All the Agents Challenge: Preface

1. Introduction

With the All the Agents Challenge at the 20th edition of the International Semantic Web Conference (ISWC) in 2021, we organisers wanted to further advance the efforts that bring two strands of work together:

1. The original vision for the Semantic Web, formulated around 2001 [1, 2], in which agents played a big part, but had not fully materialised [3, 4].
2. Recent progress in research, standardisation and technology around the Web of Things [5], the Linked Data Platform [6] and the Social Linked Data (Solid) platform [7], which can enable new use-cases and provide a substrate for agent technologies on the Semantic Web [8].

This bridging of technologies and communities has also been the topic of the Dagstuhl seminar "Autonomous Agents on the Web"[1] held in Feburary 2021. Therefore, for the 20th edition of the ISWC conference and in spirit of Jim Hendler's op-ed that asked "Where are all the agents?" [3], we posed the "All the Agents Challenge" (ATAC) to the Semantic Web community:

To build agents that do things on Linked Data

Conceptually, we assume agents that operate on a shared environment accessible via Linked Data interfaces. To allow the authors to focus on the agent part, the organisers prepared two environments, while not restricting the authors to bring their own:

- The Autonomous Maze Environment Explorer Project (AMEE) by Mike Amundsen[2] provides an environment for agents that have to escape a maze by following hypermedia affordances. AMEE can simulate mazes of different sizes. The organisers provide a RDF version of AMEE.
- The Building on Linked Data (BOLD) environment[3] is a description of a large real-world building (two floors, over 250 rooms) with simulated occupancy and sunlight. Agents can read occupancy and light level in rooms values through sensors. Agents can also actuate light switches located in the rooms.

We managed to solicit six submissions in the form of paper+video+code and subjected those submissions to the following process.

2. Screening Process

We followed a three-stage reviewing process:

[1] https://www.dagstuhl.de/en/program/calendar/semhp/?semnr=21072
[2] http://amundsen.com/talks/2021-02-dagstuhl/
[3] https://github.com/bold-benchmark

1. The organisers pre-screened the submissions whether they adhere to the hard minimum criteria (see section 2.1), and if they did, provided editorial feedback. All submissions passed this stage.
2. The organisers passed the submissions to a jury of senior scholars (see section 2.2) in the field, which reviewed the submissions and provided feedback on the content.
3. At the ISWC conference, the authors gave a brief presentation of their submission followed by question and answers. Afterwards, the judges determined the winner.

Overall, we applied following criteria and subjected the submissions to the scrutiny of a jury consisting of eminent scholars in the field.

2.1. Criteria

There are no established metrics for evaluating the performance of agents on the Semantic Web. The criteria we identified for the challenge are seperated into two groups: hard criteria to make sure the submissions are understandable and follow the Linked Data principles, and soft criteria, to be inclusive and foster creativity. Therefore, the committee evaluated submissions according to the following criteria:

- **Clarity (hard criterion)**
  As we aim to attract a diverse set of submissions, we expect an accompanying short paper and demonstrator video that describe the use-case and the agent-based system. We expect the submission to be clear and understandable to the evaluation committee.
- **Linked Data (hard criterion)**
  A hard constraint we apply is that we expect agents to operate on Linked Data on the Web, i. e. to make HTTP requests and to process RDF.
- **Degree of use of semantic technologies**
  The more a submission applies semantic technologies throughout its stack (e.g., to better cope with heterogeneous environments), the more virtual points we award.
- **Degree of dynamicity**
  During the course of the execution, an agent should exhibit adaptive behaviour (e.g., to cope with dynamic environments, if the environment is dynamic).
- **Degree of interaction and coordination**
  Next to single-agent submissions, we also appreciate agents with a form of social ability, i.e., agents that can interact and coordinate with other agents (e.g., direct message exchange, stigmergy, organisations, policies and norms, automated negotiation, interaction with people).

2.2. Jury

Our jury consisted in the following senior scholars:

- Jim Hendler, Rensselaer Polytechnic Institute, USA
- Jomi Hübner, Federal University of Santa Catarina, Brasil
- Simon Mayer, University of St Gallen, Switzerland
- Alessandro Ricci, University of Bologna, Italy
- Munindar Singh, North Carolina State University, USA

3. Challenge Submissions and Winner

The six submissions who passed the pre-screening process were:

- [9] – "Collaborative Home" by Ramparany, Trentin, Cumin, and Boissier.
- [10] – "Collaborative Route Finding in Semantic Mazes" by Beaumont, O'Neill, Bermeo, and Collier.
- [11] – "Crawl into the Dungeon with Hypermedia Agents" by Saffaf and Charpenay.
- [12] – "Integrated Planning and Execution on Read-Write Linked Data" by Aßfalg, Schneider, and Käfer.
- [13] – "Building Management using the Semantic Web and Hypermedia Agents" by O'Neill, Beaumont, Bermeo, and Collier.
- [14] – "WAT: Autonomous Hypermedia-driven Web Agents for Web of Things Devices" by Noura, Siegert, and Gaedke.

Extended versions of those submissions are compiled in this volume. A comparison of those submissions is available in Table 1. The table does not point a clear winner out by itself, which shows the importance of a collective decision made by a jury.

In terms of multi-agent platforms, the table shows a certain overlap:

- JaCaMo [16] and ASTRA [17], platforms that follow the Belief-Desire-Intention architecture, both appear in two submissions: [9, 11] and [10, 13] respectively. On top, submission [14] also uses part of the JaCaMo platform, but indirectly and at a conceptual level: although their work was not based on a pre-existing agent platform (they developed agents in Python), their multi-agent system followed the CArtAgO architectural framework [18], which is also part of JaCaMo.
- Linked-Data-Fu [19, 20], which can serve as a reactive agent platform, is used only by submission [12]. On top, two more submissions use it [11, 14], but merely as a Linked Data crawler.

Despite the overlap, these platforms were used in very different ways: the submitted systems integrate either reasoning, or planning, or learning capabilities, but no submissions uses two of these capabilities at the same time. The need for a software platform that can reliably support all capabilities at once may emerge, as Web agents find adoption.

In terms of environment platforms, the table shows a minor overlap: The two environment platforms prepared by the organizers (AMEE, BOLD) were used only once each, but environments used by other submissions share similarities with AMEE and BOLD: submission [11] relies on a maze environment similar to AMEE, while submission [9] offers functionalities that encompass those of BOLD. The main motivation for developing new environments was to provide more potential actions on the environment. The two other submissions developed a scenario in the manufacturing domain [12, 14].

From the six submissions, the jury determined the following winner:

> "WAT: Autonomous Hypermedia-driven Web Agents for Web of Things Devices"
> by Mahda Noura, Valentin Siegert, and Martin Gaedke [14].

We congratulate the winners.

Criterion		Submission					
		[9]	[10]	[11]	[12]	[13]	[14]
SEM	Semantic Web tech.	SPARQL	(RDF)	OWL	LDP	(RDF)	(RDF)
	№ data sources	4	1	1	1 + n	1	2
	Ontologies used	FOAF, SOSA	*custom*	*custom*	*custom*	Brick	TD, SOSA, SAREF
	Reasoning	-	-	✓	✓	-	✓
DYN	Dynamic sources	✓	-	-	✓	✓	✓
	Planning	-	-	-	✓	-	-
	Learning	-	✓	-	-	-	-
INT	Agent count	3	1 + n	1	1	3 + n	6
	Interaction protocol(s)	Request	Request	-	-	Subscribe	-
(Multi-)Agent platform		JaCaMo	ASTRA	JaCaMo	Linked-Data-Fu	ASTRA	*custom*
Environment platform		Home'In	AMEE	AMEE-like	JMonkey-LD	BOLD	IT'm Factory
Use case domain		Home auto.	Path finding	Path finding	Manufacturing	Building auto.	Manufacturing

Table 1: Summary of all submissions against the three soft criteria of the ATAC challenge (SEM: degree of use of semantic technologies, DYN: degree of dynamicity, INT: degree of interaction and coordination) as well as implementation details. Ontologies are considered not *custom* if there exists documented uses by third-parties. Agent interaction protocols are given as per the FIPA interaction protocol specification [15].

4. Conclusion

We were impressed by the diversity of the submissions that ATAC 2021 attracted and appreciate both the conceptual effort and the development effort of all authors to combine a diverse set of technologies and approaches at different maturity stages. We still see a gap to the original vision, where agents are able to operate in different and heterogeneous domains. The jury noted that challenges to address in possible further editions of ATAC were related to the scale of the environment, possibly geographically distributed, and the interaction with real-world objects.

We think that the first edition of ATAC showed promising steps in the right direction. Still, given the many challenges towards the vision of agents on the web, much remains to be done.

November 2021

Tobias Käfer
Andreas Harth
Andrei Ciortea
Victor Charpenay

Acknowledgements

We thank the members of our jury for serving on the panel and for the discussion on the future directions of the field. Furthermore, we thank Mike Amundsen for publishing the non-RDF version of AMEE.

References

[1] T. Berners-Lee, J. Hendler, O. Lassila, The Semantic Web, Scientific American (2001).
[2] J. A. Hendler, Agents and the Semantic Web, IEEE Intelligent Systems 16 (2001) 30–37.
[3] J. A. Hendler, Where are all the intelligent agents?, IEEE Intelligent Systems 22 (2007) 2–3.
[4] N. Shadbolt, T. Berners-Lee, W. Hall, The Semantic Web revisited, IEEE Intelligent Systems 21 (2006) 96–101.
[5] M. Kovatsch, R. Matsukura, M. Lagally, T. Kawaguchi, K. Toumura, K. Kajimoto, Web of Things (WoT) Architecture, Recommendation, W3C, 2020. Latest version available at https://www.w3.org/TR/wot-architecture10/.
[6] S. Speicher, J. Arwe, A. Malhotra, Linked Data Platform 1.0, Recommendation, W3C, 2015. Latest version available at https://www.w3.org/TR/ldp/.
[7] E. Mansour, A. V. Sambra, S. Hawke, M. Zereba, S. Capadisli, A. Ghanem, A. Aboulnaga, T. Berners-Lee, A demonstration of the Solid platform for social web applications, in: Proceedings of Demos at the 25th International Conference on World Wide Web (WWW), ACM, 2016, pp. 223–226.
[8] A. Ciortea, S. Mayer, F. Gandon, O. Boissier, A. Ricci, A. Zimmermann, A decade in hindsight: The missing bridge between multi-agent systems and the World Wide Web, in: Proceedings of the 18th International Conference on Autonomous Agents and Multi Agent Systems (AAMAS), 2019, p. 1659–1663.

[9] F. Ramparany, I. F. Trentin, J. Cumin, O. Boissier, Collaborative home, in: Proceedings of the All the Agents Challenge (ATAC 2021), 2021, pp. 7–12.
[10] K. Beaumont, E. O'Neill, N. V. Bermeo, R. Collier, Collaborative route finding in semantic mazes, in: Proceedings of the All the Agents Challenge (ATAC 2021), 2021, pp. 13–18.
[11] N. Saffaf, V. Charpenay, Crawl into the dungeon with hypermedia agents, in: Proceedings of the All the Agents Challenge (ATAC 2021), 2021, pp. 19–24.
[12] N. Aßfalg, H. Schneider, T. Käfer, Integrated planning and execution on read-write linked data, in: Proceedings of the All the Agents Challenge (ATAC 2021), 2021, pp. 25–31.
[13] E. O'Neill, K. Beaumont, N. V. Bermeo, R. Collier, Building management using the semantic web and hypermedia agents, in: Proceedings of the All the Agents Challenge (ATAC 2021), 2021, pp. 32–37.
[14] M. Noura, V. Siegert, M. Gaedke, Wat: Autonomous hypermedia-driven web agents for web of things devices, in: Proceedings of the All the Agents Challenge (ATAC 2021), 2021, pp. 38–43.
[15] S. Poslad, Specifying protocols for multi-agent systems interaction, ACM Trans. Auton. Adapt. Syst. 2 (2007) 15–es.
[16] O. Boissier, R. H. Bordini, J. F. Hübner, A. Ricci, Multi-Agent Oriented Programming – Programming Multi-Agent Systems Using JaCaMo, MIT Press, 2020.
[17] R. W. Collier, S. E. Russell, D. Lillis, Reflecting on agent programming with agentspeak(l), in: Proceedings of the 18th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA), Springer, 2015, pp. 351–366.
[18] A. Ricci, M. Viroli, A. Omicini, CArtAgO: A Framework for Prototyping Artifact-Based Environments in MAS, in: Proceedings of the 3rd International Workshop on Environments for Multi-Agent Systems (E4MAS 2006), Springer, 2007, pp. 67–86.
[19] S. Stadtmüller, S. Speiser, A. Harth, R. Studer, Data-fu: a language and an interpreter for interaction with read/write linked data, in: Proceedings of the 22nd International Conference on World Wide Web (WWW), 2013, pp. 1225–1236.
[20] T. Käfer, A. Harth, Rule-based programming of user agents for linked data, in: Proceedings of the 11th International Workshop on Linked Data on the Web (LDOW) at the 27th Web Conference (WWW), 2018.