

# On-Demand and Lightweight Knowledge Graph Generation – a Demonstration with DBpedia

Malte Brockmeier, Yawen Liu, Sunita Pateer,  
Sven Hertling<sup>[0000–0003–0333–5888]</sup>, and Heiko Paulheim<sup>[0000–0003–4386–8195]</sup>

Data and Web Science Group  
University of Mannheim, Germany  
{mbrockme,yawliu,spateer}@mail.uni-mannheim.de  
{sven,heiko}@informatik.uni-mannheim.de  
<http://dws.informatik.uni-mannheim.de/>

**Abstract.** Modern large-scale knowledge graphs, such as DBpedia, are datasets which require large computational resources to serve and process. Moreover, they often have longer release cycles, which leads to outdated information in those graphs. In this paper, we present *DBpedia on Demand* – a system which serves DBpedia resources on demand without the need to materialize and store the entire graph, and which even provides limited querying functionality.

**Keywords:** Knowledge Graph · On Demand · Live · Lightweight · Knowledge Graph as a Service · DBpedia

## 1 Introduction

Knowledge graphs on the Web – such as DBpedia [6], YAGO [9], or Wikidata [11] – are useful building blocks for intelligent applications [3]. While they differ considerably in coverage and quality [1], in many cases, timeliness of information is a crucial property in some use cases as well. Since many public knowledge graphs have rather long release cycles, sometimes up to a few months or even years [3], the information served in these knowledge graphs may be easily outdated. One approach which has been proposed to close this gap is *DBpedia Live*, which provides fresh data extracted from Wikipedia as a knowledge graph [7], and which, after a longer downtime, is available again since 2019.<sup>1</sup>

Another challenge related to such knowledge graphs is their sheer size. The creation of those graphs is often a longer running and computation-intensive process, and even serving the knowledge graph online requires quite a bit of computing power. The hosted services for hosting a static DBpedia copy or running DBpedia live cost around \$1 per hour at the time of writing this paper.<sup>2</sup>

<sup>0</sup> Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

<sup>1</sup> <https://www.dbpedia.org/blog/dbpedia-live-restart-getting-things-done/>

<sup>2</sup> <https://aws.amazon.com/marketplace/seller-profile?id=1b9b499d-15e3-479d-a644-da16d45c40a7>

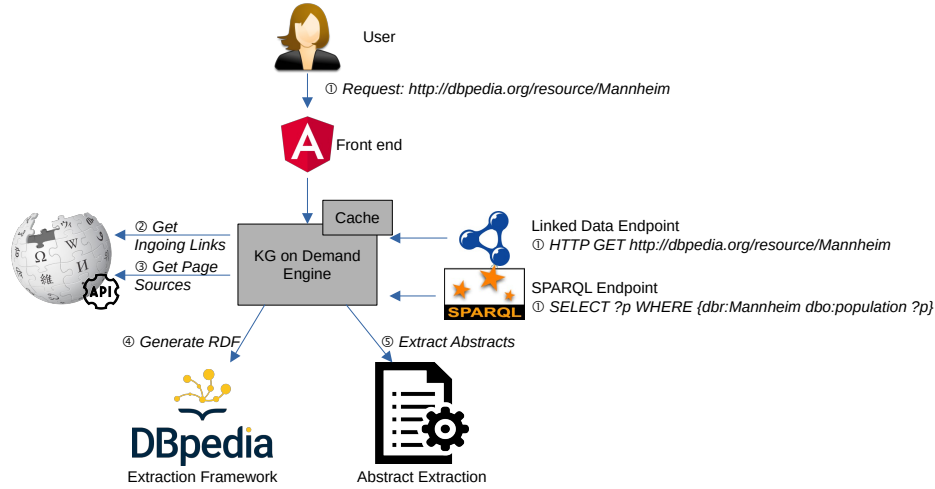


Fig. 1. Overall System Architecture of DBpedia on Demand

In this paper, we introduce *DBpedia on Demand* – an installation of DBpedia which can be run on a local machine as a service. It creates knowledge graph resources on request without materializing an entire graph, which makes its computing requirements rather minimal. It provides a Linked Data endpoint as well as limited SPARQL querying properties.<sup>3</sup>


## 2 Prototype

DBpedia is built from Wikipedia by extracting information from infoboxes and other bits of structured contents, such as categories. To that end, it uses *mappings* from infoboxes to a central ontology, which are also stored in a Wiki. When creating a release, a dump of Wikipedia is processed, and the page contents are transformed to knowledge graph instances using the mappings defined in the Mappings Wiki [6].

The idea of KG on demand is different. Instead of creating a full DBpedia copy, it creates instances upon request: when a user requests a certain resource, the corresponding Wikipedia page is downloaded and processed in real time. Mappings can be downloaded from the mappings Wiki regularly, which allows for using at least recent, although not up to date mappings. Fig. 1 shows the overall architecture of the KG on Demand engine. The process runs in five steps:

1. The URI for which to generate the RDF for is obtained, either through the user interface, a HTTP request, or by extracting it from a SPARQL query.
2. The ingoing links (i.e., Wikipedia pages pointing to the page which corresponds to the requested resource) to that resource are obtained from the Wikipedia API.

<sup>3</sup> Code available at <https://github.com/mbrockmeier/KnowledgeGraphOnDemand>



**Fig. 2.** Performance of Knowledge Graph on Demand

3. The page sources for both the requested resources and the resources of the pages linking to the resource are obtained via the Wikipedia API.
4. For all those pages, the corresponding RDF representation is generated via the DBpedia Extraction Framework.
5. Additionally, the abstract for the resource at hand is generated by extracting the first few sentences from the Wikipedia page.

In the DBpedia knowledge graph, outgoing edges from a resource are generated from the corresponding Wikipedia page’s infobox. Hence, for generating only the outgoing edges, processing one single Wikipedia page is sufficient. In contrast, ingoing edges are outgoing edges of *other* Wikipedia pages. In order to generate those, we first query the Wikipedia API for other pages that link to the one at hand, and then process those as well in order to generate the ingoing edges.

Since the abstract extraction in DBpedia requires the creation and processing of a local copy of the entire Wikipedia, we developed our own simple abstract extraction mechanism instead of using the original one.

The number of ingoing Wikipedia links is the key factor that influences the performance of the proposed approach. Fig. 2 depicts the performance for extracting resources from Wikipedia pages.<sup>4</sup> It can be observed that the system scales linearly with the number of ingoing links.<sup>5</sup>

The approach of generating resources for URIs also determines the kind of SPARQL queries that our approach can process. Generally, all queries that only consider simple ingoing and outgoing edges of a fixed resource can be handled, such as

```
SELECT ?actor WHERE { ?actor dbo:starring dbr:Lost_Highway .}
SELECT ?director WHERE { dbr:Lost_Highway dbo:director ?director .}
```

On the other hand, our approach cannot handle queries without a fixed resource, or queries with a path length greater than one, such as:

<sup>4</sup> The tests were run on an Intel Core i7-8700K processor, 32GB of DDR4 RAM, a 1TB NVMe SSD, and an internet connection with a downstream rate of 100MBit/s. The reported processing times are the averages of 10 runs.

<sup>5</sup> Since the number of ingoing links are the main driver of performance, the maximum number of ingoing links to consider can also be reduced in the configuration of our implementation. By default, all ingoing links are processed.

**Table 1.** Comparison of DBpedia Live and DBpedia on Demand

	DBpedia Live	DBpedia on Demand
Storage requirements	high	low
Online runtime	high	low
Response time	fast	slower
Latest mapping version	yes	approximate
SPARQL interface	full	limited

```

SELECT ?actor ?movie WHERE { ?actor dbo:starring ?movie .}
SELECT ?director WHERE { dbr:Tom_Cruise dbo:starring ?movie .
                          ?movie dbo:director ?director .}

```

In other words, the queries are limited to patterns with at least one resource fixed, and all query variables being at most one hop away from that resource.

### 3 DBpedia Live vs. DBpedia on Demand

DBpedia live builds a full copy of DBpedia. This copy is then updated by monitoring the update stream of Wikipedia. Whenever a change is made in Wikipedia, DBpedia live is notified and processes the change by re-computing the corresponding instance representation in the knowledge graph. Likewise, it is notified on changes in the Mappings wiki (e.g., if the mapping of an infobox to an ontology class is changed), and the affected instances can then be changed as well.

By design, DBpedia live always stores a local materialized copy of the full knowledge graph, which leads to high computational costs of serving the knowledge graph endpoint (i.e., it can be run on commodity hardware). Since all requests are processed on that local copy, they can be handled faster than in DBpedia on demand, and also all types of SPARQL queries can be answered. At the same time, DBpedia Live requires a lot of online runtime to constantly process the stream of changes from Wikipedia. Table 1 summarizes the main differences between DBpedia Live and DBpedia on Demand.

### 4 Conclusion and Outlook

In this paper, we have shown an approach which generates a knowledge graph like DBpedia in an on-demand and lightweight fashion. The approach is very resource-efficient and can be run on commodity hardware. Hence, it is an interesting building block for applications which use a knowledge graph, but do neither want to rely on a public endpoint, nor materialize an entire knowledge graph. The approach can be seen as a complement to DBpedia Live, in comparison to which it has different advantages and disadvantages.

While the current implementation is a prototype only, there are some shortcomings which we have inherited from the DBpedia Extraction Framework. First,

that framework is not capable of multithreading, which makes the current implementation suitable for a local service, rather than for setting up a public endpoint. Second, we have currently encapsulated the *entire* extraction framework, which, however, comes with a significant ramp-up time. Both issues could be addressed by branching and refactoring the extraction framework’s codebase.

While this demonstration has been based on DBpedia, it can be transferred to other approaches as well. With the same mechanism, it would be possible to use the extraction code of other Wikipedia-based knowledge graphs, such as YAGO [9] or the Linked Hypernym extraction [5], as well as to transfer the approach to other Wikis [4]. Also other refinement operators which are local to a Wikipedia page, such as extracting relations from text [2], would be applicable.

While the SPARQL capabilities in our approach are still basic, the approach could be extended towards supporting more complex queries. Possible approaches would be incremental solving of queries, or first using the Wikipedia link graph for obtaining candidates, and then running the extraction on the candidates to narrow down the solution set. For implementing a query endpoint, it would also be interesting to build a (possibly limited) linked data fragments endpoint [10] with the approach at hand, since many of the basic building blocks of those endpoints (e.g., paging) can be directly transferred to our approach.

Apart from adding querying capabilities, the endpoint could also be enriched with an on-demand approach of generation of *embedding vectors*, such as *RDF2vec Light* [8]. That way, downstream applications could leverage both explicit as well as latent representations of the entities in the knowledge graph.

## References

1. Färber, M., Bartscherer, F., Menne, C., Rettinger, A.: Linked data quality of dbpedia, freebase, opencyc, wikidata, and yago. *Semantic Web* **9**(1), 77–129 (2018)
2. Heist, N., Hertling, S., Paulheim, H.: Language-agnostic relation extraction from abstracts in wikis. *Information* **9**(4), 75 (2018)
3. Heist, N., Hertling, S., Ringler, D., Paulheim, H.: Knowledge graphs on the web-an overview. (2020)
4. Hertling, S., Paulheim, H.: Dbkwik: extracting and integrating knowledge from thousands of wikis. *Knowledge and Information Systems* **62**(6), 2169–2190 (2020)
5. Kliegr, T.: Linked hypernyms: Enriching dbpedia with targeted hypernym discovery. *Journal of Web Semantics* **31**, 59–69 (2015)
6. Lehmann, J., et al.: Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic web* **6**(2), 167–195 (2015)
7. Morsey, M., Lehmann, J., Auer, S., Stadler, C., Hellmann, S.: Dbpedia and the live extraction of structured data from wikipedia. *Program* (2012)
8. Portisch, J., Hladik, M., Paulheim, H.: Rdf2vec light—a lightweight approach for knowledge graph embeddings. *arXiv preprint arXiv:2009.07659* (2020)
9. Rebele, T., et al.: Yago: A multilingual knowledge base from wikipedia, wordnet, and geonames. In: *ISWC*. pp. 177–185. Springer (2016)
10. Verborgh, R., et al.: Triple pattern fragments: a low-cost knowledge graph interface for the web. *Journal of Web Semantics* **37**, 184–206 (2016)
11. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Communications of the ACM* **57**(10), 78–85 (2014)