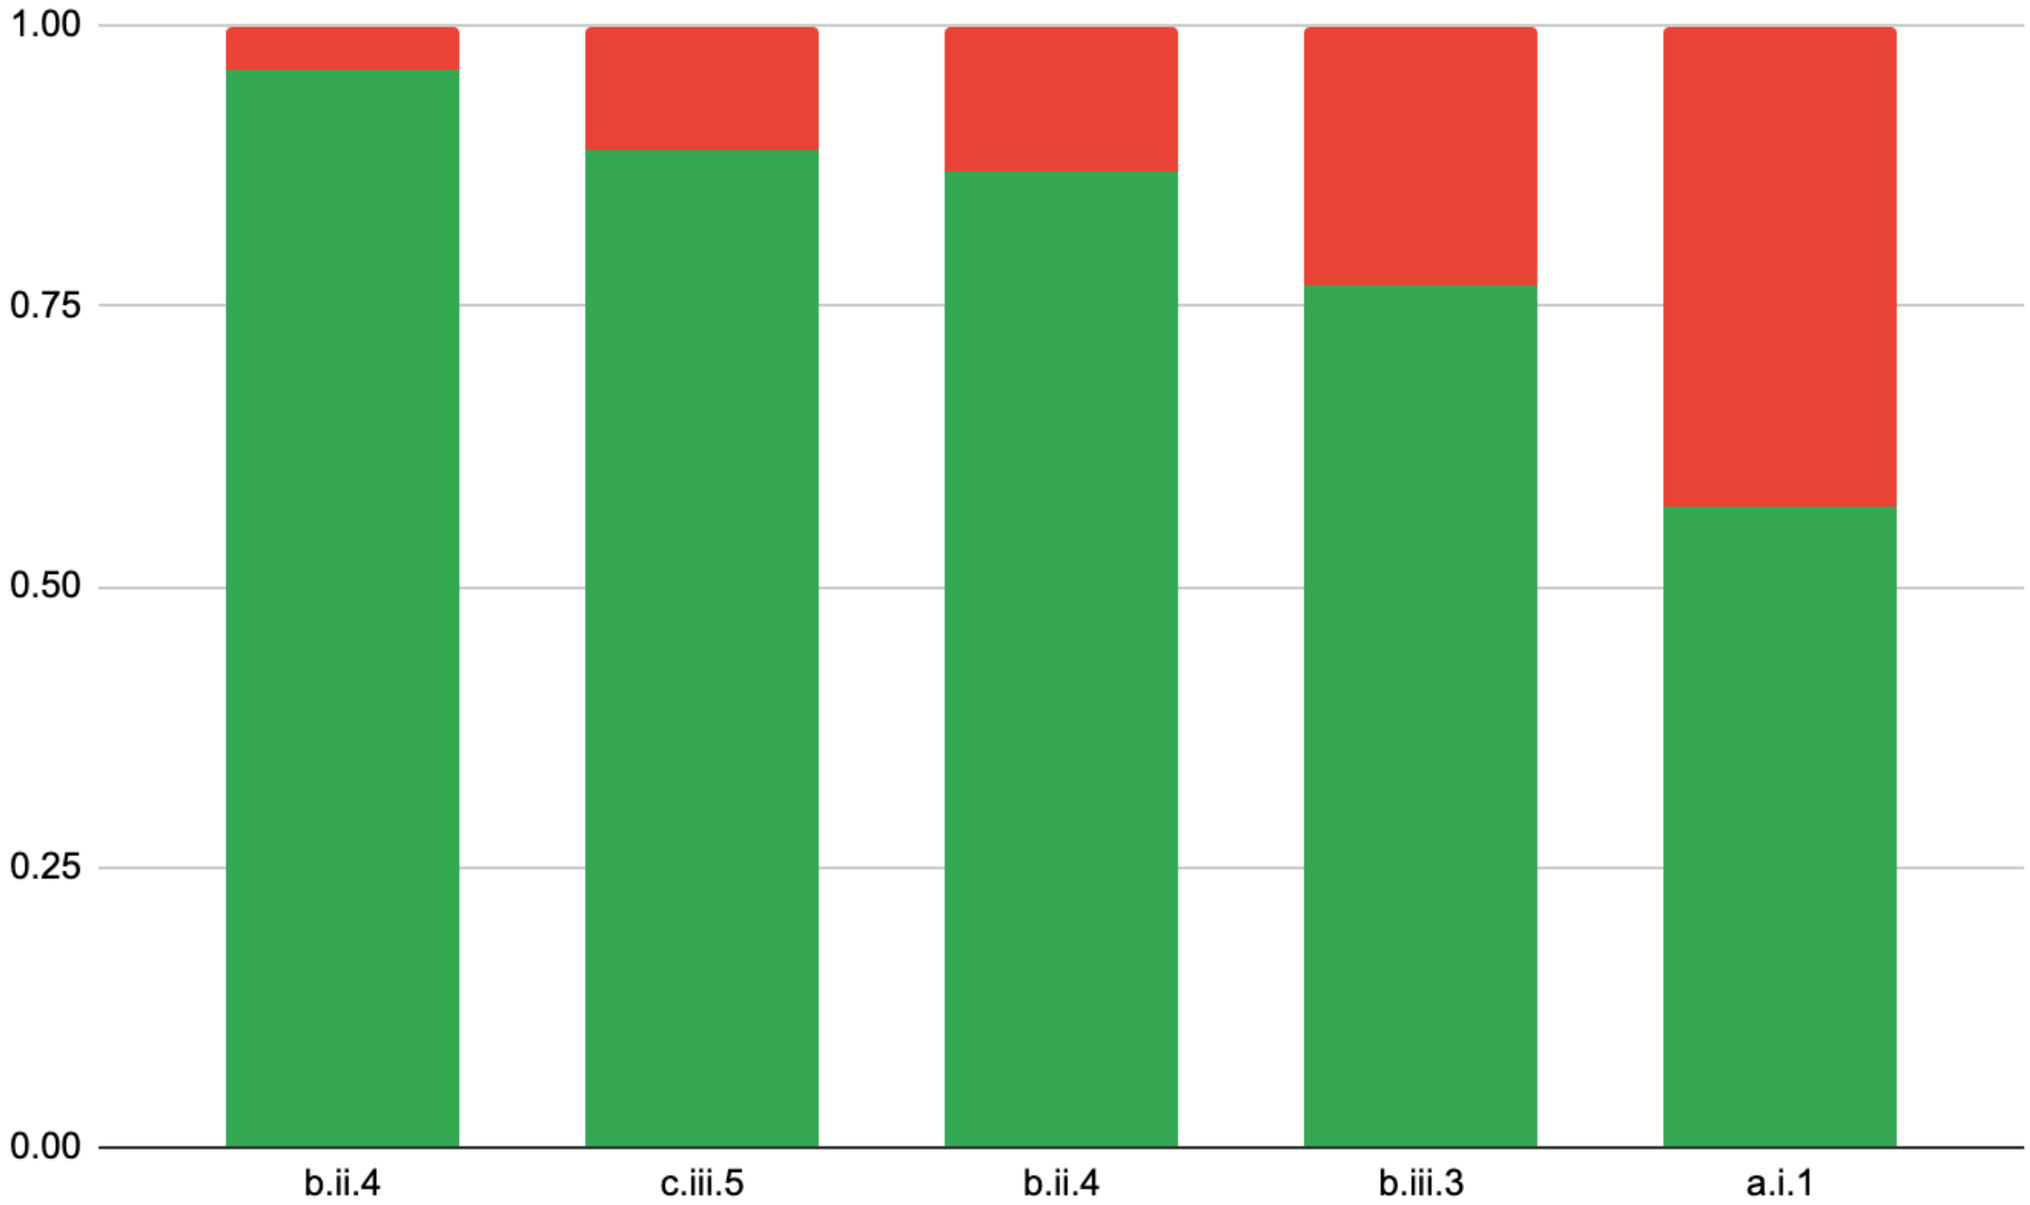

Ranking by Storage (Rf)

2- Ranking Coherence

- We opt for **Kendall index**, which counts the number of pairwise disagreements between two rank sets
- The larger the distance, the more dissimilar the rank sets are.

$$K\left(\mathcal{R}1,\mathcal{R}2\right)=\sum_{\{i,j\}\in P}\frac{\bar{K}_{i,j}\left(\mathcal{R}1,\mathcal{R}2\right)}{\mid P\mid}$$

$$\bar{K}_{i,j}\left(\mathcal{R}1,\mathcal{R}2\right)=\begin{cases}0 & \mathcal{R}1[r_i^1]=\mathcal{R}2[r_i^2]=i\ \wedge\mathcal{R}1[r_j^1]=\mathcal{R}2[r_j^2]=j\ \wedge\\ & r_i^1-r_j^1=r_i^2-r_j^2\\ 1 & \text{otherwise}\end{cases}$$

Kendall rank correlation coefficient: https://en.wikipedia.org/wiki/Kendall_rank_correlation_coefficient

100M

Rf

100M

Rp

100M

Rs

100M

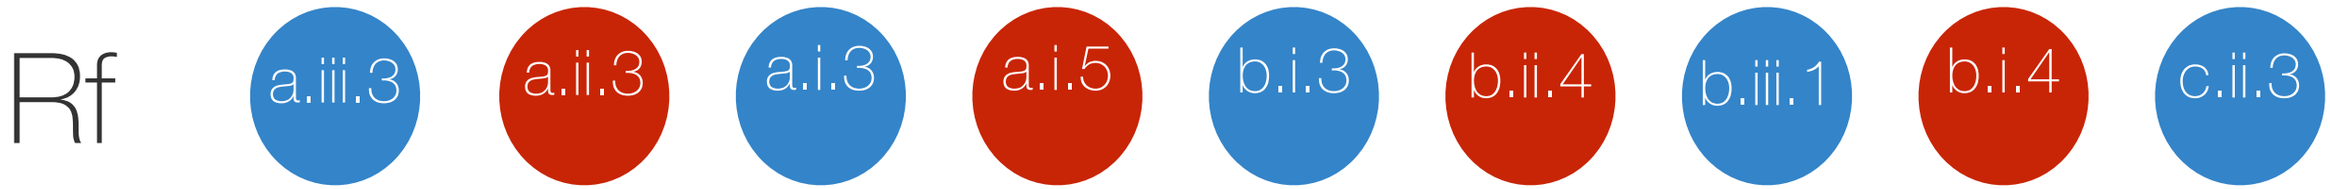

250M

500M

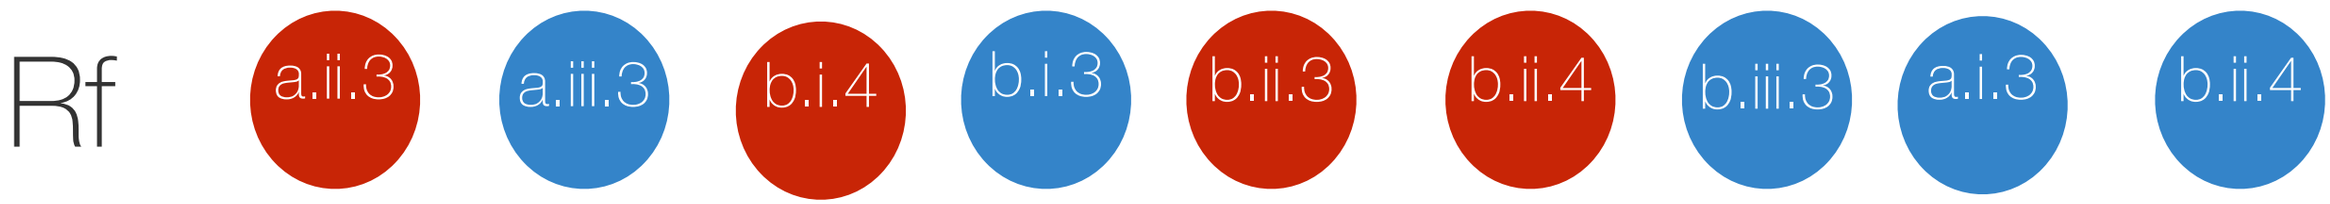

Experiments showed that Individual ranking criteria have:

- High coherence.
- Low conformance.

Bench-Ranking as a multi-Objective Optimiz. Problem

- We look at the Bench-ranking as a multi-objective problem.
- We aim to maximize the three dimensions altogether.
- We opt for the standard Pareto Front. algorithm (NSGA2).

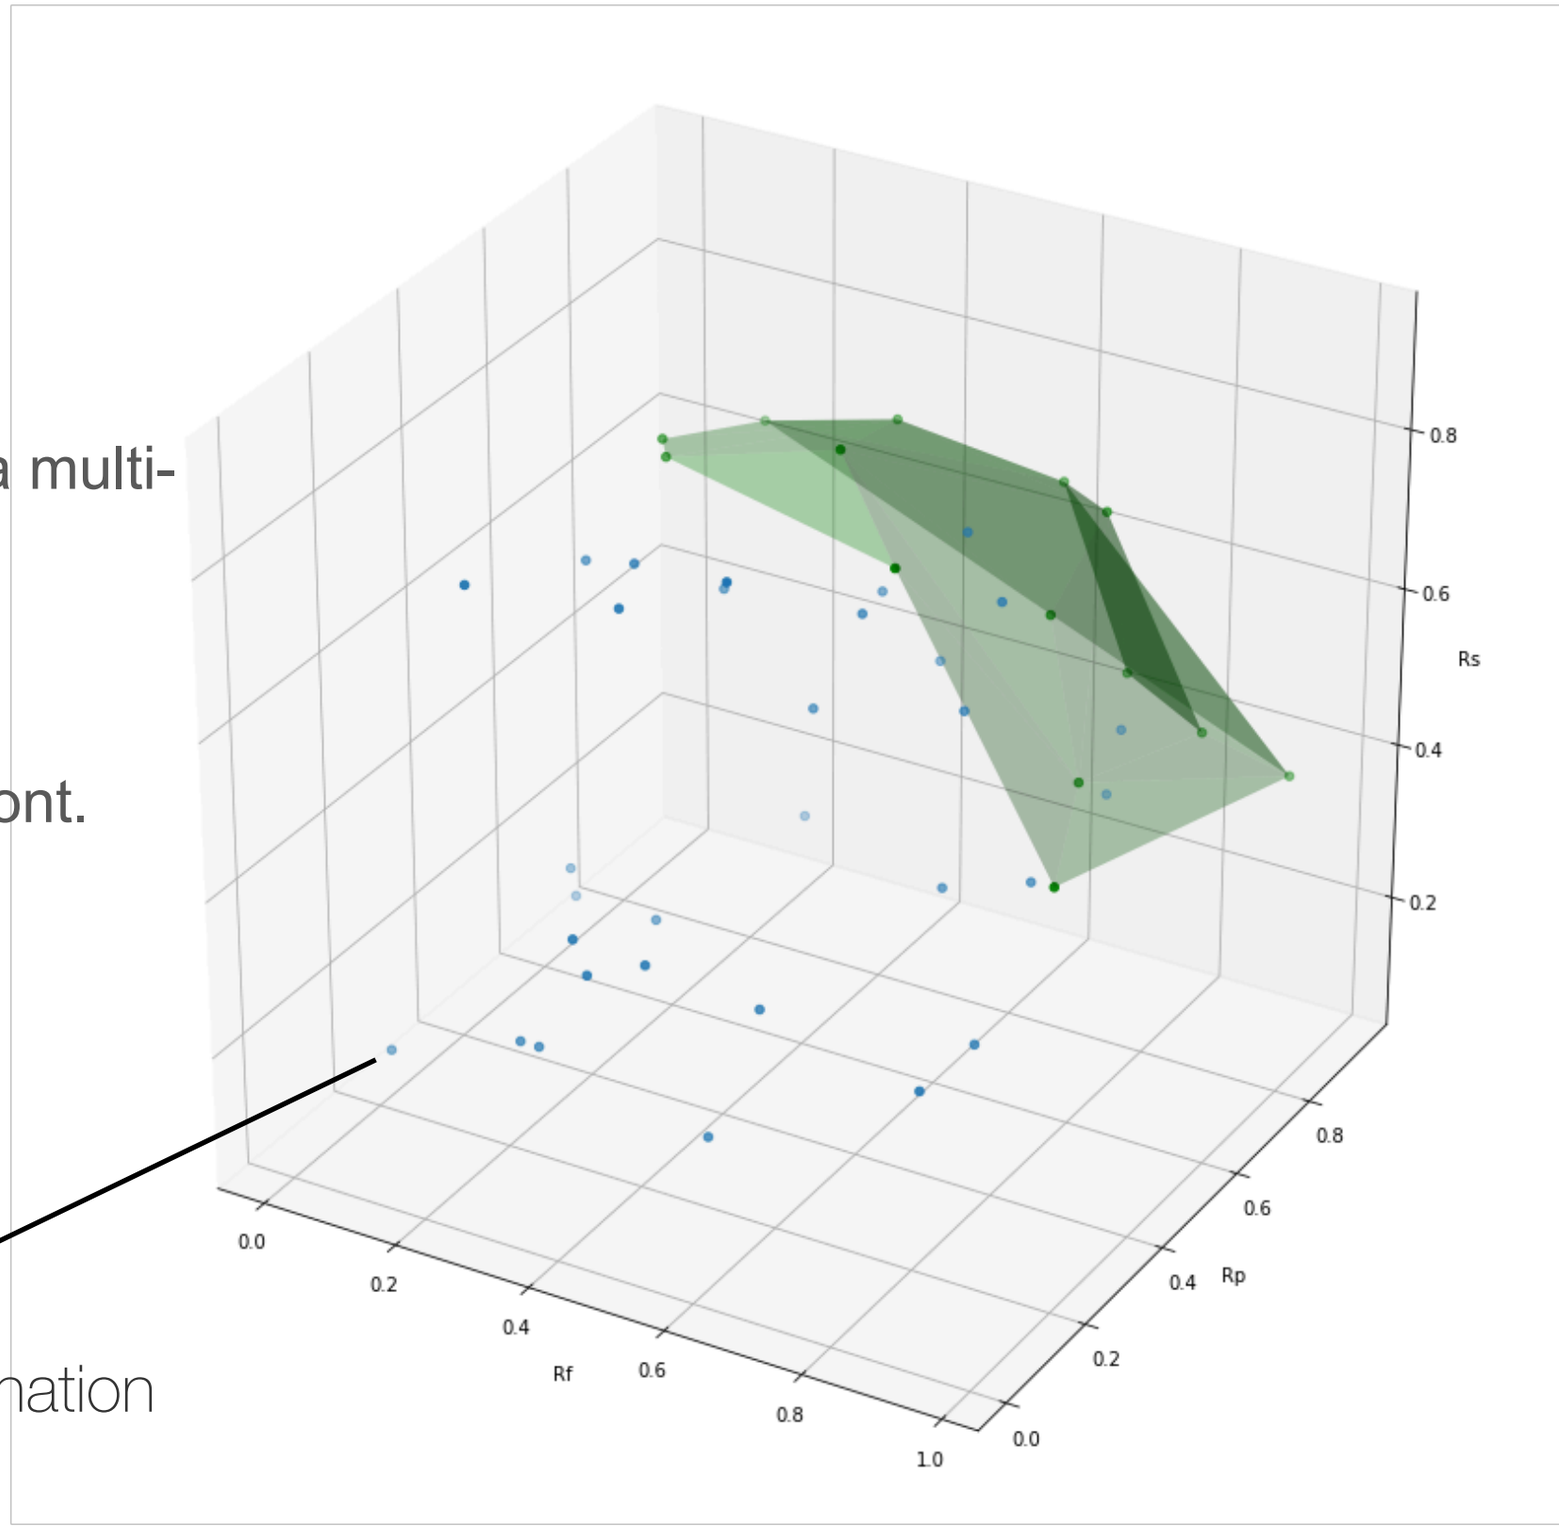

An item of configuration combination

e.g, a.i.4

A fast and elitist multiobjective genetic algorithm: NSGA-II: https://ieeexplore.ieee.org/document/996017

Conclusions

BD Prescriptive Analysis

There is still gap in BD Prescriptive analysis.

Case study

We worked the case study of Processing RDF graphs in the realm of Relational world (Apache Spark).

Bench-Ranking

Simple Yet accurate, in the cases of Selecting the best configuration combination out of this complex solution

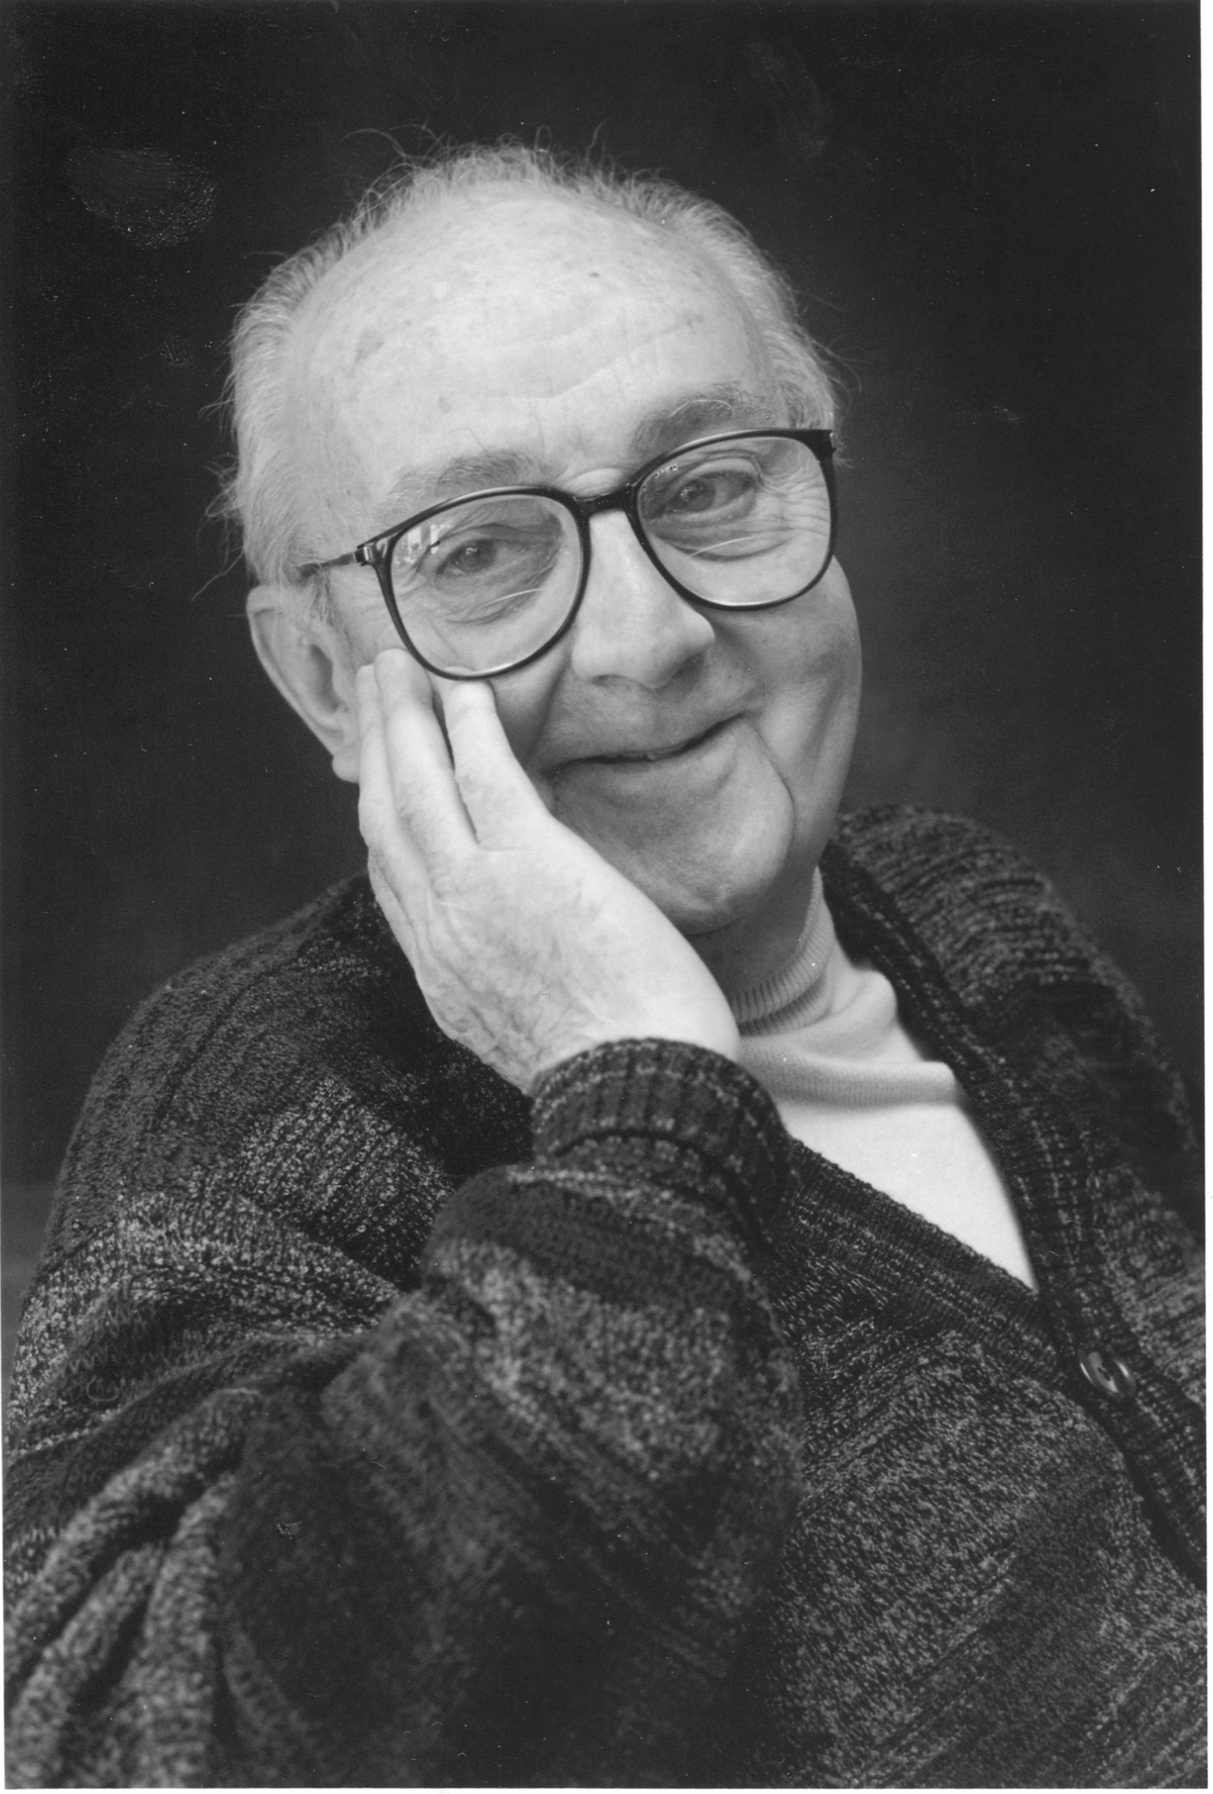

"All the ~~models~~ rankings are wrong but some of them are useful!" ~~George Box~~ Ragab

https://github.com/DataSystemsGroupUT/SPARKSQLRDFBenchmarking