Study Protocol for Combining Machine Learning and Semantic Web - A Systematic Mapping Study

ANNA BREIT, Semantic Web Company, Austria
LAURA WALTERSDORFER, TU Wien, Austria
FAJAR J. EKAPUTRA, TU Wien, Austria
MARTA SABOU, TU Wien, Austria
ANDREAS EKELHART, SBA Research, Austria
ANDREAA IANA, University of Mannheim, Germany
HEIKO PAULHEIM, University of Mannheim, Germany
JAN PORTISCH, University of Mannheim, Germany
ARTEM REVENKO, Semantic Web Company, Austria
ANNETE TEN TEIJE, VU Amsterdam, Netherlands
FRANK VAN HARMELEN, VU Amsterdam, Netherlands

Semantic Web Machine Learning Systems (SWeMLS) are the result of a cooperation of a Semantic Web (SW) knowledge structure and a Machine Learning (ML) component. The knowledge structure should fulfill the following criteria: i) represent a domain model, ii) rely on some formal logical foundations (e.g., Description Logic) , iii) be encoded in a W3C recommended language (RDF-s, OWL). This might include vocabularies, taxonomies, knowledge graphs, ontologies, and other structures based on linked data. The machine learning component may consist of any kind of machine learning model, including rule learning systems, traditional ML, as well as more recent deep learning models. SWeMLS are systems with a readily available implementation. These implementations might be of different maturity, ranging from prototypes to enterprise ready systems. This study protocol provides details about the motivation and methodology, such as research questions, search queries and literature selection procedures.

CCS Concepts: • **Information systems** → **Semantic web description languages**; • **Computing methodologies** → **Knowledge representation and reasoning**; **Machine learning**.

Additional Key Words and Phrases: semantic web, knowledge graph, machine learning, knowledge representation and reasoning, artificial intelligence, hybrid AI, neuro-symbolic integration

Authors' addresses: Anna Breit, anna.breit@semantic-web.com, Semantic Web Company, Austria; Laura Waltersdorfer, laura.waltersdorfer@tuwien.ac.at, TU Wien, Austria; Fajar J. Ekaputra, fajar.ekaputra@tuwien.ac.at, TU Wien, Austria; Marta Sabou, marta.sabou@ifs.tuwien.ac.at, TU Wien, Austria; Andreas Ekelhart, aekelhart@sba-research.org, SBA Research, Austria; Andreaa Iana, andreea@informatik.uni-mannheim.de, University of Mannheim, Germany; Heiko Paulheim, heiko@informatik.uni-mannheim.de, University of Mannheim, Germany; Jan Portisch, jan@informatik.uni-mannheim.de, University of Mannheim, Germany; Artem Revenko, artem.revenko@semantic-web.com, Semantic Web Company, Austria; Annete ten Teije, annette.ten.teije@vu.nl, VU Amsterdam, Netherlands; Frank van Harmelen, frank.van.harmelen@vu.nl, VU Amsterdam, Netherlands.

0360-0300/2021/8-ART111 $15.00
https://doi.org/10.1145/1122445.1122456

ACM Reference Format:
Anna Breit, Laura Waltersdorfer, Fajar J. Ekaputra, Marta Sabou, Andreas Ekelhart, Andreaa Iana, Heiko Paulheim, Jan Portisch, Artem Revenko, Annete ten Teije, and Frank van Harmelen. 2021. Study Protocol for Combining Machine Learning and Semantic Web - A Systematic Mapping Study. *ACM Comput. Surv.* 37, 4, Article 111 (August 2021), 10 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Recent years have seen an increased interest and amazing developments in both data-driven, sub-symbolic machine learning techniques (Deep Learning etc) and techniques that make use of explicit, symbolic knowledge representations. Knowledge representation and reasoning has been classically researched in the Semantic Web community, starting from early efforts like Cyc to modern large-scale knowledge graphs like Wikidata, and including all flavours of reasoning. Indeed, arguably, knowledge representation techniques have been popularised by the Semantic Web research community in the last two decades (ontologies, thesauri, Linked Data sets) leading to a great interest in and uptake of knowledge graphs. At the same time, exciting developments happen at the cross-roads of these technology areas in terms of emerging neuro-symbolic AI systems that combine technology elements from both fields. These are now broadly applied in a range of application areas. To keep the scope of the study manageable, we focus on a subset of neuro-symbolic AI systems, those that rely on Semantic Web resources and Machine Learning components (SWeML).

SWeMLS are of interest not only the core research areas involved, but also tertiary areas that adapt such systems to address/solve problems in their domains. Firstly, researchers from the areas of Semantic Web and Machine learning (but also from broader fields of symbolic and subsymbolic AI) want to gain an overview on combination methods and patterns, and get an insight on the ingredients of SWeMLS, get an overview of the main trends and underrepresented topics that are indication of future work. Similarly, researchers from domains adopting SWeMLS would benefit from an overview of the trends in the field in terms of tasks that have been solved as well as in engineering aspects of these systems (which datasets and methods are used for a certain task and how are they combined?) Finally, authors of SWeML would benefit from a structured way to describe their system and its key characteristics. Newcomers, would benefit from a structured way of intepreting such systems.

1.1 Goals and expected results

Targeted research in the direction of Semantic Web Machine Learning System is quite limited. However, existing works show that a wide variety of research efforts have already been made that would fall under the term of SWeML. To this end, we are aiming to collect different existing systems in order to give an overview over the research field. More precisely, we will provide

- insight on how, when, and in which areas SWeMLS are currently used, as well as an overview about the development and the maturity of the field as a whole.
- a typology of SWeMLS with regard to the characteristics of their ML and SW module, as well as their patterns of connection.

2 METHODOLOGY

As our goal is to gain an overview of existing research efforts that would fall under the term of Semantic Web Machine Learning Systems, conducting a **Systematic Mapping Study (SMS)** [15] forms an appropriate approach in order to structure this broad research area. The Systematic Mapping Study consists of three consecutive phases, i.e., (1) *Study Planning* phase, (2) *Study Execution* phase, and (3) *Analysis and Reporting* phase . The first phase of the study, planning focuses on *scoping the study* and, accordingly, *proposing the methodology* for each step of the study as documented in the *Study Protocol.* Study scoping includes positioning the planned work in the

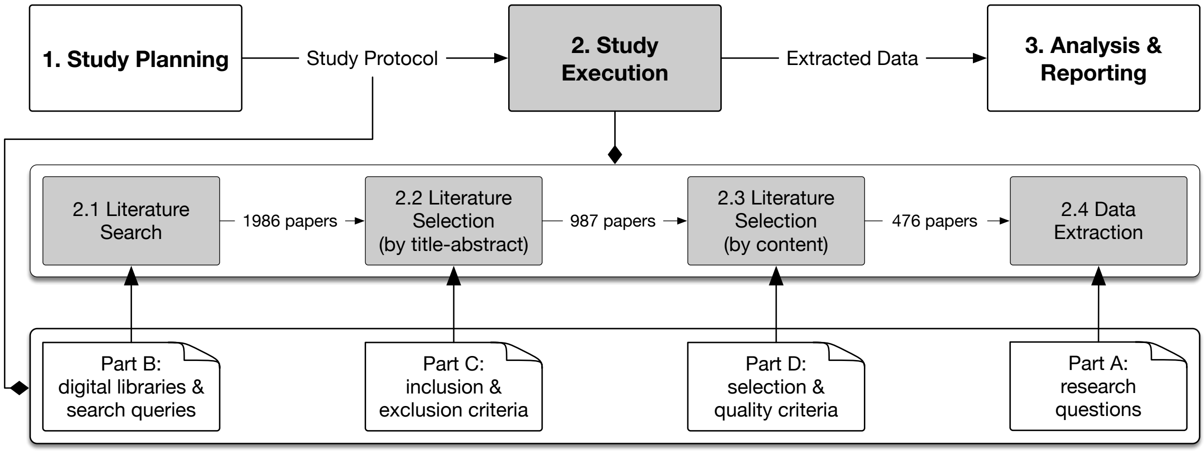

Fig. 1. Overview of the Semantic Mapping Study Process

context of related research areas and related work in terms of similar surveys. This forms a basis for deriving pertinent Research Questions (Part A in cf. Figure 1, Section 2.1), which are then translated into appropriate Search Queries (Part B, Section 2.1), a number of paper selection criteria (Part C & D in Figure 1, Section 2.3.1) used to identify relevant papers and a Data Extraction form (Part E in cf. Figure 1) that facilitates the objective and unbiased extraction of data. The methodological details captured in the Study Protocol aim to make the study process transparent and reproducible.

In the following sections, we will present details concerning study execution (cf. Figure 1). *Research Questions* (Section 2.1), *Literature Search* (Section 2.2) and *Literature Selection* (Section 2.3), and *Data Extraction* (Section 2.4) procedures will be specified.

2.1 Research Queries

As existing surveys targeting Semantic Web Machine Learning Systems are limited, a wide variety of questions regarding these systems have not been yet answered. In the following, we will present the selection of research questions which we answer in the course of this systematic survey:

RQ1 Bibliographic characteristics of the System. How is the geographic and temporal distribution?
- a. How is the geographic and temporal distribution?
- b. How are the systems positioned, which keywords are used to describe them?

RQ2 Patterns of Connection. What is the processing flow of the systems in terms of inputs/outputs and the order or processing units?

RQ3 Application Area. What is the application area of the systems?
- a. What kind of tasks are solved? (e.g., text analysis, information extraction)
- b. In which domains are SWeML systems applied? (e.g., natural sciences, general)

RQ4 Characteristics of the ML Module. What are the characteristics of the machine learning model(s) used in the SWeML system?
- a. What training type(s) are used in the model(s)? (e.g., unsupervised, supervised, self-supervised)
- b. What ML model category can be identified? (e.g., decision trees, convolutional neural network)
- c. Which ML components can be identified? (e.g., attention, convolutions)
- d. How can ML model types be classified? (e.g., classical ML, Deep Learning)

RQ5 Characteristics of the SW Module. What are the characteristics of the Semantic Web knowledge structure used in the SWeML system?
  a. What type of semantic web structure is used? (e.g., taxonomy, ontology)
  b. What is the degree of semantic exploitation? (e.g., only labels are used, hierarchical relations are used)
  c. What is the size of the resource? (number of triples)
  d. In what formalism is the resource represented? (e.g., RDF-S, OWL)
  e. Does the system integrate semantic processing modules? (e.g., reasoner, ontology matcher)

RQ6 Maturity, Transparency and Auditability
  a. What is the level of maturity of the systems (prototype, beta, stable)?
  b. How transparent are the systems? (i.e., whether software, infrastructure and details of evaluation setup are explained in the papers)
  c. For systems with explicit provenance mechanisms, what are the goals of having such mechanism? (e.g., explainability, validation) What specific context information is captured to support the goals? (e.g., user consent, stakeholder involvements)

While answering RQ1, RQ5 and RQ6 will provide a better understanding of the how and where SWeML systems are currently used, as well as an overview about the development and the maturity of the field as a whole, RQ2, RQ3 and RQ4 will provide a basis of different forms SWeML systems can take.

2.2 Literature Search

2.2.1 Digital Libraries. For our systematic mapping study, we selected the following resources as basis for our query-based search:

- WebOfScience[1]
- ACM Digital Library [2]
- IEEE Xplore[3]
- Scopus[4]

Our goal was to retrieve important conferences and journals, the aforementioned which are also referred to as good sources for software engineering according to [5, 15].

2.2.2 Search Strings. The search query was derived from the study research questions and iteratively refined to obtain a high number of relevant papers while keeping the number of retrieved papers manageable. The query consists of three sub-queries targeting the *SW module* (Q1), the *ML module* (Q2), and the *system* aspect (Q3) of SWeMLS, respectively. The search strings contained in these queries are presented in Table 1. Each sub-query consists of a union of the listed search terms, the final query used for study search is an intersection of the three sub-queries.

Q1 - SW Module. Q1 keywords include (1) the name of the overall field of interest (semantic web); (2) the most frequent terms to refer to semantic structures (ontolog*, linked data, knowledge graph) as well as (3) a number of W3C standard names that are likely to be used when implementing semantic web structures. During the keyword selection a number of keywords were tested but not included in the final query as they all lead to very large result sets containing a lot of false positive hits. These include: semantic, semantic model, vocabulary.

Q2 - ML Module. The field of machine learning is extremely wide, which is why the task of finding the keywords that most accurately describe the ML module of SWeMLS is highly challenging. To overcome this

[1]http://www.webofknowledge.com/
[2]https://dl.acm.org/
[3]https://ieeexplore.ieee.org/
[4]https://www.scopus.com/

Table 1. Search strings used in the search queries.

Sub-Query	Used Search Strings
Q1 (SW module)	knowledge graph, linked data, semantic web, ontolog*, RDF, OWL, SPARQL, SHACL
Q2 (ML module)	deep learning, neural network, embedding, representation learning, feature learning, language model, language representation model, rule mining, rule learning, rule induction, genetic programming, genetic algorithm, kernel method
Q3 (system)	Natural Language Processing, Computer Vision, Information Retrieval, Data Mining, Information integration, Knowledge management, Pattern recognition, Speech recognition

problem and to avoid bias in selecting, an intersection of keywords present in multiple sources was chosen. These sources were the ACM Computing Classification System[5], where we considered all concepts narrower than *CCS→Computing methodologies→Artificial intelligence* and *CCS→Computing methodologies→Machine learning*, topics extracted by *Microsoft Academic*[6] that were child topics to *Machine Learning*[7] or *Artificial Intelligence*[8], and subcategories and pages of the Wikipedia categories *Artificial Intelligence*[9] and *Machine Learning*[10]. A keyword was considered if it appeared in all three resources.

To reduce the amount of keywords considered and to increase their quality, we removed: (i) terms that described a specific approach or model such as "BERT" or "PCA"; (ii) terms that were too unspecific and did not show a clear connection to ML or AI without knowing their broader terms (e.g., "visual inspection"); (iii) terms whose substrings were already included in the search query, for example, "deep artificial neural network" would be discarded if "neural network" is included. The final keywords for Q2 are shown in Table 1.

Q3 - System. A separate query is introduced to assure that retrieved papers present systems aiming to solve specific tasks. This query focuses on application fields to avoid bias by including specific tasks as search terms. Herefore, an intersection of all relevant children of (1) ACM Computing Classification System[5] concepts *CCS→Computing methodologies→Artificial intelligence*, *CCS→Computing methodologies→ Machine learning*, and *CCS→Information systems*, (2) *Microsoft Academic*[6] Topics *Machine Learning*[7] and *Artificial Intelligence*[8], and (3) Wikipedia categories *Artificial Intelligence*[9] and *Machine Learning*[10] were taken into account, where children were considered 'relevant' if they represent an application area. The final keywords for Q3 can be found in Table 1.

Since query results from different digital libraries often contain duplicates, we performed multiple rounds of de-duplication both automatically and manually on the result set.

2.3 Literature Selection

After conducting the search based on the keywords and resources defined in Section 2.2, retrieved papers need to be filtered as the keyword-based search will generate several false positive regarding the relevance for this SMS. In order to conduct this filtering transparently, a set of study selection and quality criteria was defined. Based on

[5]https://dl.acm.org/ccs, visited Oct 2nd 2020
[6]https://academic.microsoft.com/topics
[7]https://academic.microsoft.com/topics/41008148,119857082, visited Oct 2nd 2020
[8]https://academic.microsoft.com/topics/41008148,154945302, visited Oct 2nd 2020
[9]https://en.wikipedia.org/wiki/Category:Artificial_intelligence, visited Oct 2nd 2020
[10]https://en.wikipedia.org/wiki/Category:Machine_learning, visited Oct 2nd 2020

Table 2. Inclusion and Exclusion Criteria for Study Selection

Criteria		Inclusion Criteria (IC)	Exclusion Criteria (EC)
C1	Publication Date	Papers published between 2010 and 2020.	Papers published before 2010 or after 2020.
C2	Language	Papers written in English.	Papers written in a language other than English.
C3	Type of Publication	Primary studies subject to peer review including journal papers, papers as part of conference or workshop proceedings, book chapters.	Non-peer-reviewed papers such as technical reports, theses, books, abstracts, presentations, tutorials, guidelines or summaries of conferences/editorials. Also, secondary studies such as systematic literature reviews or mapping studies, and surveys.
C4	Accessibility	Papers, which can be accessed from a major technical university (TU Wien) without additional paywall.	Papers, which cannot not be accessed from a major technical university (TU Wien) without additional paywall.
C5	Duplicates	If multiple publications of the same study exist presenting the same analysis, the latest version (i.e., most complete study report) will be included.	Studies for which a newer / more complete version exists.
C6	SW and ML Interconnection	The SW and ML modules do interact and are both used to tackle the provided task, e.g. (but not limited to), SW knowledge is input to ML model, or ML model produces SW knowledge.	The SW and ML modules do not interact and do not aim at the same task.
C7	System	Papers that present an implemented system that is used to solve a specific task.	Papers only providing theoretical systems, or presenting an implementation without evaluating the system on a specific task.
C8	Quality	(C8a) The use of English is sufficiently proficient to allow understanding the details of the system (C8b) Good scientific quality, all key information present.	(C8a) English language issues hamper understanding paper contributions and system details, (C8b) Low scientific quality (e.g., shallow analysis, key information missing).

these criteria we selected applicable studies in multiple rounds. First, reviewers focused only on metadata such as title, abstract, keywords, publication venue type and year. In a second and third round entire paper contents were analyzed to decide on inclusion or exclusion.

2.3.1 Study Selection Criteria. For each criterion, an *inclusion criterion (IC)*, and a complementary *exclusion criterion (EC)* is given. This improves the specificity of the criteria. Inclusion criteria (IC) 1-5 concern metadata of the publications, such as publication data (C1), language (C2), publication type (C3), accessibility (C4) and duplicates (C5). C6 and C7 refer to our to our SWeMLS definition whether described systems have an interconnection between SW and ML component (C6), and whether the system solves a task (C7). In later selection rounds, two quality criteria were added which could lead to the rejection of papers: 1) English language issues making it difficult to understand contributions and overall workings of described systems, 2) scientific quality in case important depth in analysis was missing or other missing key information.

2.4 Data Extraction

After completion of literature selection, data extraction was conducted with the help of a data extraction form.

2.4.1 Data Extraction Form. The preparation of a data extraction form prior to study analysis is done to reduce researcher bias and allow multiple researchers to extract data objectively [15]. The data extraction form is designed to define how and which data is collected regarding the surveyed papers. Table 3 provides an overview of the data items to be collected to answer the research questions within this study context. First part of the form is focused on bibliographic information, such as publication title, year, venue etc.

Furthermore, data was extracted concerning the second part of the form concentrating on the research question defined in Section 2.1. We derived these data items based on the research questions, as it can be observed from the right column, indicating the mapping.

Table 3. Fields of the Data Extraction Form

ID	Data Item	Description	RQ
Bibliographic Information			
D1	Publication Title	Title of paper	
D2	Publication Year	Calendar year	
D3	Publication Type	Journal, conference, workshop, book chapter	
D4	Publication Venue	Conference name, book title, journal title	
D5	Author Country	Country of the Affiliation associated with the author	RQ1 a
D6	Keywords	Keywords assigned to the publication by the authors	RQ1 b
D7	Paper summary	Short summary of the paper (with regard to this SMS)	
Study Information			
D8	Level of Maturity	with regard to stability: prototype, beta, or stable	RQ6 a
D9	Targeted Tasks	e.g., KG completion, question answering	RQ3 a
D10	Application Domain	e.g., health care, entertainment	RQ3 b
D11	Training Type of ML Model	e.g., supervised, reinforcement learning	RQ4 a
D12	ML Models	ML models used in the system	RQ4 b- d
D13	Semantic Resource	e.g., DBpedia	RQ5 a
D14	Author SW Type	Type of of Semantic Resource as stated in the paper	RQ5 a
D15	Our SW Type	Type of of Semantic Resource as defined in this study	RQ5 a
D16	Use of SW Module	part of the semantics which is actually used	RQ5 b
D17	Size of SW Module	e.g. number triples	RQ5 c
D18	SW Formalism	e.g. RDF-S, OWL	RQ5 d
D19	Type of Semantic Processing	Presence and type of semantic processing modules e.g., reasoner, SPARQL query engine	RQ5 e
D20	Processing Flow	processing flow of the system in terms of inputs/outputs and the order or processing units	RQ2
D21	Provenance Capturing	Presence of input/output data provenance collection: yes, no	RQ6 c
D22	Infrastructure Documentation	Presence of documentation on used infrastructure: yes, no	RQ6 b
D23	Software Documentation	Presence of documentation on used software and libraries: yes, no	RQ6 b
D24	Dataset Documentation	Presence of documentation on used data sets: yes, no	RQ6 b
D25	Data-Split documentation	Presence of documentation on used training, development and test set used in evaluation: yes, no	RQ6 b
D26	Processing Steps Documentation	Presence of documentation on performed processing steps, such as pre-processing, cross-validation: yes, no	RQ6 b
D27	Metrics Documentation	Presence of documentation on used metrics in evaluation: yes, no	RQ6 b

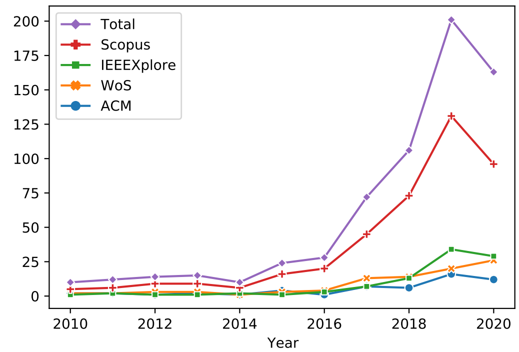

Fig. 2. Publications in search engines per year (not exclusive)

3 DATA ANALYSIS

3.1 General analysis

A total of 476 papers have been selected for inclusion in the survey following the aforementioned data acquisition process. Figure 2 shows the non-exclusive distribution per year of publications retrieved from various search engines.[11] We observed two trends in the publication count over the years. On the one hand, starting from 2016 there has been a surge in the number of papers published in all databases. On the other hand, a large portion of the selected papers were retrieved from Scopus. Between 2010 and 2016, the published papers account yearly for less than 5% of the total number of selected publications. From 2016 onward, 15-20% were retrieved yearly, increasing in 2019 and 2020 to over 35% of all publications selected for data extraction. An important aspect we take into account in the remainder of the data analysis is that the decrease from 2019 might be a side effect of the fact that the set of papers from 2020 is incomplete.[12]

Furthermore, we found that 67% of the academic venues hosted less than 2% of the selected papers. The percentage, as well as the absolute number of publications published at the remaining 33% of the venues is listed below:

- 13% (62) in Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence & Lecture Notes in Bioinformatics)
- 5% (25) in CEUR Workshop Proceedings
- 4% (17) in IEEE Access
- 3% (15) in ACM SIGKDD International Conference on Knowledge Discovery & Data
- 3%(13) in EMNLP - Conference on Empirical Methods in Natural Language Processing
- 3%(12) in Communications in Computer Information Science
- 3%(12) in SIGIR- International ACM SIGIR Conference on Research Development in Information Retrieval

Figure 3 illustrates the regional distribution of the authors of selected publications. We found three major clusters dominating the countries of authors publishing in the domain of SWeMLs. More specifically, 43% of the surveyed papers have an author affiliated with an institution in Asia, approximately 29% one affiliated in Europe and nearly 19% an author based in North America. Among the Asian countries, in 71% of the cases the author is based in China, while in North America the authors are located in the US in 86% of the cases. The geographical

[11]Note for Figure 2 that several papers were counted multiple times in the graph due to being available in more than one digital library.
[12]The search for papers was performed in October 2020 and many digital libraries have delays of several months for indexing publications.

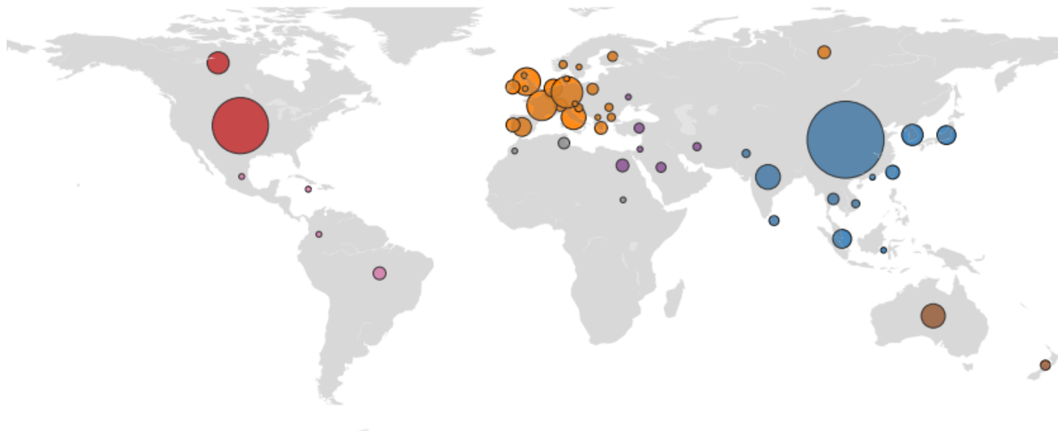

Fig. 3. Regional distribution of paper authors. The circles represent countries, their color denotes geographical regions, whereas their area indicates the amount of papers with authors affiliated at an institution in the corresponding country.

distribution in Europe is less skewed, with Germany, France, the United Kingdom, and Italy, being the most frequent countries of affiliation of the authors, each in over 10% of the cases. In only 1% of the cases publications have an author from Africa or from South and Central America. Similarly, the Middle East and Oceania are also underrepresented, as in only 2.5% to 3.5% of the cases, respectively, publications have an author affiliated with an institution from one of the two regions.

REFERENCES

[1] ACM. 2017. Statement on Algorithmic Transparency and Accountability. *Association for Computing Machinery US Public Policy Council (USACM)* January 12 (2017), 1–2.

[2] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* 58 (2020), 82 – 115. https://doi.org/10.1016/j.inffus.2019.12.012

[3] Tarek R Besold, Artur d'Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pascal Hitzler, Kai-Uwe Kühnberger, Luis C Lamb, Daniel Lowd, Priscila Machado Vieira Lima, et al. 2017. Neural-symbolic learning and reasoning: A survey and interpretation. *arXiv preprint arXiv:1711.03902* (2017).

[4] Tarek R. Besold and Oliver Kutz (Eds.). 2017. *First International Workshop on Comprehensibility and Explanation in AI and ML 2017.* Number 994 in CEUR Workshop Proceedings. http://ceur-ws.org/Vol-994

[5] Pearl Brereton, Barbara A Kitchenham, David Budgen, Mark Turner, and Mohamed Khalil. 2007. Lessons from applying the systematic literature review process within the software engineering domain. *Journal of systems and software* 80, 4 (2007), 571–583.

[6] Claudia D'Amato. 2020. Machine Learning for the Semantic Web: Lessons learnt and next research directions. *Semantic Web* 11, 1 (jan 2020), 195–203. https://doi.org/10.3233/SW-200388

[7] Derek Doran, Sarah Schulz, and Tarek R. Besold. 2017. What Does Explainable AI Really Mean? A New Conceptualization of Perspectives, See [4]. http://ceur-ws.org/Vol-2071/CExAIIA_2017_paper_2.pdf

[8] A. Garcez, K. Broda, D. Gabbay, et al. 2002. *Neural-symbolic learning systems: foundations and applications.* Springer.

[9] Artur S. d'Avila Garcez, Dov M. Gabbay, and Krysia B. Broda. 2002. *Neural-Symbolic Learning System: Foundations and Applications.* Springer-Verlag, Berlin, Heidelberg.

[10] Google. 2018. *Artificial Intelligence at Google - Our Principles.* Technical Report. Google. https://ai.google/principles

[11] Pascal Hitzler, Federico Bianchi, Monireh Ebrahimi, and Md Kamruzzaman Sarker. 2020. Neural-symbolic integration and the Semantic Web. *Semantic Web* 11, 1 (2020), 3–11. https://doi.org/10.3233/SW-190368

[12] Pascal Hitzler, Federico Bianchi, Monireh Ebrahimi, and Md Kamruzzaman Sarker. 2020. Neural-symbolic integration and the Semantic Web. *Semantic Web* 11, 1 (2020), 3–11.

[13] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. 2020. A Survey on Knowledge Graphs: Representation, Acquisition and Applications. *arXiv* (2020), 1–25. arXiv:2002.00388 http://arxiv.org/abs/2002.00388

[14] Henry Kautz. [n.d.]. The Third AI Summer, AAAI Robert S. Engelmore Memorial Lecture, 34th AAAI, 2020.

[15] Barbara Kitchenham, Stuart Charters, et al. 2007. Guidelines for performing systematic literature reviews in software engineering version 2.3. *Engineering* 45, 4ve (2007), 1051.

[16] Maria Angela Pellegrino, Abdulrahman Altabba, Martina Garofalo, Petar Ristoski, and Michael Cochez. 2020. GEval: A Modular and Extensible Evaluation Framework for Graph Embedding Techniques. In *European Semantic Web Conference*. Springer, 565–582.

[17] Achim Rettinger, Uta Lösch, Volker Tresp, Claudia D'Amato, and Nicola Fanizzi. 2012. Mining the semantic web: Statistical learning for next generation knowledge bases. *Data Mining and Knowledge Discovery* 24, 3 (feb 2012), 613–662. https://doi.org/10.1007/s10618-012-0253-2

[18] Petar Ristoski and Heiko Paulheim. 2016. Semantic Web in data mining and knowledge discovery: A comprehensive survey. *Journal of Web Semantics* 36 (jan 2016), 1–22. https://doi.org/10.1016/j.websem.2016.01.001

[19] R. Sapna, H. G. Monikarani, and Shakti Mishra. 2019. Linked data through the lens of machine learning: An Enterprise view. In *Proceedings of 2019 3rd IEEE International Conference on Electrical, Computer and Communication Technologies, ICECCT 2019*. Institute of Electrical and Electronics Engineers Inc. https://doi.org/10.1109/ICECCT.2019.8869283

[20] Md Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, and Pascal Hitzler. 2021. Neuro-Symbolic Artificial Intelligence Current Trends. *arXiv preprint arXiv:2105.05330* (2021).

[21] Arne Seeliger, Matthias Pfaff, and Helmut Krcmar. 2019. Semantic web technologies for explainable machine learning models: A literature review. *Proceedings of the 1st Workshop on Semantic Explainability co-located with the 18th International Semantic Web Conference (ISWC 2019)* 2465 (2019), 30–45.

[22] Satyaveer Singh and Mahendra Singh. 2018. Semantic Web Mining: Survey and Analysis. *Journal of Web Engineering & Technology* 5 (01 2018), 20–31.

[23] Dezhao Song, Frank Schilder, Shai Hertz, Giuseppe Saltini, Charese Smiley, Phani Nivarthi, Oren Hazai, Dudi Landau, Mike Zaharkin, Tom Zielund, et al. 2017. Building and querying an enterprise knowledge graph. *IEEE Transactions on Services Computing* 12, 3 (2017), 356–369.

[24] Gerd Stumme, Andreas Hotho, and Bettina Berendt. 2006. Semantic Web Mining. State of the art and future directions. *Journal of Web Semantics* 4 (06 2006), 124–143. https://doi.org/10.1016/j.websem.2006.02.001

[25] Michael van Bekkum, Maaike de Boer, Frank van Harmelen, André Meyer-Vitali, and Annette ten Teije. 2021. Modular design patterns for hybrid learning and reasoning systems. *Applied Intelligence* (2021), 1–19.

[26] Frank Van Harmelen and Annette Ten Teije. 2019. A boxology of design patterns for hybrid learning and reasoning systems. *Journal of Web Engineering* 18, 1-3 (2019), 97–124. https://doi.org/10.13052/jwe1540-9589.18133 arXiv:1905.12389

[27] Laura von Rueden, Sebastian Mayer, Katharina Beckh, Bogdan Georgiev, Sven Giesselbach, Raoul Heese, Birgit Kirsch, Julius Pfrommer, Annika Pick, Rajkumar Ramamurthy, et al. 2019. Informed Machine Learning–A Taxonomy and Survey of Integrating Knowledge into Learning Systems. *arXiv preprint arXiv:1903.12394* (2019).

[28] Jeroen Voogd, Paolo de Heer, Kim Veltman, Patrick Hanckmann, and Jeroen van Lith. 2019. Using Relational Concept Networks for Explainable Decision Support. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 78–93.

[29] Changchang Yin, Rongjian Zhao, Buyue Qian, Xin Lv, and Ping Zhang. 2019. Domain knowledge guided deep learning with electronic health records. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 738–747.