Intelligent Agents: The Vision Revisited

Sabrina Kirrane[1] and Stefan Decker[2]

[1] Vienna University of Economics and Business, Vienna, Austria
[2] RWTH Aachen University, Germany, Germany

Abstract. As early as the mid sixties, motivated by the ever growing body of scientific knowledge, scholars identified the need for data to be organised in a manner that is more intuitive for humans to digest. Additionally, they envisioned a future where intelligent systems would be able to make sense of vast amounts of data and alleviate humans from performing complex analytical tasks. Although Semantic Web technologies have demonstrated great potential in this regard, the vision has yet to be realised. In this position paper, we examine the status quo in terms of making data available as Linked Data and highlight some of the challenges experienced by Linked Data publishers and consumers. Following on from this we revisit the original vision of the Semantic Web and argue for additional research to support interaction between intelligent software agents constrained via goals, preferences, norms and usage restrictions, in a manner that fosters trustworthiness in the services delivered.

1 Introduction

The idea to use graphs to represent knowledge, which can be automatically actioned upon by machines, has been around since the early 60's. Both, Engelbart [9] (in 1962) and Lickleder [18, 23] (in 1965) imagined a future where machines would be able to automatically process and reason over data represented in knowledge graphs. Almost forty years later the seminal Semantic Web paper by Berners-Lee et al. [3] described their vision of a Semantic Web, whereby the existing web infrastructure could be used to represent data in a manner that could be automatically actioned upon by intelligent software agents.

Roughly five years after the seminal paper both Shadbolt et al. [25] and Feigenbaum et al. [10] reflected on the state of the art at the time and concluded that although intelligent agents were still far from being realised the technology was steadily gaining traction especially as a means of data integration. More recently, Glimm and Stuckenschmidt [13] and Bernstein et al. [4] confirm, that approximately 17 years on, the vision has yet to become a reality. The authors observe that although the primary focus of the Semantic Web community was initially on knowledge representation, reasoning and querying, in recent years there has been a broadening beyond pure Semantic Web topics to include knowledge extraction, discovery, search and retrieval. However, according to Bernstein et al. [4] there are still a number of issues concerning heterogeneity both in terms of representation and semantics, and diversity in terms of web data quality. Additionally, the authors identify new challenges that arise with increasing data volume and publishing velocity.

Although Bernstein et al. [4] point to several challenges that the community will face in the future, there is a distinct lack of focus on topics that are important from an intelligent agents perspective, yet remain under represented within the community. For instance, technologies, techniques and protocols that enable agents to interact with other agents in order to carry out their activities according to constraints in the form of goals, preferences, and usage limitations in a trustworthy manner.

In order to fill this gap, this position paper revisits the original vision of the Semantic Web and highlights several open research questions that are important if we hope to one day have intelligent agents that are able to act on our behalf. We start by examining the status quo in terms of existing Linked Data management practices. In particular, we discuss data management through the lens of the FAIR[3] (**F**indable, **A**ccessible, **I**nteroperable and **R**eusable) data principles, and highlight several open research challenges in terms of persistent identifiers, indexing, and usage constraints. Following on from this we argue for adapting and extending these FAIR principles to guide the development of FAIR ICT Agents, whereby ICT denotes **I**nteractive intelligent agents that are **C**onstrained via goals, preferences, norms and usage restrictions, in a manner that fosters **T**rustworthiness.

The remainder of the paper is structured as follows: *Section* 2 introduces the FAIR data principles and discusses some of the limitations of current Linked Data management practices. *Section* 3 presents several challenges and opportunities that need to be overcome for FAIR ICT Agents to become a reality. Finally *Section* 4 concludes the paper and identifies several open research questions.

2 Making Linked Data FAIR

The Semantic Web enables things, otherwise known as resources represented using the Resource Description Framework (RDF) data model, to be linked using Internationalised Resource Identifiers (IRIs) in a similar way to how web documents are linked using the HyperText Markup Language (HTML) hypertext reference (HREF) attribute. Linked Data is a related concept, which refers to a set of best practices for publishing and connecting structured data on the Web [19]. In recent years, we have seen significant advances in the technology used to both publish and consume Linked Data, however a recent article by Beek et al. [2] claims that the existing Semantic Web is neither traversable nor machine-processable, and consequently argues that the Semantic Web needs centralisation. In this position paper, we argue for treating the root cause of the problem (i.e., highlighting existing data management challenges and calling for best practices guidelines and research to address them) rather than the symptoms (i.e., developing centralised solutions on top of distributed web data that

[3] FAIR data principles, https://www.force11.org/node/6062

address the inherit limitations of the existing infrastructure). In terms of the former we argue that a necessary first step is to provide additional guidelines for data publishers that go beyond the original Linked Data principles and the well known 5 star rating system[4].

An emerging best practice in terms of scientific knowledge dissemination is the adoption of FAIR data principles [29], whereby researchers strive to ensure that their research objects (papers, datasets, code etc...) are **F**indable, **A**ccessible, **I**nteroperable and **R**eusable. Although the FAIR principles were devised to provide guidance for managing scholarly assets, we believe that said principles could be adapted to provide guidance to Linked Data publishers in order to improve the findablility, accessibility, interoperability and reusability of machine readable data available on the Web.

2.1 FAIR data

The core objective of the FAIR data principles is to provide guidance to scholarly data publishers in terms of making their data reusable by both humans and machines. The four foundational principles can be summarised as follows:

- To be deemed **F**indable, data should be uniquely identifiable via persistent identifiers, these identifiers should be used to associate descriptive metadata with the data, and both data and metadata should be indexed in a manner that is easy to search.
- In order to make data **A**ccessible it should be possible to retrieve the data via common protocol(s), that are open, free, universally implementable and can support usage constraints where desirable.
- Making (meta)data[5] **I**nteroperable is primarily concerned with the representation of (meta)data in a manner that facilitates integration e.g. using common/standard ontologies and vocabularies.
- Finally, (meta)data is **R**eusable if it is richly described in terms of relevant attributes, contains relevant provenance information and is compatible with domain specific standards.

2.2 FAIR Linked Data

There is clearly a strong connection between said principles and Semantic Web technologies and Linked Data principles. Both **R**eusability and **I**nteroperablity are at the core of the Resource Description Framework (RDF) data model. By using RDF to describe resources, it is possible to describe complex relations between resources in a machine readable format. Ontologies provide for a shared understanding of things and how they are related, that can easily be reused and extended. Data is linked to other data using HyperText Transfer Protocol (HTTP) IRIs that can be used to identify things (papers, datasets, code etc...).

[4] https://www.w3.org/DesignIssues/LinkedData.html
[5] In order to improve readability in this paper we use (meta)data to denote to data and metadata.

Although the RDF data model and Linked Data principles are good starting points in terms of making Linked data FAIR, there are still a number of open research challenges. In terms of **F**indablility, according to FAIR **(meta)data should be identifiable via persistent identifiers**. Despite a push by the community to use persistent identifiers, for instance for resources submitted to the International Semantic Web Conference (ISWC) resources track[6], they are still not widely used in practice. Another key aspect of **F**indablility is the **indexing of (meta)data in a manner that is easy to search**. Although there have been a number of proposals (cf. [2, 11]), given that indexing is done in a centralised manner existing proposals suffer from data freshness issues. From an **A**ccessibility perspective when it comes to **usage constraints that describe how the data should be used** there is a large body of work on access control specification and enforcement strategies for RDF [21] and licensing [14, 15, 16, 27] proposals for data exposed as Linked Data (cf. Section 3 for additional details). The challenge here is the fact that existing usage control strategies (where used) are still very primitive.

Although FAIR was devised to provide guidance in terms of effective scholarly data management, we argue that by adapting the FAIR data principles for Linked Data it will be possible not only to identify existing challenges in terms of data management, which we only touch upon in this article, but also to provide a best practice guide for dealing with these challenges.

3 Towards Intelligent Agents

Returning to the original visions by Berners-Lee et al. [3], whereby intelligent agents are able to make sense of Web data and alleviate humans from performing complex analytical tasks, it is clear that FAIR principles alone are not enough as they focus on data management without considering how intelligent agents might make use of this data. In this respect, we identify the need for FAIR ICT Agents, whereby ICT denotes **I**nteractive intelligent agents that are **C**onstrained via goals, preferences, norms and usage restrictions, in a manner that fosters **T**rustworthiness.

3.1 Interactive intelligent agents

When it comes to intelligent agents the services offered by each agent need to be designed in a manner such that multiple agents can interact (and possibly even collaborate) in order to complete tasks and solve problems. Each agent needs to maintain a list of services that it is capable of executing based on the (meta)data in its knowledge graph (including descriptive attributes, constraints and provenance data). Ideally, the list of services should grow organically with the data and as the agent uncovers new insights based on incremental analysis of its knowledge graph.

[6] http://iswc2018.semanticweb.org/call-for-resources-track-papers/

Unlike traditional web services, semantic web services use formal ontology-based annotations to describe the service in a manner that can be automatically interpreted by machines. In the early years of the Semantic Web there were several standardisation initiatives, namely the Web Ontology Language for Web Services (OWL-S)[7], the Web Service Modeling Language (WSML)[8], the W3C standard Semantic Annotations for WSDL and XML Schema (SAWSDL)[9]. A survey conducted by Klusch et al. [22] provides a summary of existing work and describes the various semantic web service search architectures (i.e. centralised and decentralised directory based, and decentralised directoryless). The authors conclude that **research into decentalised semantic service search is lagging far behind its centralised counterpart**. When it comes to semantic web services the big question is how do we support **adaptive discovery and composition of semantic services**? Other open research challenges are concerned with enabling **interoperability between policy aware agents**, and **dealing with agents joining and leaving the network at will**.

3.2 Constrained via goals, preferences, norms and usage restriction

Berners-Lee et al. [3] originally envisioned a system, where intelligent agents were capable of acting on behalf of humans. One of the key components of such a system is the policy language that is capable of capturing the constraints under which the agents operate. During the early days of the Semantic Web the development of general policy languages that leverage semantic technologies (such as KAoS [7], Rei [20] and Protune [6]), was an active area of research. General policy languages cater for a diverse range of functional requirements (e.g., access control, query answering, service discovery, negotiation, to name but a few). Considering that the policy language needs to be interpreted by machines, **formal semantics is important as it allows for the verification of correctness**. However, research into general semantic policy languages seems to have reduced considerably in recent years and **the suitability of existing general policy languages towards the intelligent agents vision** is still an open research question.

In terms of specific policy languages access control is a topic that has received a lot of attention over the years. Kirrane et al. [21] provide a detailed survey of the various access control models, standards and policy languages, and the different access control enforcement strategies for RDF. Although there have been several different proposals over the years, there is still no standard access control strategy for Linked Data. Considering the array of access control specification and enforcement mechanisms proposed to date, a necessary first step towards ensuring that intelligent agents have the ability to decide with whom they share information is to develop **a framework that can be used to evaluate existing access control offerings** in terms of expressivity, correctness and completeness. When it comes to usage control in the form

[7] https://www.w3.org/Submission/OWL-S/
[8] https://www.w3.org/Submission/WSML/
[9] https://www.w3.org/TR/sawsdl/

of licensing, research topics range from using Natural Language Processing to extract license rights and obligations [8] to licenses compatibility validation and composition [14, 15, 16, 27]. More recently, the Open Digital Rights Language (ODRL)[10], which became a W3C recommendation in February 2018, provides a promising first step towards the general adoption of machine understandable licenses, however **it remains to be seen if data publishers embrace the new standard** and if **license aware data querying and processing mechanisms become common practice**.

Another important research direction that remains underdeveloped is the use of policies to specify **societal norms and personal values that would enable agents to understand the constraints of the environment** in which they operate. Also, there are also several open research questions in terms of the suitability of the existing languages to **deal with the volume, velocity, variety and veracity of data** we are faced with today, the ability to **balance expressivity and computational complexity**, and ensuring that the intelligent agent ecosystem can deal with the **policy interoperabiliy needs of collaborating agents**.

3.3 Fostering trustworthiness

Artz and Gil [1] conducted a comprehensive survey of trust mechanisms in computer science in general and the Semantic Web in particular. The authors highlight that traditional approaches focused primarily on authentication via assertions by third parties, however in later years the topic evolved to include historical interaction data, the transfer of trust from trusted entities, and decentralised trust mechanisms (e.g. voting mechanisms or other consensus decision making mechanisms). Although there is a large body of computer science literature relating to trust **the effectiveness of existing trust mechanisms in the context of intelligent agents** has yet to be determined.

In an intelligent agent ecosystem local provenance chains could be used by agents to provide explanations for decisions made, while global provenance chains could be used to provide transparency with respect to collaborating agents or the distributed system as a whole. These provenance chains could also be used to record and retrieve historical data and to build trust between agents. Although there has been a number of proposals for representing provenance events (cf. [12, 17]). To date the focus has been on recording where the data came from or capturing the source of the data or changes to data over time. In this regard there have been several standardisation initiatives, such as *PROV*[11] and *OWL-Time*[12] ontologies, that can be used to represent *provenance* and *temporal* information respectively. In the context of intelligent agents **there is a need to record provenance with respect to both data and processing in a manner that can be easily digestible**.

[10] https://www.w3.org/TR/odrl-model/
[11] PROV,https://www.w3.org/TR/prov-overview/
[12] OWL-Time,https://www.w3.org/TR/owl-time/

From a provenance chains perspective there are two distinct avenues that could be leveraged, one built on top of existing web protocols [24, 28] and another based on blockchain technologies [30]. Weitzner et al. [28] present their vision of a policy-aware architecture for the Web, which includes three basic components: policy-aware audit logging, a policy language framework, and accountability reasoning tools. Specifically, they discuss how transparency and accountability can be achieved via distributed accountability appliances that communicate using existing web protocols. Seneviratne and Kagal [24] build on this idea by proposing a distributed accountability platform known as Accountable Hyper Text Transfer Protocol (HTTPA) that allows data producers to express usage restrictions and data consumers to express usage intentions. Unfortunately the authors only touch upon the required features and **the proposed accountability platform has yet to be assessed from both a functional or a non-functional requirements perspective**. Alternative distributed architectures for transparent personal data processing are discussed by Bonatti et al. [5], however the authors simply describe the opportunities and challenges, and the concrete implementation is left to future work. Zyskind et al. [30] discuss how blockchain technology could be extended to keep track of both data and access transactions. One of the primary drawbacks of the work is the fact that the authors focus on how to repurpose the blockchain as an access-control moderator as opposed to exploring the suitability of the proposed architecture for data transparency and governance. Another related avenue of research by Third and Domingue [26] proposes a semantic index for distributed ledgers.

Although, Blockchain platforms such as Ethereum[13] and Hyperledger Fabric[14] have the capability to support policy aware service provision, via smart contracts and chaincode, **the suitability of blockchain platforms in terms of both functional and non functional requirements** remains an open research question. In addition, there are a variety of societal challenges that also need to be considered, such as **the right to be forgotten, algorithmic biases, fake news, filter bubbles**, to name but a few.

4 Conclusion and Future Work

In this paper, we revisit the original vision of the Semantic Web whereby software agents are able to perform complex computational tasks on behalf of humans [3]. Inspired by recent surveys [4, 10, 13, 25] that analyse the evolution of Semantic Web technologies over almost two decades, we strive to shed light on important research topics that are necessary for the development of intelligent agents however are currently under represented at the predominant Semantic Web publishing venues. In order to frame the discussion we started by examining existing Linked Data publishing and consumption practices through the lens of the scientific FAIR data principles. From a data management perspective, we identified several open research challenges in terms of persistent identifiers,

[13] https://www.ethereum.org/
[14] https://www.hyperledger.org/projects/fabric

indexing, and usage constraints. From an application perspective, we argued for additional research to support interaction between agents that are constrained via goals, preferences, norms and usage restrictions, in a manner that fosters trustworthiness in the services delivered. From a best practices perspective, a potential first step is to adapt and extend the FAIR data principles such that they can serve as a best practice guide for FAIR ICT Agents. We do not claim that this is an exhaustive list of challenges, but rather with this position paper we hope to rejuvenate interest in these under represented topics with a view to bringing us closer to making the intelligent agent vision a reality.

Acknowledgments. Supported by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) DALICC project https://www.dalicc.net.

References

[1] D. Artz and Y. Gil. A survey of trust in computer science and the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 2007.

[2] W. Beek, L. Rietveld, S. Schlobach, and F. van Harmelen. Lod laundromat: Why the semantic web needs centralization (even if we don't like it). *IEEE Internet Computing*, 2016.

[3] T. Berners-Lee, J. Hendler, O. Lassila, et al. The semantic web. *Scientific american*, 2001.

[4] A. Bernstein, J. Hendler, and N. Noy. A new look at the semantic web. *Communications of the ACM*, 2016.

[5] P. Bonatti, S. Kirrane, A. Polleres, and R. Wenning. Transparent personal data processing: The road ahead. In *International Conference on Computer Safety, Reliability, and Security*, 2017.

[6] P. A. Bonatti and D. Olmedilla. Rule-based policy representation and reasoning for the semantic web. In *Proceedings of the Third International Summer School Conference on Reasoning Web*. Springer-Verlag, 2007.

[7] J. M. Bradshaw. *Software agents*. MIT press, 1997.

[8] E. Cabrio, A. Palmero Aprosio, and S. Villata. These Are Your Rights. In *The Semantic Web: Trends and Challenges*. Springer International Publishing, 2014.

[9] D. C. Engelbart. Augmenting human intellect: A conceptual framework. *Stanford Research Institute. Retrieved March*, 1962.

[10] L. Feigenbaum, I. Herman, T. Hongsermeier, E. Neumann, and S. Stephens. The semantic web in action. *Scientific American*, 2007.

[11] J. D. Fernández, W. Beek, M. A. Martínez-Prieto, and M. Arias. Lod-a-lot: A queryable dump of the lod cloud. In *The Semantic Web – ISWC 2017*. Springer/Verlag, 2017.

[12] G. Fu, E. Bolton, N. Q. Rosinach, L. I. Furlong, V. Nguyen, A. Sheth, O. Bodenreider, and M. Dumontier. Exposing provenance metadata using different rdf models. *arXiv preprint arXiv:1509.02822*, 2015.

[13] B. Glimm and H. Stuckenschmidt. 15 years of semantic web: An incomplete survey. *KI-Künstliche Intelligenz*, 2016.

[14] G. Governatori, A. Rotolo, S. Villata, and F. Gandon. One license to compose them all. In *International Semantic Web Conference*, 2013.

[15] G. Governatori, H.-P. Lam, A. Rotolo, S. Villata, G. A. Atemezing, and F. L. Gandon. Live: a tool for checking licenses compatibility between vocabularies and data. In *International Semantic Web Conference*, 2014.

[16] G. Guido, L. Ho-Pun, R. Antonino, V. Serena, and G. Fabien. Heuristics for Licenses Composition. *Frontiers in Artificial Intelligence and Applications*, 2013.

[17] O. Hartig. Provenance information in the web of data. *LDOW*, 2009.

[18] J. R. Hauben. Vannevar bush and jrc licklider: Libraries of the future 1945–1965. *The Amateur Computerist*, 2005.

[19] T. Heath and C. Bizer. *Linked data: Evolving the web into a global data space*. 2011.

[20] L. Kagal and T. Finin. A policy language for a pervasive computing environment. In *Proceedings POLICY 2003. IEEE 4th International Workshop on Policies for Distributed Systems and Networks*, 2003.

[21] S. Kirrane, A. Mileo, and S. Decker. Access control and the resource description framework: A survey. *Semantic Web*, 2017. URL `http://www.semantic-web-journal.net/system/files/swj1280.pdf`.

[22] M. Klusch, P. Kapahnke, S. Schulte, F. Lecue, and A. Bernstein. Semantic web service search: a brief survey. *KI-Künstliche Intelligenz*, 2016.

[23] J. C. R. Licklider. Libraries of the future. 1965.

[24] O. Seneviratne and L. Kagal. Enabling privacy through transparency. In *Privacy, Security and Trust (PST), 2014 Twelfth Annual International Conference on*, 2014.

[25] N. Shadbolt, T. Berners-Lee, and W. Hall. The semantic web revisited. *IEEE intelligent systems*, 2006.

[26] A. Third and J. Domingue. Linked data indexing of distributed ledgers. In *Proceedings of the 26th International Conference on World Wide Web Companion*, 2017.

[27] S. Villata and F. Gandon. Licenses compatibility and composition in the web of data. In *Third International Workshop on Consuming Linked Data (COLD2012)*, 2012.

[28] D. J. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. Hendler, and G. J. Sussman. Information accountability. *Communications of the ACM*, 2008.

[29] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, et al. The fair guiding principles for scientific data management and stewardship. *Scientific data*, 2016.

[30] G. Zyskind, O. Nathan, et al. Decentralizing privacy: Using blockchain to protect personal data. In *Security and Privacy Workshops (SPW), 2015 IEEE*, 2015.