

The Semantic Web: Two Decades On

Aidan Hogan

IMFD; Department of Computer Science, University of Chile, Santiago, Chile

E-mail: ahogan@dcc.uchile.cl

Abstract. More than two decades have passed since the establishment of the initial cornerstones of the Semantic Web. Since its inception, opinions have remained divided regarding the past, present and potential future impact of the Semantic Web. In this paper – and in light of the results of over two decades of development on both the Semantic Web and related technologies – we reflect on the current status of the Semantic Web, the impact it has had thus far, and future challenges. We first review some of the external criticism of this vision that has been put forward by various authors; we draw together the individual critiques, arguing both for and against each point based on the current state of adoption. We then present the results of a questionnaire that we have posed to the Semantic Web mailing list in order to understand respondents’ perspective(s) regarding the degree to which the original Semantic Web vision has been realised, the impact it can potentially have on the Web (and other settings), its success stories thus far, as well as the degree to which they agree with the aforementioned critiques of the Semantic Web in terms of both its current state and future feasibility. We conclude by reflecting on future challenges and opportunities in the area.

Keywords: Semantic Web, Ontologies, Linked Data, Knowledge Graphs

1. Introduction

Arguably the first concrete milestones towards realising the Semantic Web were the 1998 release of the initial versions of the Resource Description Framework (RDF) [13] and RDF Schema (RDFS) specifications [48]. In 2001, Berners-Lee et al. [8] would position RDF as a key technology for realising their vision of what they called the “Semantic Web”, which would “bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users”. A slew of developments were to follow, culminating in the release of numerous standards, such as OWL, SPARQL, SKOS, RIF, RDB2RDF, SHACL, ShEx, as well as a variety of updates to existing standards. Each standard has received varying degrees of attention and acceptance from researchers, developers, and publishers alike. We refer the reader to the survey by Gandon [24] for further details on the developments and trends in the Semantic Web research area spanning the first two decades.

More than two decades on there are varying opinions on the extent to which the original vision of Berners-Lee et al. [8] has been realised—or indeed, the extent to which it *can* or *should* be realised.

Within the Semantic Web community, there has long been a consensus that while the vision has yet to be fully translated into reality, it was a question of *when*, not *if*. In 2006, Shadbolt et al. [72], while admitting that the Semantic Web wasn’t “yet with us on any scale”, argued that it soon would be once the “standards are well established”. In 2007, Horrocks [42], while likewise admitting that “fully realising the Semantic Web still seems some way off”, argued that OWL had “already been very successful” and had “become a de facto standard for ontology development in fields as diverse as geography, geology, astronomy, agriculture, defence and the life sciences”. The years that followed were marked by optimism with regard to Linked Data, with authors claiming an exponential growth of data published following these principles [19, 45, 55, 62]. Optimism was further expressed with the selective adoption of Semantic Web technologies by household names, including the BBC [46], the New York Times [69], Oracle [84], Facebook [89], Google [10, 29], Wikimedia [83], Amazon [90], and so forth. More recent announcements of the development of knowledge graphs by Google [74], LinkedIn [33], Bing [73], eBay [65], Amazon [47], Airbnb [16], etc.,

1 have also been viewed as a win for the Semantic Web
2 community.

3 The Semantic Web has not only had numerous pro-
4 ponents down through the years, but also numerous vo-
5 cal opponents. As early as 2001, impassioned criticism
6 of the vision of the Semantic Web began to emerge,
7 with Doctrow’s often cited “*Metacrap*” essay [22]
8 laying out the seven “*insurmountable obstacles*” that
9 made the Semantic Web vision “*a pipe-dream*” in his
10 view; in summary, he criticises the naivety of expect-
11 ing users to create high-quality structured content, and
12 of expecting domain ontologies to be globally agreed-
13 upon given the many possible interpretations on how
14 a particular domain may be described. Various other
15 online articles and blog posts criticising the Semantic
16 Web emerged through the years. Here we summarise a
17 number of recent, prominent examples (found through
18 web searches for Semantic Web-related terms com-
19 bined with negative terms such as “fail”, “dead”, etc.,
20 further following hyperlinks to related articles):
21

- 22 – In 2013, ter Heide [79] suggested that the Seman-
23 tic Web had “*failed*” mainly due to: not catering
24 to a typical user’s interests, not considering new
25 streams of information such as messages, and ex-
26 pecting users to pull complex information rather
27 than being pushed content relevant to them.
- 28 – In 2014, Rothkind [67] discusses a thread on
29 Hacker News, asking “*is the Semantic Web still a*
30 *thing?*”, critiquing in particular the lack of incen-
31 tive for publishers to invest in publishing Linked
32 Data versus publishing the data in its native for-
33 mat; he highlights the lack of clear business mod-
34 els for doing so, noting that the infrastructure to
35 exploit Linked Data had “*not really materialized,*
36 *and it’s hardly clear that it will*”.
- 37 – In 2016, Cagle [15] suggested that the Semantic
38 Web had “*failed*”, primarily because it is hard to
39 understand, and it does not fit with other famil-
40 iar paradigms (citing Object Oriented Program-
41 ming), arguing for more lightweight semantics
42 (taxonomies) to alleviate the burden on users.
- 43 – In 2017, Cabeda [14] suggested that the rapid
44 advancement in Machine Learning techniques
45 “*leaves the Semantic Web in the dust*”, and con-
46 cluded that it “*needs to evolve and integrate its*
47 *ideas with artificial intelligence*”.
- 48 – In 2018, Target [78] – while giving a brief his-
49 tory on the major developments of the Seman-
50 tic Web – suggests that it has “*threatened to re-*
51 *cede as an idea altogether*”, observing that “*work*

1 *on the Semantic Web seems to have petered out*”;
2 while he acknowledges adoption in settings such
3 as the Open Graph Protocol and schema.org, and
4 commends technologies such as JSON-LD, he ul-
5 timately concludes that there are many “*engineer-*
6 *ing and security issues*” to be addressed before
7 the original decentralised vision of the Semantic
8 Web can be meaningfully realised.
9

10 These critiques of the Semantic Web raise a number of
11 important issues in terms of the feasibility of realising
12 its original vision and should be carefully considered
13 in the context of the Semantic Web community: while
14 the community is perhaps generally aware of such po-
15 tential criticisms, it is not always clear what (if any-
16 thing) should be done to address them.

17 Some such critiques have been addressed by mem-
18 bers of the community, both formally and informally.
19 In a 2013 keynote, Hendler [35] counters a number of
20 criticisms of the Semantic Web – such as the lack of
21 need for ontologies, the inability of the relevant tech-
22 nologies to scale, etc. – while ultimately concluding
23 that there are open challenges to face, particularly in
24 terms of uniting Ontologies and Linked Data, and de-
25 veloping practical reasoning methods for the Web. In
26 a 2017 keynote, Mika [58] provides a brief history of
27 the Semantic Web, noting a “*chicken and egg*” prob-
28 lem in the early days of applications requiring data and
29 applications being needed to incentivise the publica-
30 tion of data, but discussing how more and more incen-
31 tives are available for publishing data through initia-
32 tives such as Linking Open Data, schema.org, etc.; he
33 further discusses some application domains – Seman-
34 tic Search, eCommerce, Social Web – in which Seman-
35 tic Web concepts are being deployed.

36 Given the differing opinions that yet exist two
37 decades on, we believe it to be a fitting moment to un-
38 derstand the varying perspectives within the Seman-
39 tic Web community itself regarding its impact thus far,
40 the aforementioned critique, and the opportunities pre-
41 sented and challenges faced when looking to the fu-
42 ture. Along these lines, in this paper:
43

- 44 § 2 we first review external critique of the Semantic
45 Web, synthesising the primary criticisms raised,
46 presenting an argument both for and against each;
- 47 § 3 we present the results of a questionnaire posed to
48 the Semantic Web mailing list, aiming to ascertain
49 the various perspectives of respondents re-
50 garding the extent to which Berners-Lee et al.’s
51 original vision of the Semantic Web has been re-

1 alised or can be realised, the level of perceived
2 impact that the Semantic Web has had thus far on
3 the current Web, the success stories of the Semantic
4 Web, as well as opinions of the main points of
5 critique resulting from the previous analysis;

6 § 4 we summarise the main success stories, opportu-
7 nities, and challenges found regarding the past,
8 present and future of the Semantic Web.
9

10 2. Critique of the Semantic Web

11 Based on the previous critiques of the Semantic
12 Web, we now distil ten main criticisms paraphrased
13 from these articles [14, 15, 22, 67, 78, 79]; though the
14 list of issues should not be considered comprehensive,
15 it covers the main points in the articles found. We first
16 summarise the point of criticism, providing references
17 for sources that inspire its inclusion; we then argue
18 both for and against each point in turn to better under-
19 stand its implications.¹ The criticism presented stems
20 from authors with different roles and perspectives, rep-
21 resenting diverse points of view. We further categorise
22 each criticism according to: (i) *economic*: relating to
23 financial costs, incentives, market, etc.; (ii) *human*: re-
24 lating to individual users in terms of usability, accessi-
25 bility, etc.; (iii) *social*: relating to groups of people in
26 terms of agreement, social trends, etc.; and (iv) *techni-
27 cal*: relating to issues such as computational cost, dif-
28 ficulty to implement, etc. Later we will use these crit-
29 icisms to form a questionnaire posted to the Semantic
30 Web mailing list in order to gain insights into the
31 perspectives of experts from the community on these
32 issues.
33
34

35 2.1. The Semantic Web addresses a niche 36 problem [79]

37 *Categories:* Human, Social

38 *Critique:* Scenarios used to motivate the Semantic
39 Web are fact-based and often overly specific and com-
40 plex. The majority of users are only interested in
41 finding individual webpages with simple facts, opin-
42 ions, social recommendations, etc., rather than solving
43 complex queries on factual content involving multiple
44 sources. The current Web, with the help of search en-
45 gines like Google, thus covers (and will continue to
46 cover) the needs of the vast majority of users.
47
48
49

50 ¹We do so in the style of a debate, meaning that the author does
51 not necessarily hold the point-of-view being argued for/against.

1 *For:* Search engines such as Google, Bing, Yandex,
2 etc., have improved considerably over the years, where
3 finding information on the Web is now easier than
4 ever. In a July 2014 analysis of organic Google click-
5 through rates, Petrescu [64] estimated that users click
6 on a result listed on the first page for 71.3% of searches
7 and on a later page for 5.6% of searches; these fig-
8 ures do not account for users clicking paid results,
9 finding answers directly on the results page, refining
10 their search, etc. With current search engines, most
11 user searches can be quickly and easily resolved. Aside
12 from search, use-cases relating to modelling complex
13 domains using ontologies, data integration in enter-
14 prises, etc., are not tangible for ordinary web users.
15

16 *Against:* There are many niche problems of impor-
17 tance to society with which the Semantic Web can
18 help, including, for example, drug discovery in the
19 case of rare diseases [43]. However, the Semantic
20 Web is not limited to niche use-cases. Search engines
21 themselves have been adopting Semantic Web con-
22 cepts to enable semantic search; for example, through
23 schema.org [29], Knowledge Graphs [73, 74], etc.
24 On the other hand, while current search engines are
25 excellent for finding individual webpages, the Se-
26 mantic Web vision addresses more complex types of
27 queries that require drawing information from multi-
28 ple sources on the Web. While current searches gener-
29 ally appear to be resolved quickly (e.g., are answered
30 by a single high-ranking result), users may not be *cur-
31 rently* issuing more “complex” searches as they know
32 search engines will not offer useful results. Searches
33 requiring cross-referencing multiple webpages are not
34 necessarily niche, but may rather be *personalised* [8];
35 for example, finding the *closest store open now selling
36 aspirin* does not appear to be niche, and could be bet-
37 ter automated with Semantic Web techniques. Regard-
38 ing users’ interests, the Semantic Web does not only
39 address encyclopaedic data, nor does it only address
40 search; for example, its graph-based data model can be
41 used to integrate and find novel connections within so-
42 cial data [12]. Regarding other use-cases, though users
43 may not know of the use of Semantic Web techniques
44 within specific domains or enterprises, this does not
45 prevent them from benefiting from such technologies.
46

47 2.2. The Semantic Web will be made redundant by 48 advances in Machine Learning before it has a 49 chance to take off [14].

50 *Category:* Technical
51

Critique: The Semantic Web assumes that the current (HTML-based) Web is poorly machine-readable. However, advances in Machine Learning are increasingly undermining this assumption. By the time the Semantic Web could reach enough maturity to have major impact on the Web, Machine Learning will have advanced to a point where such technologies for publishing/consuming structured content are made redundant.

For: Advances in areas such as Deep Learning have led to results that previously seemed unachievable in the short term. Machines can now perform more “human-like” tasks with increasing precision and recall. These advances, combined with developments in Information Extraction, increasingly blur the lines between human-readable and machine-readable content [54]; as a relevant example, in the TAC–KBP “Cold Start” challenge, which requires systems to extract knowledge-bases from scratch from text, systems improved their F_1 scores from 0.48 to 0.58 on English texts in the space of a year (from 2016 to 2017) [25]. The need for a specialised machine-readable Web becomes more tenuous as machines succeed in processing our natural language with increasing fidelity.

Against: Techniques like Deep Learning are still applied as a form of specialised Artificial Intelligence, requiring extensive training data to build models for one particular task. Though impressive gains are being made, the aforementioned F_1 score of 0.58 [25] still leaves much to be desired. Addressing the tasks discussed by Berners-Lee et al. [8] on the current Web – without structured content – would require a general form of Artificial Intelligence as yet without precedent (sometimes referred to as *AI-complete* tasks [85]). Many of the prominent data-driven AI-style applications found in practice – such as digital assistants (Siri, Alexa, etc.) – in fact already rely on Semantic Web resources to provide structured content [52]. While the Semantic Web undoubtedly stands to benefit from Machine Learning, so too can applications using Machine Learning benefit from advances in the Semantic Web.

2.3. The Semantic Web depends too much on reliable publishers [15, 22]

Categories: Human, Technical

Critique: The Semantic Web is founded on the idea that machines will automatically process structured content on the Web. Such processing is particularly brittle in the face of both indeliberate errors and delib-

erate deception due to unreliable publishers (as commonplace on the Web).

For: Automatically solving complex tasks on the Semantic Web involves processes such as inferencing to integrate information. Such processes work by assuming input data to be held true and computing other entailments that then follow; this assumption is clearly naive for Web data. Even small errors in the input data (e.g., inconsistent claims) can lead to nonsensical entailments; in previous work we found 301 thousand RDFS/OWL inconsistencies in a crawl of 4 million RDF documents (294 thousand relating to datatypes, 7 thousand relating to instances of disjoint classes) [11]. More complex tasks require more complex chains of inferencing, where each step accumulates a higher probability of error. Such processes could then be easily manipulated by deceptive agents.

Against: The Semantic Web community recognises that publishers are not always reliable, and though the issue of data quality is a major challenge, it is one that the community has been addressing [87]. Much like on the Web, rather than assume all information to be trustworthy, two elements are required: reliable sources of data, and methods to accurately estimate the reliability of sources. Specifically regarding inferencing, methods such as paraconsistent reasoning [51] are more robust to noisy inference, while methods such as authoritative and quarantined reasoning [66] select more trustworthy sources for inferencing based on link analysis. Finally – as acknowledged by the original vision paper [8] – users should not blindly trust results, but can rather be provided details (on-demand) of how these results were achieved, refining criteria as required.

2.4. The Semantic Web depends too much on ontological agreement [15, 22]

Categories: Social, Technical

Critique: There is no single way to model a domain using an ontology. There is no global truth. Different stakeholders in the domain may consider different semantics for terms or even hold contradictory claims. The Semantic Web is brittle to differing views.

For: Is a tomato a “fruit” or a “vegetable”? Is Pluto a “planet”? Is Sherlock Holmes a “person”? The answer to each such question depends, either due to a lack of consensus, or ambiguity on what terms like “fruit”, “person”, etc., mean. While we might define in an ontology that all mayors are people, Bosco the Dog

was elected mayor of Sunol, California while Duke the Dog was elected mayor of Cormorant, Minnesota. The real-world is messy and hosts innumerable perspectives on what is true, or what “truth” even means. Edit wars on Wikipedia evidence such disagreement [86]. These ambiguities and conflicts are the true underlying cause of interoperability issues, and rather than solving them, ontologies (particularly expressive ones), require them to be have been solved beforehand; doing so at the scope of the Web presupposes either a utopian (global agreement reached) or a dystopian (global agreement enforced) view of society.

Against: To be more precise, the Semantic Web *benefits from* – rather than *requires* – ontological agreement. The fact that full agreement cannot always be reached does not preclude the utility of formally capturing the agreement that can be reached. While agreement on detailed domain definitions is costly, ontologies such as SNOMED CT [49] show that it can be achieved with sufficient will and organisation. For the broader Web, initiatives such as schema.org [29] show that agreement is possible on lightweight semantic definitions (given sufficient incentives). The impact of collaboratively-edited datasets such as Wikidata [52, 83] further exemplify ways in which (partial) agreement can be fostered in an emergent way. Considerable attention has been given by the Semantic Web literature to resolving inconsistencies reflecting different views [11], to inferencing over contextual data reflecting different versions of truth [30], and so forth. Furthermore, ontologies are defined in a decentralised way [81], where stakeholders can adopt their preferred ontology or define their own, giving rise to an emergent agreement; exemplifying this, Schmachtenberg et al. [70] found that FOAF and Dublin Core were used by 69% and 56% of the 1,014 RDF datasets that they crawled. In the case of multiple competing ontologies, mappings can be computed or defined to enable interoperability by bridging the concepts on which they agree [23]; along these lines Vandebussche et al. [81] find over 5 thousand links between different vocabularies in their collection.

2.5. Publishing Semantic Web content on the Web has a prohibitively high cost [15]

Categories: Economic, Technical

Critique: Given data in a legacy format, a relational database, JSON, CSV, etc., there is a prohibitively high cost associated with publishing the data using the Semantic Web standards.

For: Publishing Semantic Web content in a suitable way – e.g., following Linked Data principles [34] – requires expertise. Where data are available in a structured format, conversion to RDF is far from straightforward, especially when issues such as offering dereferenceable IRIs, adding links, etc., are considered [41]. While certain types of data are easily conceptualised as RDF graphs, others require various forms of indirection (e.g., reification [36]) to be properly represented.

Against: Most websites are now based on data stored in databases. Standards have been developed to reduce the cost of publishing RDF from legacy data, key amongst which are the RDB2RDF mappings [5, 20] for generating RDF data from relational databases, and JSON-LD for lifting JSON to an RDF-style data model [76]. Tools have been developed to help with tasks such as linking, most prominently Silk [82] and LIMES [61]. Exporters built into commonly-used platforms such as Drupal allow thousands of websites to begin publishing RDF quickly and easily [17]. Work continues to better support more and more types of data, such as the standardisation of the RDF Data Cube vocabulary for representing statistical data [18].

2.6. There are too few incentives for adopting Semantic Web technologies on the Web [67]

Categories: Economic, Social

Critique: Aside from the costs of using Semantic Web technologies on the Web, there is little incentive to do so, due in part to the fact that the infrastructure for publishing and/or exploiting such content on the Web has not been adequately developed or adopted.

For: The Semantic Web has long faced a chicken-and-egg problem [58]: incentives for publishing data require infrastructure to exploit those data, while infrastructure for exploiting data cannot develop without data. While the Linked Data community partially resolved this dilemma by successfully convincing various stakeholders to publish data on the (implicit) promise that applications would arrive to justify the cost, these applications did not emerge, and as a result, many datasets and related services went offline [4, 40]; for example, Aranda et al. [4] estimated, in 2013, that around 29% of the 427 public SPARQL services they found had gone offline. The dearth of Linked Data applications hint at an important lesson: publishing data independently of a particular application implies higher costs for leveraging that data in that application;

publishing data independently of *any* application then implies higher costs for *all* applications. Finally, one of the main incentives for publishing on the current Web is advertising revenue, where it is not clear how advertising would work on the Semantic Web where software agents, rather than humans, access websites [28].

Against: In the case of schema.org [29], publishers are incentivised to embed structured data in their webpages by the promise of “rich snippets”: having the data – denoting images, ratings, etc. – displayed in search engine results, offering a more eye-catching result summary that attracts more clicks; as a result, schema.org has been widely adopted on the Web, where Meusel et al. [56] found more than 700 thousand pay-level-domains (websites) hosting schema.org content in the 2014 WebDataCommons dataset. Such examples show that incentives do exist for Web publishers to provide more structured content: offering such content can, in the context of certain applications, help direct traffic back to a website or increase demand for a particular product or service it describes, which can drive new business models that replace traditional advertising revenues [28]. The varied use of datasets such as Wikidata [52, 83] – whose SPARQL service received over over 3.8 million queries per day in the first quarter of 2018 [52] – show that a variety of applications – including some not originally envisaged – can benefit from the increasing availability of structured content offered by the Semantic Web.

2.7. The Semantic Web standards are too verbose [15, 78]

Category: Human

Critique: The Semantic Web standards are (unnecessarily) long, complex and difficult to understand. This creates a major barrier for attracting new adopters. More concise standards would have been better.

For: Most of the Semantic Web standards have been designed by committee, anticipating use-cases that had yet to arrive or be fully understood, sometimes focusing on academic rather than practical issues. The resulting standards are difficult to understand, with much of their complexity dedicated to relatively niche issues; as a result, we can find various calls to simplify the standards, with, e.g., Berners-Lee calling for the deprecation of various features in the RDF standard in 2010 [7]. In the same way that JSON has become more popular than its more complex XML cousin, simpler

standards that suffice for common needs will tend to win out versus complex standards that (additionally) address more niche need; along these lines, for example, Meusel et al. [57] found over five times more Microdata/Microformats statements than RDFa in the 2013 Common Crawl dataset; in previous works, we found that (pure) RDFS is much more prevalently used than OWL in Web data [26]; and so forth.

Against: When speaking of verbose standards, one should not overlook the SQL:2016 standard [88], which has 1,732 pages—yet the core of SQL is broadly adopted and understood. One does not need to understand the entire standard in order to profitably use parts of it. Along the same lines, one does not need to understand the model theoretic definitions of RDF to describe data in RDF, nor does one need to understand the semantic conditions defined for OWL to use it to describe an ontology, etc.; rather practitioners can start with a simple system based on the parts of the standards important for them, extending their use of the standards – as needs arise – towards building more complex (and powerful) systems that work for them. Simpler standards that arise can also be mapped to more complex standards; for example, Microdata and Microformats are directly convertible to RDF. More modern Semantic Web standards – such as JSON-LD [76] – have also had success in terms of adoption.²

2.8. The Semantic Web will not scale [6]

Category: Technical

Critique: Consuming data published using the Semantic Web standards requires algorithms with poor scalability and/or performance. Current implementations exhibit poor scalability and/or performance.

For: Even the most common tasks that one might consider over (most of) the Semantic Web standards are intractable. Deciding if two RDF graphs have been parsed from the same document, potentially with different blank node labels (aka. *RDF isomorphism*), is GI-complete [38]. SPARQL query evaluation is PSPACE-hard (PSPACE-complete for the original standard [63]). Entailment is undecidable for OWL (2) Full and N2EXPTIME-complete for OWL 2 DL [60]; infamously even the OWL “Lite” fragment of the original OWL standard – motivated as a more

²See <https://w3techs.com/technologies/details/da-jsonld/all/all>; retrieved 2019-09-29: JSON-LD is used by 26.5% of websites.

1 terse fragment permitting more efficient reasoning –
 2 was later found to have EXPTIME-complete entail-
 3 ment. Other experimental works have shown Semantic
 4 Web query engines to be considerably outperformed
 5 by relational databases; for example, with the Berlin
 6 SPARQL Benchmark, Bizer and Schultz [9] show that,
 7 in some cases, MySQL can execute 13 times more
 8 queries in a given time period than the best SPARQL
 9 store tested (Sesame) considering comparable queries.

10 *Against:* Such complexity results are not particular
 11 to Semantic Web proposals, where for example the
 12 complexity of SPARQL query evaluation is analogous
 13 to that for SQL [63]. More generally, worst-case com-
 14 plexity results rarely tell the whole story: the fact that
 15 there exists at least one input for which a task is dif-
 16 ficult tells us little about how efficient solutions might
 17 be for practical inputs (see, e.g., [38]). Achieving scale
 18 and efficiency often requires trade-offs, where by trad-
 19 ing in completeness, OWL reasoning has been shown
 20 to scale to billions of triples [66, 80]; along similar
 21 lines, a variety of tractable profiles of OWL 2 have
 22 been defined that trade expressivity for efficiency of
 23 reasoning tasks [60]. More practically speaking, a poor
 24 implementation does not refute its underlying idea.
 25 With this aside, some more recent benchmarks show,
 26 for example, SPARQL engines being capable of out-
 27 performing graph databases and relational databases
 28 for more complex graph patterns [37]. Anecdotally, we
 29 can also point to Wikidata’s decision to use Semantic
 30 Web technologies (RDF, SPARQL, etc.) to publish and
 31 manage its content, with positive (performance) re-
 32 sults [52]. Adoption of the Semantic Web standards by
 33 major vendors – such as Oracle [84] and Amazon [90]
 34 – further help to (anecdotally) refute this criticism.

35 2.9. *The Semantic Web lacks usable systems &* 36 *tools [15]*

37 *Categories:* Human, Technical

38 *Critique:* Practitioners who are initially interested
 39 in adopting Semantic Web technologies are quickly
 40 alienated by a lack of usable tools for their use-cases.

41 *For:* While one may argue that end-users need not
 42 understand the Semantic Web to benefit from it – that
 43 the Semantic Web is something “under the hood” pow-
 44 ering end-user applications – such an argument still
 45 supposes the availability of systems, tools, etc., for
 46 building these applications. While many systems and
 47 tools have been developed for the Semantic Web, the

1 bulk have been created in an academic context for
 2 the purposes of proving a concept described in a pa-
 3 per. Systems often go offline after the paper is pub-
 4 lished; tools may rather be of a more prototypical na-
 5 ture; few resources are tested in terms of usability [44];
 6 etc. On the other hand, newer competing technolo-
 7 gies with more usable, developer-friendly resources
 8 are seeing more adoption, including formats such as
 9 JSON/Microdata/Microformats being more popular
 10 than RDF [57], the Neo4j graph database being far
 11 more popular than its closest SPARQL rival³, Face-
 12 book’s GraphQL [32] being widely adopted for public
 13 query interfaces (versus SPARQL/Linked Data), etc.
 14 The Semantic Web is thus left in the wake of alterna-
 15 tive, more lightweight, more usable technologies.

16 *Against:* While the Semantic Web could always ben-
 17 efit from having more (usable) systems and tools, most
 18 standards have a variety of mature implementations to
 19 choose from (including from well-known vendors such
 20 as Oracle [84], Amazon [90], etc.). On the other hand,
 21 the adoption of similar, competing technologies is an
 22 *opportunity* for the Semantic Web, as in the case of
 23 JSON-LD [76] successfully leveraging the popularity
 24 of JSON to help (implicitly) bridge the gap between
 25 developers and the Semantic Web. Along similar lines,
 26 various works have looked at making property graphs
 27 – the model underlying many graph databases [3] – and
 28 RDF graphs interoperable [21, 31]. The same story is
 29 borne out with proposals such as GraphQL-LD [77],
 30 this time bridging GraphQL and SPARQL. What we
 31 see, then, is increasing adoption of the core concepts
 32 underlying the Semantic Web: structured data formats,
 33 graph-based data modelling, public query APIs, etc.;
 34 with some syntactic glue, these advances can be lever-
 35 aged as advances, in turn, for the Semantic Web.⁴

36 2.10. *The Semantic Web advocates decentralisation,* 37 *which is too costly [78]*

38 *Categories:* Economic, Social, Technical

39 ³<https://db-engines.com/en/ranking> ranks graph databases (in-
 40 cluding SPARQL engines) in terms of popularity, where as of
 41 2019/05/25, Neo4j is ranked first (49.46 points), while the highest-
 42 ranked SPARQL engine – Virtuoso – is ranked fifth (2.73 points).

43 ⁴In a signed public comment in the questionnaire described later,
 44 Staab refers to this as a “hijacking strategy” (e.g., JSON-LD “hijack-
 45 ing” JSON, adding a core Semantic Web principle), expressing the
 46 opinion that it is an excellent way forward.

Critique: The original vision of the Semantic Web is a decentralised one (where, e.g., individual health care providers host their own web-site with their own structured content). On the other hand, on the current Web, centralisation has become the predominant paradigm (considering Google, Facebook, etc.). Decentralising the Semantic Web is too costly.

For: Berners-Lee et al. [8] talk about individual providers (doctors, physical therapists, etc.) hosting their own websites and agents, giving a decentralised setting for the Semantic Web. However, the Web has tended more and more towards centralisation, with individual providers rather collecting on central, specialised websites. For example, rather than hosting personal websites, most people rather host profiles on social networks. Likewise success stories sometimes quoted for the Semantic Web have involved some level of centralisation: Wikidata [83] centralises data creation and curation, schema.org [29] centralises the schema/ontology, and so forth. Decentralisation incurs significant conceptual and practical costs in terms of design, performance, etc. In terms of querying, for example, Schmidt et al. [71] demonstrate that local query processing is often orders of magnitude more efficient than federated querying over endpoints, even when statistics about remote data are made available for optimisation purposes. More generally, no precedent exists in the Semantic Web setting for the type of decentralised infrastructure envisaged by Berners-Lee [8].

Against: There is an emergent public awareness of the problems associated with growing centralisation in terms of users' privacy, control of data, etc. Along these lines, the recently standardised Linked Data Platform [75], along with projects such as Solid [53], not only further a decentralised vision of the Semantic Web, but also position the Semantic Web as a path towards a more decentralised Web. Abstractly, the benefits of centralisation versus decentralisation are mostly technological—benefits that will inevitably shrink as technology continues to improve. Conversely, the benefits of decentralisation versus centralisation are mostly social, be they upholding privacy, avoiding hegemony and monopoly, averting censorship, etc.—benefits that will at least remain constant, or more likely grow, over time. Asymptotically speaking, the relative benefits of decentralisation will thus, over time, increasingly dominate those of centralisation.

3. Questionnaire

We have, thus far, presented ten points critiquing the Semantic Web, arguing both for and against each individual point; the goal in each case was not to reach a verdict, but rather to understand possible arguments on both sides. We are now rather interested to see what members of the Semantic Web community, more broadly, think of the current state of adoption of the Semantic Web, what impact it could have in future, what they view as the main success stories thus far, and finally, what they think of the previously raised points of critique. We are particularly interested in the perspectives of experts in the Semantic Web who have read and worked extensively on the topic and can thus offer a more informed opinion; it is important to keep in mind, however, that targeting experts in this way may in turn lead to a pro-Semantic Web bias.

We designed a questionnaire for these issues and sent it to the W3C Semantic Web mailing list⁵ soliciting responses. All questions were left optional. The questionnaire was open to responses from May 12th to May 25th, 2019, in which time 113 responses were collected. In this section we present the details of the questionnaire and the responses received. Additional material is available online for the purposes of further analysis, including details of the questionnaire design, individual responses, public comments, keywords of success stories, and word clouds in SVG format [39].

3.1. Expertise of Participants

The questionnaire began with two questions to ascertain the self-assessed level of expertise of the respondent in terms of Semantic Web topics. The first question asked respondents to select one of the following options regarding their own level of expertise:

- Zero expertise (e.g., I have not read about the topic nor worked on the topic)
- Some expertise (e.g., I have read about the topic but not worked on the topic)
- Considerable expertise (e.g., I have read about the topic and worked occasionally on the topic)
- Strong expertise (e.g., I have read and worked extensively on the topic)

The results are shown in Figure 1, indicating strong expertise on the Semantic Web amongst respondents, as was the goal of the questionnaire: to target experts.

⁵semantic-web@w3.org; we also asked that members share the list with others who might be interested.





Fig. 1. Self-reported expertise of respondents

Fig. 2. Type of expertise of respondents in terms of *Academia*, *Industry*, *Other*, and combinations thereof.

We were further interested to know if respondents' expertise was mainly relating to academia, industry, or other settings; we thus asked respondents to select all that applied to them from the following:

- I have worked on the Semantic Web in academia (more than 1 year of experience).
- I have worked on the Semantic Web in industry (more than 1 year of experience).
- I have worked on the Semantic Web outside of both academia and industry (more than 1 year of experience).
- None of the above.

The results shown in Figure 2 reveal that 63.7% of respondents have an academic background, while 42.5% (also) have an industrial background.⁶

3.2. Realisation and Impact

In order to understand to what extent the respondents believe that the original vision of the Semantic Web has been already realised, to what extent they believe it can be realised in future, the impact it has had thus far and the impact it will have (in terms of both the Web and other settings), we posed the questions shown in Figure 3 to the participants. The results are shown in Figure 4, displaying the distribution of votes, as well

⁶We highlight a possible ambiguity in the question for what students should choose (noticed after posting the questionnaire).

as the mean \bar{x} and p -value computed using a two-tailed z -test (recalling that $n = 113$) for the null hypothesis of the results being approximately normally distributed with $\mu = 3$ (i.e., the mean being the neutral value: 3).

From these results, we observe the following:

Q1: regarding the original vision of the Semantic Web, the majority of respondents believe that it remains mostly or completely unrealised;

Q2: regarding the potential for realising the original vision of the Semantic Web in future, while 10 respondents believe it is completely unfeasible to realise, 14 believe it is completely feasible to realise; other responses were weighted towards believing it is mostly feasible to realise;

Q3: regarding current impact on the Web, responses were weighted towards the centre: that while Semantic Web technologies play some role on the Web, they do not play a key role;

Q4: regarding future impact on the Web, responses were weighted towards an optimistic view, with 76 respondents indicating their belief that Semantic Web technologies will play a significant or key role on the future Web;

Q5: regarding current impact in settings other than the Web, responses were weighted towards the centre: that while Semantic Web technologies play some role, they do not play a key role;

Q6: regarding future impact in settings other than the Web, responses were again weighted towards optimism, with 76 respondents again indicating their belief that Semantic Web technologies will play a significant or key role in the future.

While respondents tend to be reserved about the extent to which the Semantic Web has been realised and the impact that related technologies have had thus far, they tend to be much more positive regarding the future; per Q2, however, the bright future they envisage for the Semantic Web does not necessarily depend on completely realising the original vision.

We also performed a two-tailed z -test to look for statistically-significant differences in mean responses for each of the six questions between the 40 respondents who have only worked in academy (i.e., the *Acad (Only)* group of Figure 2) and all other respondents; with $\alpha = 0.05$, we found no such statistically-significant differences. The most notable difference was for Q6, with means of 3.615 and 3.931, respectively, for the academic-only group and the other group, suggesting that the latter group may believe

1 more in the potential impact of the Semantic Web for
 2 settings not directly involving the Web than the former
 3 group; however, as aforementioned, with $p \approx 0.131 <$
 4 α , the difference was not statistically significant.

6 3.3. Success Stories

8 We next asked respondents to list success stories
 9 they associate with the Semantic Web; specifically:

10 What are the main success stories that you would associate with
 11 the Semantic Web thus far (if any)? Please specify one per line;
 12 you may use simple keywords referring to the name of a technol-
 13 ogy, system, standard, dataset, project, etc.

14 A text field was provided below the question.

15 A total of 90 non-empty responses were collected.
 16 In order to summarise the main success stories men-
 17 tioned, the raw responses required some manual cura-
 18 tion. While some respondents provided keywords
 19 on individual lines, others rather answered with full
 20 sentences or paragraphs of free text; in these cases,
 21 we manually extracted a list of keywords from such
 22 text. While some responses referred to concrete stan-
 23 dards, datasets, initiatives, etc., other responses rather
 24 referred to more general concepts and domains. Re-
 25 garding the latter cases, distinct but related terms –
 26 such as biology, bioinformatics, life sciences, etc. –
 27 were used by different respondents, potentially “split-
 28 ting the vote”; in such cases, we manually selected and
 29 mapped related terms to a canonical term (e.g., in the
 30 previous case, we selected bioinformatics). A total of
 31 394 occurrences of 136 unique keywords were found.

32 Figure 5 illustrates the main success stories ref-
 33 erenced in the responses, with schema.org [29] be-
 34 ing the most referenced project. Knowledge Graphs
 35 (e.g., [16, 33, 47, 65, 73, 74]), Wikidata [83] and DB-
 36 pedia [50] fill the next positions, followed by two key-
 37 words often mentioned side-by-side: Bioinformatics
 38 and Ontologies. Linked Data was next, followed by
 39 a sequence of three standards: RDF, JSON-LD and
 40 SPARQL. Informally, we noticed a number of clus-
 41 ters of responses: (1) those focused on the Web and
 42 Public Datasets, including search engines, embedded
 43 meta-data, Wikidata, DBpedia; etc.; (2) those focused
 44 on Semantics, including the use of ontologies in spe-
 45 cific domains, particularly bioinformatics; (3) those fo-
 46 cused on Enterprises, particularly relating to Knowl-
 47 edge Graphs, Data Integration and Data Governance,
 48 etc.; and (4) those focused on the Public Sector, includ-
 49 ing relevant initiatives within governments, libraries,
 50 museums, etc.
 51

1 3.4. Reaction to Critique

2 The next part of the questionnaire sought feedback
 3 on the ten points of critique presented previously. More
 4 specifically, we presented the title and description of
 5 each point of critique as given in Section 2 without the
 6 associated arguments for or against. We then asked re-
 7 spondents to indicate the extent to which they agreed
 8 with the stated critique, both in terms of the current
 9 state of the Semantic Web, as well as how significant
 10 an obstacle it might pose to future development and
 11 adoption of the Semantic Web. In the cases of points
 12 (7) verbose standards, (8) does not scale and (9) lacks
 13 usable tools, we further ask respondents to indicate
 14 the standards they believe to be most problematic re-
 15 garding the highlighted issue (if any), selecting zero-
 16 to-many from RDF (data model), RDFS, OWL and
 17 SPARQL. By way of example, Figure 6 shows the
 18 question issued for point (8); the same structure was
 19 followed for other points, with C1 and C2 posed for all
 20 points, and C3 posed for points (7–9).
 21

22 The results for C1 – level of agreement with re-
 23 spect to the current state of the Semantic Web – are
 24 summarised for all ten points of critique in Figure 7,
 25 again showing the mean \bar{x} and p -value computed as
 26 before. We see that respondents were most in agree-
 27 ment with the critiques regarding (9) a lack of usable
 28 tools, (6) a lack of incentives, and (3) a lack of toler-
 29 ance to unreliable publishers. On the other hand, they
 30 mostly disagreed with the idea that (2) advances in Ma-
 31 chine Learning render the Semantic Web redundant,
 32 and that (10) decentralisation is too costly. Other cri-
 33 tiques rather saw a balance of responses.
 34

35 Looking to the future, the results for C2 are pre-
 36 sented in Figure 8; while in general we see few re-
 37 sponses indicating that the presented issue is insur-
 38 mountable (option 1), we see many responses indica-
 39 tive of major obstacles (option 2) to be overcome.
 40 More generally, the most critical challenges that the
 41 Semantic Web must face in future according to the re-
 42 spondents are (6) a lack of incentives, and (3) a lack
 43 of tolerance to unreliable publishers; when compared
 44 to responses for the current state of the Semantic Web,
 45 respondents are slightly more optimistic regarding (9)
 46 a lack of usable tools. Conversely, respondents do not
 47 see (2) redundant w/ML or (5) costly publishing as
 48 posing major challenges relative to other issues.
 49

50 With respect to the four categories previous dis-
 51 cussed (Economic, Human, Social, Technical), from
 Figures 7 and 8, we observe that the criticisms for
 which there was most agreement related predomi-

In a 2001 Scientific American article, Berners-Lee et al. outlined their vision of the Semantic Web as follows: "The Semantic Web will bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users." They quote an example where a user's software agent automatically accesses the individual websites of different health care providers and returns those within a certain distance from the user's home offering treatments covered by the user's insurance plan (according to the insurer's website), with excellent ratings (according to a trusted third-party website), and with open appointments matching the user's availability (based on the user's local calendar).

Q1. As of May 2019, to what extent do you believe that Berners-Lee et al.'s 2001 vision of the Semantic Web has been realised in practice?

Completely unrealised 1 2 3 4 5 Completely realised

Q2. Independent of the current state, to what extent do you believe that Berners-Lee et al.'s 2001 vision of the Semantic Web is feasible to realise, in practice, within the short-to-medium term?

Completely unfeasible 1 2 3 4 5 Completely feasible (or it has already been realised)

Q3. To what extent do you believe that Semantic Web technologies play an important role on the current Web?

They play no role 1 2 3 4 5 They play a key role

Q4. To what extent do you believe that Semantic Web technologies will play an important role on the future Web?

They will play no role 1 2 3 4 5 They will play a key role

Q5. To what extent do you believe that Semantic Web technologies currently play an important role in settings not directly involving the Web?

They play no role 1 2 3 4 5 They play a key role

Q6. To what extent do you believe that Semantic Web technologies will play an important future role in settings not directly involving the Web?

They will play no role 1 2 3 4 5 They will play a key role

Fig. 3. Realisation and impact section of the questionnaire

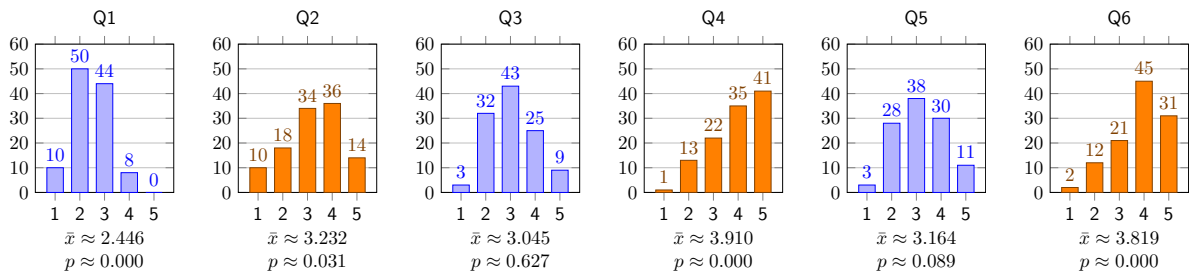


Fig. 4. Responses to realisation and impact section of the questionnaire (shown in Figure 3)

nantly to Human issues: (3) unreliable publishers, (7) verbose standards, and (9) lacks usable tools; on the other hand, respondents tended to disagree with purely Technical issues: (2) redundant w/ML, and (8) won't scale. These results suggest that respondents see more pressing issues relating to the human aspect of Semantic Web technologies rather than the technical aspect.

Within these results, we again look at the difference between the respondents with only academic experience and other respondents; we find statistically-significant differences for three questions:

– C2: *Unreliable publishers* The means were different with statistical significance ($p \approx 0.002$) –

2.450 for the academic-only group, versus 2.986 for the other group – indicating that the academic-only group tends to view noisy/deceptive data as a greater future obstacle for the Semantic Web.

– C1: *Won't Scale* The means were different with statistical significance ($p \approx 0.007$) – 2.900 for the academic-only group, versus 3.507 for the other group – indicating that the academic-only group tends to agree more that scalability is currently a significant issue facing the Semantic Web.

– C2: *Won't Scale* The means were different with statistical significance ($p \approx 0.002$) – 2.775 for the academic-only group, versus 3.411 for the other




Fig. 9. Responses to C3 for critiques 7–9 indicating the standards believed to be most problematic with respect to the highlighted issue

group – indicating that the academic-only group tends to view a lack of scalability as a greater future obstacle for the Semantic Web.

Finally, regarding (7) verbose standards, (8) problems with scale, and (9) a lack of usable tools, Figure 9 presents the results of C3 indicating the standards that respondents feel most problematic. We see that OWL, followed by SPARQL, have the most responses in terms of being problematic for each of the three highlighted issues. Notably, the OWL 2 standard defines three tractable profiles [60] that aim to address issues (7) and (8), and a number of non-standard proposals such as RDFS⁺ [1] or OWL-LD [26] have also been put forward; despite these proposals, the responses show that the majority of the respondents view these issues as unresolved for OWL. Of the three critiques, (9) a lack of usable tools is the one identified as most universally affecting the standards according to respondents, 34 of whom identified all four standards as being problematic with respect to this issue.

3.5. Comments

The questionnaire ended with a comments section, where respondents could indicate both public and private comments. These comments varied in content.

Some comments, both positive and negative, spoke directly of the questionnaire. Aside from individual comments relating to the questionnaire being too long, the way in which options were ordered, and the lack of a “don’t know” option (rather each question was optional) a number of public comments suggested other issues not raised, specifically relating to: social aspects, shared vocabularies, complex information modelling, agility of standardisation, RDF syntaxes, semantic modelling, lack of high-level abstractions, etc.

Other comments expressed more detailed opinions on the overall theme of the questionnaire, on specific critiques, or on their outlook for the Semantic Web. Some comments related to being less focused on adoption of Semantic Web standards and more focused on

the adoption of its concepts and best practices (even if not using RDF(S), SPARQL, OWL, etc.); how incentives may be bootstrapped; a lack of focus on how data are used; key use-cases such as data maintenance and research data management (under FAIR principles); the need for new/improved standards; the difficulty of modelling certain data in RDF; the need for more dog-fooding, education and marketing; problems with the Semantic Web being driven primarily by academia; etc. Other comments rather took a more pessimistic view, noting that if the Semantic Web were useful we should have seen more of it by now, that the Web of “walled gardens” looks set to continue, etc. We refer to the public comments online for more details [39].

4. Discussion

Two decades on, the general consensus in the Semantic Web community appears to be that there is still a long way to go before the original vision of the Semantic Web is realised. On the other hand, the consensus is that Semantic Web technologies are presently having some impact on both the Web and in non-Web settings, and will continue to have more impact looking to the future. Along these lines, respondents to our survey cite success stories such as schema.org, Knowledge Graphs, Wikidata, DBpedia, Biomedical Ontologies, etc., as examples where the Semantic Web has had most impact thus far. On the other hand, a lack of usable tools, a lack of incentives, a lack of robustness for unreliable publishers, and overly verbose standards, in particular, are widely acknowledged as valid criticisms of the Semantic Web in its current state.

Looking to the future, the general consensus is that while none of the highlighted issues are insurmountable, many do pose non-trivial obstacles to the further adoption and development of the Semantic Web. A theme widely recognised as a key obstacle for the Semantic Web is the lack of availability of usable tools; such issues are known with the community and have been discussed, for example, by Karger et al. [44].

Part of the reason for the lack of usable tools may also be due to the largely academic nature of the Semantic Web, where work on such tools is difficult to publish (seen as “engineering” rather than “science”), while the community perhaps lacks expertise in areas such as Human Computer Interaction (HCI) relating to conducting and publishing usability studies. Another major issue is the lack of incentives, which, with some exceptions such as schema.org [29], remains a general challenge; while some authors have begun to tackle this issue from a more general point-of-view [28], more work is called for. The results of the questionnaire also highlight the need for more work on data quality [87] and methods to ensure robustness in the presence of unreliable publishers [51, 66]. The results further reveal issues relating to the (perceived) verbosity of the core standards, particularly OWL, perhaps suggesting the need for (further [1, 26, 60]) work to better understand and address this issue. A more transversal theme is implicit in the responses: the Semantic Web needs more contributors from other research communities and from outside academia.

Among all of the mentioned issues, one that stands out, in particular, relates to the *usability of Semantic Web technologies and their accessibility to newcomers*. We thus call for more work on this particular topic—work that may take a number of directions. First, we require more work on tools and interfaces that reduce the cognitive load and expertise required for users to benefit from Semantic Web technologies; ideally, the design of such tools and interfaces should be guided by usability studies with target end-users. Second, we require more work on making the Semantic Web standards more accessible and appealing to newcomers; this may involve simplifying standards, creating more lightweight profiles of existing standards, creating interactive primers to motivate and introduce the standards in a more engaging manner, and so forth.

Along similar lines, we further call for more works that *bridge the Semantic Web with other technologies having similar goals*, particularly those that gain (or have gained) considerable traction. This may take a number of forms. In the case of languages, mappings can be created to make the technologies interoperable, as was done in the case of OBO and OWL [27], SPARQL and SQL [20], and so forth; more work can be done to align new query languages like Cypher [59], Gremlin [68] or GraphQL [32] with SPARQL [2, 21, 77], thus more closely aligning the graph database/NoSQL/Web developer community with the Semantic Web. A second option is to

take existing technologies, and extend them to support Semantic Web concepts; this has worked particularly well for JSON-LD [76], taking a familiar concept for developers – JSON – and adding some additional syntactic sugar to create an RDF-compatible data format.

The results presented herein highlight that the original vision of the Semantic Web still eludes us, though major strides have been made in recent years. No matter how elusive, however, the Semantic Web vision remains an alluring one (at least to some, including the present author). We are all intimately aware of how the Web has revolutionised society, where the Semantic Web has the potential to further propel the Web to a new stage, marked by unprecedented levels of automation and convenience for users. Unlike twenty years ago, we now have the benefit of many years of experience and research on the topic, as well as established success stories like schema.org, Wikidata, Biomedical Ontologies, etc., to further build upon. Even a partial realisation of the Semantic Web vision will serve (and arguably is serving) as a great boon to society, much like how A.I. is finding more and more applications without ever having surpassed the Turing test. Part of the criticism, perhaps, stems from comparing the Semantic Web with the Web: a technological development to which almost anything else would pale in comparison; while the Semantic Web has not seen the same level of rapid growth and penetration as the Web, this does not devalue the (sometimes quiet) impact that the Semantic Web community can point to, while still hinting at the vast impact it *could* potentially have. Two decades on, it is thus still a vision that merits patient pursuit, even if – or perhaps even *especially given that* – there is much work left to be done before the Semantic Web holds the sorts of conclusive answers that might satisfy even its most ardent critics.

Acknowledgements We thank the respondents to the questionnaire and the reviewers for their helpful comments. This work was funded in part by the Millennium Institute for Foundational Research on Data (IMFD) and Fondecyt, Grant No. 1181896.

References

- [1] D. Allemang and J.A. Hendler, *Semantic Web for the Working Ontologist - Effective Modeling in RDFS and OWL, Second Edition*, Morgan Kaufmann, 2011. ISBN 978-0-12-385965-5.
- [2] R. Angles, H. Thakkar and D. Tomaszuk, RDF and Property Graphs Interoperability: Status and Issues, in: *Alberto Mendelzon International Workshop (AMW)*, Vol. 2369, CEUR-WS.org, 2019.

- [3] R. Angles, M. Arenas, P. Barceló, A. Hogan, J.L. Reutter and D. Vrgoc, Foundations of Modern Query Languages for Graph Databases, *ACM Comput. Surv.* **50**(5) (2017), 68:1–68:40. doi:10.1145/3104031.
- [4] C.B. Aranda, A. Hogan, J. Umbrich and P. Vandenbussche, SPARQL Web-Querying Infrastructure: Ready for Action?, in: *International Semantic Web Conference*, 2013, pp. 277–293. doi:10.1007/978-3-642-41338-4_18.
- [5] M. Arenas, A. Bertails, E. Prud’hommeaux and J. Sequeda, A Direct Mapping of Relational Data to RDF, 2012, <https://www.w3.org/TR/rdb-direct-mapping/>.
- [6] M. Bergman, Scalability of the Semantic Web (2006), <http://www.mkbergman.com/227/scalability-of-the-semantic-web/>.
- [7] T. Berners-Lee, The Future of RDF, 2010, <https://www.w3.org/DesignIssues/RDF-Future.html>.
- [8] T. Berners-Lee, J. Hendler and O. Lassila, The Semantic Web, *Scientific American* **284**(5) (2001), 34–43.
- [9] C. Bizer and A. Schultz, The Berlin SPARQL Benchmark, *Int. J. Semantic Web Inf. Syst.* **5**(2) (2009), 1–24. doi:10.4018/jswis.2009040101.
- [10] K.D. Bollacker, C. Evans, P. Paritosh, T. Sturge and J. Taylor, Freebase: a collaboratively created graph database for structuring human knowledge, in: *ACM SIGMOD International Conference on Management of Data*, 2008, pp. 1247–1250. doi:10.1145/3178876.3186602.
- [11] P.A. Bonatti, A. Hogan, A. Polleres and L. Sauro, Robust and scalable Linked Data reasoning incorporating provenance and trust annotations, *J. Web Semant.* **9**(2) (2011), 165–201. doi:10.1016/j.websem.2011.06.003.
- [12] J.G. Breslin, A. Passant and D. Vrandečić, Social Semantic Web, in: *Handbook of Semantic Web Technologies*, Springer, 2011, pp. 467–506. doi:10.1007/978-3-540-92913-0_12.
- [13] D. Brickley, R.V. Guha and A. Layman, Resource Description Framework (RDF) Schemas, 1998, <https://www.w3.org/TR/1998/WD-rdf-schema-19980409/>.
- [14] J. Cabeda, Semantic Web is Dead, Long live the AI!!!, 2017, <https://hackernoon.com/semantic-web-is-dead-long-live-the-ai-2a5ea0cf6423>.
- [15] K. Cagle, Why the Semantic Web Has Failed, 2016, <https://www.linkedin.com/pulse/why-semantic-web-has-failed-kurt-cagle>.
- [16] S. Chang, Scaling Knowledge Access and Retrieval at Airbnb, 2018, <https://medium.com/airbnb-engineering/scaling-knowledge-access-and-retrieval-at-airbnb-665b6ba21e95>.
- [17] S. Corlosquet, R. Delbru, T. Clark, A. Polleres and S. Decker, Produce and Consume Linked Data with Drupal!, in: *International Semantic Web Conference*, 2009, pp. 763–778. doi:10.1007/978-3-642-04930-9_48.
- [18] R. Cyganiak, D. Reynolds and J. Tennyson, The RDF Data Cube Vocabulary, 2014, <https://www.w3.org/TR/vocab-data-cube/>.
- [19] A. Dadzie and E. Pietriga, Visualisation of Linked Data - Reprise, *Semantic Web* **8**(1) (2017), 1–21. doi:10.3233/SW-160249.
- [20] S. Das, S. Sundara and R. Cyganiak, R2RML: RDB to RDF Mapping Language, 2012, <https://www.w3.org/TR/r2rml/>.
- [21] S. Das, J. Srinivasan, M. Perry, E.I. Chong and J. Banerjee, A Tale of Two Graphs: Property Graphs as RDF in Oracle, in: *International Conference on Extending Database Technology (EDBT)*, OpenProceedings.org, 2014, pp. 762–773. doi:10.5441/002/edbt.2014.82.
- [22] C. Doctorow, Metacrap: Putting the torch to seven straws of the meta-utopia (2001), <https://people.well.com/user/doctorow/metacrap.htm>.
- [23] J. Euzenat and P. Shvaiko, *Ontology Matching, Second Edition*, Springer, 2013. ISBN 978-3-642-38720-3.
- [24] F. Gandon, A Survey of the First 20 Years of Research on Semantic Web and Linked Data, 2018, [ff10.3166/ISI.23.3-4.11-56ff.fhah-01935898](https://arxiv.org/abs/1803.01935).
- [25] J. Getman, J. Ellis, Z. Song, J. Tracey and S.M. Strassel, Overview of Linguistic Resources for the TAC KBP 2017 Evaluations: Methodologies and Results, in: *Text Analysis Conference (TAC)*, NIST, 2017.
- [26] B. Glimm, A. Hogan, M. Krötzsch and A. Polleres, OWL: Yet to arrive on the Web of Data?, in: *Linked Data on the Web (LDOW)*, CEUR-WS.org, 2012.
- [27] C. Golbreich, M. Horridge, I. Horrocks, B. Motik and R. Shearer, OBO and OWL: Leveraging Semantic Web Technologies for the Life Sciences, in: *International Semantic Web Conference (ISWC)*, Springer, 2007, pp. 169–182. doi:10.1007/978-3-540-76298-0_13.
- [28] T. Grubenmann, A. Bernstein, D. Moor and S. Seuken, Financing the Web of Data with Delayed-Answer Auctions, in: *World Wide Web Conference (WWW)*, ACM, 2018, pp. 1033–1042. doi:10.1145/3178876.3186602.
- [29] R.V. Guha, D. Brickley and S. Macbeth, Schema.org: evolution of structured data on the web, *CACM* **59**(2) (2016), 44–51. doi:10.1145/2844544.
- [30] R.V. Guha, R. McCool and R. Fikes, Contexts for the Semantic Web, in: *International Semantic Web Conference*, Springer, 2004, pp. 32–46. doi:10.1007/978-3-540-30475-3_4.
- [31] O. Hartig, Foundations of RDF* and SPARQL*, in: *Alberto Mendelzon International Workshop on Foundations of Data Management and the Web*, CEUR-WS.org, 2017.
- [32] O. Hartig and J. Pérez, Semantics and Complexity of GraphQL, in: *World Wide Web Conference (WWW)*, ACM, 2018, pp. 1155–1164. doi:10.1145/3178876.
- [33] Q. He, B.-C. Chen and D. Agarwal, Building The LinkedIn Knowledge Graph, 2016, <https://engineering.linkedin.com/blog/2016/10/building-the-linkedin-knowledge-graph>.
- [34] T. Heath and C. Bizer, *Linked Data: Evolving the Web into a Global Data Space*, Synthesis Lectures on the Semantic Web, Morgan & Claypool Publishers, 2011. doi:10.2200/S00334ED1V01Y201102WBE001.
- [35] J. Hendler, “Why the Semantic Web will never work”, in: *Extended Semantic Web Conference (ESWC)*, Keynote, 2013, http://videolectures.net/eswc2011_hendler_work/.
- [36] D. Hernández, A. Hogan and M. Krötzsch, Reifying RDF: What Works Well With Wikidata?, in: *International Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS)*, Vol. 1457, CEUR-WS.org, 2015, pp. 32–47.
- [37] D. Hernández, A. Hogan, C. Riveros, C. Rojas and E. Zerega, Querying Wikidata: Comparing SPARQL, Relational and Graph Databases, in: *International Semantic Web Conference*, 2016, pp. 88–103. doi:10.1007/978-3-319-46547-0.
- [38] A. Hogan, Canonical Forms for Isomorphic and Equivalent RDF Graphs: Algorithms for Learning and Labelling Blank Nodes, *TWEB* **11**(4) (2017), 22:1–22:62. doi:10.1145/3068333.
- [39] A. Hogan, Responses to “Semantic Web: Perspectives” Questionnaire, 2019, <https://doi.org/10.5281/zenodo.3229401>. doi:10.5281/zenodo.3229401.

- [40] A. Hogan, P. Hitzler and K. Janowicz, Linked Dataset description papers at the Semantic Web journal: A critical assessment, *Semantic Web* **7**(2) (2016), 105–116. doi:10.3233/SW-160216.
- [41] A. Hogan, J. Umbrich, A. Harth, R. Cyganiak, A. Polleres and S. Decker, An empirical survey of Linked Data conformance, *J. Web Semant.* **14** (2012), 14–44. doi:10.1016/j.websem.2012.02.001.
- [42] I. Horrocks, Semantic Web: the story so far, in: *International Cross-Disciplinary Conference on Web Accessibility (W4A)*, 2007, pp. 120–125. doi:10.1109/ICTAI.2007.181.
- [43] S. Kanza and J.G. Frey, A new wave of innovation in Semantic Web tools for drug discovery, *Expert Opinion on Drug Discovery* **14**(5) (2019), 433–444. doi:10.1080/17460441.2019.1586880.
- [44] D.R. Karger, The Semantic Web and End Users: What’s Wrong and How to Fix It, *IEEE Int. Comp.* **18**(6) (2014), 64–70. doi:10.1109/MIC.2014.124.
- [45] M. Karnstedt, K. Sattler and M. Hauswirth, Scalable distributed indexing and query processing over Linked Data, *J. Web Semant.* **10** (2012), 3–32. doi:10.1016/j.websem.2011.11.010.
- [46] G. Kobilarov, T. Scott, Y. Raimond, S. Oliver, C. Sizemore, M. Smethurst, C. Bizer and R. Lee, Media Meets Semantic Web - How the BBC Uses DBpedia and Linked Data to Make Connections, in: *European Semantic Web Conference (ESWC)*, 2009, pp. 723–737. doi:10.1007/978-3-642-02121-3.
- [47] A. Krishnan, Making search easier: How Amazon’s Product Graph is helping customers find products more easily, 2018, <https://blog.aboutamazon.com/innovation/making-search-easier>.
- [48] O. Lassila and R.R. Swick, Resource Description Framework (RDF) Model and Syntax Specification, 1998, <https://www.w3.org/TR/1998/WD-rdf-syntax-19981008/>.
- [49] D. Lee, N. de Keizer, F. Lau and R. Cornet, Literature review of SNOMED CT use, *Journal of the American Medical Informatics Association* **21**(e1) (2013), e11–e19. doi:10.1136/amiajnl-2013-001636.
- [50] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer and C. Bizer, DBpedia - A large-scale, multilingual knowledge base extracted from Wikipedia, *Semantic Web* **6**(2) (2015), 167–195. doi:10.3233/SW-140134.
- [51] F. Maier, Y. Ma and P. Hitzler, Paraconsistent OWL and related logics, *Semantic Web* **4**(4) (2013), 395–427. doi:10.3233/SW-2012-0066.
- [52] S. Malyshev, M. Krötzsch, L. González, J. Gonsior and A. Bielefeldt, Getting the Most Out of Wikidata: Semantic Technology Usage in Wikipedia’s Knowledge Graph, in: *International Semantic Web Conference*, 2018, pp. 376–394. doi:10.1007/978-3-030-00668-6_23.
- [53] E. Mansour, A.V. Samba, S. Hawke, M. Zereba, S. Capadisli, A. Ghanem, A. Aboulnaga and T. Berners-Lee, A Demonstration of the Solid Platform for Social Web Applications, in: *International Conference on World Wide Web (WWW)*, 2016, pp. 223–226. doi:10.1145/2872518.2890529.
- [54] J.L. Martinez-Rodriguez, A. Hogan and I. Lopez-Arevalo, Information Extraction meets the Semantic Web: A Survey, *Semantic Web* (2020), To Appear; <http://www.semantic-web-journal.net/content/information-extraction-meets-semantic-web-survey-0>. doi:10.3233/SW-180333.
- [55] P.N. Mendes, H. Mühleisen and C. Bizer, Sieve: linked data quality assessment and fusion, in: *PJoint EDBT/ICDT Workshops*, 2012, pp. 116–123. doi:10.1145/2320765.2320803.
- [56] R. Meusel, C. Bizer and H. Paulheim, A Web-scale Study of the Adoption and Evolution of the schema.org Vocabulary over Time, in: *International Conference on Web Intelligence, Mining and Semantics (WIMS)*, ACM, 2015, pp. 15:1–15:11. doi:10.1145/2797115.2797124.
- [57] R. Meusel, P. Petrovski and C. Bizer, The WebDataCommons Microdata, RDFa and Microformat Dataset Series, in: *International Semantic Web Conference*, 2014, pp. 277–292. doi:10.1007/978-3-319-11964-9_18.
- [58] P. Mika, What Happened To The Semantic Web?, in: *ACM Conference on Hypertext and Social Media (HYPERTEXT)*, 2017, p. 3. doi:10.1145/3078714.3078751.
- [59] J.J. Miller, Graph Database Applications and Concepts with Neo4j, in: *Southern Association for Information Systems Conference (SAIS)*, AIS eLibrary, 2013. doi:10.1002/bult.2010.1720360610.
- [60] B. Motik, B.C. Grau, I. Horrocks, Z. Wu, A. Fokoue and C. Lutz, OWL 2 Web Ontology Language Profiles (Second Edition), 2012, <https://www.w3.org/TR/owl2-profiles/>.
- [61] A.N. Ngomo and S. Auer, LIMES - A Time-Efficient Approach for Large-Scale Link Discovery on the Web of Data, in: *International Joint Conference on Artificial Intelligence (IJCAI)*, 2011, pp. 2312–2317. doi:10.5591/978-1-57735-516-8/IJCAI11-385.
- [62] E. Oren, R. Delbru, M. Catasta, R. Cyganiak, H. Stenzhorn and G. Tummarello, Sindice.com: a document-oriented lookup index for open linked data, *IJMSO* **3**(1) (2008), 37–52. doi:10.1504/IJMSO.2008.021204.
- [63] J. Pérez, M. Arenas and C. Gutiérrez, Semantics and complexity of SPARQL, *ACM Trans. Database Syst.* **34**(3) (2009), 16:1–16:45. doi:10.1145/1567274.1567278.
- [64] P. Petrescu, Google Organic Click-Through Rates in 2014, 2014, <https://moz.com/blog/google-organic-click-through-rates-in-2014>.
- [65] R. Pittman, A. Srivastava, S. Hewavitharana, A. Kale and S. Mansour, Cracking the Code on Conversational Commerce, 2017, <https://www.ebayinc.com/stories/news/cracking-the-code-on-conversational-commerce/>.
- [66] A. Polleres, A. Hogan, R. Delbru and J. Umbrich, RDFS and OWL Reasoning for Linked Data, in: *Reasoning Web*, Springer, 2013, pp. 91–149. doi:10.1007/978-3-642-39784-4.
- [67] J. Rochkind, Is the semantic web still a thing?, 2014, <https://bibwild.wordpress.com/2014/10/28/is-the-semantic-web-still-a-thing/>.
- [68] M.A. Rodriguez, The Gremlin graph traversal machine and language, in: *Symposium on Database Programming Languages (DBPL)*, ACM, 2015, pp. 1–10. doi:10.1145/2815072.2815073.
- [69] E. Sandhaus, Build Your Own NYT Linked Data Application, 2010, <https://open.blogs.nytimes.com/2010/03/30/build-your-own-nyt-linked-data-application/>.
- [70] M. Schmachtenberg, C. Bizer and H. Paulheim, Adoption of the Linked Data Best Practices in Different Topical Domains, in: *International Semantic Web Conference (ISWC)*, Springer, 2014, pp. 245–260. doi:10.1007/978-3-319-11964-9.
- [71] M. Schmidt, O. Görlitz, P. Haase, G. Ladwig, A. Schwarte and T. Tran, FedBench: A Benchmark Suite for Federated Semantic Data Query Processing, in: *International Seman-*

- 1 *tic Web Conference (ISWC)*, Springer, 2011, pp. 585–600.
2 doi:10.1007/978-3-642-25073-6.
- 3 [72] N. Shadbolt, T. Berners-Lee and W. Hall, The Semantic Web
4 Revisited, *IEEE Intelligent Systems* **21**(3) (2006), 96–101.
5 doi:10.1109/MIS.2006.62.
- 6 [73] S. Shrivastava, Bring rich knowledge of people,
7 places, things and local businesses to your apps, 2017,
8 [https://blogs.bing.com/search-quality-insights/2017-
9 07/bring-rich-knowledge-of-people-places-things-and-local-
10 businesses-to-your-apps](https://blogs.bing.com/search-quality-insights/2017-07/bring-rich-knowledge-of-people-places-things-and-local-businesses-to-your-apps).
- 11 [74] A. Singhal, Introducing the Knowledge Graph: things,
12 not strings, 2012, [https://www.blog.google/products/search/
13 introducing-knowledge-graph-things-not/](https://www.blog.google/products/search/introducing-knowledge-graph-things-not/).
- 14 [75] S. Speicher, J. Arwe and A. Malhotra, Linked Data Platform
15 1.0, 2015, <https://www.w3.org/TR/ldp/>.
- 16 [76] M. Sporny, D. Longley, G. Kellogg, M. Lanthaler and N. Lind-
17 ström, JSON-LD 1.0 – A JSON-based Serialization for Linked
18 Data, 2014, <https://www.w3.org/TR/json-ld/>.
- 19 [77] R. Taelman, M.V. Sande and R. Verborgh, GraphQL-LD:
20 Linked Data Querying with GraphQL, in: *International Se-
21 mantic Web Conference – Posters & Demos (ISWC P&D)*,
22 CEUR-WS.org, 2018.
- 23 [78] S. Target, Whatever Happened to the Semantic Web?, 2018,
24 <https://twobithistory.org/2018/05/27/semantic-web.html>.
- 25 [79] D. ter Heide, Three reasons why the Semantic Web has failed,
26 2013, [https://gigaom.com/2013/11/03/three-reasons-why-the-
27 semantic-web-has-failed/](https://gigaom.com/2013/11/03/three-reasons-why-the-semantic-web-has-failed/).
- 28 [80] J. Urbani, S. Kotoulas, J. Maassen, F. van Harmelen and
29 H.E. Bal, WebPIE: A Web-scale Parallel Inference En-
30 gine using MapReduce, *J. Web Semant.* **10** (2012), 59–75.
31 doi:10.1016/j.websem.2011.05.004.
- 32 [81] P. Vandenbussche, G. Ateazing, M. Poveda-Villalón and
33 B. Vantant, Linked Open Vocabularies (LOV): A gateway to
34 reusable semantic vocabularies on the Web, *Semantic Web* **8**(3)
35 (2017), 437–452. doi:10.3233/SW-160213.
- 36 [82] J. Volz, C. Bizer, M. Gaedke and G. Kobilarov, Silk - A Link
37 Discovery Framework for the Web of Data, in: *Linked Data on
38 the Web (LDOW)*, 2009.
- 39 [83] D. Vrandečić and M. Krötzsch, Wikidata: a free col-
40 laborative knowledgebase, *CACM* **57**(10) (2014), 78–85.
41 doi:10.1145/2629489.
- 42 [84] Z. Wu, G. Eadon, S. Das, E.I. Chong, V. Kolovski, M. An-
43 namalai and J. Srinivasan, Implementing an Inference Engine
44 for RDFS/OWL Constructs and User-Defined Rules in Ora-
45 cle, in: *International Conference on Data Engineering*, 2008,
46 pp. 1239–1248. doi:10.1109/ICDE.2008.4497533.
- 47 [85] R.V. Yampolskiy, AI-Complete, AI-Hard, or AI-Easy – Classi-
48 fication of Problems in AI, in: *Midwest Artificial Intelligence
49 and Cognitive Science Conference (MAICS)*, CEUR Workshop
50 Proceedings, 2012, pp. 94–101.
- 51 [86] T. Yasseri, R. Sumi, A. Rung, A. Kornai and J. Kertész, Dy-
namic of Conflicts in Wikipedia, *PLOS ONE* **7**(6) (2012).
doi:10.1371/journal.pone.0038869.
- [87] A. Zaveri, A. Rula, A. Maurino, R. Pietrobon, J. Lehmann and
S. Auer, Quality assessment for Linked Data: A Survey, *Se-
mantic Web* **7**(1) (2016), 63–93. doi:10.3233/SW-150175.
- [88] Information technology – Database languages – SQL, 2016,
SQL:2016; ISO/IEC 9075:2016.
- [89] The Open Graph Protocol, 2017, <http://ogp.me/>.
- [90] Amazon Neptune: Fast, reliable graph database built for the
cloud, 2019, <https://aws.amazon.com/neptune/>.