

ABSTRACT: More and more applications use the RDF framework as their data model and RDF stores to index and retrieve their data. Many of these applications require both structured queries as well as fulltext search. SPARQL addresses the first requirement in a standardized way, while fulltext search is provided by store-specific implementations. RDF benchmarks enable developers to compare structured query performance of different stores, but for fulltext search on RDF data no such benchmarks and comparisons exist so far. In this paper, we extend the LUBM benchmark with synthetic scalable fulltext data and corresponding queries for fulltext-related query performance evaluation. Based on the extended benchmark, we provide a detailed comparison of fulltext search features and performance of the most widely used RDF stores. Results show interesting RDF store insights for basic fulltext queries (classic IR queries) as well as hybrid queries (structured and fulltext queries). Our results are not only valuable for selecting the right RDF store for specific applications, but also reveal the need for performance improvements for certain kinds of queries.