

Preface of MEPDaW, SeWeBMeDA and SWeTI 2018*

Oya Beyan¹, Jeremy Debattista², Stefan Decker¹, Javier D. Fernández³, Ali Hasnain⁴, Muhammad Intizar Ali⁴, Pankesh Patel⁵, Dietrich Rebholz-Schuhmann⁴, Amit Sheth⁶, Dhavalkumar Thakker⁷, Jürgen Umbrich³, and Maria-Esther Vidal⁸

¹ Fraunhofer FIT, RWTH Aachen University
{beyan, decker}@dbis.rwth-aachen.de

² Adapt Centre, Trinity College Dublin, Ireland
debattij@tcd.ie

³ Vienna University of Economics and Business, Vienna, Austria
{javier.fernandez, juergen.umbrich}@wu.ac.at

⁴ Insight Centre for Data Analytics, National University of Ireland, Galway, Ireland
{ali.hasnain, ali.intizar, rebholz}@insight-centre.org

⁵ Fraunhofer CESE, US
Appatel@fc-md.umd.edu

⁶ Kno.e.sis – Wright State University, Ohio, USA
amit.sheth@wright.edu

⁷ University of Bradford, UK
D.Thakker@bradford.ac.uk

⁸ Technische Informationsbibliothek (TIB)
maria.vidal@tib.eu

Abstract. This joint volume of proceedings gathers together papers from the 4th Workshop on Managing the Evolution and Preservation of the Data Web (MEP-DaW), the 2nd Workshop on Semantic Web solutions for large-scale biomedical data analytics (SeWeBMeDA), and the Workshop on Semantic Web of Things for Industry 4.0 (SWeTI), held on the 3rd and 4th of June of 2018 during the 15th ESWC conference in Heraklion, Crete, Greece.

1 Managing the Evolution and Preservation of the Data Web

The 4th edition of this workshop has targeted one of the emerging and fundamental problems in the Semantic Web, specifically the preservation of evolving linked datasets. This topic is of particular relevance to the Semantic Web community since it raises awareness of the many research challenges for preserving and managing dynamic linked datasets. Fostering active usage of such evolving datasets requires further research advances on topics such as storage, synchronisation, change representation and querying over evolving graphs.

This year, we accepted three papers, we invited a keynote speaker, and we discussed on future steps of the community, which we describe in brief.

In this year's contributions we see a focus on the management of data versioning and the preservation of evolving knowledge. Singh et al. [6] present DELTA-LD, a change

* Joint proceedings are publicly available in [1].

detection mechanisms for linked datasets. DELTA-LD focuses on detecting changes at both resource level (creation, removal, update, movement, or renewal of a resource) and triple level (deleting or adding a triple). To do so, the approach considers (i) the extraction of features from the linked datasets in order to detect changes and identify similar representations in different versions (i.e. moved resources), and (ii) a classification of the changes and a representation of the change model using a provided ontology.

Pandit et al. [5] investigate on how to represent changes in consents and activities regarding the novel General Data Protection Regulation (GDPR). In their position paper, they first discuss the use of PROV to represent the provenance of activities and ODRL to represent the consent, and identify the influence of consent changes. Then, they discuss on detecting and representing change in activities and how to link and use the changes to demonstrate the compliance w.r.t DDPR obligations.

Laajimi et al. [4] focus on evaluating the performance of archiving engines. In particular, they propose and evaluate the use of the SPARK distributed system to archive RDF data. Thus, authors represent RDF data and changes in SPARK dataframes, while archiving queries are resolved via SPARK SQL. Then, the performance of different versioning approaches (e.g. fully materialized version versus representing only the changing triples in each version) are evaluated, with particular attention to measuring the different performance of star- and chain queries.

Furthermore, in this workshop, Miel Vander Sande keynote⁹ provides an in-depth review on the current state of affairs, lessons and challenges when preserving Linked Data. His talk shows, first of all, the importance of preservation (e.g. for link preservation and content and concept drift management) and its roots in Web archiving. In this regards, current strategies can be seen from an observational point of view (discrete snapshots) or being perceived as a continuous flow where the changes can be provided/detected by different ways (e.g. via versioning system or notification-based). In the talk, Miel Vander Sande shows that, although there are (increasingly popular) technical solutions (i.e. products), the Linked Data archiving process (with some building blocks such as interfaces, change detection, publishing, crawling and querying) is still a challenge from the technical but also infrastructural and societal point of view. Finally, in spite of current challenges (e.g. who will be responsible for archiving, parallel truths, robust links), the talk looks at the bright side showing that Linked Data is, in essence, feasible to archive (e.g. it is raw, self-contained and machine processable) and there are many technologies than can support archiving (e.g. Memento, HDT, Triple Pattern Fragments), although we need archiving to be add to the discussion, and stimulate archiving.

We close the workshops with a discussion on open topics and future directions. The main topic was the lack of clear guidelines to archive/preserve linked data, which is preventing further adoption by publishers and consumers. To solve this, we promote the collection of best practices and tools. In addition, we highlight that current systems do not consider the versioned datasets in their query planning strategies. Including this topic in the discussion would foster and boost the performance of current RDF archiving engines.

⁹ See slides at: <https://mielvds.github.io/MEPDaW2018/>

2 Semantic Web solutions for large-scale biomedical data analytics (SeWeBMeDA-2018)

The second edition of SeWeBMeDA-2018 workshop invited papers for life sciences and biomedical data processing, as well as the amalgamation with Linked Data and Semantic Web technologies for better data analytics, knowledge discovery and user-targeted applications.

This workshop at the Extended Semantic Web Conference (ESWC) targeted original contributions describing theoretical and practical methods and techniques that present the anatomy of large scale linked data infrastructure, which covers: the distributed infrastructure to consume, store and query large volumes of heterogeneous linked data; using indexes and graph aggregation to better understand large linked data graphs, query federation to mix internal and external data-sources, and linked data visualisation tools for health care and life sciences. It will further cover topics around data integration, data profiling, data curation, querying, knowledge discovery, ontology mapping / matching / reconciliation and data / ontology visualisation, applications / tools / technologies / techniques for life sciences and biomedical domain. SeWeBMeDA aims to provide researchers in biomedical and life science, an insight and awareness about large scale data technologies for linked data, which are becoming increasingly important for knowledge discovery in the life sciences domain.

This year, we accepted three papers, we invited a keynote speaker, organised a short hackathon and also discussed on current issues along with future steps for large scale data in biomedical domain.

Keynote talk was given by Maria-Esther Vidal who is the head of the Scientific Data Management group at TIB Leibniz Information Centre for Science and Technology, Germany and a full professor (on-leave) at Universidad Simón Bolívar (USB) Venezuela. Her interests include Big data and knowledge management, knowledge representation, and semantic web with more than 130 peer-reviewed papers in Semantic Web, Databases, Bioinformatics, and Artificial Intelligence. The title of her talk was "Synthesizing Big Data into Actionable Knowledge", where she discussed the role of Big data in promoting emerging scientific and interdisciplinary research by enabling decision-making. She described that knowledge-driven approach is capable to ingest Big data sources and integrate them into a knowledge graph that represents not only the meaning of the entities published by these data sources, but also that provides the basis for the discovery of unknown patterns and associations between these entities. The features of this knowledge-driven framework are shown in the context of the EU funded project iASiS (<http://project-iasis.eu/>), where it is used to pave the way for personalized diagnosis and treatments. The presentation slides are available at: (<https://goo.gl/aH92pM>).

As mentioned we had three paper presentations:

Gleim et al [3], proposes an automated schema extraction approach compatible with existing Semantic Web-based technologies. The extracted schema enables ad-hoc query formulation against privacy sensitive data sources without requiring data access, and successive execution of that request in a secure enclave under the data provider's control. The developed approach permit user to extract structural information from non-uniformed resources and merge it into a single schema to preserve the privacy of each

data source. Initial experiments show that this approach overcomes the reliance of previous approaches on agreeing upon shared schema and encoding a priori in favor of more exible schema extraction and introspection.

Hasnain et al [2], assess the FAIR principles against the LOD principles to determine, to which degree, the FAIR principles reuse LOD principles, and to which degree they extend the LOD principles. This assessment helps to clarify the relationship between both schemes and gives a better understanding, what extension FAIR represents in comparison to LOD. This publication concludes, that LOD gives a clear mandate to the openness of data, whereas FAIR asks for a stated license for access and thus includes the concept of reusability under consideration of the license agreement. Furthermore, FAIR makes strong reference to the contextual information required to improve reuse of the data, e.g., provenance information. According to the LOD principles, such meta-data would be considered interoperable data as well, however, the requirement of extending of data with meta-data does indicate that FAIR is an extension of the LOD (in contrast to the inverse).

Nayak et al [7], propose that the use of topic modeling, specifically non-negative matrix factorization (NMF), as a first step towards dimensionality reduction when dealing with large amounts of data. In this position paper, as a use case, author applied NMF to the BioSamples metadata and present preliminary results.

At the end of the workshop we organised a short Hackathon title "Privacy-Preserving Information Extraction with Bloom Filters". At the beginning of the hackathon, we provided a short introduction to the prerequisites, such as bloom filters, general privacy issues and frameworks that can be used (Python or KNIME). Then, each team involved in the hackathon was given a unique Knowledge Graph onto which they could apply information retrieval techniques to build up some experience with the given framework. Next, the the Bloom Filters were applied and discussed the suitable metrics for valuing an unseen knowledge graph based on a query response that may contain false positives. Finally, each team formulated queries for estimating the worth of an unseen Knowledge Graph and ultimately made a decision about which other teams Knowledge Graph complements their own Knowledge Graph the best.

3 Semantic Web of Things for Industry 4.0

insert your preface here

Acknowledgments

We would like to thank the authors for their contribution and active participation in the workshops, and all the program committee members for reviewing the submissions and provide valuable feedback. We are also grateful to the organisers of the ESWC 2018 conference for their support, and our keynote speakers, Miel Vander Sande from imec-Ghent University and Maria-Esther Vidal (for keynote speaker at SeWeBMeDA-2018) who is the head of the Scientific Data Management group at TIB Leibniz Information Centre for Science and Technology, Germany and a full professor (on-leave) at Universidad Simón Bolívar (USB) Venezuela [insert your keynote speaker here](#).

The MEPDaW workshop was co-organised by members funded by the EU’s Horizon 2020 research and innovation programme: grant 731601 (SPECIAL) and the Austrian Research Promotion Agency’s (FFG) program “ICT of the Future”: grant 861213 (CitySpin).

SeWeBMeDA-2018 workshop was co-organised by Insight Centre for Data Analytics NUI Galway and Fraunhofer FIT, RWTH Aachen University. This workshop has been supported in part by Science Foundation Ireland under Grant Number SFI/12/RC/2289. [insert your ACKS here](#)

References

1. J. Debattista, J. D. Fernández, J. Umbrich, M.-E. Vidal, A. Hasnain, O. Beyan, S. Decker, and D. Rebbholz-Schuhmann, editors. *Joint proceedings of the 4th Workshop on Managing the Evolution and Preservation of the Data Web (MEPDaW), the 2nd Workshop on Semantic Web solutions for large-scale biomedical data analytics (SeWeBMeDA), and the Workshop on Semantic Web of Things for Industry 4.0 (SWeTI)*, number XXXX in CEUR Workshop Proceedings, Aachen, 2018.
2. A. Hasnain and D. Rebbholz-Schuhmann. Assessing fair data principles against the 5-star open data principle. In *2nd workshop on Semantic Web solutions for large-scale biomedical data analytics (SeWeBMeDA)*, 2018.
3. L. Z. O. K. H. S. S. D. Lars C. Gleim, Md. Rezaul Karim and O. Beyan. Using schema extraction for query design without data access to enable privacy maintaining processing of sensitive data. In *2nd workshop on Semantic Web solutions for large-scale biomedical data analytics (SeWeBMeDA)*, 2018.
4. A. B. Meriem Laajimi and N. Y. Ayadi. Evaluation of rdf archiving strategies with spark. In *4th Workshop on Managing the Evolution and Preservation of the Data Web (MEPDaW)*, 2018.
5. H. J. Pandit, D. O’sullivan, and D. Lewis. Gdpr-driven change detection in consent and activity metadata. In *4th Workshop on Managing the Evolution and Preservation of the Data Web (MEPDaW)*, 2018.
6. A. Singh, R. Brennan, and D. O’sullivan. Delta-ld: A change detection approach for linked datasets. In *4th Workshop on Managing the Evolution and Preservation of the Data Web (MEPDaW)*, 2018.
7. A. Z. Stuti Nayak and M. Dumontier. Quality assessment of biomedical metadata using topic modeling. In *2nd workshop on Semantic Web solutions for large-scale biomedical data analytics (SeWeBMeDA)*, 2018.