

alised. Key to this was the publication of the *5★'s of Linked Open Data*, which outlined how applying the Linked Data principles can enhance the reusability of structured data published on the Web. This led to the emergence of Linked Data Portals in various domains. Governments – spurred on by the growing political demand for transparency – were a major adopter, with *data.gov.uk* and *data.gov* publishing multitudinous datasets in LD for the British and U.S. governments, respectively; other public organisations have since followed suit, with, e.g., the *EU Data Portal* adopting LD technology. Other portals emerged focussing on scientific data, with *Bio2RDF* and *OpenPhacts* offering an entry-point onto various key biomedical datasets; the *RKBExplorer* portal aggregating bibliographical meta-data from various sources; and so forth. Other initiatives emerged from the cultural heritage domain, where the *Europeana* portal offers multi-lingual meta-data on items found in museums, libraries and other cultural institutions. Other Data Portals emerged to catalogue datasets available across domains, including *DataHub*, *LODStats*, and *LOD Laundromat*.

Standardisation. As LD began to achieve more and more adoption, new requirements were emerging for the Semantic Web standards, leading to the recommendation of new versions. An important achievement was the standardisation of *SPARQL* in 2008: the first standard query language for RDF; this was extended with a variety of features in 2013, very much informed by emerging requirements for LD. *OWL* received an update in 2009, with new features and new tractable profiles emerging to deal with larger-scale applications. *RDF* later received a major update in 2014, key to which was the standardisation of various syntaxes – such as *Turtle* – that had found favour in the LD community; various flavours of the *JSON-LD* and *RDFa* syntaxes were also standardised. Another major development was the 2015 standardisation of the *Linked Data Platform (LDP)*, which paved the cowpaths worn by the LD community. Other standardisation efforts included the recommendation of *RDB2RDF* methods and languages for exporting relational data to RDF; *SKOS* for describing taxonomic information; *PROV* for capturing provenance; and others besides. Standardisation continues, with (just for example) the *RDF Data Shapes* Working Group soon to recommend a language for specifying structural constraints on RDF graphs.

LD Tools and Applications. As new standards emerged, so too did implementations thereof, with a wide range of parsers, query engines, reasoners, extractors, etc., being released (mostly open source). Aside from standards, however, other core challenges quickly became clear in LD adoption. In particular, the question of how to automatically link datasets – how to realise the fourth LD principle at large-scale – emerged as key for the LD vision; various proposals and tools for linking datasets thus emerged, with *SILK* and *LIMES* becoming popular frameworks amongst LD publishers. On the flip side, another key question was how to consume LD: how could end-users search and browse the collective data that the Web of Data provides? A number of data search engines emerged, key amongst which was *Sindice* and the related *Sig.ma* interface, which allowed for semantic search. Other applications, such as *Tabulator* or *Haystack*, allow for live browsing of Linked Data. The use of LD technologies – albeit to various degrees of adoption – likewise found use within major companies; with the BBC building LD-powered websites for various major sporting events; *IBM*

Watson incorporating LD knowledge-bases for its Jeopardy-playing engine; Facebook proposing its *Open Graph* protocol to decentralise certain social features; and Google offering novel semantic search features powered by *schema.org* and the *Google Knowledge Graph*.

Research. A variety of key issues have been tackled in the LD literature, relating to linking, decentralisation, data quality, amongst others. The resulting literature has been key to realising the aforementioned achievements and others besides. When LDOW began, LD research was a niche topic, but a decade on, such research can be found published not only in Semantic Web and Web venues, but also in areas such as AI, Databases, Data Mining, etc. There are still many open questions left to tackle, and the LD literature is sure to continue to expand in the next decade and onward.

3. ONWARD TO THE NEXT DECADE

Aside from the presentation of eight technical papers on a variety of emerging Linked Data research topics, this year's workshop will also include an open session to discuss the trends that are currently observed concerning the publication and consumption of structured data from the Web as well as the challenges that arise from these trends. We expect the discussion to touch upon points such as:

- the roles of the different data publication mechanisms in the context of the Web of Data, ranging from data portals publishing CSV files, over websites annotating simple data structures using *schema.org*, to networks of data providers setting RDF links in order to ease the integration of complex structured data.
- the relation between corporate knowledge graphs and the effort that is invested into cleansing them in order to present users with a single truth; and the decentralized nature of the Web of Data, which allows everybody to make statements about everything leading to higher heterogeneity and a more complex concept of truth on the one side, but multiple perspectives on the other side.
- mechanisms and incentives for increasing the degree of interlinking in structured data on the Web, especially, e.g., *schema.org* data.
- an increased focus on the role of, and support for, graph-based data (such as RDF) in popular big data processing frameworks.
- the role of Linked Data in enabling data discovery on the Web, irrespective of data model or serialisation.

Within the open session, we will also discuss the next steps for the LDOW workshop series to better understand its role and influence as part of the broader Web Science community, as well as initiatives that emerge around the workshop. One of the emerging initiatives is proposed as "Pioneering the Linked Open Research Cloud", which is intended to encourage researchers to apply techniques and best practices in Linked Data to scholarly communication. The intention of this initiative is to ensure that research articles as well as peer-reviews for example are better discoverable, accessible, interactive, and more social through the use of native Web technologies and standards.