

# SoFOS: Demonstrating the Challenges of Materialized View Selection on Knowledge Graphs

Georgia Troullinou  
FORTH-ICS  
troulin@ics.forth.gr

Matteo Lissandrini  
Aalborg University  
matteo@cs.aau.dk

Haridimos Kondylakis  
FORTH-ICS  
kondylak@ics.forth.gr

Davide Mottin  
Aarhus University  
davide@cs.au.dk

## ABSTRACT

Analytical queries over RDF data are becoming prominent as a result of the proliferation of knowledge graphs. Yet, RDF databases are not optimized to perform such queries efficiently, leading to long processing times. A well known technique to improve the performance of analytical queries is to exploit *materialized views*. Although popular in relational databases, view materialization for RDF and SPARQL has not yet transitioned into practice, due to the non-trivial application to the RDF graph model. Motivated by a lack of understanding of the impact of view materialization alternatives for RDF data, we demonstrate SoFOS, a system that implements and compares several cost models for view materialization. SoFOS is, to the best of our knowledge, the first attempt to adapt cost models, initially studied in relational data, to the generic RDF setting, and to propose new ones, analyzing their pitfalls and merits. SoFOS takes an RDF dataset and an analytical query for some facet in the data, and compares and evaluates alternative cost models, displaying statistics and insights about time, memory consumption, and query characteristics.

## CCS CONCEPTS

• **Computing methodologies** → **Semantic networks**; • **Information systems** → **Database views**.

### ACM Reference Format:

Georgia Troullinou, Haridimos Kondylakis, Matteo Lissandrini, and Davide Mottin. 2021. SoFOS: Demonstrating the Challenges of Materialized View Selection on Knowledge Graphs. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD '21)*, June 20–25, 2021, Virtual Event, China. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3448016.3452765>

## 1 INTRODUCTION

Companies of all types and sectors, such as Amazon, Google, Bosh, and Zalando, use the graph model to represent and store their enterprise knowledge bases [10]. Moreover, large knowledge repositories

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*SIGMOD '21, 2021, Virtual Event, China*

© 2021 Association for Computing Machinery.  
ACM ISBN 978-1-4503-8343-1/21/06...\$15.00  
<https://doi.org/10.1145/3448016.3452765>




Figure 1: An example Knowledge Graph.

are now available with a wide range of information in many different domains – DBpedia and WikiData are two notable examples. Most of this knowledge is available as RDF datasets [12] through SPARQL endpoints [2], organized as *knowledge graphs* (KGs). In KGs like the one in Figure 1, nodes represent entities and edges represent relationships and attributes. KGs allow storing a wide range of heterogeneous, factual, and statistical information that forms a valuable asset for businesses, organizations, and individuals.

As more data is stored in KGs, there is an increasing need to enable quick exploration of those sources [15, 16] and to answer more complex queries [10, 14]. However, in SPARQL query processing, the research mainly focuses on queries that identify nodes and edges satisfying some specific conditions (e.g., entities by name, friends of friends, or product categories) [1, 2, 5].

*Example 1.1.* Consider a KG like DBpedia or WikiData storing for each country the list of official languages and the number of people speaking that language in that country. This data can be used to answer analytical queries like “in how many countries is French an official language?” or “what is the total amount of French-speaking population in the American continent?”.

Given the growing importance of KGs as knowledge repositories, there is a need for effective *analytical query* answering to extract relevant insights from the data [3, 8, 14].

The study of analytical queries (i.e., OLAP) over relational systems has attracted substantial attention in the past decades [9] and recently, different methodologies have also been proposed in the context of KGs [3, 6]. Nonetheless, obtaining answers to analytical queries is usually time-consuming and prohibitively expensive for most RDF data-management systems [14]. A technique to improve

the performance of analytical queries is view materialization [7]. View materialization precomputes and stores the results of analytical queries offline to serve new incoming queries faster. Nonetheless, this requires the system to select which views to materialize. In addition, the intricacies of the RDF model, e.g., complex schema, entailment, and blank nodes, further complicate the direct adoption of techniques proposed for the relational data.

A recent work [8] applies an approach designed for relational OLAP [7] to RDF data. Yet, since existing approaches are adaptations of relational techniques, there is no understanding of their appropriateness to knowledge graphs. We shed a light on the use of multiple alternative approaches over KGs by showcasing Sofos, a system that compares various cost models for view materialization. A cost model is the main building block for selecting the views to materialize, as it provides an estimate of the time for querying a database with and without the materialized views.

**Contributions.** Sofos proposes, evaluates, and compares a variety of existing cost models for view selection, adapted for the RDF setting. It allows users to run a set of queries on the materialized views and inspect the performance in executing the query workload. The goal of this prototype is to identify strengths and limitations of multiple cost estimation techniques for view selection on RDF data. In summary, Sofos (1) addresses the problem of providing fast query answering for analytical queries on KGs, (2) provides a generic solution to be deployed on any RDF triple store with SPARQL query processing, and (3) highlights possible limitations of six alternative approaches. Given a KG, a facet over the KG, and a constraint on the number of views to materialize, Sofos generates a set of views to answer aggregated queries over the provided facet.

## 2 RELATED WORKS

KGs gained traction in the last few years, due to the proliferation of Linked Open Data [2, 13, 17] and proprietary enterprise knowledge graphs [10]. Recently, companies and researchers require to perform complex analytics on the data in the form of aggregate queries. In the following, we highlight how existing methods have tried to adapt techniques for relational data to the graph model. In this demonstration, we present a system that can showcase the limitations of these adaptations.

**Data cube analysis.** In relational data, *data cubes* [7] conveniently represent aggregates over multiple data dimensions. That is, they model data as a set of observations, each carrying one or more measures, and a set of dimensions across which the measures of the observations can be aggregated (e.g., consider the population recorded for each city in each country, which can be aggregated across time, regions and continents, or language spoken in order to retrieve, for instance, the amount of population per country speaking each language). Analyses in such data cubes are notoriously computationally expensive since they involve the processing of large portions of the dataset. Therefore, a common approach is that of employing *materialized views* so that queries can be executed over a smaller portion of pre-processed data, significantly reducing query time [7, 9]. For instance, one can pre-aggregate population across countries, languages, and years, so that a query asking for the total amount of people speaking German during 2020 can be computed by processing the pre-aggregated results instead of the

whole data for each city. Yet, given a data-cube with many different dimensions, there are multiple ways in which data could be aggregated (e.g., across cities and regions, or languages and years, and so on). Materializing views for all these combinations is expensive both in terms of processing time as well as in terms of space occupation on disk. Therefore, *view selection* techniques have been proposed for the case of relational databases [7, 9]. These techniques estimate the benefit that materializing a specific view can provide. Such benefit is estimated as a linear function of the size of the materialized view compared against the size of the data from which such a view should be derived. For instance, a view aggregating daily records into yearly records provides an expected reduction factor of  $\sim 350$ , and one would expect a proportional improvement in processing speed when using the view for querying, instead of the daily data.

For the case of RDF data, instead, the state of the art approaches simply set-out to adapt solutions from the relational model to the graph model. Yet, the research on relational data cannot be directly applied on graphs, as the structure and the schema is not known a-priori in such datasets.

**OLAP approaches for RDF.** The MARVEL system [8], belonging to this line of work, implements view materialization for optimizing query answering of OLAP SPARQL queries [4]. MARVEL employs a cost model, a view selection algorithm, and an algorithm for rewriting SPARQL queries using the available materialized views. Although the approach is the first to tackle the challenges of answering analytical queries on KGs through view materialization, the input data should actually adopt a data cube model (in particular the QB4OLAP [4]) and the cost model simply considers the number of edges (triples) in each view.

Other approaches have investigated the need for enabling complex aggregate queries in SPARQL [3, 14]. In particular, the Analytical schema model [3] enables different views on generic KGs. Yet, this model does not tackle the problems of view materialization for RDF data, instead, they propose to map the data to a relational model and exploit traditional optimizations for relational queries. Finally, a distinct approach for RDF analytics [14] converts a complex aggregate query to a set of smaller, approximate, queries. Yet, this approach has the sole goal to diminish the load for the database answering the query, and not to speed up query processing.

Therefore, to date, no solution has explored in detail the case of view materialization for KGs as a graph-centric problem. Instead, existing solutions, simply resort to map the data to a relational model. Sofos aims at systematically analyzing view materialization by shedding a light on existing methods to pave the road to a native graph-aware model for answering analytical queries on KGs.

## 3 THE SOFOS SYSTEM

The Sofos system *implements, adapts and compares several cost models for view selection on RDF KGs*. The system, given an initial analytical *facet* of the graph to analyze, materializes a set of views based on a cost model, and then it measures the performance, in terms of storage cost and query response-time, of the selected views. Sofos comprises of two main modules: ① an *offline module* for **selective view materialization** (Section 3.1), and ② an

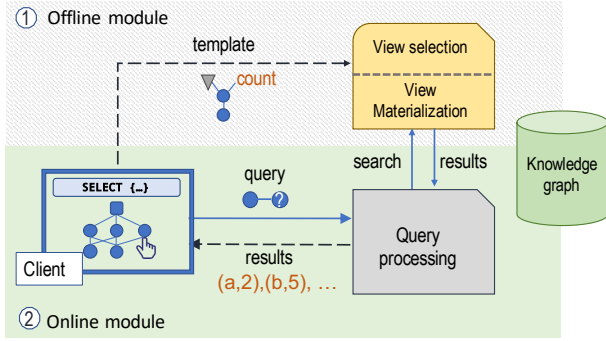


Figure 2: The Sofos system.

online module for query execution and performance comparison (Section 3.2). Figure 2 shows its main components.

**Background & problem:** At its core, the Sofos system takes a knowledge graph  $G$  and an analytical facet  $F$ , which describes the information that should be aggregated in different views, and materializes a set of  $k$  views  $\mathcal{V}_1, \dots, \mathcal{V}_k$  based on  $F$ . Then, given any query  $Q$  targeting  $F$ , the system either answers  $Q$  querying one of the  $k$  materialized views, or accesses the graph  $G$  if none of the views can be used to compute the required answer.

In Sofos, a *knowledge graph*  $G$  is represented as a set of RDF triples  $(s, p, o) \in (\mathcal{I} \cup \mathcal{B}) \times (\mathcal{I}) \times (\mathcal{I} \cup \mathcal{B} \cup \mathcal{L})$ , where  $\mathcal{I}$  is a set of entity identifiers,  $\mathcal{B}$  is a set of “blank” nodes with no identifier, and  $\mathcal{L}$  is a set of literals. A *query*  $Q$  on an RDF graph is a set of *triple patterns*, that is, a set of triples in which some of the triple’s components  $s, p$ , or  $o$  are variables from a set  $\mathcal{X}$ , and is expressed in the SPARQL query language. An *answer* to a query  $Q$  is computed based on the matchings in  $G$  of the triple patterns in the query and the values corresponding to instances of the variables in the query. We denote as  $Q(G)$  the set of query answers on the knowledge graph  $G$ . Here, we focus on *analytical queries* of the kind `SELECT  $\vec{X}$  agg( $u$ ) WHERE  $P$  GROUP BY  $\vec{X}$` , in which  $\vec{X} \subseteq \mathcal{X}$  are grouping variables, i.e., a subset of the variables appearing in  $P$ ,  $u \in \mathcal{X}$  is the specific variable over which the aggregation is computed, and *agg* is an aggregation expression in  $\{\text{SUM, AVG, COUNT, MAX, MIN}\}$ .

The Sofos system builds on analytical *facets* that determine the triples of the graph that are the target of some queries and hence provide the conditions to construct a set of views. A *facet* has the same form of an analytical query and is then identified by the triple  $F = \langle \vec{X}, P, \text{agg}(u) \rangle$ . Finally, a *view* from a facet  $F$  is a query  $\mathcal{V} = \langle \vec{X}', P', \text{agg}(u) \rangle$ , where  $P'$  is derived from  $P$ , and  $\vec{X}' \subseteq \vec{X}$  aggregates over just a subset of variables in  $\vec{X}$ . Therefore, the facet  $F$  induces a *lattice* of views  $\mathcal{V}(F)$ , in which different subsets of variables are used for aggregation and results are represented at different levels of granularity. Moreover, in Sofos, a materialized view is also an RDF graph that contains an encoding of only the answers to the query used to generate it. Analytical queries targeting a facet  $F$  also contain a subset of  $\vec{X}$  and  $P$  but can be further specialized by also introducing additional `FILTER` conditions.

Given a query  $Q$ , *view materialization* allows for answering the query by exploiting the contents of a precomputed view  $\mathcal{V}_i$ , avoiding in this way the need to query the underlying graph  $G$ . Materializing the entire lattice would allow to always select the best view

$\mathcal{V}_i$  for any query. Nonetheless, materializing the entire lattice is impractical from the memory consumption standpoint. As such, Sofos explores different strategies to select a subset  $\mathcal{V}_1, \dots, \mathcal{V}_k$  of views from the lattice. In the relational case, the system would always select the smallest possible view to answer  $Q$ , since there is a linear correlation between number of tuples and running time [7]. This linear correlation does not trivially hold in the case of knowledge graphs, because a graph is not defined in terms of tuples. As such, we need a cost function  $C : \mathcal{V}(F) \rightarrow \mathbb{R}^+$  predicting the running time of any query  $Q$  if the view  $\mathcal{V}_i$  is materialized.

In practice, to select the best set of views, we adopt a greedy approach [7]. Given a set of selected views, the greedy approach exploits the estimated time from the cost function and compares the expected running time of a set of queries with and without including the candidate view  $\mathcal{V}_i$  in the set of views. While, in the relational case, the cost is derived directly from the number of tuples in the view, Sofos proposes a comparison among different cost functions to select  $k$  views from a facet  $F$ , and shows the advantages and shortcoming of each of them when tested against a specific set of queries. We opt for a budget representing the number of views  $k$  to allow for a more straightforward comparison on memory and time consumption. However, note that this budget can be adapted to regulate the space consumption on the selected views as well, i.e., instead of selecting  $k$  views, select up to  $k$  views within a certain memory budget.

### 3.1 Selective view materialization

Sofos performs two offline operations, (a) **view selection** that decides on the best views to materialize given the cost function  $C$ , and (b) **view materialization** that augments the graph with extra information to store aggregation values.

**View selection.** Sofos supports six cost models: (1) a random baseline, (2) a direct adaptation of tuple counting for relational data, two RDF-based cost models, namely (3) the number of aggregated values and (4) the number of nodes, (5) a learned cost model, and (6) a user-defined one.

- (1) **Random:** This cost function is constant  $C(\mathcal{V}_i) = 1$ , for each view  $\mathcal{V}_i \in \mathcal{V}(F)$ , i.e., this will output a random  $k$ -size subset of  $\mathcal{V}(F)$ .
- (2) **Number of triples:** This cost function is analogous to the number of tuples in relational databases. On a knowledge graph, this cost corresponds to the number of RDF triples in the corresponding graph  $G_{\mathcal{V}_i}$ ,  $C(\mathcal{V}_i) = |G_{\mathcal{V}_i}|$ .
- (3) **Number of aggregated values:** This corresponds to the number of results of the query representing the view, i.e.,  $C(\mathcal{V}_i) = |\mathcal{V}_i(G)|$ .
- (4) **Number of nodes:** This cost corresponds to the number of node values in the view  $\mathcal{V}_i$ , i.e.,  $C(\mathcal{V}_i) = |\mathcal{I}_i \cup \mathcal{B}_i \cup \mathcal{L}_i|$ .
- (5) **Learned cost:** For comparison, we adapt a cost estimate from a learned deep regression model  $f : \mathcal{V}(F) \rightarrow \mathbb{R}$  [11]. We encode a query into a vector representing the relationships, the attributes, and the type of aggregates in the query, along with statistics about the relationship frequency and the attribute frequency. In the offline training phase, the model takes the encoding of either a given workload or randomly generated queries and their running time. In the online phase, the model receives the

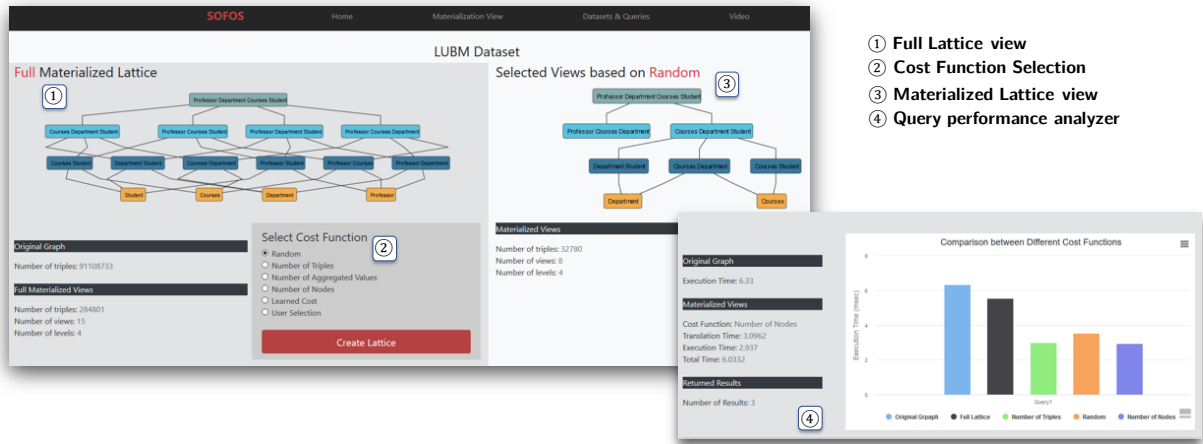


Figure 3: The GUI of Sofos system.

encoding of a query (i.e., view)  $V_i$  and outputs the estimated running time, such that  $C(V_i) = f(V_i)$ .

- (6) **User defined:** The user acts as a cost function, selecting  $k$  views from the lattice.

**View materialization.** View materialization in Sofos consists of generating a new graph for each view  $V_i$ . Each graph contains a set of extra blank nodes to which is attached the value of the aggregation of different bindings for the subset of the template variables in  $\bar{X}$ . This materialization procedure is a generalization of the standard techniques adopted in MARVEL [8]. The result of view materialization is hence an *expanded RDF graph*  $G^+$ .

### 3.2 Query Performance Comparison

After materialization of a specific subset of views, the system runs a set of queries randomly generated from the facet  $F$  against the expanded graph  $G^+$  and measures the performance of each query.

When answering a query, Sofos identifies the best view to adopt and translates the input query  $Q$  into a query  $Q'$  in the expanded RDF graph  $G^+$  targeting the data of the selected view. In practice, the translation straightforwardly substitutes aggregate variables with the blank nodes representing the aggregation and reformulates triples patterns accordingly.

Therefore, Sofos allows running any set of queries on different sets of materialized views for each cost function. The user can then compare the relative performance of each view selection method and hence the appropriateness of different cost models.

## 4 DEMONSTRATION SCENARIO

The goal for the demonstration is to show, through experiments, the challenges involved in materialized view selection on knowledge graphs, exploring various alternative cost models. A screenshot of our system is shown in Figure 3. The demonstration will start by guiding the participants through the different design choices in Sofos. We will then walk them through the following steps:

**Configuration:** In this step, the three datasets used for our demonstration (i.e., the LUBM, the DBpedia, and the Semantic Web Dogfood datasets) will be presented along with the corresponding

query facets for these datasets. Each query facet will be accompanied by a high-level description and a corresponding SPARQL query template, enabling the active exploration of the data available each time. For each dataset we will propose a query workload composed of different parametrized queries for a given query template.

**Exploration of the Full Lattice:** By selecting a specific combination of dataset and facet, the full materialized lattice will be presented to the users, explaining why such a large structure is required, precomputing at the various levels, the aggregations that the query template might ask. By selecting a node (view) in the lattice the user will be able to check the data that are stored for this specific node.

**Exploring Cost Models:** Using the full lattice as input, the various view selection algorithms (and the accompanied cost models) will be explained to the participants and demonstrated in practice. In each case, the trade-off in query execution and storage amplification will be shown, enabling users to understand which cost model is better in each case.

**User Selected Views:** Besides exploiting an already existing view selection algorithm, the users will be able to select individual nodes from the lattice to be materialized and see the impact of their choices on the query execution time. Each time the space amplification and the query execution time will be contrasted, enabling users to explore the sweet-spot where space amplification is minimized and query execution time is improved.

**“Hands-on” Challenge:** In this phase, conference participants would be challenged, so that given a specific query and budget, to optimally select the views to be materialized for optimizing query execution. The participant that will make the best selection will receive a Sofos-related small prize.

## ACKNOWLEDGMENTS

Matteo Lissanrini is supported by the EU’s H2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 838216. This research project was supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “2nd Call for H.F.R.I. Research Projects to support Post-Doctoral Researchers” (iQARuS Project No 1147).

## REFERENCES

- [1] Günes Aluç, Olaf Hartig, M. Tamer Özsu, and Khuzaima Daudjee. 2014. Diversified Stress Testing of RDF Data Management Systems. In *ISWC*.
- [2] Angela Bonifati, Wim Martens, and Thomas Timm. 2019. An analytical study of large SPARQL query logs. *The VLDB Journal* (aug 2019).
- [3] Dario Colazzo, François Goasdoué, Ioana Manolescu, and Alexandra Roatiş. 2014. RDF analytics: lenses over semantic graphs. In *WWW*. ACM, 467–478.
- [4] Lorena Etcheverry and Alejandro A Vaisman. 2012. QB4OLAP: a new vocabulary for OLAP cubes on the semantic web. In *COLD*, Vol. 905. 27–38.
- [5] Yuanbo Guo, Zhengxiang Pan, and Jeff Hefflin. 2005. LUBM: A benchmark for OWL knowledge base systems. *JWS* 3, 2-3 (2005), 158–182.
- [6] Nurefsan Gür, Jacob Nielsen, Katja Hose, and Torben Bach Pedersen. 2017. GeoSemOLAP: Geospatial OLAP on the Semantic Web Made Easy. In *WWW Companion*.
- [7] Venky Harinarayan, Anand Rajaraman, and Jeffrey D Ullman. 1996. Implementing data cubes efficiently. *SIGMOD Rec.* 25, 2 (1996), 205–216.
- [8] Dilshod Ibragimov, Katja Hose, Torben Bach Pedersen, and Esteban Zimányi. 2016. Optimizing aggregate SPARQL queries using materialized RDF views. In *ISWC*. 341–359.
- [9] Tapio Niemi, Jyrki Nummenmaa, and Peter Thanisch. 2001. Constructing OLAP cubes based on queries. In *DOLAP*. 9–15.
- [10] Natasha Noy, Yuqing Gao, Anshu Jain, Anant Narayanan, Alan Patterson, and Jamie Taylor. 2019. Industry-scale knowledge graphs: Lessons and challenges. *ACM Queue* 17, 2 (2019), 48–75.
- [11] Jennifer Ortiz, Magdalena Balazinska, Johannes Gehrke, and S Sathya Keerthi. 2019. An empirical analysis of deep learning for cardinality estimation. *arXiv preprint arXiv:1905.06425* (2019).
- [12] W3C RDF Working Group. 2014. *Resource description framework*. W3C. <http://www.w3.org/RDF/>.
- [13] Andy Seaborne and Eric Prud'hommeaux. 2006. *SPARQL query language for RDF*. Technical Report. W3C.
- [14] Arnaud Soulet and Fabian M. Suchanek. 2019. Anytime Large-Scale Analytics of Linked Open Data. In *ISWC*. 576–592.
- [15] Georgia Troullinou, Haridimos Kondylakis, Evangelia Daskalaki, and Dimitris Plexousakis. 2015. RDF Digest: Ontology Exploration using Summaries. In *ISWC*.
- [16] Georgia Troullinou, Haridimos Kondylakis, Kostas Stefanidis, and Dimitris Plexousakis. 2018. Exploring RDFS KBs Using Summaries. In *ISWC*. 268–284.
- [17] Marcin Wylot, Manfred Hauswirth, Philippe Cudré-Mauroux, and Sherif Sakr. 2018. RDF Data Storage and Query Processing Schemes: A Survey. *ACM Comput. Surv.* 51, 4 (2018), 84:1–84:36.