A map without a legend: the semantic web and knowledge evolution

Jérôme Euzenat

► To cite this version:

Jérôme Euzenat. A map without a legend: the semantic web and knowledge evolution. Semantic Web – Interoperability, Usability, Applicability, 2020, 11 (1), pp.63-68. 10.3233/SW-190376 . hal-02984945

HAL Id: hal-02984945

<https://hal.science/hal-02984945>

Submitted on 1 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A map without a legend

The semantic web and knowledge evolution

Jérôme Euzenat

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France

E-mail: Jerome.Euzenat@inria.fr

Abstract. The current state of the semantic web is focused on data. This is a worthwhile progress in web content processing and interoperability. However, this does only marginally contribute to knowledge improvement and evolution. Understanding the world, and interpreting data, requires knowledge. Not knowledge cast in stone for ever, but knowledge that can seamlessly evolve; not knowledge from one single authority, but diverse knowledge sources which stimulate confrontation and robustness; not consistent knowledge at web scale, but local theories that can be combined. We discuss two different ways in which semantic web technologies can greatly contribute to the advancement of knowledge: semantic eScience and cultural knowledge evolution.

Keywords: semantic web, linked data, big data, open data, knowledge, knowledge representation, ontology, machine learning, web of data, reproducible research, eScience, cultural evolution

1. Life before the semantic web

Many animals are able to learn from their environment. This can be achieved through perceiving the environment and experimenting with it: acting and learning from the results of actions. Some of these animals learn from others. Most of them do it by mere imitations, such as birds learning songs [1]. However, it is known that some monkeys are able to signal danger and even that some other species can learn the signal from others [2].

The resulting behaviour shared among a population is called culture. Culture encompasses know-how, knowledge, languages, traditions, among other things, and does not have to be explicitly expressed.

Human beings are special among these species because they can express the knowledge¹ that they learn and communicate it to other human beings.

Bees can communicate how to reach food sources in an articulated language, but this stunning capability seems to be fixed. Direct knowledge communication is a very powerful mechanism because it leads to transmit knowledge without having to relearn it from experience. This capability grew slowly and led to teaching, monasteries, universities, and conferences.

Besides articulated expression, written expression and communication have allowed to get rid of time and space in knowledge transmission. Through tablets, rolls, books, libraries, or scholar journals it is possible to build on someone else knowledge without even meeting.

¹In this paper, knowledge is not opposed to belief, but to data, as customary in the field of knowledge representation. Thus, we are not focusing on the epistemological status of statements, but on their form. Knowledge is concerned with general statements; data with particular statements.

This led to a variety of ways to acquire knowledge: by observing and experimenting, by imitating, by communicating (talking, writing, reading). These key features have provided a selective advantage to the human species. This has been more eloquently storytold [3].

The worldwide web, facilitating cultural exchange, is a culminating point in this story, so far. One initial achievement of the web has been to make knowledge more easily and readily accessible across the planet. Knowledge was expressed in a not-very different way than before: through natural language texts and pictures, later through movies, interpretable by human beings. Those who have experienced the transition of the world before and after the web can only be grateful. The ease of use of the web pushed knowledgeable people all over the world to provide worldwide access to their knowledge. In this respect, wikipedia and stackoverflow in which millions of individuals care about articulating what they know for others, are just two but wonderful and precious successes.

Hence, a semantic web [4, 5], allowing machines to grasp this knowledge, is a tremendous idea.

2. The data advent: a little semantics went a long way... but may remain stuck here

The semantic web could be characterised by one of its early slogans: a web for machines. Not only for humans, but also for machines. The web was already processed by machine, but its content was aimed at being interpreted by human beings. More precisely it meant that the *content* of the web would be more machine processable.

However, there may be various understanding of ‘content’. It could be metadata: a set of attributes that help web search and classification. It could be precise data that is interpreted by specific programs, e.g. calendars or software dependencies. It could be relatively general knowledge representations of the content, such as causal relations or universal statements. Technologies developed by the W3C Semantic web activity covered these three aspects.

The semantic web did not immediately took off. This is not particularly surprising, but in our world in which everything has to change fast, this is considered a shortcoming. In particular, though enthusiastic scholars had been eager to develop sophisticated representations or ontologies, these were instantiated with little data. Hence of little *practical* use.

It was necessary to provide data to this semantic web because no one was interested by a bunch of theorem provers. Fortunately, by that time, the web was not short of data. Linked data [6], building on semantic web technologies, was welcome to benefit from these data and semantic web technologies. The linked data trend was spot on as it met other trends such as openness, big data, and later, data flow from the internet of things. Governments supported the process to connect data sources of administrations. Companies competed to acquire giant meshes of data on which they could run cleverly engineered pieces of software to learn what was there.

Moreover, data is the fuel for the knowledge mill. The availability of data, and computing power, unleashed the amazing capabilities of data analytics and machine learning from these data. This made some sense on top of raw data, that which can lead to impressive applications. Hence, it effectively contributes to make the web more intelligible to machines [7].

Nowadays, web users are not expected to provide knowledge, nor to access it. It seems that they are mere data provider, mostly through their actions, e.g. click, buy, like. These data are machine processable, but not open. They are kept secret, in silos, to the exclusive exploitation of a single organisation. They are processed by corporations which eventually learn knowledge from that data. But this knowledge, in turn, is not shared nor even prone to be communicated because not necessarily expressed in an articulated language. Instead, it is directly actioned. Hence, knowledge does not improve.

3. For knowledge!

After twenty years, the semantic web field is mostly focused on data, even when it is made of so-called knowledge graphs.

Of course, schemata and vocabularies are used. But, they are most of the time used for helping machines to parse data. Most of them simply express the structure and eventual relations between entities, but not what we know about them. Hence, the interpretation of data relies on the system after parsing. Knowledge is not made explicit nor shared; it is not available for scrutiny, reuse or evolution.

A semantic web without knowledge is for machines like a map without a legend for humans. The lack of a legend does not prevent to distinguish a road from a river, a bridge from a boat shuttle. But it does not help understanding how these can be used to travel. It does not enable the explorer to plan her trip, by reasoning on elevation or humidity.

Hence, this seems like a regression which brings us down the knowledge evolution ladder: humanity may eventually be fed with data, but any knowledge will have to be learnt again and again.

This is why the grand goal of formally expressing knowledge on the web must be rehabilitated. One ambition of the semantic web was to make knowledge available to machines, so that at least they can help us work with this knowledge, not only to find it out but to process it, to improve it and to communicate it to the world [7], as the next knowledge medium [8]. This could have led to a further knowledge ecosystem in which knowledge is elaborated while it is used for providing services. Like the web can be seen as a data washing machine [9], all the knowledge of the web, may be used to clean up and to improve knowledge. Reconsidering knowledge sharing and evolution at web scale would empower humans and machines alike with knowledge that can be both used to provide services and jointly refined and elaborated.

The goal is not to build the all-encompassing ontology with which everyone will agree. Knowledge does not have to be centralised: diversity is source of disputation and robustness. Knowledge is not cast in stone for ever, but it has to seamlessly evolve. Knowledge need not be consistent at web scale, it can come in local theories that can be combined.

The semantic web initiative has already provided a good basis for expressing and sharing knowledge. First, universal and well-defined languages such as RDF, RDFS, OWL, or SPARQL have been designed and reliable implementations exist. Then many specialised ontologies using these technologies have been developed from the most abstract, e.g. work, provenance, to the more concrete, e.g. proteins, scientific articles.

In the remainder, two different aspects are considered: how explicitly sharing formalised knowledge contributes to improve scientific practices, hence our shared knowledge; how knowledge, shared or not, may be seamlessly evolved and how this can be studied in an effective way. Of course, these directions are not without connection and not the only possible ones. However, they are in a totally different position with respect to the current state of the semantic web.

4. eScience: an example of web of knowledge

Let us take a typical example of collaborative knowledge elaboration, sharing and evolution: scientific research. eScience is a paramount example of knowledge expression and evolution.

Science has not been left aside of the professional computerisation: it has been very active from this standpoint. Statistical packages, plotting and editing notebooks are now commonplace and *in silico*

simulations are accepted. This facilitates the production, analysis and dissemination of knowledge by and for humans helped by machines. Meanwhile, little is made for helping machines to have a grasp at the content.

Semantic eScience is a further step in that direction [10]. It exploits semantic web technologies to provide an interoperable and machine-interpretable infrastructure for scientific enquiry. Within the past 10 years, it has been a continuing source of attention.

Necessary technologies are already in place for metadata and identification: IRIs provide a natural way to identify entities such as researchers (ORCID) or documents (DOI). Support has been provided to express bibliographic data in RDF [11]. This allows expressing metadata about scientific literature.

On the content side, i.e. the objects of scientific statements, there are already many resources to express these with semantic web technologies. In life and health sciences, Gene ontology, OBO Foundry, biportal, and Bio2RDF were already there 10 years ago and continuously improved [12]. On the mathematical side, efforts have been made to provide ways to offer a semantic encoding [13].

It is also possible to deal with the methodological aspect of research. Some work have proposed the expression of hypotheses [14]. On the experimental side: the researchobject project has provided a way to describe protocols [15]. Attempts are made to describe evaluation methods [16]. Finally, efforts have been made to address open science requirements to publish data sets. Google dataset search [17] offers search among data sets described with schema.org dataset and DCAT. Relations between such elements—which process produced which data, what hypothesis is it supposed to support, in which paper has it been published—can be kept track of through provenance assertions [18].

All this allows for expressing scientific knowledge on the web, not only the results of scientific enquiry but the whole process that led to establish such results. Hence, ingredients for semantic scientific research are available. What is missing is that scientific practices embrace semantic web technologies. This can come from more integration in work habits and new applications providing benefits. Here are a few example of these.

Scientific knowledge expressed with semantic web technologies would clearly facilitate searching results [19]. However, more can be obtained from it with the help of machines. Knowledge expressed formally on the web could lead to formal scientific collaboratories [20]. For instance, data analysis workflows can be exploited for continuous reevaluation of hypotheses by updating data analyses when new data is available [21]. It may be checked, beforehand, that an experiment is prone to refute a claim. Papers whose results contradict one another could be identified. On the mathematical side, proof checking, now performed at the scale of one theory may be attempted at larger scales. Experiments testing a particular hypothesis, or a more specific one, can be found to avoid duplication. A protocol may be analysed to pinpoint what should be changed to affect the results of experiments. Literature can be exploited to predict adverse effects [22]. More simply, machine learning could help cleaning and mining result as well as suggesting interesting tracks. These applications require knowledge and cannot rely on data alone.

In addition to contributing to open research, semantic expression of research processes may help addressing reproducibility issues. Providing accurate descriptions of computer-based experiments can allow a computer to reproduce them. Tools may be developed to help (debugging and) peer-reviewing experiment preregistrations. They may also help to identify missing information in descriptions and so even facilitate off-line reproducibility.

This is where full knowledge can play its role, by enabling machines to help us improving our knowledge by confronting it to existing data, by finding contradiction with other pieces of knowledge, by learning from data and knowledge. Machines could then join the course of knowledge evolution adding value to knowledge and massively confronting it.

5. Evolving knowledge

Knowledge, whether expressed on the web or not, whether generated through eScience or not, must evolve. Knowledge that does not evolve, and systems based on it, are at risk of obsolescence. If the semantic web is to be more than a parenthesis in supporting the knowledge of humanity, efforts should be made for ensuring its smooth and continuous evolution.

Semantic eScience can contribute to this, but is not sufficient. Human knowledge evolves independently from scientific research, which is a relatively recent way to deal with knowledge. Contrary to science, people do not necessarily elaborate knowledge for itself, but evolve it through its use. However, both processes contribute to ‘improve’ knowledge and evolutionary interpretations of science are not new [23].

Natural selection can be thought of as a simple control mechanism based on variation, selection and transmission [24]. This can be implemented in computers as had already been done for genetic programming [25]. Concerning human knowledge, anthropology identified and studied cultural evolution [26]. Thus one way of addressing the problem of knowledge evolution is through understanding how cultural evolution techniques may be adapted to the semantic web.

Experimental cultural evolution has been successfully applied to natural language [27]. It is perfectly reasonable to apply experimental cultural evolution to knowledge [28]. Knowledge generates individual behaviour, which is subject to selective pressure from the environment and thus spreads differentially. Knowledge evolution can indeed be implemented as a mechanism that makes knowledge evolve seamlessly while it is used.

Experiments have been performed with knowledge-carrying agents adapting minimally their knowledge when it reveals inadequate. The behaviour of such agents becomes more coherent as more interactions are performed. Like in cultural evolution, knowledge transmission does not necessarily happen through inheritance. On the contrary, agents may transmit knowledge by cooperating or by directly exchanging it. Such minimal distributed social mechanisms for evolving and shaping knowledge are desirable features for social machines [29].

Experimental cultural evolution has been adapted to evolve abstract cultures [30], natural language features [27], and, closer to the semantic web, ontology alignments. In particular, it can be used to repair alignments better than blind logical repair [31], to create alignments based on entity descriptions [32], to learn alignments from dialogues framed in interaction protocols [33, 34], or to correct alignments until no error remains and to start with no alignment [35, 36]. Each study provides new insights and opens directions on the effects of local reactive adaptation on the resulting knowledge.

There are two important challenges to provide evolvability to the semantic web knowledge: understanding cultural knowledge evolution and embedding it within the semantic web. Contrary to the situation described for eScience: the task is not to merely use and spread semantic web technologies, but to properly extend them.

Cultural knowledge evolution principles may be studied by following the path of experimental cultural evolution presented above or developing more theoretical approaches. Many challenging questions are yet unanswered. For instance, the formation and evolution of common knowledge, i.e. proper culture, must be studied. Another relevant issue is the co-evolution of knowledge learnt from the environment and knowledge acquired through communication.

Concerning web integration, we are back in the read-write web debate: as long as agents are reduced to knowledge consumers, only top-down evolution is possible; as soon as they become proper learners,

i.e. able to evolve and communicate the knowledge they learn, a more lively evolution is possible. This requires these agents, human or software applications, to expose and adapt the knowledge they use.

6. Conclusions

The semantic web endeavour has so far provided impressive outcomes in terms of linked data, which can be interpreted by machines, as well as representation languages and ontologies. This is only the beginning of the journey towards a semantic web contributing to the knowledge of humanity. A solid basis is available that should be pushed further.

Human beings shape their culture and knowledge progressively through evolution. The current emphasis of the semantic web towards data undermines this enterprise. Knowledge learnt from data is not made explicit nor communicated. Hence, it cannot properly evolve, but has to be relearnt.

For the semantic web to take its full part in knowledge advancement, it has to be complemented by explicit knowledge expression and sharing. This would unleash the capability to properly evolve knowledge as illustrated by work on two directions: (a) scientific knowledge elaboration processes may be improved by expressing them semantically, and (b) knowledge on the web may be evolved smoothly through evolutionary techniques.

References

- [1] R. Dawkins, *The selfish gene*, Oxford University Press, 1976.
- [2] M. Hauser, *The evolution of communication*, The MIT Press, 1997.
- [3] Y.N. Harari, *Sapiens: a brief history of humankind*, Penguin Random House, 2011.
- [4] T. Berners-Lee, What the semantic web can represent, *Design issues*, W3C, 1998. <https://www.w3.org/DesignIssues/RDFnot.html>.
- [5] T. Berners-Lee, J. Hendler and O. Lassila, The Semantic Web, *Scientific American* **284**(5) (2001), 34–43.
- [6] T. Berners-Lee, Linked data, *Design issues*, W3C, 2006. <https://www.w3.org/DesignIssues/LinkedData.html>.
- [7] K. Janowicz, F. van Harmelen, J. Hendler and P. Hitzler, Why the data train needs semantic rails, *AI Magazine* **36**(1) (2015), 5–14, ISSN 0738-4602.
- [8] M. Stefik, The next knowledge medium, *AI Magazine* **7**(1) (1986), 34–46, ISSN 0738-4602.
- [9] S. Auer and J. Lehmann, Creating knowledge out of interlinked data, *Semantic Web* **1**(1–2) (2010), 97–104. doi:10.3233/SW-2010-0019.
- [10] B. Brodaric and M. Gahegan, Ontology use for semantic e-Science, *Semantic Web* **1**(1–2) (2010), 149–153. doi:10.3233/SW-2010-0013.
- [11] P. Ciccarese, D. Shotton, S. Peroni and T. Clark, CiTO + SWAN: The web semantics of bibliographic records, citations, evidence and discourse relationships, *Semantic Web* **5**(4) (2014), 295–311. doi:10.3233/SW-130098.
- [12] M. Dumontier, Building an effective Semantic Web for health care and the life sciences, *Semantic Web* **1**(1–2) (2010), 131–135. doi:10.3233/SW-2010-0018.
- [13] M. Kohlhase and F. Rabe, Semantics of OpenMath and MathML3, *Mathematics in Computer Science* **6**(3) (2012), 235–260. doi:10.1007/s11786-012-0113-x.
- [14] D. Garijo, Y. Gil and V. Ratnakar, The DISK Hypothesis Ontology: Capturing Hypothesis Evolution for Automated Discovery, in: *Proceedings of Workshops and Tutorials of the 9th International Conference on Knowledge Capture (K-CAP2017)*, Austin, Texas, USA, December 4th, 2017, I. Tiddi, G. Rizzo and . Corcho, eds, CEUR Workshop Proceedings, Vol. 2065, CEUR-WS.org, 2017, pp. 40–46. <http://ceur-ws.org/Vol-2065/paper08.pdf>.
- [15] K. Hettne, H. Dharuri, J. Zhao, K. Wolstencroft, K. Belhajjame, S. Soiland-Reyes, E. Mina, M. Thompson, D. Cruickshank, L. Verdes Montenegro, J. Garrido, D. de Roure, O. Corcho, G. Klyne, R. van Schouwen, P. t Hoen, S. Bechhofer, C. Goble and M. Roos, Structuring research methods and data with the research object model: genomics workflows as a case study, *Journal of biomedical semantics* **5**(1) (2014), 41, ISSN 2041-1480. doi:10.1186/2041-1480-5-41.

- [16] M. Stocker, M. Prinz, F. Rostami and T. Kempf, Towards research infrastructures that curate scientific information: A use case in life sciences, in: *Data Integration in the Life Sciences - 13th International Conference, DILS 2018, Hannover, Germany, November 20-21, 2018, Proceedings*, S. Auer and M. Vidal, eds, Lecture Notes in Computer Science, Vol. 11371, Springer, 2018, pp. 61–74. ISBN 978-3-030-06015-2. doi:10.1007/978-3-030-06016-9_6.
- [17] N. Noy, Making it easier to discover datasets, *Google blog*, 2018. <https://www.blog.google/products/search/making-it-easier-discover-datasets/>.
- [18] T. Lebo, S. Sahoo and D. McGuinness (eds.), PROV-O: The PROV Ontology, Recommendation, W3C, 2013. <https://www.w3.org/TR/prov-o/>.
- [19] W. Hu, H. Qiu, J. Huang and M. Dumontier, BioSearch: a semantic search engine for Bio2RDF, *Database* **2017** (2017), 059. doi:10.1093/database/bax059.
- [20] J. Euzenat, Building consensual knowledge bases: context and architecture, in: *Towards very large knowledge bases*, N. Mars, ed., IOS press, Amsterdam (NL), 1995, pp. 143–155.
- [21] Y. Gil, D. Garijo, V. Ratnakar, R. Mayani, R. Adusumilli, H. Boyce, A. Srivastava and P. Mallick, Towards Continuous Scientific Data Analysis and Hypothesis Evolution, in: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, S.P. Singh and S. Markovitch, eds, AAAI Press, 2017, pp. 4406–4414.
- [22] A. Noor, A. Assiri, S. Ayvaz, C. Clark and M. Dumontier, Drug-drug interaction discovery and demystification using Semantic Web technologies, *Journal of the American Medical Informatics Association* **24**(3) (2017), 556–564. doi:10.1093/jamia/ocw128.
- [23] D. Hull, *Science as a process: an evolutionary account of the social and conceptual development of science*, University of Chicago Press, 1988.
- [24] C. Darwin, *On the origin of species by means of natural selection, or the preservation of favoured races in the struggle for life*, John Murray, 1859, 6th edition: 1872.
- [25] A. Eiben and J. Smith, From evolutionary computation to the evolution of things, *Nature* **521** (2015), 476–482. doi:10.1038/nature14544.
- [26] A. Mesoudi, A. Whiten and K. Laland, Towards a unified science of cultural evolution, *Behavioral and brain sciences* **29**(4) (2006), 329–383. doi:10.1017/S0140525X06009083.
- [27] L. Steels (ed.), *Experiments in cultural language evolution*, John Benjamins, 2012.
- [28] J. Euzenat, Replicator-interactor in experimental cultural knowledge evolution, in: *Proceedings of the 2nd JOWO workshop on Interaction-Based Knowledge Sharing (WINKS), Graz (AT), K. Adrian, J. Euzenat, D. Gromann, E. Jiménez-Ruiz, M. Schorlemmer and V. Tamma, eds, CEUR Workshop Proceedings, Vol. 2518, CEUR-WS.org, 2019. <http://ceur-ws.org/Vol-2518/paper-WINKS2.pdf>.*
- [29] J. Hendler and T. Berners-Lee, From the Semantic Web to social machines: A research challenge for AI on the World Wide Web, *Artificial intelligence* **174**(2) (2010), 156–161. doi:10.1016/j.artint.2009.11.010.
- [30] R. Axelrod, The dissemination of culture: a model with local convergence and global polarization, *Journal of conflict resolution* **41**(2) (1997), 203–226. doi:10.1177/0022002797041002001.
- [31] J. Euzenat, First experiments in cultural alignment repair (extended version), in: *The Semantic Web: ESWC 2014 Satellite Events - ESWC 2014 Satellite Events, Anissaras, Crete, Greece, May 25-29, 2014, Revised Selected Papers*, V. Presutti, E. Blomqvist, R. Troncy, H. Sack, I. Papadakis and A. Tordai, eds, Lecture Notes in Computer Science, Vol. 8798, Springer, 2014, pp. 115–130. doi:10.1007/978-3-319-11955-7_10.
- [32] M. Anslow and M. Rovatsos, Aligning experientially grounded ontologies using language games, in: *Graph Structures for Knowledge Representation and Reasoning - 4th International Workshop, GKR 2015, Buenos Aires, Argentina, July 25, 2015, Revised Selected Papers*, M. Croitoru, P. Marquis, S. Rudolph and G. Stapleton, eds, Lecture Notes in Computer Science, Vol. 9501, Springer, 2015, pp. 15–31. ISBN 978-3-319-28701-0. doi:10.1007/978-3-319-28702-7_2.
- [33] M. Atencia and M. Schorlemmer, An interaction-based approach to semantic alignment, *Journal of Web Semantics* **13**(1) (2012), 131–147. doi:10.1016/j.websem.2011.12.001.
- [34] P. Chocron and M. Schorlemmer, Vocabulary alignment in openly specified interactions, in: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*, K. Larson, M. Winikoff, S. Das and E.H. Durfee, eds, ACM, 2017, pp. 1064–1072.
- [35] P. Chocron and M. Schorlemmer, Attuning ontology alignments to semantically heterogeneous multi-agent interactions, in: *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, G.A. Kaminka, M. Fox, P. Bouquet, E. Hüllermeier, V. Dignum, F. Dignum and F. van Harmelen, eds, Frontiers in Artificial Intelligence and Applications, Vol. 285, IOS Press, 2016, pp. 871–879. ISBN 978-1-61499-671-2. doi:10.3233/978-1-61499-672-9-871.
- [36] J. Euzenat, Interaction-based ontology alignment repair with expansion and relaxation, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, C. Sierra, ed., ijcai.org, 2017, pp. 185–191. ISBN 978-0-9992411-0-3. doi:10.24963/ijcai.2017/27.