



RESEARCH ARTICLE

AIDA: A knowledge graph about research dynamics in academia and industry

Simone Angioni¹, Angelo Salatino², Francesco Osborne²,
Diego Reforgiato Recupero¹, and Enrico Motta²

¹Department of Mathematics and Computer Science, University of Cagliari (Italy)

²Knowledge Media Institute, The Open University, Milton Keynes (UK)

an open access  journal



Citation: Angioni, S., Salatino, A., Osborne, F., Recupero, D. R., & Motta, E. (2021). AIDA: A knowledge graph about research dynamics in academia and industry. *Quantitative Science Studies*, 2(4), 1356–1398. https://doi.org/10.1162/qss_a_00162

DOI:
https://doi.org/10.1162/qss_a_00162

Corresponding Author:
Simone Angioni
simone.angioni@unica.it

Keywords: bibliographic data, knowledge graph, research dynamics, scholarly data, scholarly ontologies, topic detection

ABSTRACT

Academia and industry share a complex, multifaceted, and symbiotic relationship. Analyzing the knowledge flow between them, understanding which directions have the biggest potential, and discovering the best strategies to harmonize their efforts is a critical task for several stakeholders. Research publications and patents are an ideal medium to analyze this space, but current data sets of scholarly data cannot be used for such a purpose because they lack a high-quality characterization of the relevant research topics and industrial sectors. In this paper, we introduce the Academia/Industry DynAmics (AIDA) Knowledge Graph, which describes 21 million publications and 8 million patents according to the research topics drawn from the Computer Science Ontology. 5.1 million publications and 5.6 million patents are further characterized according to the type of the author's affiliations and 66 industrial sectors from the proposed Industrial Sectors Ontology (INDUSO). AIDA was generated by an automatic pipeline that integrates data from Microsoft Academic Graph, Dimensions, DBpedia, the Computer Science Ontology, and the Global Research Identifier Database. It is publicly available under CC BY 4.0 and can be downloaded as a dump or queried via a triplestore. We evaluated the different parts of the generation pipeline on a manually crafted gold standard yielding competitive results.

1. INTRODUCTION

Academia and industry share a complex, multifaceted, and symbiotic relationship. Their collaboration and exchange of ideas, resources, and persons (Anderson, 2001a) are conducive to the production of new knowledge that will ultimately shape the society of the future. Analyzing the knowledge flow between academia and industry, understanding which directions have the biggest potential, and discovering the best strategies to harmonize their efforts is thus a critical task for several stakeholders (Salatino, Osborne, & Motta, 2020a). Governments and funding agencies need to regularly assess the potential impact of research areas and technologies to inform funding decisions. Commercial organizations have to monitor research developments and adapt to technological advancements. Researchers must keep up with the latest trends and be aware of complementary research efforts from the industrial sector.

The relationship between academia and industry has been analyzed from several perspectives in the literature, focusing for instance on the characteristics of direct collaborations (Ankrah & Omar, 2015), the influence of industrial trends on curricula (Weinstein, Kellar, &

Copyright: © 2021 Simone Angioni, Angelo Salatino, Francesco Osborne, Diego Reforgiato Recupero, and Enrico Motta. Published under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.



Hall, 2016), and the quality of the knowledge transfer (Ankrah, Burgess et al., 2013). However, most of the quantitative studies on this relationship were limited to small-scale data sets or focused on very specific research questions (Anderson, 2001a; Bikard, Vakili, & Teodoridis, 2019).

Research articles and patents are an ideal medium to analyze the knowledge generated and developed by academia and industry (Ankrah & Omar, 2015; Ankrah et al., 2013). Today, we have several large-scale knowledge graphs which describe research papers according to their titles, abstracts, authors, organizations, and other metadata. Examples include Microsoft Academic Graph¹ (Wang, Shen et al., 2020), Scopus², Semantic Scholar³, AMiner (Zhang, Zhang et al., 2018), CORE (Knoth & Zdrahal, 2012), OpenCitations (Peroni & Shotton, 2020), and others. Other resources, such as Dimensions⁴, the United States Patent and Trademark Office (USPTO)⁵, the Espacenet data set⁶, and the PatentScope corpus⁷, offer a similar description of patents. However, these data sets cannot be directly used to analyze the research dynamics of academia and industry as they lack a high-quality characterization of the relevant research topics and industrial sectors.

In particular, they suffer from three main limitations. First, current solutions do not allow us to easily discriminate if a document (research paper or patent) is from academia or industry. Second, they typically offer a coarse-grained characterization of research topics, which are usually represented only as a list of terms chosen by the authors or extracted from the abstract. This purely syntactic solution is unsatisfactory (Osborne & Motta, 2015), as it fails to distinguish research topics from other generic keywords; to deal with situations where multiple labels exist for the same research area; and to model and take advantage of the semantic relationships that hold between research areas. For instance, we want to be able to infer that all documents tagged with the topic Neural Network are also about Machine Learning and Artificial Intelligence. This richer representation would allow us to retrieve all the publications that address the concept Artificial Intelligence, even if the metadata does not contain the specific string “artificial intelligence.” A third issue is that current scholarly data sets do not characterize companies according to their sectors. Therefore, it is not possible to measure the impact of a topic (e.g., sentiment analysis, deep learning, semantic web) on different types of industry (e.g., automotive, financial, energy).

These limitations affect also the performance of machine learning systems, typically based on neural networks, for predicting the impact of research trends and forecasting patents (Choi & Jun, 2014; Marinakis, 2012; Ramadhan, Malik, & Sjafrizal, 2018; Zang & Niu, 2011). These solutions typically work with limited features, such as the number of patents associated with a topic for each year, as current data sets do not integrate articles and patents, lack a granular representation of research topics, and cannot distinguish whether a document was produced by academia or industry. We hypothesize that considering a richer characterization of this space would ultimately yield better performance in comparison to state-of-the-art approaches.

In this paper, we introduce the Academia/Industry DynAmics (AIDA) Knowledge Graph, which describes 21 million publications and 8 million patents in the field of *Computer*

¹ Microsoft Academic Graph: <https://aka.ms/microsoft-academic>

² Scopus: <https://www.scopus.com/>

³ Semantic Scholar: <https://www.semanticscholar.org/>

⁴ Dimensions: <https://www.dimensions.ai/>

⁵ USPTO: <https://www.uspto.gov/>

⁶ Espacenet dataset: <https://worldwide.espacenet.com/>

⁷ PatentScope: <https://patentscope.wipo.int/>

Science. Papers and patents are associated to the research topics in the Computer Science Ontology (CSO). In addition, 5.1 million publications and 5.6 million patents are also characterized according to the type of the author's affiliations (e.g., academia, industry) and 66 industrial sectors (e.g., automotive, financial, energy, electronics) from the Industrial Sectors Ontology (INDUSO). AIDA is also linked to several other knowledge bases, including MAKG, Dimensions, Google Patents, GRID, DBpedia, and Wikidata.

AIDA is available at <https://w3id.org/aida/>. It can be downloaded as a dump or queried via a Virtuoso triplestore at <https://w3id.org/aida/sparql/>. We plan to release a new version of AIDA every 6 months, to regularly update the publications, the topics, and the industrial sectors.

AIDA was generated using an automatic pipeline that integrates data from Microsoft Academic Graph (MAG)⁸, Dimensions, English DBpedia, the Computer Science Ontology (CSO), and the Global Research Identifier Database (GRID), respectively containing information about 242 million research papers, 38 million patents, 4.58 million entities, 14,000 research topics, and 97,000 organizations.

The resulting knowledge base enables analyzing the evolution of research topics across academia and industry and studying the characteristics of several industrial sectors. For instance, it enables detecting the research trends most interesting for the automotive sector or which prevalent industrial topics were recently adopted by academia. It can thus be utilized by a variety of deep learning methods for predicting the impact of research trends on industry and academia (Chung & Sohn, 2020; Ramadhan et al., 2018; Zang & Niu, 2011). It can also be used to characterize authors, citations, countries, and several other entities in MAG according to their topics and industrial sectors. This makes it possible to study further dynamics, such as the migration of researchers and the citation flow between academia and the industry.

We evaluated the different parts of the pipeline for generating AIDA on manually crafted gold standards yielding competitive results. We also report an evaluation of the impact of AIDA on forecasting systems for predicting the impact of research topics on the industry. Specifically, we tested five classifiers on 17 combinations of features and found that the forecaster based on Long Short-Term Memory neural networks and exploiting the full set of features from AIDA obtain significantly better performance ($p < 0.0001$) than alternative methods.

A preliminary version of AIDA which included a smaller data set and a limited number of semantic relations was previously discussed in a short workshop paper (Angioni, Salatino et al., 2020). The current paper greatly expands on that work by presenting a novel and up-to-date version of AIDA (including about 5 million additional articles), an improved version of the pipeline for generating AIDA, a more extensive ontological schema, and a comprehensive evaluation of AIDA.

In summary, our main contributions include the following:

- the first official release of AIDA, a knowledge graph for studying the research dynamics of academia and industry;
- a pipeline for automatically generating AIDA based on a robust semantic model and a state-of-the-art topic detection approach;
- a detailed discussion of AIDA schema, content, and links to other knowledge graphs;
- an evaluation of the AIDA pipeline and its ability to classify documents in terms of research topics and industrial sectors;

⁸ We used the dump released in April 2020.

- an illustrative overview of the Computer Science domain according to the data in AIDA;
- a discussion of AIDA possible usage that summarizes some research efforts that adopted preliminary versions of AIDA;
- an analysis of the current limitations of the AIDA pipeline and a sustainability plan developed in collaboration with Springer Nature for replacing MAG with a combination of Dimensions and DBLP, after MAG will be decommissioned at the end of 2021; and
- an appendix detailing several exemplary SPARQL queries in order to support the reuse of AIDA.

The rest of the paper is organized as follows. In Section 2, we review the literature on methods and data sets for studying and quantifying the relationship between academia and industry. In Section 3, we describe the pipeline to generate AIDA, give an overview of the resulting knowledge graph, and discuss our strategy for releasing new versions. Section 4 presents the evaluation of the different parts of the AIDA pipeline and the experiments showing that AIDA can support effectively deep learning approaches for predicting the impact of research topics. In Section 5 we focus on the usage of AIDA and report three exemplary research efforts that adopted preliminary versions of AIDA: a bibliometric analysis of the research dynamics across academia and industry; a study of the main research trends in two main venues of Human-Computer Interaction; and a new web application that we developed to support Springer Nature editors in assessing the quality of scientific conferences. Section 6 describes the main limitations of the proposed pipeline and how we will address them going forward. Finally, in Section 7 we summarize the main conclusions and outline future directions of research.

2. LITERATURE REVIEW

In this section, we review the current state of the art regarding knowledge graphs describing research papers and patents (Section 2.1) and approaches for analyzing the relationships between industry and academia (Section 2.2).

2.1. Knowledge Graphs of Research Articles and Patents

Knowledge graphs are graph of data intended to accumulate and convey knowledge of the real world, whose nodes represent entities of interest and whose edges represent relations between these entities (Hogan, Blomqvist et al., 2021). Such descriptions have formal semantics allowing both computers and people to process them efficiently and unambiguously. Knowledge graphs about research articles and patents typically describe the relevant actors (e.g., authors, organizations) and entities (e.g., topics, tasks, technologies), as well as any other contextual information (e.g., project, funding) in an interlinked manner.

In recent years we have seen the emergence of several knowledge graphs describing research publications and their metadata.

Microsoft Academic Graph (MAG) (Wang et al., 2020) is a heterogeneous knowledge graph that contains the metadata of more than 248 million scientific publications, including citations, authors, institutions, journals, conferences, and fields of study. Microsoft Academic Knowledge Graph (MAKG)⁹ (Färber, 2019) is a large RDF data set based on MAG that also provides entity embeddings for the research papers.

⁹ MAKG: <https://makg.org/>

The Semantic Scholar Open Research Corpus¹⁰ (Ammar, Groeneveld et al., 2018) is a data set of about 185 million publications released by Semantic Scholar, an academic search engine provided by the Allen Institute for Artificial Intelligence (AI2). The OpenCitations Corpus (Peroni & Shotton, 2020) is released by OpenCitations, an independent infrastructure organization for open scholarship dedicated to the publication of open bibliographic and citation data with semantic technologies. The current version includes 55 million publications and 655 million citations. Scopus is a well-known data set curated by Elsevier, which includes about 70 million publications and is often used by governments and funding bodies to compute performance metrics. The AMiner Graph (Zhang et al., 2018) is the corpus of more than 200 million publications generated and used by the AMiner system¹¹. AMiner is a free online academic search and mining system that also extracts researchers' profiles from the Web and integrates them into the metadata. The Open Academic Graph (OAG)¹² is a large knowledge graph integrating Microsoft Academic Graph and AMiner Graph. The current version contains 208 million papers from MAG and 172 million from AMiner. CORE (Knoth & Zdrahal, 2011)¹³ is a repository that integrates 24 million open access research outputs from repositories and journals worldwide. The Dimensions corpus is a data set produced by Digital Science that integrates and interlinks 109 million research publications, 5.3 million grants, and 40 million patents. Publications and citations are freely available for personal, noncommercial use.

DBLP (Ley, 2009) is a very well-curated bibliographic database of conferences, workshops, and journals in Computer Science. It currently covers 5.7 million articles, 5,443 conferences, and 1,773 journals. The ACL Anthology Reference Corpus (Bird, Dale et al., 2008) is a digital archive of conference and journal papers in natural language processing and computational linguistics, which aims to serve as a reference repository of research results. UnarXive (Saier & Färber, 2020) is a data set including over one million publications from arXiv.org for which it provides the full text and in-text citations annotated via global identifiers. AceKG (Wang, Yan et al., 2018) is a large-scale knowledge graph that provides 3 billion triples of academic facts about papers, authors, fields of study, venues, and institutes, as well as the relations among them. It was designed as a benchmark data set for challenging data mining tasks, including link prediction, community detection, and scholar classification. DOI-boost (La Bruzzo, Manghi, & Mannocci, 2019) provides an enhanced version of Crossref¹⁴ that integrates information from Unpaywall, ORCID, and MAG, such as author identifiers, affiliations, organization identifiers, and abstracts. It is periodically released on Zenodo¹⁵.

Several other knowledge graphs and resources focus specifically on patents (Schwartz & Sichelman, 2019). For instance, the European Patent Office (EPO) curates the Espacenet data set, which currently covers about 110 million patents from all over the world. Similarly, the United States Patent and Trademark Office produces a corpus that includes more than 14 million US patents. The World Intellectual Property Organization (WIPO) offers the PatentScope data set, which contains 84 million patent documents, including 4 million international patent applications.

Deng, Huang, and Zhu (2019) propose a method based on conditional random fields for automatically generating knowledge graphs describing technologies extracted from a set of

¹⁰ ORC: <https://s2-public-api-prod.us-west-2.elasticbeanstalk.com/corpus/>

¹¹ AMiner: <https://www.aminer.cn/>

¹² Open Academic Graph: <https://www.openacademic.ai/oag/>

¹³ CORE: <https://core.ac.uk/>

¹⁴ Crossref: <https://www.crossref.org/>

¹⁵ DOIboost laster release: <https://zenodo.org/record/3559699>

patents. However, the approach was only tested on about 5,000 patents and the resulting knowledge base was not made available. TechNet (Sarica, Luo, & Wood, 2019)¹⁶ is a semantic networks which includes 4 million terms extracted from 5.8 million patents in the US patents database. Specifically, the authors created an NLP approach to mine generic engineering terms and used their word embeddings to assess their semantic similarity.

Another category of knowledge graphs offers a semantic representation of the content of scientific articles. The Semantic Web community has been working for a while on this direction, fostering the Semantic Publishing paradigm (Shotton, 2009), creating bibliographic repositories in the Linked Data Cloud (Nuzzolese, Gentile et al., 2016), generating knowledge bases of biological data (Belleau, Nolin et al., 2008), formalizing research workflows (Wolstencroft, Haines et al., 2013), implementing systems for managing nano-publications (Groth, Gibson, & Velterop, 2010; Kuhn, Chichester et al., 2016) and micropublications (Schneider, Ciccarese et al., 2014), and developing a variety of ontologies to describe scholarly data (e.g., SWRC¹⁷, BIBO¹⁸, BiDO¹⁹, FABIO²⁰, SPAR²¹ (Peroni & Shotton, 2018), and SKGO²² (Fathalla, Auer, & Lange, 2020)).

A recent example is the Open Research Knowledge Graph (ORKG) (Jaradeh, Auer et al., 2019)²³, which aims to describe research papers in a structured manner to make them easier to find and compare.

Several of these knowledge bases focus on describing the research areas of scientific publications. These include the Medical Subject Heading (MeSH)²⁴ in Biology, Mathematics Subject Classification (MSC)²⁵ in Mathematics, Physics Subject Headings (PhySH)²⁶ in Physics, and many others.

In the field of Computer Science, the best-known taxonomies of research areas are the ACM Computing Classification System²⁷ and the Computer Science Ontology (CSO) (Salatino, Thanapalasingam et al., 2018b). The first one is developed and maintained by the Association for Computing Machinery (ACM). It contains around 2,000 concepts and it is manually curated. Conversely, CSO is automatically generated from a large collection of publications by the Open University and includes about 14,000 research areas. We adopted CSO for AIDA because it is one order of magnitude larger than the alternatives and it comes with the CSO Classifier (Salatino, Osborne et al., 2019b; Salatino, Thanapalasingam, & Mannonci, 2019c), which is a tool for automatically annotating documents with CSO topics. Hence, it allows us to easily generate a granular representation of all the documents integrated from MAG and Dimensions.

Currently, there are no data sets that enable the study of fine-grained research topics and their relation with industrial sectors across research papers and patents.

For this reason, we decided to undertake this new endeavor and develop AIDA.

¹⁶ TechNet: <https://www.tech-net.org/>

¹⁷ SWRC: <https://ontoware.org/swrc>

¹⁸ BIBO: <https://bibliontology.com>

¹⁹ BiDO: <https://purl.org/spar/bido>

²⁰ FABIO: <https://purl.org/spar/fabio>

²¹ SPAR: <https://www.sparontologies.net/>

²² SKGO: <https://github.com/saidfathalla/Science-knowledge-graph-ontologies>

²³ ORKG: <https://www.orkg.org/orkg/>

²⁴ Medical Subject Heading: <https://www.ncbi.nlm.nih.gov/mesh>

²⁵ Mathematics Subject Classification: <https://mathscinet.ams.org/msc>

²⁶ Physics Subject Headings: <https://physh.aps.org/>

²⁷ ACM Classification System: <https://www.acm.org/publications/class-2012>

We decided to adopt MAG over the alternatives knowledge graphs of articles for two main reasons. First, it appears to be the most comprehensive among the publicly available data sets of publications (Visser, van Eck, & Waltman, 2021). Second, it associates articles with DOIs and organizations with GRID identifiers and therefore can be easily integrated with other knowledge graphs.

For patents, we chose Dimensions because of its comprehensiveness and also because it identifies organizations with GRID IDs, allowing us to easily integrate them with MAG affiliations.

After the first version of this manuscript was written, Microsoft announced that MAG will be decommissioned in 2022. For this reason, we formulated a plan in collaboration with Springer Nature for using a combination of Dimensions and DBLP as our source for research publications in the following versions of AIDA. This plan is presented in Section 6.

2.2. Relationship Between Academia and Industry

Academia and industry typically tend to influence each other by exchanging ideas, resources, and researchers (Powell & Snellman, 2004). Analyzing their relationship allows us to understand their role within the whole knowledge economy (Anderson, 2001b): from production, towards adoption, enrichment, and ultimately deployment as a new commercial product or service. In some cases, academia and industry engage in collaborations as an opportunity for a more productive division of tasks: academia focusing on scientific insights, and industry on commercialization (Bikard et al., 2019). Stilgoe (2020) discusses the main drivers of scientific innovation and focuses on the central role of the industry sector in pushing innovation by constantly deploying new technologies. However, it can be argued that innovation advances also through a more complex route, which involves the birth of a new scientific area, the development of its theoretical framework, and the creation of innovative products that capitalize on the new knowledge (Kuhn, 1962).

The knowledge transfer between academia and industry has been studied according to both qualitative (Grimpe & Hussinger, 2013; Michaudel, Ishihara, & Baran, 2015) and quantitative methods (Huang, Yang, & Chen, 2015; Larivière, Macaluso et al., 2018). A good example of the first category is Michaudel et al. (2015), who share their personal experience on how the collaboration between industry and academia impacted their research program. Similarly, Grimpe and Hussinger (2013) perform a survey-based analysis to understand the innovation performance associated with collaborations between universities and German manufacturers. In the category of quantitative approaches, Larivière et al. (2018) employ both research papers and patents to understand the primary interests of both sides in this symbiosis. Huang et al. (2015) also take a quantitative approach and analyze 20,000 research papers and 8,000 patents in the area of fuel cells to assess the direct benefits of collaborations between academia and industry.

Hanieh, AbdElall et al. (2015) argue that a partnership agreement between industry and academia aims at enhancing economic prosperity, social equity, and environmental protection. This partnership includes also carrying out scientific research activities and solving industrial problems. In their paper, the authors analyze the state of affairs in Palestine, showing that such cooperation is weak, and hence they advocate improving this partnership. Also, they suggest to develop curricula by including sustainability concepts and improving teaching methods.

However, these approaches focus on relatively narrow areas of science and do not use a granular characterization of research areas. Conversely, AIDA allows researchers to analyze the interaction of research topics and industrial sectors across millions of documents. The

resulting data can support a variety of studies that are not feasible with current knowledge bases. For instance, AIDA makes it possible to analyze how industrial sectors (e.g., automotive) contribute to specific research fields (e.g., AI, Robotics) and how certain research lines lead to the development of concrete commercial services. It also enables us to quantify the impact of a field on industry across the years, in order to better assess the concrete fallback of scientific research.

3. AIDA: ACADEMIA INDUSTRY DYNAMICS KNOWLEDGE GRAPH

The Academia/Industry DynAmics (AIDA) Knowledge Graph includes about 1.3 billion triples that describe a large collection of publications and patents in Computer Science according to their research topics, industrial sectors, and author's affiliations (academia, industry, or collaborative). Specifically, 21 million publications from MAG and 8 million patents from Dimensions are classified according to the research topics drawn from the Computer Science Ontology (CSO). On average, each publication is associated with 27 ± 19 topics and each patent with 33 ± 14 ²⁸.

The 5.1 million publications and 5.6 million patents that were associated with GRID IDs in the original data are also classified according to the type of the author's affiliations (e.g., academia, industry) and 66 industrial sectors (e.g., automotive, financial, energy, electronics) drawn from the Industrial Sectors ontology (INDUSO)²⁹, which was specifically designed to support AIDA.

Because these annotations require at least an affiliation of the authors of the document to be associated with a GRID ID (as detailed in Section 3.1), they are currently restricted only to the document linked to GRID by Microsoft Academic Graph and Dimensions.

About 4.5 million articles and 4.9 million patents were also typed with the three main categories of our schema: academia, industry, and collaboration (between academia and industry). We also included additional affiliation categories from GRID, such as "Government," "Facility," "Healthcare," and "Nonprofit."

AIDA was generated and will be regularly updated by an automatic pipeline that integrates and enriches data from Microsoft Academic Graph (MAG), Dimensions, English DBpedia, the Global Research Identifier Database (GRID), CSO, and INDUSO.

Table 1 shows the number of publications and patents from academia, industry, and collaborative efforts. Note that only the documents associated with a GRID ID (about 5.1 million publications and 5.6 million patents) can be classified as academia, industry, collaborative, or any other additional category from GRID.

When considering the affiliation types, most publications (69.8%) are written by academic institutions. However, industry contributes to a good number of them (15.3%). The situation is reversed when considering patents: 84% of them are from industry and only 2.3% from academia. Another interesting finding is that the collaborative efforts are limited, involving only 2.6% of the publications and 0.2% of the patents. These numbers require further analysis but may suggest that we need to improve the mechanisms to support and fund collaborative works.

The data model of AIDA builds on AIDA Schema, Schema.org, FOAF, OWL, CSO, and others. We created AIDA Schema to define all the specific relations that could not be reused from state-of-the-art ontologies. It is available at <https://w3id.org/aida/ontology>.

²⁸ With $x \pm y$ we refer to x being the average and y the standard deviation.

²⁹ INDUSO: <https://w3id.org/aida/downloads/induso.ttl>

Table 1. AIDA—Affiliation types

	Publications	Patents
Academia	3,906,131	122,390
Industry	834,443	4,760,614
Collaborative	133,781	16,806
Additional categories in GRID	627,179	747,618
Documents with GRID ID	5,133,171	5,639,252
Total documents	20,850,710	7,940,034

Figure 1 depicts the full data model of AIDA KG, including both the relations that we defined within AIDA Schema and those we imported from external schemas. It focuses on six types of entities (light blue boxes in Figure 1): papers, patents, authors, affiliations, industrial sectors, and DBpedia categories. To be compatible with other knowledge graphs in this space (e.g., MAG, Scopus, DBLP, Semantic Scholar), papers are identified according to their Digital Object Identifier (DOI) and patents according to their World Intellectual Property Organization (WIPO) ID. We also retain the original MAG IDs for papers and authors as additional identifiers. These are used to link AIDA to MAKG and to identify articles that lack a DOI. In addition, affiliations are identified with GRID IDs. Industrial sectors and DBpedia categories are identified according to the instances available within INDUSO.

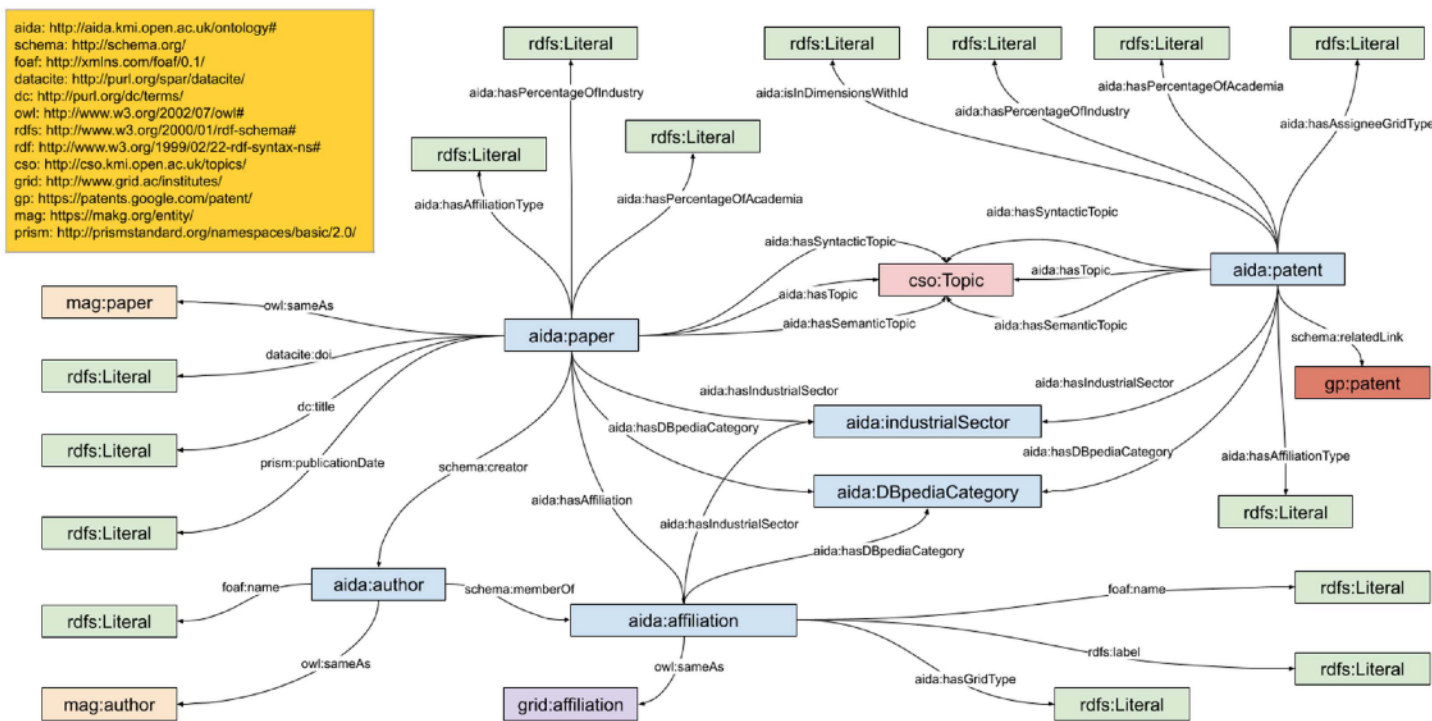


Figure 1. AIDA KG data model. For an enlarged version, visit <https://w3id.org/aida#aidaschema>.

The main information about papers and patents is given by means of the following semantic relations:

- *hasTopic*, which associates with the documents all their relevant topics drawn from CSO.
- *hasIndustrialSector*, which associates with documents and affiliations the relevant industrial sectors drawn from INDUSO.
- *hasAffiliationType*, which associates with the documents the three categories (academia, industry, or collaborative) describing the affiliations of their authors.

AIDA Schema includes also some additional relationships which support more complex queries:

- *hasSyntacticTopic* and *hasSemanticTopic*, which indicate, respectively, all the topics extracted using the syntactic module and the semantic module of the CSO Classifier (Salatino, Osborne et al., 2019b). The first set is composed of topics that are explicitly mentioned in the documents. It has high precision but low recall and may be used by applications for which precision is paramount. The second one consists of topics that do not directly appear in the text but were inferred using word embeddings.
- *hasAffiliation*, which identifies the affiliations of a paper.
- *hasPercentageOfAcademia* and *hasPercentageOfIndustry*, which link to articles and patents the percentage of authors from academia and industry. It may be used to generate analytics that are needed to further segment the *collaborative* category.
- *hasGridType*, *hasAssigneeGridType*, which associate the eight categories of organizations described in GRID (Education, Healthcare, Company, Archive, Nonprofit, Government, Facility, and Other) with affiliations and patents.
- *hasDBpediaCategory*, which associates with papers the industrial categories found in DBpedia (through the *About:Property* and *About:Industry*).
- *isInDimensionsWithId*, which identifies the patent id used within the Dimensions database.

As already mentioned, the AIDA knowledge graph also adopts several relations from external sources:

- <https://schema.org/creator>, which links documents to authors and authors to affiliations.
- <https://schema.org/memberOf>, which links authors to affiliations.
- <https://www.w3.org/1999/02/22-rdf-syntax-ns#type>, which defines the type of the entity.
- <https://www.w3.org/2000/01/rdf-schema#label>, which indicates the label of an affiliation.
- <https://purl.org/dc/terms/title>, which indicates the title of a paper.
- <https://purl.org/spar/datacite/doi>, which indicates the DOI of a paper.
- <https://xmlns.com/foaf/0.1/name>, which indicates the name of an author or an affiliation.
- <https://schema.org/relatedLink>, which states the related link of a patent (typically a Google Patent URL).
- <https://prismstandard.org/namespaces/basic/2.0/publicationDate>, which indicates the year of publication of a paper.
- <https://www.w3.org/2002/07/owl/sameAs>, which links papers, authors, or affiliations to their representations on external knowledge bases.

Table 2 reports the number of triples available in the current version of AIDA for each relation. AIDA includes about 1.3 billion triples: 1.2 billion with object properties and 98 million

Table 2. Number of triples for each relation in AIDA

Provenance	Relation	Triples
AIDA	https://aida.kmi.open.ac.uk/ontology#hasTopic	847,931,791
AIDA	https://aida.kmi.open.ac.uk/ontology#hasSemanticTopic	159,711,581
AIDA	https://aida.kmi.open.ac.uk/ontology#hasSyntacticTopic	70,349,962
AIDA	https://www.w3.org/2000/01/rdf-schema#type	54,839,960
AIDA	https://www.w3.org/2002/07/owl#sameAs	46,950,925
AIDA	https://aida.kmi.open.ac.uk/ontology#hasIndustrialSector	12,006,596
AIDA	https://aida.kmi.open.ac.uk/ontology#hasAffiliationType	9,774,165
AIDA	https://aida.kmi.open.ac.uk/ontology#hasDBpediaCategory	9,691,511
AIDA	https://aida.kmi.open.ac.uk/ontology#isInDimensionWithId	7,940,034
AIDA	https://schema.org/relatedLink	7,940,034
AIDA	https://aida.kmi.open.ac.uk/ontology#hasPercentageOfAcademia	4,179,108
AIDA	https://aida.kmi.open.ac.uk/ontology#hasPercentageOfIndustry	5,745,644
MAG	https://schema.org/creator	53,647,155
MAG	https://xmlns.com/foaf/0.1/name	26,048,450
MAG	https://purl.org/dc/terms/title	20,850,710
MAG	https://prismstandard.org/namespaces/basic/2.0/publicationDate	20,850,710
MAG	https://purl.org/spar/datacite/doi	5,636,401
MAG	https://schema.org/memberOf	4,828,260
MAG	https://aida.kmi.open.ac.uk/ontology#hasAffiliation	6,613,216
GRID	https://aida.kmi.open.ac.uk/ontology#hasAssigneeGridType	5,056,426
GRID	https://aida.kmi.open.ac.uk/ontology#hasGridType	13,171
GRID	https://www.w3.org/2000/01/rdf-schema#label	13,171

with datatype properties. Here, we distinguish the provenance of the triples to highlight which ones are directly generated by the AIDA pipeline (described in Section 3.1) and which ones are reused from other knowledge graphs. Overall, 1.18 billion triples (89.1 % of the total) were generated by our pipeline, while 185 million were derived from MAG and 7 million from GRID. We reused some relations from MAG, because they enable several kinds of useful queries involving, for instance, the years of publication of the articles and the names of the authors. In the set of triples generated by the AIDA pipeline, 1.08 billion (82.6%) regard the three main contributions of AIDA. Specifically, 1.07 billion triples regard the topics (*hasSyntacticTopic*, *hasSemanticTopic*, *hasTopic*), 19.6 million the affiliation types (*hasAffiliationType*, *hasPercentageOfAcademia*, *hasPercentageOfIndustry*), and 12.0 million the industrial sectors (*hasIndustrialSector*).

Table 3. Links of AIDA with external knowledge bases

Knowledge base	Type	Distinct entities	Total triples
CSO	Topic	11,091	1,077,993,334
MAKG	Author	26,035,279	26,035,279
MAKG	Paper	20,850,710	20,850,710
INDUSO	Industrial Sector	66	12,007,438
Dimensions	Patent	7,940,034	7,940,034
Google Patents	Patent	7,940,034	7,940,034
GRID	Affiliation	13,171	13,171
DBpedia	Organization	13,171	13,171
DBpedia	Concept	3,864	3,864
Wikidata	Concept	3,842	3,842

Table 3 reports the number of triples linking AIDA to external knowledge bases and the number of relevant distinct entities. For instance, AIDA includes more than 1 billion triples having as object a topic in CSO and overall links to 11,000 unique topics. AIDA is mostly linked to MAKG (the RDF version of MAG), including *owl:sameAs* relationships for 21 million papers and 25 million authors. It also links to Dimensions (8 million patents), Google Patents (8 million patents), GRID (13,000 affiliations), and DBpedia (3,864 concepts and 13,000 affiliations), and Wikidata (3,842 concepts). It should be noted that we cannot link directly to MAG, as it is not available online. However, as we use MAG IDs for papers and authors, mapping MAG and AIDA is trivial.

AIDA includes also the most recent mappings between CSO and DBpedia and between CSO and Wikidata, which implicitly links the documents in AIDA to 3,864 DBpedia entities and 3,842 Wikidata entities. Currently, those statements are not materialized for reason of space. However, materializing these links would yield an additional 460 million triples linking papers and patents to DBpedia entities (e.g., https://dbpedia.org/resource/Machine_learning) and 450 million triples linking them to Wikidata entities (e.g., <https://www.wikidata.org/entity/Q2539>). Alternatively, the user can explore these links by formulating SPARQL queries that take advantage of the *owl:sameAs* relationship between CSO, DBpedia, and Wikidata (see example in the Appendix).

The online documentation of AIDA Schema is available at <https://w3id.org/aida#aidaschema>.

AIDA is accessible via a Virtuoso triplestore at <https://w3id.org/aida/sparql>. The user can click the “help” button in the upper right of the web page for instructions on how to use the endpoint and some exemplary queries. The full dump of the last versions of AIDA is available at <https://w3id.org/aida/>. The dumps of the previous versions are available at <https://w3id.org/aida/downloads.php#datasets>.

AIDA is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0), meaning that everyone is allowed to copy and redistribute the material in any medium or format; and remix, transform and build upon the material for any purpose, even commercially.

In the following subsections, we will describe the pipeline for the automatic generation of AIDA (Section 3.1) and present an overview of the data (Section 3.2).

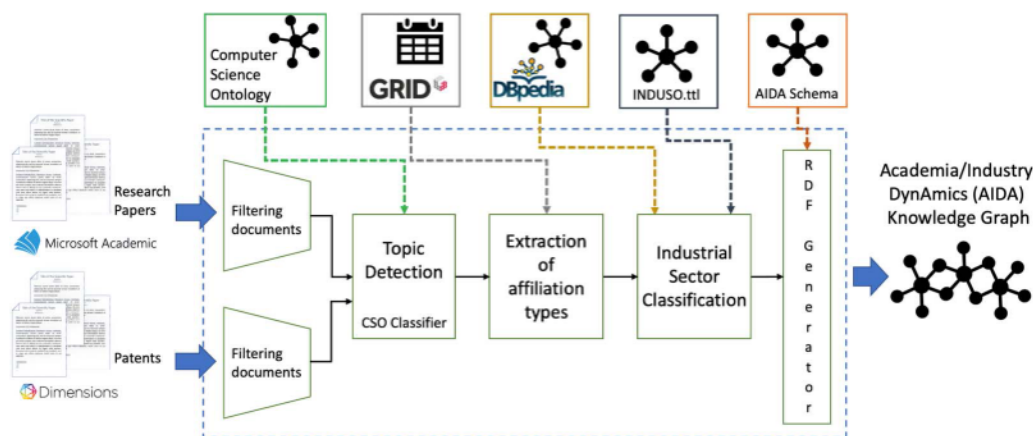


Figure 2. Workflow for the generation of AIDA.

3.1. AIDA Generation

The automatic pipeline for generating AIDA works in three steps: topic detection, integration of affiliation types, and industrial sector classification, as shown in Figure 2.

In the following, we will describe each phase of the process (Sections 3.1.1–3.1.3), discuss the scalability (Section 3.1.4), and present our plan for producing new versions (Section 3.1.5).

3.1.1. Topic detection

We first collect all the publications and patents from MAG and Dimensions within the Computer Science domain. In particular, we extract the papers from MAG classified as “Computer Science” in their Field of Science (FoS) (Sinha, Shen et al., 2015), an in-house taxonomy of research domains developed by Microsoft. Similarly, the patents in Dimensions are classified according to the International Patent Classification (IPC) and the fields of research (FoR) taxonomy, which is part of the Australian and New Zealand Standard Research Classification (ANZSRC). To extract only the patents from the Computer Science domain, we select those with the following IPC classification: “Computing, Calculating or Counting” (G06), “Educating, Cryptography, Display, Advertising, Seals” (G09), “Information Storage” (G11), “Information and Communication Technology” (G16), and others (G99). We also select those having the following field of research: “Information and Computing Science” (08) and “Technology” (10).

In the current version, the resulting data set includes 21 million publications and 8 million patents. The publications (21 million) and authors (25 million) extracted from MAG are also linked (*owl:sameAs*) to the relevant entities in MAKG. The patents obtained from Dimensions (8 million) are linked (*schema:relatedLink*) to the relevant patents in Google Patents.

Because the fields of study in MAG and fields of research in Dimensions are not specific enough for a detailed analysis of the knowledge flow, we then annotate each document with the research topics from the Computer Science Ontology (CSO) (Salatino et al., 2018b). CSO is an automatically generated ontology of research topics in the field of *Computer Science*. We used the current version (3.2), which includes 14,000 research topics and 159,000 semantic relationships. The CSO data model³⁰ is an extension of SKOS³¹ and the main semantic

³⁰ CSO Schema: <https://cso.kmi.open.ac.uk/schema/cso>

³¹ Simple Knowledge Organization System: <https://www.w3.org/2004/02/skos/>

relationships are *superTopicOf*, which is used to define the hierarchical relations within the field of Computer Science (e.g., *<artificial intelligence, superTopicOf, machine learning>*) and *relatedEquivalent*, which is used to define alternative labels for the same topic (e.g., *<ontology matching, relatedEquivalent, ontology alignment>*).

We adopted CSO because it offers a much more granular characterization of research topics than standard classification schemas (e.g., the ACM Classification) and generic knowledge graphs (e.g., DBpedia, Wikidata). For instance, a recent analysis (Salatino, Thanapalasingam et al., 2020b) reported that less than 37% of the topics in CSO are covered by DBpedia.

CSO was officially released in 2019 and has been already adopted by several major organizations, including Springer Nature. In the last 2 years, CSO supported the creation of many innovative applications and technologies, including ontology-driven topic models (e.g., CoCoNoW [Beck, Rizvi et al., 2020]), recommender systems for articles (e.g., SBR [Thanapalasingam, Osborne et al., 2018]) and video lessons (Borges & dos Reis, 2019), visualization frameworks (e.g., ScholarLensViz [Löffler, Wesp et al., 2020], ConceptScope [Zhang, Chandrasegaran, & Ma, 2021]), temporal knowledge graphs (e.g., TGK [Rossanez, dos Reis, & da Silva Torres, 2020]), NLP frameworks for entity extraction (Dessi, Osborne et al., 2021), tools for identifying domain experts (e.g., VeTo [Vergoulis, Chatzopoulos et al., 2020]), and systems for predicting academic impact (e.g., ArtSim [Chatzopoulos, Vergoulis et al., 2020a]). It was also used for several large-scale analyses of the literature (e.g., Cloud Computing [Lula, Dospinescu et al., 2021], Software Engineering [Chicaiza & Reátegui, 2020], and Ecuadorian publications [Chicaiza & Reátegui, 2020]).

We annotated publications and patents using the CSO Classifier (Salatino et al., 2019b), an open-source Python tool³² that we developed for annotating documents with research topics from CSO (Salatino et al., 2019c).

The CSO Classifier was initially developed in the context of a collaboration with Springer Nature, with the aim of automatically classifying scientific volumes according to a granular set of research areas. In this context, it supported Smart Topic Miner (Salatino et al., 2019a), a web application for assisting the Springer Nature editorial team in annotating conference proceedings in Computer Science, such as LNCS, LNBIP, CCIS, IFIP-AICT, and LNICST. This solution brought a 75% cost reduction and dramatically improved the quality of the annotations, resulting in 12 million additional downloads over 3 years from the SpringerLink portal³³.

The CSO Classifier is an unsupervised method that operates in three phases. First the syntactic module finds all topics in the ontology that are explicitly mentioned in the paper. Secondly, a semantic module identifies further semantically related topics using part-of-speech tagging and similarity over word embeddings. Finally, the CSO Classifier enriches the resulting set by including the superareas of these topics according to CSO.

Specifically, in the *syntactic* module, the text is split into unigrams, bigrams, and trigrams. Each n-gram is then compared with concepts labels in CSO using the Levenshtein similarity. As result, it returns all matched topics having similarities greater than or equal to the pre-defined threshold.

The *semantic* module takes advantage of a pretrained Word2Vec word embedding model which captures semantic properties of words (Mikolov, Sutskever et al., 2013). We trained this

³² CSO Classifier: <https://pypi.org/project/cso-classifier/>

³³ SpringerLink: <https://link.springer.com/>

model using the titles and abstracts of over 4.6 million English publications in the field of Computer Science from MAG. We preprocessed this data by replacing spaces with underscores in all n-grams matching the CSO topic labels (e.g., “semantic web” became “semantic_web”). We performed also a collocation analysis to identify frequent bigrams and trigrams (e.g., “highest_accuracies,” “highly_cited_journals”). This solution allows the CSO Classifier to better disambiguate concepts and treat terms such as “deep_learning” and “e-learning” as completely different words. The model parameters are: *method* = skipgram, *embedding-size* = 128, *window-size* = 10, *min-count-cutoff* = 10, *max-iterations* = 5. The semantic module based on these embeddings identifies candidate terms composed of a combination of nouns and adjectives using a part-of-speech tagger. Then, it splits these candidate terms into unigrams, bigrams, and trigrams. For each n-gram we retrieve its most similar word from the Word2Vec model and we compute their cosine similarity with the topic labels in CSO. For bigrams and trigrams, we first check in the model their glued version, creating one single word (e.g., “semantic_web”). If this word is not available within the model vocabulary, the classifier uses the average of the embedding vectors of all its tokens. Then, for each identified topic, the CSO Classifier computes the relevance score as the product between the number of times it was identified (frequency) and the number of unique n-grams that helped it to be inferred (diversity). Finally, it uses the elbow method (Satopaa, Albrecht et al., 2011) for selecting the set of most relevant topics.

Finally, the resulting set of topics is enriched by including all their supertopics in CSO up to the root: *Computer Science*. For instance, a paper tagged as *neural network* is also tagged with *machine learning*, *artificial intelligence*, and *computer science*. This solution yields an improved characterization of high-level topics that are not directly referred to in the documents.

The CSO ontology contains nine levels of topics. When we detect a specific topic (e.g., Neural Networks) we also infer all the super topics in the CSO taxonomy (Machine Learning, Artificial Intelligence, Computer Science). The user can choose to just use the topics directly mentioned in the paper (*hasSyntacticTopic*), those inferred by using word embeddings (*hasSemanticTopic*), or the full set of topics that also includes the supertopics (*hasTopic*).

More details about the CSO Classifier are available in Salatino et al. (2019b).

We also import in AIDA the mapping between CSO and DBpedia, which is a set of 3,864 *owl:sameAs* relationships aligning the two knowledge bases and the mapping between CSO and Wikidata, which includes 3,842 *owl:sameAs* relationships. This allows us to establish several implicit links between documents in AIDA and concepts in DBpedia and Wikidata, which can be materialized with a reasoner or queried using SPARQL (see example in the Appendix).

3.1.2. Integration of affiliation types

In the second step, we classify papers and patents according to the nature of the relevant organizations in the GRID database. Both MAG and Dimensions link organizations to their GRID IDs. In turn, GRID associates each ID with geographical location, date of establishment, alternative labels, external links, and type of institution (e.g., Education, Healthcare, Company, Archive, Nonprofit, Government, Facility, Other). In total, 5.1 million articles and 5.6 million patents were associated with GRID IDs. We leverage this last field to tag 4.5 million articles and 4.9 million patents as “academia,” “industry,” or “collaborative.” A document is assigned an “academia” type if all the authors or original assignees have an academic affiliation (“Education” in GRID), an “industry” type if they have an industrial affiliation (“Company” in GRID), and a “collaborative” type if there is at least one creator from academia and one from industry. AIDA includes also the other categories from GRID through the relation *hasGridType*.

3.1.3. Industrial sector classification

To characterize the industrial sectors addressed by each document, we designed the Industrial Sector Ontology (INDUSO), which is a two-level taxonomy describing 66 sectors and their relationships. INDUSO was created using a bottom-up method that took into consideration the large collection of publications and patents from MAG and Dimensions. Specifically, for each affiliation described in the documents with a GRID ID, we extracted from DBpedia the objects of the properties *About:Purpose* and *About:Industry*. This resulted in a noisy and redundant set of 699 sectors. We then applied a bottom-up hierarchical clustering approach for merging similar sectors. For instance, the industrial sector “Computing and IT” was derived from categories such as “Networking hardware,” “Cloud Computing,” and “IT service management.”

This structure was used as a starting point by a team of ontology engineers from the Open University and the University of Cagliari and domain experts from Springer Nature, who manually revised these categories and arranged the resulting sectors in a two-level taxonomy.

For example, the first level sector “energy” includes “nuclear power,” “oil and gas industry,” and “air conditioning.” Specifically, the INDUSO ontology contains the following properties:

- the *skos:broader* property, which links the first level sectors to the second level sectors.
- the *prov:wasDerivedFrom* property, which associates each of the 66 industrial sectors to the original 699 sectors that were derived from DBpedia.
- the *rdf:type* property, which is used to define the 66 sectors as *:industrialSector* and the original 699 sectors as *:DBpediaCategory*.

To tag a document with INDUSO, we identify its affiliations on DBpedia using the link between GRID and DBpedia and then retrieve the objects of the properties *About:Purpose* and *About:Industry*. We then use the previously defined mapping between DBpedia and INDUSO to obtain the industrial sectors.

For instance, a document with an author affiliation described in DBpedia as “natural gas utility” is tagged with the second level sector “Oil and Gas Industry” and the first level sector “Energy.”

3.1.4. Scalability

The pipeline currently runs on a server with 128 Gbyte of RAM, CPU Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40 GHz. Typically, one single paper requires 0.83 seconds to be processed and classified according to the CSO, Academia/Industry, and INDUSO classifications. Therefore, considering the 29 million documents (21 million papers and 8 million patents), and using a multithread programming style (we used 10 threads), it takes about 27 days to perform the classification of the entire data set.

For each following update, we only need to include new documents and update the citations of existing papers. This operation is much faster than processing the entire data set and we plan to run it periodically. For instance, considering a typical amount of new papers for 3 months in 2020, equal to about 350,000, the update will take around 8 hours.

3.1.5. Generation of updates

We plan to periodically release new versions of AIDA, which will include the most recent publications and patents, as well as the latest versions of CSO and INDUSO. Specifically, we will run the pipeline described in this section – and depicted in Figure 2 – over a new dump of documents every 6 months. In addition, we also plan to release a new version whenever a significant new version of CSO or INDUSO is produced.




Figure 3. Distribution of the main topics.

During the writing of this paper, Microsoft decided to decommission the MAG project after 2021. We have formulated a plan to switch to other sources that is discussed in Section 6.

3.2. AIDA Overview

In this section, we present an overview of AIDA and discuss some exemplary analytics supported by this resource.

Figure 3 shows the 16 high-level topics (direct subtopics of Computer Science in CSO) associated with most research articles in AIDA and reports the relevant percentage of academic publications, industrial publications, academic patents, and industrial patents.

These figures were computed by normalizing the number of documents associated with a topic in a category (e.g., academic publications) with the total number of documents in the same category. It should be noted that the percentages do not add to 100% because documents can be associated with multiple topics.

Some topics, such as Artificial Intelligence and Theoretical Computer Science, are mostly addressed by academic publications. Other (e.g., Computer Security, Computer Hardware, and Information Retrieval) attract a stronger interest from the industry. The topics which are mostly associated with patents are Computer Networks, Internet, and Computer Hardware.

Figure 4 shows the percentage of publications from academia (A) and industry (I) for the same 16 topics across three windows of time (1991–2000, 2001–2010, and 2011–2020). The split into three intervals of 10 years is useful to highlight the trend of each topic across the years. Some evident trends include the sharp growth of Computer Security, Information




Figure 4. Distribution of the topics in publications across time.




Figure 5. Distribution of the main industrial sectors.

Retrieval, Computer Network, and Internet. Some other topics, such as Software Engineering and Computer Aided Design, appear to have become less prolific in recent years.

Figure 5 (Main Industrial Sectors I and Main Industrial Sectors II) shows the 16 industrial sectors associated with most research articles and reports their percentage of publications and patents in AIDA.

Because AIDA mainly covers Computer Science, the most popular sectors (e.g., Technology, Computing and IT, Electronics, and Telecommunications, and Semiconductors) are linked to this field. However, we can also appreciate the solid presence of sectors such as Financial, Health Care, Transportation, Home Appliance, and Editorial.

AIDA also enables us to analyze how these sectors have a different composition with regard to research topics. Table 4 highlights the key topics of a set of exemplary sectors by reporting the difference between the normalized number of publications in a sector and overall. The darker cells mark the main topics for each sector. For instance, the publications written by authors from the Semiconductor sector refer to the topics Computer Aided Design 90% more frequently than the average publication.

The industrial sectors have a very distinct composition, even when considering just the high-level topics in the table. For instance, the Automotive sector focuses mainly on Robotics, Software Engineering, and Artificial Intelligence; the Telecommunications sector mainly focuses on Computer Network, Internet, and Computer Hardware; and the Photography sector on Information Retrieval, Computer Vision, and Artificial Intelligence.

AIDA can also be queried via triplestore using SPARQL³⁴. The ontological schema of AIDA allows users to formulate queries about topics, industrial sectors, and affiliation types associated with articles and patents. In the Appendix we report a selection of sample queries that can be run on our SPARQL endpoint.

4. EVALUATION

To show that AIDA is both correct and useful, we performed two evaluations. In the first, reported in Section 4.1, we measured the precision and recall of the three components of the pipeline that produce the data about topics, the academia/industry classification, and the industrial sectors. In the second, presented in Section 4.2, we evaluated the ability of AIDA to support the task of predicting the impact of a research topic on industry. Specifically, we ran

³⁴ AIDA triplestore: <https://w3id.org/aida/sparql>

Table 4. Topic composition of some prominent industrial sectors. Bold indicates the highest value for each row

	Computing and IT		Telecommunications		Electronics	Semiconductor	Inf. Technology		Photography	Automotive	Financial
Artificial Intelligence	9%	5%	9%	-17%	0%	22%	8%	-6%			
Computer Aided Design	-21%	-27%	-2%	90%	1%	-5%	2%	-36%			
Computer Hardware	-7%	7%	-7%	31%	-5%	-12%	-9%	-17%			
Computer Network	-3%	17%	-9%	11%	-9%	-18%	-15%	-8%			
Computer Programming	18%	-19%	-1%	12%	52%	-31%	-16%	-32%			
Computer Security	6%	-1%	-2%	-27%	-1%	9%	-35%	21%			
Computer Systems	1%	1%	-3%	1%	4%	-2%	-12%	-10%			
Computer Vision	-7%	-1%	21%	-16%	-29%	44%	-7%	52%			
Data Mining	28%	-25%	12%	-35%	49%	-18%	-34%	-17%			
Human-computer Inter.	14%	-9%	8%	-41%	9%	-21%	-6%	32%			
Information Retrieval	6%	-16%	14%	-55%	-6%	71%	-37%	29%			
Information Technology	20%	-15%	-5%	-41%	55%	13%	-41%	-20%			
Internet	4%	13%	-6%	-1%	1%	-19%	-24%	-4%			
Operating Systems	14%	-40%	-8%	1%	61%	-24%	-55%	-30%			
Robotics	3%	-1%	16%	-14%	-9%	-18%	322%	15%			
Software Engineering	22%	16%	6%	2%	55%	-24%	20%	-31%			

several classifiers on different combination of features and found that the richer representation of topics in AIDA was conducive to significantly better performance than alternative solutions.

4.1. Evaluation of AIDA Generation

The following subsections describe the evaluations performed for assessing the topic classification, the academia/industry classification, and the industrial sector classification.

4.1.1. Topic classification

We compared the CSO Classifier, which we use to annotate documents according to their topics, against 13 unsupervised approaches using a gold standard made of 70 most cited papers (Salatino et al., 2019b) within the fields of Natural Language Processing (23 papers), Semantic Web (23), and Data Mining (24). We chose the most cited papers because this solution offers a simple, deterministic, and not arbitrary selection criterion. The 70 papers were annotated by 21 human experts. Each human expert annotated 10 papers; each paper was annotated by three human experts, resulting in 210 annotations overall. The 21 experts were researchers working in different areas of Computer Science with over 5 years of experience. They were asked to read title, abstract, and keywords and assign all the relevant topics from the CSO ontology so as to emulate the classifier's task. Each paper was associated with 14 ± 7.0 topics using the majority voting strategy.

The interannotator agreement was 0.45 ± 0.18 according to Fleiss' Kappa, resulting in a moderate interrater agreement.

It should be noted that this range of agreement is normal when using a large number of granular categories, such as the 14,000 topics in CSO.

In Table 5 we report the values of precision, recall, and F1 of all tested classifiers.

Table 5. Values of precision, recall, and f-measure. Bold indicates the best results

Classifier	Description	Prec.	Rec.	F1
TF-IDF	TF-IDF	16.7%	24.0%	19.7%
TF-IDF-M	TF-IDF mapped to CSO concepts	40.4%	24.1%	30.1%
LDA100	LDA with 100 topics	5.9%	11.9%	7.9%
LDA500	LDA with 500 topics	4.2%	12.5%	6.3%
LDA1000	LDA with 1,000 topics	3.8%	5.0%	4.3%
LDA100-M	LDA with 100 topics mapped to CSO	9.4%	19.3%	12.6%
LDA500-M	LDA with 500 topics mapped to CSO	9.6%	21.2%	13.2%
LDA1000-M	LDA with 1,000 topics mapped to CSO	12.0%	11.5%	11.7%
W2V-W	W2V on windows of words	41.2%	16.7%	23.8%
STM	Classifier used by STM	80.8%	58.2%	67.6%
SYN	Syntactic module	78.3%	63.8%	70.3%
SEM	Semantic module	70.8%	72.2%	71.5%
INT	Intersection of SYN and SEM	79.3%	59.1%	67.7%
CSO-C	The CSO Classifier	73.0%	75.3%	74.1%

The first eight classifiers are based on TF-IDF and Latent Dirichlet Allocation (LDA) (Blei, Ng, & Jordan, 2003), and their performance did not exceed an F1 of 30.1%. For each paper, TF-IDF returns a ranked list of words according to their TF-IDF score. The TF-IDF-M classifier, instead, returns the set of CSO topics having Levenshtein similarity higher than 0.8 with the words with the best TF-IDF score. This threshold was set empirically, because it yielded the best performance for the baselines.

LDA100, LDA500, and LDA1000 are three LDA classifiers, respectively trained on 100, 500, and 1,000 topics. These three classifiers select all LDA topics with a probability of at least j and return all their words with a probability of at least k . The best values of j and k were found by performing a grid search. In a similar way, we trained LDA100-M, LDA500-M, and LDA1000-M, but the resulting keywords are then mapped to the CSO topics, as for TF-IDF-M.

W2V-W processes the input document with a 10-words sliding window, and uses the word2vec model to identify CSO topics that are semantically similar to the embedding of the window. The embedding of the window are obtained by averaging the embeddings of the single tokens.

STM is the classifier originally adopted by Smart Topic Miner (Osborne, Salatino et al., 2016), the application used by Springer Nature for classifying proceedings within the Computer Science domain. It detects exact matches between the terms extracted from the text and the CSO topics. SYN represents the syntactic module of the CSO classifier, introduced in Salatino, Thanapalasingam et al. (2018a). SEM consists of the semantic module of the CSO classifier. INT represents a hybrid version that returns the intersection of the topics produced by the SYN and SEM modules. Finally, CSO-C is the default implementation of the CSO Classifier which produces the union of the topics returned by the two modules. The overall values of precision and recall for a given classifier are computed as the average of the values of precision and recall obtained over the papers.

The data produced in the evaluation, the Python implementation of the approaches, and the word embeddings are available at <https://w3id.org/cso/cso-classifier>.

Note that TF-IDF-M, LDA100-M, LDA500-M, LDA1000-M, W2V-W, STM, SYN, SEM, INT, and CSO-C are all general algorithms that classify a text according to the categories from an input taxonomy. Therefore, no method is specifically biased towards CSO.

The LDA500-M and TF-IDF-M approaches performed poorly with an f-measure of 30.1%. STM and SYN yielded very good precision of, respectively, 80.8% and 78.3%. These methods were able to find topics explicitly mentioned in the text, which tend to be very relevant. However, they suffered from low recall, 58.2% and 63.8% respectively, as they failed to identify more subtle topics. SEM had lower precision than SYN but higher recall and f-measure, suggesting that it can identify further topics that do not directly appear in the paper. INT generated higher precision (79.3%) compared to SYN and SEM (78.3% and 70.8%), but it did not yield good recall, dropping to 59.1%. Finally, CSO-C outperformed all the other methods in terms of both recall (75.3%) and f-measure (74.1%).

It should be noted that F1 in the 70%–75% range is remarkably good, given the granularity of the topics in the benchmark, and consistent with the results of other studies that used large classification schemas (e.g., MeSH [Costa, Rei et al., 2021]).

Indeed, the agreement (computed with Fleiss' Kappa) among the three annotators which created the gold standard was 0.451 ± 0.177 , indicating a moderate interrater agreement (Landis & Koch, 1977). When adding the CSO Classifier as a fourth annotator the agreement lowers only slightly to 0.392 ± 0.144 . The difference from human annotators may completely disappear

Table 6. Agreement between annotators (including the CSO classifier) and average agreement of each annotator according to the evaluation in Osborne et al. (2019). Bold indicates the best agreements for each annotator

	CSO-C	User1	User2	User3	User4	User5	User6
CSO-C	–	56%	68%	64%	64%	76%	64%
User1	56%	–	40%	56%	36%	48%	44%
User2	68%	40%	–	64%	52%	76%	64%
User3	64%	56%	64%	–	52%	64%	68%
User4	64%	36%	52%	52%	–	64%	52%
User5	76%	48%	76%	64%	64%	–	72%
User6	64%	44%	64%	68%	52%	72%	–
Av. agreement	66%	45%	58%	59%	51%	63%	60%

when considering a simpler classification schema. A recent experiment using the CSO Classifier for assisting systematic reviews (Osborne, Muccini et al., 2019) reported that its performance were not statistically significantly different from the ones of six senior researchers ($p = 0.77$) when classifying 25 papers according to five main subtopics of Software Architecture. We report in Table 6 the degree of agreement between the annotator (including also CSO-C), computed as the ratio of papers which were tagged with the same category by both annotators.

Since its introduction, in 2019, the CSO Classifier has been adopted by several applications and research efforts (Chatzopoulos, Vergoulis et al., 2020b; Dörpinghaus & Jacobs, 2020; Jose, Jagathy Raj, & George, 2021; Vergoulis, Chatzopoulos et al., 2020). For instance, Dörpinghaus and Jacobs (2020) used it for annotating the articles from the DBLP computer science library. Chatzopoulos et al. (2020b) integrated it in ArtSim, an approach for predicting the popularity of new research papers. Vergoulis et al. (2020) classified 1.5 million papers and use such topical representation for identifying experts that share similar publishing habits. Finally, Jose et al. (2021) developed an ontology-based framework that integrates CSO and the CSO Classifier for retrieving journal articles from academic repositories and dynamically expanding the ontology with new research areas.

4.1.2. Academia/industry and industrial sector classifications

To evaluate the quality of the academia/industry classification in AIDA we randomly selected 100 papers: 33 academic papers, meaning that all the authors of each paper are reported with academic affiliations only; 33 industry papers, whose authors are reported with affiliation in the industry only; and 34 collaborative papers, meaning that each paper in this set includes authors with affiliations from academia and authors with affiliations from the industry.

We then asked three independent researchers to manually annotate each paper as “academic,” “industrial,” or “collaborative” according to the classification above. They were allowed to check online whether a certain institution was academic or industrial. The average agreement score of the three experts was 92.6%. We generated a gold standard by using a majority voting strategy. That is, if a paper was considered an academic paper by at least two researchers, it was labeled as such. There were no cases where a paper was annotated with three different classes by the researchers.

The resulting gold standard perfectly matched the automatic classification.

Table 7. Performance of industrial sector classification task

<i>Industrial sector</i>	<i>Precision</i>	<i>Recall</i>	<i>F1 score</i>
Automotive	1.000	1.000	1.000
Healthcare	0.894	0.894	0.894
Computing and IT	0.850	0.809	0.829
Electronic	0.700	0.777	0.736
Telecommunication	0.944	0.894	0.918
<i>Macro Average</i>	0.877	0.875	0.875
<i>Weighted Average</i>	0.879	0.875	0.877

To evaluate the accuracy of our approach for identifying the industrial sectors of a document, we selected 100 organizations, equally divided (20 per each industrial sector) among telecommunication companies, healthcare companies, automotive companies, computing and information technology companies, and electronic companies. We then asked three independent experts (three senior researchers working within ICT companies and with a computer science background) to annotate each organization among the five classes above (or the *other* category if none of the previous categories was appropriate). The average agreement score of the experts was 84.0%.

We created a gold standard using a majority voting strategy. For instance, if a company was classified as healthcare by at least two experts, then its label was “healthcare.” Note that for each company, at least two experts always gave the same label. We then performed a precision-recall analysis of the categories forecasted by our approach and, for each category, we obtained the performance shown in Table 7.

It is interesting to note that, while the performance of our approach is overall quite good, it can differ according to the category. For example it is quite easy to recognize organizations in the “Automotive” sector, but much less so to identify the ones in “Electronic.” The same issues also affected human annotators. An analysis of the results seem to suggest that some categories (e.g., Electronic) are potentially more ambiguous according to both human annotators and the linked categories on DBpedia. Conversely, some other categories are more well defined and relatively easy to identify.

In conclusion, the evaluation substantiated that our approaches for classifying documents work remarkably well, performing similarly to human annotators.

4.2. Impact Forecasting

In this section, we present an evaluation of the ability of AIDA to support machine learning forecasters for predicting the impact of research topics on the industry, which is a typical task in the study of academia/industry relationship (Altuntas, Dereli, & Kusiak, 2015; Choi & Jun, 2014; Marinakis, 2012; Ramadhan et al., 2018; Zang & Niu, 2011). The impact of research topics on the industry has been traditionally quantified using the number of relevant patents. For instance, in AIDA the topic *wearable sensors* was granted only two patents during 2009. In the following years, a lot of commercial organizations started to invest in this area and submitted several patents, ultimately producing 135 patents in 2018. Predicting these dynamics is very advantageous for companies that need to stay at the forefront of innovation and anticipate new technologies.

The literature proposes a range of approaches to patent and technology prediction through patent data, using for instance weighted association rules (Altuntas et al., 2015), Bayesian clustering (Choi & Jun, 2014), and various statistical models (Marinakos, 2012) (e.g., Bass, Gompertz, Logistic, and Richards). In the last few years, we saw the emergence of several approaches based on Neural Networks (Ramadhan et al., 2018; Zang & Niu, 2011), which lately have obtained the most competitive results. However, most of these tools focus only on patents, as they are limited by current data sets that do not typically integrate research articles nor can they distinguish between documents produced by academia or industry. We thus hypothesized that a knowledge graph such as AIDA, which integrates a lot of information about publications and patents and their origin, should offer a richer set of features, ultimately yielding a better performance in comparison to approaches that rely solely on the number of publications or patents (Choi & Jun, 2014; Marinakis, 2012; Ramadhan et al., 2018; Zang & Niu, 2011).

To test this hypothesis, we generated a gold standard that associates with each topic in AIDA all the time frames of 5 years in which the topic had not yet emerged (fewer than 10 patents). These samples were labeled as *True* whenever the topic produced more than 50 industrial patents (PI) in the following 10 years and *False* otherwise. We then associated to each sample six time series composed respectively of the number of research articles (R), the number of patents (P), the number of research articles from academia (RA), research articles from industry (RI), patents from academia (PA), patents from industry (PI). For instance, the sample involving the topic *wearable sensors* in 2005–2009 contains the six series (R, P, RA, RI, PA, PI) describing the number of documents in each category during those 5 years and was labeled as *True*, as *wearable sensors* produced more than 50 industrial patents (PI) in the following years. The resulting data set includes 9,776 labeled samples.

We trained five machine learning classifiers on this gold standard: Logistic Regression (LR), Random Forest (RF), AdaBoost (AB), Convolutional Neural Network (CNN), and Long Short-term Memory Neural Network (LSTM). LR, RF, and AB use the standard implementation of scikit-learn 0.22. CNN and LSTM were implemented using Tensorflow and Keras. CNN was composed of two Convolution1D/MaxPooling1D layers and one output layer computing the softmax function. LSTM uses one LSTM hidden layer of 128 units and one output layer computing the softmax function. We used both binary cross-entropy as loss functions and trained them over 50 epochs. For the LSTM, we used 32, 64, 128, 256, and 512 units, and 128 performed the best. Moreover, after 50 epochs the accuracy started dropping.

We ran each of the classifiers on research papers (R), patents (P), and the 15 possible combinations of the other four-time series (RA, RI, PA, PI) to assess which set of features would yield the best results. We performed 10-fold cross-validation of the data and measured the performance of the classifiers by computing the average precision (P), recall (R), and F1 (F). The data set, the results of experiments, the parameters, the implementation details, and the best models are available at <https://w3id.org/aida/downloads>.

Table 8 shows the results of our experiment. LSTM outperforms all the other solutions, yielding the highest F1 for 12 of the 17 feature combinations and the highest average F1 (73.7%). CNN (72.8%) and AB (72.3%) also produce competitive results. The reader notices that our main goal was to show that the combination of the four time series (number of papers from academia, number of papers from industry, number of patents from academia, and number of patents from industry) improves the performance of all the predictors. This proves that the granular representation of documents in AIDA yields significant advantages to these systems.

We can observe that using the combination (RA-RI-PI) significantly ($p < 0.0001$) outperforms (F1: 84.7%) the version which uses only the number of patents by companies

Table 8. Performance of the five classifiers on 17 combinations of time series. Bold indicates the best F1 (F) for each combination. The table and the experiments were previously reported in Salatino et al. (2020b)

	LR			RF			AB			CNN			LSTM		
	P%	R%	F%	P%	R%	F%	P%	R%	F%	P%	R%	F%	P%	R%	F%
RA	70.8	45.2	55.2	63.3	55.8	59.2	66.0	58.4	61.9	64.1	66.3	65.0	65.2	64.2	64.6
RI	83.5	67.1	74.4	78.9	69.8	74.0	80.0	73.1	76.4	79.2	75.1	77.0	79.1	74.8	76.9
PA	58.3	15.3	24.2	60.4	15.4	24.5	59.3	16.0	25.2	60.5	15.7	24.9	60.8	15.6	24.8
PI	76.5	69.0	72.5	73.9	68.4	71.0	75.6	71.8	73.6	73.7	76.6	75.0	74.1	76.6	75.2
R	73.7	48.8	58.7	65.5	59.7	62.5	68.6	63.1	65.6	67.6	69.2	68.3	67.2	69.4	68.2
P	76.5	68.6	72.3	72.8	67.6	70.0	74.4	71.6	73.0	73.2	76.1	74.6	73.1	76.6	74.8
RA, RI	85.7	70.9	77.6	80.5	76.0	78.2	82.6	76.6	79.5	78.9	75.1	76.8	82.2	79.3	80.7
RA, PA	70.3	47.0	56.3	63.1	55.5	59.0	66.5	59.3	62.6	64.5	65.1	64.5	65.4	64.2	64.6
RA, PI	79.6	73.7	76.5	77.2	74.3	75.7	79.1	76.5	77.7	75.2	76.3	75.7	77.4	81.9	79.5
RI, PA	83.3	67.0	74.3	77.9	70.8	74.1	79.6	73.0	76.1	78.6	75.6	77.0	79.1	75.2	77.1
RI, PI	83.4	77.3	80.2	81.0	77.3	79.1	82.7	78.6	80.6	82.0	78.6	80.2	81.7	81.2	81.4
PA, PI	76.7	68.6	72.4	74.2	69.0	71.5	75.9	71.5	73.6	71.1	70.8	70.9	73.8	76.7	75.2
RA, RI, PA	85.2	71.4	77.7	80.8	75.4	78.0	82.5	77.0	79.6	82.6	78.1	80.3	82.6	78.2	80.3
RA, RI, PI	85.4	79.8	82.5	84.5	80.5	82.4	84.6	81.2	82.9	83.8	84.7	84.2	84.1	85.4	84.7
RA, PA, PI	79.6	73.9	76.6	77.5	74.4	75.9	79.2	76.5	77.8	78.9	78.6	78.6	77.4	81.4	79.2
RI, PA, PI	83.6	77.5	80.4	81.1	78.0	79.5	82.7	78.6	80.6	82.2	80.9	81.5	81.1	81.0	81.1
RA, RI, PA, PI	85.4	79.8	82.5	83.8	80.0	81.8	84.6	81.2	82.9	84.7	81.3	82.9	83.2	86.1	84.6

(74.8%). PA (academic patents) is the weakest of all the indicators, probably because there is a very small number of academic patents. Considering the origin (academia and industry) of the publications and the patents also increases performance: RA-RI (80.7%) significantly ($p < 0.0001$) outperforms R (68.2%) and PA-PI (75.2%) is marginally better than P (74.8%). This confirms that the most granular representation of the document origin in AIDA can increase the forecaster performance.

Another interesting outcome is that, when considering only one of the time series, the number of publications from industry (RI) is a significantly ($p = 0.004$) better indicator than patents from industry (PI), yielding an F1 of 76.9%, followed by RA, and PA. The best combination of two time series is RI-PI (81.4%), while the best combination of three time series is RA-RI-PI (84.7%).

In conclusion, the experiments substantiate the hypothesis that the granular representation of publications and patents in AIDA can effectively support deep learning approaches for

forecasting the impact of research topics on the industrial sector. It also validates the intuition that including features from research articles can be very useful when predicting industrial trends.

5. AIDA USAGE

To test AIDA's ability to generate advanced analytics, in the last year we generated preliminary versions of AIDA for analyzing the research trends in Computer Science. The feedback collected during these studies was used to improve the semantic schema of AIDA and the scalability of its pipeline. We summarize here the main results of these research efforts. Specifically, in Section 5.1 we report a study about topic dynamics across publications and patents from academia and industry (Salatino et al., 2020b) that used an initial version of AIDA focused on the main 5,000 topics in Computer Science. In Section 5.2 we present an analysis of the main research trends among papers published in two main venues of Human-Computer Interaction (HCI) (Mannocci, Osborne, & Motta, 2019). To further showcase AIDA ability to support tools for analyzing the research landscape, in Section 5.3 we describe the *AIDA Dashboard*, a new web application based on AIDA that we developed to support Springer Nature editors in assessing the quality of scientific conferences.

5.1. Analyzing Academia Industry Relationship

Monitoring the research trends across articles and patents can lead to a deeper understanding of the knowledge flow between academia and industry. In our recent study (Salatino et al., 2020b), we used an initial version of AIDA to represent a set of 5,000 topics in CSO according to four time series reporting the time frequency of papers from academia; papers from industry; patents from academia; and patents from industry. We then analyzed the resulting time series to identify insightful patterns.

Figure 6 shows the distribution of these topics in a bidimensional diagram according to two indexes: academia-industry (horizontal axis) and papers-patents (vertical axis). The papers-patents index of a certain topic t is the difference between the number of research papers R_t ,




Figure 6. Distribution of the most frequent 5,000 topics according to their academia-industry and publication-papers indexes (Salatino et al., 2020b).

and patents P_t related to t , over the whole set of documents $(R_t + P_t)$: $(R_t - P_t)/(R_t + P_t)$. If this index is positive, a topic tends to be associated with a higher number of publications, while if it is negative, with a higher number of patents. On the other hand, the academia-industry index for a certain topic t is the difference between the documents in academia A_t and industry I_t , over the whole set of documents $(R_t + P_t)$: $(A_t - I_t)/(R_t + P_t)$. If this index is positive, a topic tends to be mostly associated with academia, if it is negative, with industry.

As we can observe from Figure 6, topics are tightly distributed around the bisector: the ones that attract more interest from academia are prevalently associated with publications (top-right quadrant), while the ones in industry are mostly associated with patents (bottom left quadrant).

We also performed an analysis of the emergence of topics across the four time series. In particular, we determined when a topic emerges in all time series, and compared the time elapsed between each pair of them. To avoid false positives, we considered a topic as “emerged” when it was associated with at least 10 documents. Our results showed that 89.8% of the topics first emerged in academic publications, 3.0% in industrial publications, 7.2% in industrial patents, and none in academic patents. On average, publications from academia preceded publications from industry by 5.6 ± 5.6 years, and in turn, the latter preceded patents from industry by 1.0 ± 5.8 years, as showed in Figure 7. Publications from academia also preceded by 6.7 ± 7.4 years patents from industry. This outcome is consistent with previous studies which identified academia as the main creator of new knowledge (Larivière et al., 2018), but it is able to quantify much more accurately when specific research topics emerge. More details about this analysis are available in Salatino et al. (2020b).

5.2. Detecting Research Trends

A preliminary version of AIDA focusing only on publications in Human-Computer Interaction (HCI) in 1969–2018 was used to perform an analysis of the field that was published on the special issue of the *International Journal of Human-Computer Studies* celebrating 50 years of the journal (Mannocei et al., 2019). The analysis focuses on two main venues of HCI: the *International Journal of Human-Computer Studies* (IJHCS) and the Conference on Human Factors in Computing Systems (CHI). The resulting data reporting the evolution of topics were analyzed with the help of domain experts to detect the most prominent topics in various time frames and the most significant trends in the last 10 years. We briefly report the main results as they are an excellent example of the bibliographic analyses that AIDA can support.

Figure 8 compares the percentage of publications tagged with the main topics in IJHCS (blue) and CHI (orange). It was created by computing the percentage of publications associated with the same research topics in the preliminary version of AIDA. The two top venues in HCI tend to address a similar set of topics but also present some intriguing differences. For

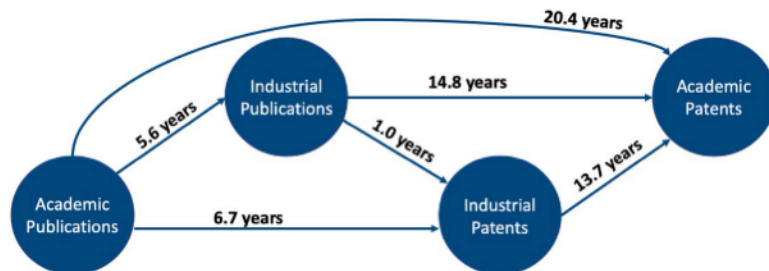


Figure 7. Average time lags when analyzing the emergence of topics through their four time series.




Figure 8. Comparison of the main research topics of IJHCS and CHI during 1960–2018.

instance, IJHCS has a more interdisciplinary focus, and in particular, it addresses several topics related to Artificial Intelligence such as Knowledge-Based Systems, Knowledge Management, Formal Languages, and Natural Language Processing. This outcome was also confirmed by the editors of IJHCS.

Figure 9 shows the main emerging topics in the two venues under analysis. These were the topics that experienced the steepest improvement in terms of the number of associated articles

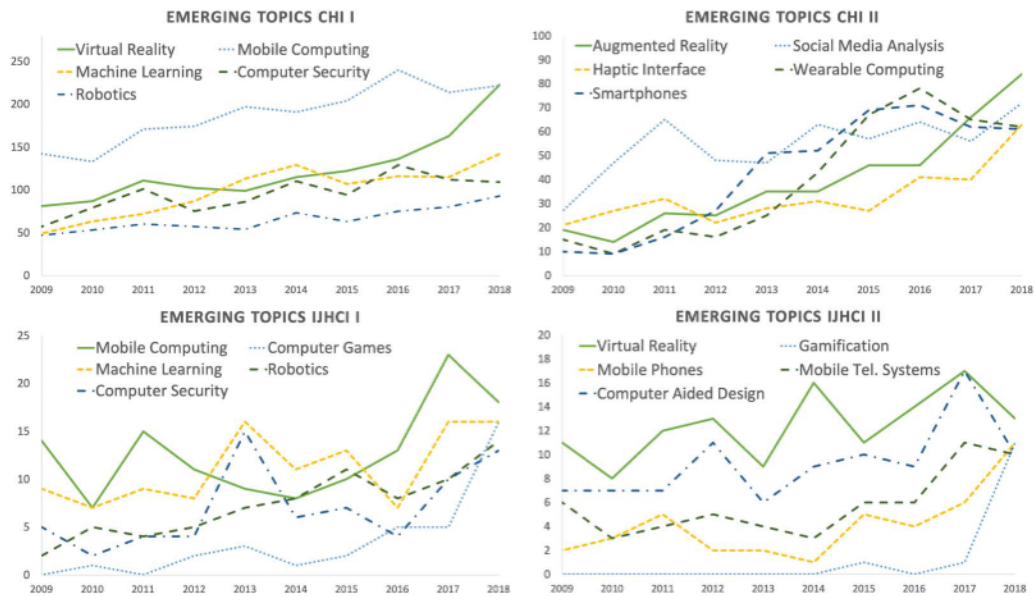


Figure 9. Emerging topics in IJHCS and CHI during 2009–2018.

in the decade 2009–2018. AIDA allows users to compute these analytics by simply querying and aggregating the relevant data. In this instance, we can easily detect that the emerging research trends of HCI in the last years include Virtual Reality, Mobile Computing, Robotics, Haptic Interfaces, Social Media Analysis, and Gamifications. A more comprehensive analysis of these trends is available in Mannocci et al. (2019).

5.3. The AIDA Dashboard: Assessing Scientific Conferences

Scientific conferences play a crucial role in the field of Computer Science by offering high-quality venues for research articles, promoting new collaborations, and connecting research efforts from academia and industry. Understanding and monitoring conferences is thus a crucial task for researchers, editors, funding bodies, and other users in this space. While several academic search engines (e.g., Microsoft Academic Graph, Semantic Scholar, Scopus) provide basic information about conferences, they do not offer advanced analytics to rank and compare them, assess their main trends, or study their involvement with specific industrial sectors.

To address these limitations, we created the *AIDA Dashboard*, a new web application that takes advantage of AIDA for supporting users in analyzing scientific conferences. The AIDA Dashboard was developed in collaboration with Springer Nature, with the primary objective of supporting their team in assessing the quality of a conference in order to inform editorial decisions. However, the analyses supported by the AIDA Dashboard can assist several other stakeholders, including researchers and funding bodies. Specifically, the AIDA Dashboard introduces three novel features that state-of-the-art systems currently lack. First, it characterizes conferences according to the granular representation of topics from AIDA, hence providing high-quality analytics about their research trends over time. Second, it enables users to easily compare conferences in the same fields according to several bibliometrics. Third, it allows users to assess the involvement of commercial organizations in a conference by offering analytics about the academia/industry collaborations and the relevant industrial sectors.

The AIDA Dashboard describe each conference according to eight tabs: *Overview*, *Citation Analysis*, *Organizations*, *Countries*, *Authors*, *Topics*, *Similar Conferences*, and *Industry*. The *Overview* tab (see Figure 10) summarizes the most important information with the aim of allowing the user to immediately understand what the conference is about and how it has performed in the last few years. The *Citation Analysis* tab reports several citation-based bibliometrics and highlights how the conference ranks in its main research areas. The *Authors*, *Organizations*, and *Countries* tabs enable users to analyze the actors that produced the articles at different levels of granularity (researchers, institutions, and geographical locations). The *Topics* tab allows users to inspect the main research topics and analyze their trends in time. The *Similar Conferences* tab compares the conference under analysis with all the other conferences in the same fields according to different bibliometrics. Finally the *Industry* tab reports the percentage of articles and citations from academia, industry, and collaborative efforts as well as the frequency of the industrial sectors from AIDA.

The AIDA Dashboard is still under development and we aim to release a first stable version in the second part of 2021. A demo of the current prototype is available at <https://aida.kmi.open.ac.uk/dashboard/>.

To showcase the functionalities of the AIDA Dashboard, Figures 10–14 illustrate some of the analytics generated for one of the main conferences in the field of Neural Networks: the *Neural Information Processing Systems Conference (NeurIPS)*.

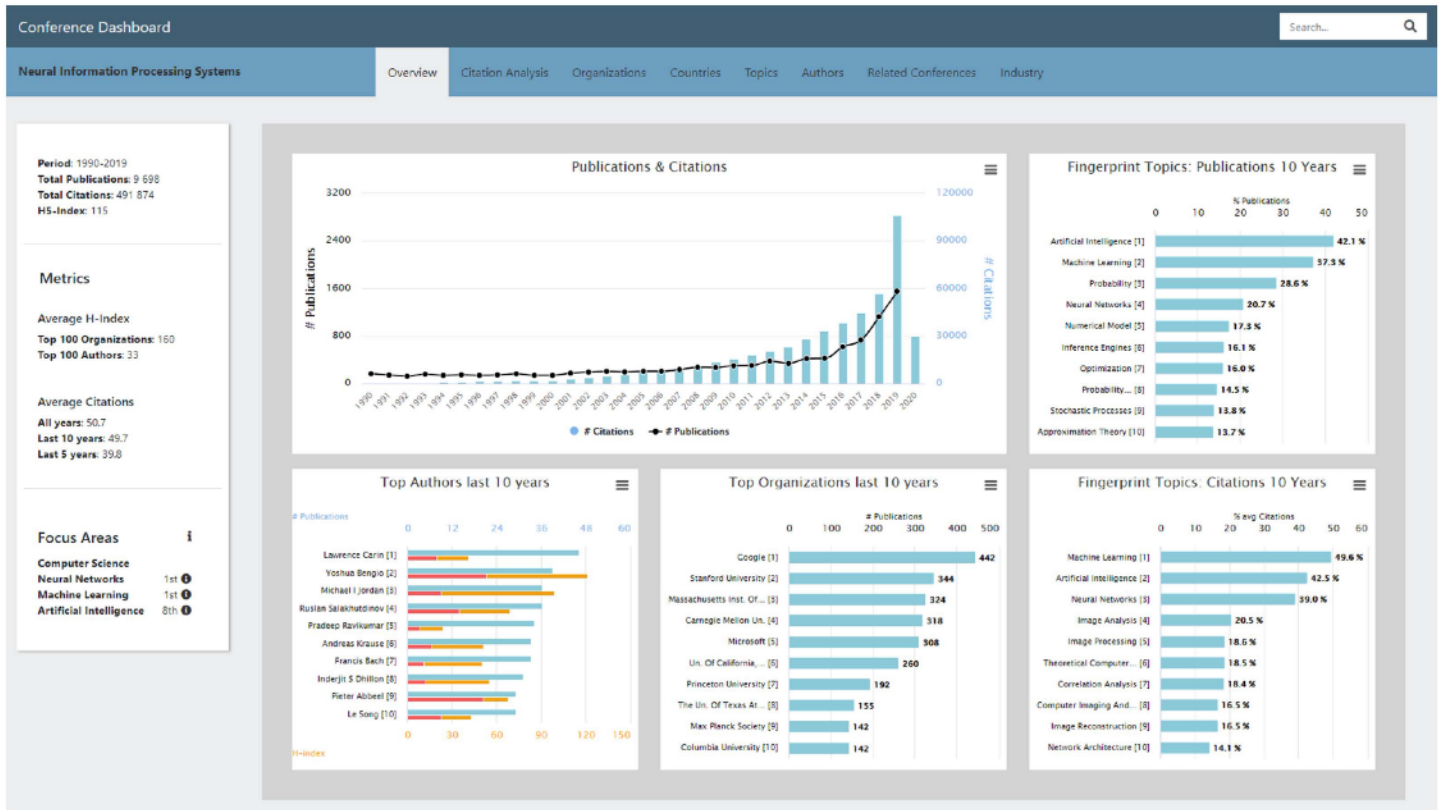


Figure 10. The overview page of the NeurIPS conference according to the AIDA Dashboard.




Figure 11. The rank of NeurIPS in its three main focus areas (neural networks, machine learning, artificial intelligence) across time. The conferences are ranked according to their average citations per article.

The users can search any conference from the main page. After they select a conference (e.g., NeurIPS) they are redirected to its *Overview* tab. Figure 10 shows the *Overview* tab of NeurIPS, which displays several pieces of high-level information, including basic bibliometrics and the main authors, organizations, and topics. We can note the presence of organizations such as Google, Stanford, and MIT and of a Turing Award winner (Yoshua Bengio) and many world-leading researchers in neural networks in the main authors. At the bottom left side, the AIDA Dashboard reports the *focus areas* of NeurIPS: *Neural Networks*, *Machine Learning*, and *Artificial Intelligence*. These are high-level fields used to categorize and compare conferences. They are computed automatically by analyzing the topic distribution of the conference in AIDA.

The line chart in Figure 11, from the *Citation Analysis* tab, shows how NeurIPS ranks in terms of average citations per paper in the three focus areas. In the last 10 years, NeurIPS has always been rippling between the first and second position in the fields of neural networks and machine learning.

The plot in Figure 12 is from in the *Topics* tab and shows the topics that received most citations in the conference. In addition to the focus areas of the conference (Neural Networks, Machine Learning, Artificial Intelligence) we can see many other relevant high-level topics (e.g., Mathematics, Probability, Signal Processing) as well as some important domains of application (e.g., Image Processing, Human Computer Interaction).

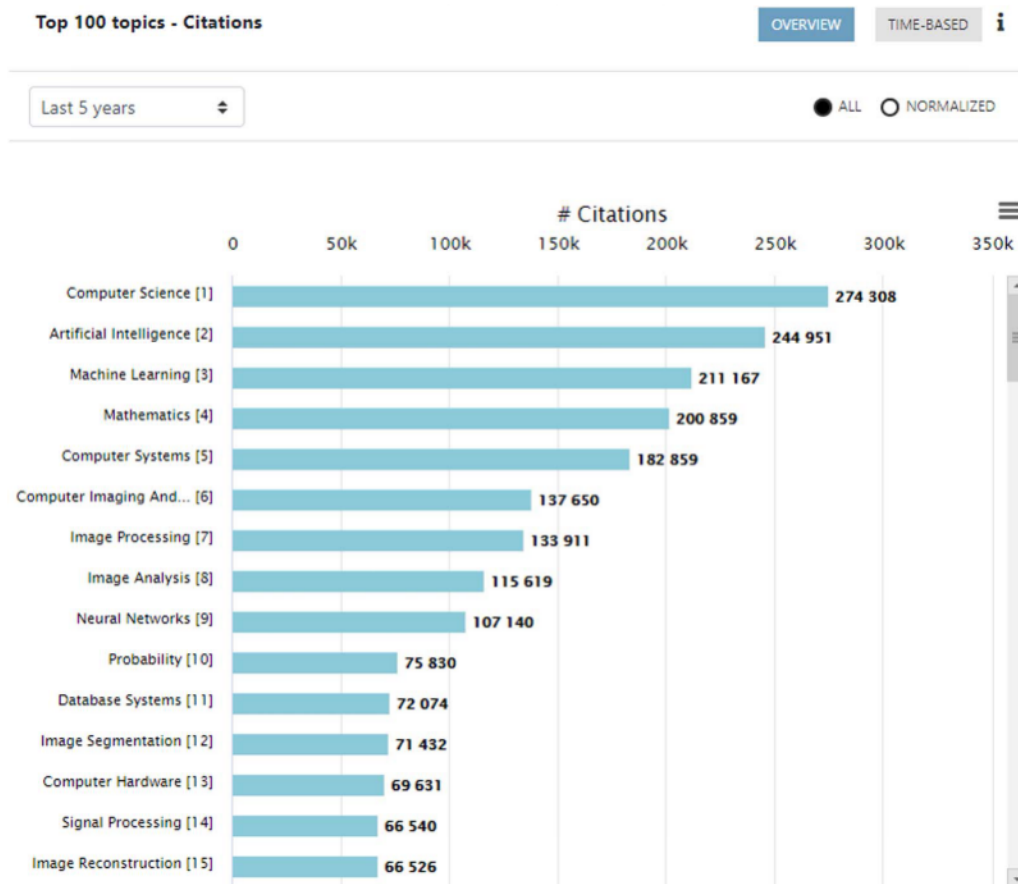


Figure 12. The most cited topics in NeurIPS during the last 5 years.

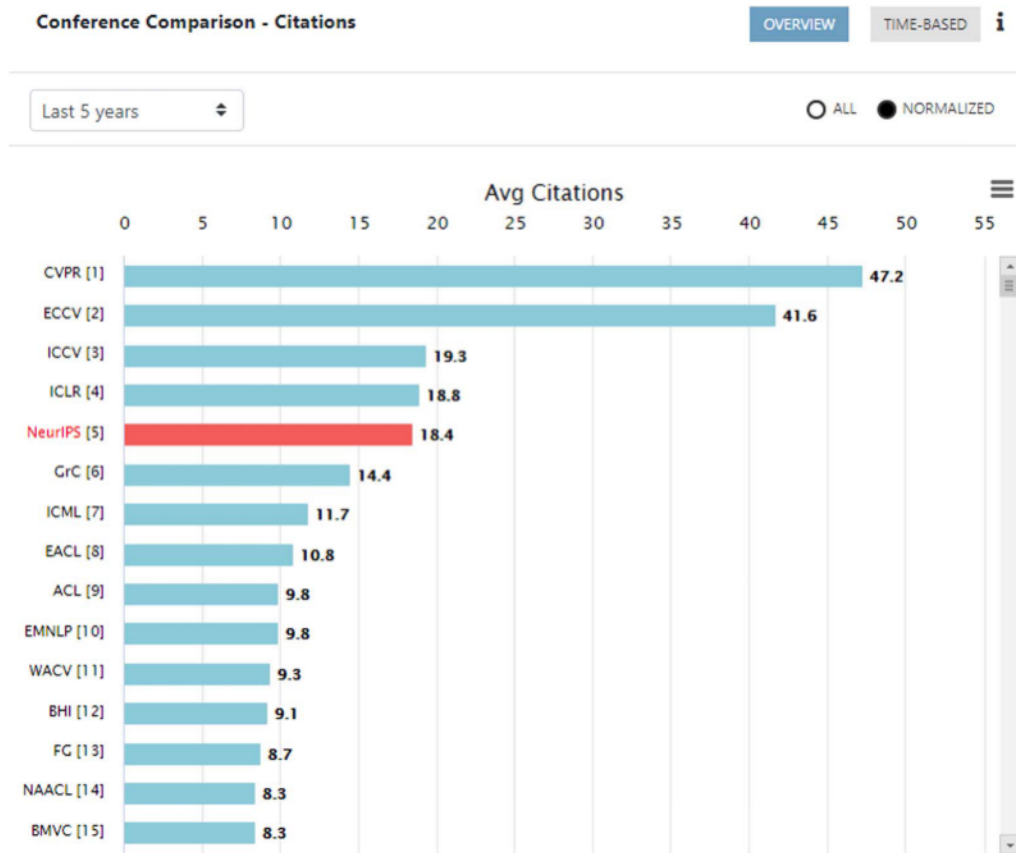


Figure 13. The best Artificial Intelligence conferences in terms of average citations in the last 5 years. NeurIPS is in fifth position, highlighted in red.

Figure 13, from the *Related Conferences* tab, shows the comparison between NeurIPS and all the other conferences in Artificial Intelligence in terms of average citations in the last 5 years. As we can see, NeurIPS ranks fifth with an average of 18.4 citations for article.

Finally, the bar chart in Figure 14, from the *Industry* tab, shows the percentages of the published articles relevant to several industrial sectors from the INDUSO ontology. For NeurIPS, 96.3% of the articles are from Computing and IT, 27% from Electronics, 9.7% from Information Technology, and so on. The *Industry* tab also shows the frequencies of articles published by authors exclusively from academia; authors exclusively from industry; and from a joint collaboration of authors from both academia and industry. In Table 9 we report the percentage of articles based on their affiliation. While most articles are from academia, the percentage of industrial and collaborative articles is significantly higher in the last 5 years, suggesting a growing interest by commercial organizations. The overview page, shown in Figure 10, shows some of the companies involved in this shift. The user can also use the *Organizations* tab to display in a line chart the growing number of publications associated with commercial organizations such as Google, Microsoft, IBM, and Facebook.

6. LIMITATIONS

In this section, we discuss some limitations of the current pipeline, and describe our plans to address them in the future.

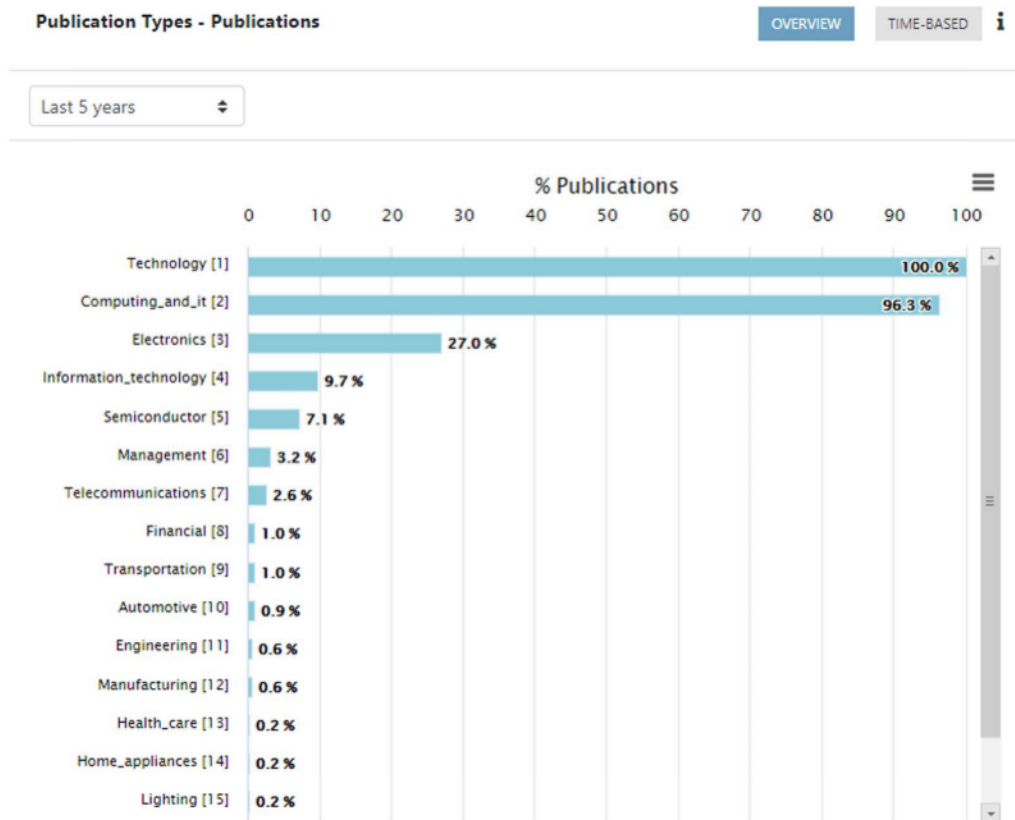


Figure 14. Most frequent industrial sectors in NeurIPS during the last 5 years.

A first challenge regards improving the scalability. A significant bottleneck of the current version is that it uses the DBpedia REST API for identifying industrial sectors. This solution relies on REST requests on the web and therefore it is quite slow. We plan to switch to a local DBpedia instance to solve this issue. In addition, we are currently working on a new version of the CSO Classifier that uses a smarter cache in the semantic module to improve scalability. We believe that these changes may be able to cut the computational time by half or more.

A second limitation regards the fact that only a subset of the documents (5.1 million articles and 5.6 million patents) are mapped to GRID and can thus be assigned with the types of affiliations and industrial sectors. We plan to address this issue from different directions. First, we intend to directly map the names of the organizations to DBpedia and knowledge bases of companies using entity-linking solutions. We are also working on link prediction techniques for graph completion that can be used to automatically classify the affiliations according to contextual information in the knowledge graph. An interesting challenge in this regard is that AIDA contains several N to M relations with $N \gg M$. Given a triple (h, r, t) , this situation arises

Table 9. Percentages of articles written by Academia/Industry/Collaborative in NeurIPS

	All years	Last 5 years
Academia	80.48%	71.59%
Industry	5.40%	6.61%
Collaborative	14.11%	21.79%

when the cardinality of the entities in the head position (h) for a certain relation (r) is much higher than the one of the entities in the tail position (t). This is actually the case for most scholarly knowledge graphs (Ammar et al., 2018; Knoth & Zdrahal, 2011; Peroni & Shotton, 2020; Wang et al., 2020; Zhang et al., 2018) that usually categorize millions of documents (e.g., papers, patents) according to a relatively small set of categories (e.g., topics, countries, chemical compounds). Another important requirement is the scalability of these methods, because we need to be able to process million of entities. We are thus focusing on the creation of link prediction approaches that perform well in this space. The first output of this research line was Trans4E (Nayyeri, Cil et al., 2021), a scalable model which tackles these issues by providing a very large number of possible vectors ($8^d - 1$, where d is the embedding dimension) to be assigned to entities involved in N to M relations.

A final important limitation is that the current version of the pipeline uses MAG as source for research articles. Unfortunately, during the writing of this paper, Microsoft decided to decommission the MAG project after 2021³⁵. To react in a timely manner, we worked on this issue with Springer Nature data science team and devised a strategy to obtain the article metadata from Dimensions. We chose this knowledge graph due to its wide coverage of Computer Science and low cost of integration (AIDA already uses Dimensions for patents). As Dimensions does not disambiguate conferences, we also plan to leverage the conference representation of DBLP, which currently includes 5,438 conferences in Computer Science. Preliminary experiments show that most conferences available in MAG are also covered by DBLP. We plan to integrate Dimensions and DBLP using the paper DOIs. For the few conferences and workshops that do not assign DOIs to articles, we will map the papers across the two data sets by computing the string similarity of their titles and authors, after applying filters that normalize, uniform cases, and remove punctuation. We will also leverage additional fields, such as the year of publication and the proceedings title, to reduce the number of papers to compare and provide further confirmation of the alignments. We plan to switch to this new solution before the end of 2021.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we have introduced AIDA, the Academic/Industry DynAmics Knowledge Graph. This resource characterizes 21 million publications and 8 million patents according to the research topics drawn from the Computer Science Ontology (CSO). 5.1 million publications and 5.6 million patents are also classified according to the type of the author's affiliations and industrial sectors. To characterize documents according to their industrial sectors, we designed the Industrial Sectors Ontology (INDUSO), which describes 66 sectors in a two-level taxonomy.

AIDA was generated using an automatic pipeline that merges and integrates information from Microsoft Academic Graph, Dimensions, DBpedia, the Computer Science Ontology, and the Global Research Identifier Database. It allows researchers to analyze the evolution of research topics across academia and industry as well as to understand their dynamics within several industrial sectors. It can be used to identify the research trends of different industries and how and when academia and/or industry tackle these in particularly significant ways, thus facilitating a granular analysis of the interaction between these two worlds. Moreover, AIDA can also be employed to investigate authors, citations, countries, and other entities already present in Microsoft Academic Graph.

³⁵ Next Steps for Microsoft Academic – Expanding into New Horizons: <https://www.microsoft.com/en-us/research/project/academic/articles/microsoft-academic-to-expand-horizons-with-community-driven-approach/>

To showcase how AIDA can be used by the wider community, we also presented some exemplary studies that take advantage of AIDA for producing advanced bibliometric analysis and introduced the AIDA Dashboard, a novel tool that aims to support Springer Nature editors in assessing the quality of scientific conferences.

The process for producing AIDA is general and can be applied to other domains of science. In this case, the CSO Classifier, which is the main computer science-specific portion of our pipeline, needs to be tailored to the new field. To do so, it is necessary to replace CSO with a different domain ontology and retrain the word2vec model with a corpus of documents that fits the new domain. This procedure is detailed in <https://doi.org/10.5281/zenodo.3459286>.

We evaluated different parts of the pipeline using a manually created gold standard and obtaining very competitive results. We also evaluated the impact of AIDA on forecasting systems for predicting the impact of research trends on the industry. In particular, we found that a forecaster based on LSTM neural networks and exploiting the full representation of articles and patents from AIDA yielded significantly better performance ($p < 0.0001$) than alternative methods. In addition, the version of this classifier using the full set of features (84.6%) gained almost 10% in terms of F1 in comparison with the one using only the number of patents across time (74.8%). This substantiates the hypothesis that adopting a more granular representation of articles and patents is critical for this task.

The resource presented in this paper opens up several interesting directions of work. First, we will produce a comprehensive analysis of AIDA and the most significant research trends in academia and industry. We also intend to use AIDA to support systems for predicting the impact of specific areas of industry research.

We plan to further improve AIDA using graph completion and link prediction techniques. As many state-of-the-art solutions in this space may suffer when dealing with knowledge graphs that categorize a very large number of entities (e.g., research articles, patents, persons), we are currently investigating new scalable approaches that can deal with this situation (Nayyeri et al., 2021). We are also exploring the possibility of using other knowledge graphs, such as Wikidata and BabelNet, to further improve the performance of graph completion techniques on AIDA.

We plan to explore the application of our pipeline to other fields, such as Biology and Engineering. To this purpose we intend to develop a new version of our classifier, testing also a range of recent word embeddings solutions, such as BERT and SciBERT. One more direction regards a further classification of papers into peer reviewed and not peer reviewed.

As far as the dashboard is concerned, we are currently performing a comprehensive evaluation with different kinds of users and will make available the results in a future paper. Finally, we are going to employ AIDA for human-robot interaction and develop a robot that can answer questions about the scholarly domain in natural language.

ACKNOWLEDGMENTS

We gratefully acknowledge Springer Nature for funding this research. We also thank Dimensions for sharing their large data set of patents. Last but not least, we gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research.

AUTHOR CONTRIBUTIONS

Simone Angioni: Data curation, Formal analysis, Writing—original draft. Angelo Salatino: Data curation, Formal analysis, Writing—original draft. Francesco Osborne: Formal analysis,

Project administration, Writing—original draft. Diego Reforgiato Recupero: Formal analysis, Project administration, Validation, Writing—original draft. Enrico Motta: Project administration, Supervision, Writing—review & editing.

COMPETING INTERESTS

The authors have no competing interests.

FUNDING INFORMATION

We gratefully acknowledge Springer Nature for funding this research. We also thank NVIDIA Corporation for the donation of the Titan X GPU used for this research.

DATA AVAILABILITY

The whole AIDA dataset can be downloaded from <http://w3id.org/aida> under CC BY 4.0 license.

REFERENCES

- Altuntas, S., Dereli, T., & Kusiak, A. (2015). Analysis of patent documents with weighted association rules. *Technological Forecasting and Social Change*, *92*, 249–262. <https://doi.org/10.1016/j.techfore.2014.09.012>
- Ammar, W., Groeneveld, D., Bhagavatula, C., Beltagy, I., Crawford, M., ... Etzioni, O. (2018). Construction of the literature graph in semantic scholar. *arXiv preprint arXiv:1805.02262*.
- Anderson, M. S. (2001a). The complex relations between the academy and industry: Views from the literature. *The Journal of Higher Education*, *72*(2), 226–246. <https://doi.org/10.1080/00221546.2001.11778879>
- Anderson, M. S. (2001b). The complex relations between the academy and industry: Views from the literature. *The Journal of Higher Education*, *72*(2), 226–246. Retrieved from <https://www.jstor.org/stable/2649323>. <https://doi.org/10.1080/00221546.2001.11778879>
- Angioni, S., Salatino, A. A., Osborne, F., Recupero, D. R., & Motta, E. (2020). Integrating knowledge graphs for analysing academia and industry dynamics. In *ADBIS, TPDL, and EDA 2020 Common Workshops and Doctoral Consortium* (pp. 219–225). https://doi.org/10.1007/978-3-030-55814-7_18
- Ankrah, S., & Omar, A.-T. (2015). Universities–industry collaboration: A systematic review. *Scandinavian Journal of Management*, *31*(3), 387–408. <https://doi.org/10.1016/j.scaman.2015.02.003>
- Ankrah, S. N., Burgess, T. F., Grimshaw, P., & Shaw, N. E. (2013). Asking both university and industry actors about their engagement in knowledge transfer: What single-group studies of motives omit. *Technovation*, *33*(2–3), 50–65. <https://doi.org/10.1016/j.technovation.2012.11.001>
- Beck, M., Rizvi, S. T. R., Dengel, A., & Ahmed, S. (2020). From automatic keyword detection to ontology-based topic modeling. In *International Workshop on Document Analysis Systems* (pp. 451–465). https://doi.org/10.1007/978-3-030-57058-3_32
- Belleau, F., Nolin, M.-A., Tourigny, N., Rigault, P., & Morissette, J. (2008). Bio2RDF: Towards a mashup to build bioinformatics knowledge systems. *Journal of Biomedical Informatics*, *41*(5), 706–716. <https://doi.org/10.1016/j.jbi.2008.03.004>, PubMed: 18472304
- Bikard, M., Vakili, K., & Teodoridis, F. (2019). When collaboration bridges institutions: The impact of university–industry collaboration on academic productivity. *Organization Science*, *30*(2), 426–445. <https://doi.org/10.1287/orsc.2018.1235>
- Bird, S., Dale, R., Dorr, B. J., Gibson, B., Joseph, M. T., ... Tan, Y. F. (2008). The ACL anthology reference corpus: A reference dataset for bibliographic research in computational linguistics.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, *3*(null), 993–1022.
- Borges, M. V. M., & dos Reis, J. C. (2019). Semantic-enhanced recommendation of video lectures. In *2019 IEEE 19th International Conference on Advanced Learning Technologies (ICALT)* (Vol. 2161, pp. 42–46). <https://doi.org/10.1109/ICALT.2019.00013>
- Chatzopoulos, S., Vergoulis, T., Kanellos, I., Dalamagas, T., & Tryfonopoulos, C. (2020a). ArtSim: Improved estimation of current impact for recent articles. In *ADBIS, TPDL, and EDA 2020 Common Workshops and Doctoral Consortium* (pp. 323–334). https://doi.org/10.1007/978-3-030-55814-7_27
- Chatzopoulos, S., Vergoulis, T., Kanellos, I., Dalamagas, T., & Tryfonopoulos, C. (2020b). ArtSim: Improved estimation of current impact for recent articles. In L. Bellatreche et al. (Eds.), *ADBIS, TPDL, and EDA 2020 Common Workshops and Doctoral Consortium* (pp. 323–334). Cham: Springer. https://doi.org/10.1007/978-3-030-55814-7_27
- Chicaiza, J., & Reátegui, R. (2020). Using domain ontologies for text classification. A use case to classify computer science papers. In *Iberoamerican Knowledge Graphs and Semantic Web Conference* (pp. 166–180). https://doi.org/10.1007/978-3-030-65384-2_13
- Choi, S., & Jun, S. (2014). Vacant technology forecasting using new bayesian patent clustering. *Technology Analysis & Strategic Management*, *26*(3), 241–251. <https://doi.org/10.1080/09537325.2013.850477>
- Chung, P., & Sohn, S. Y. (2020). Early detection of valuable patents using a deep learning model: Case of semiconductor industry. *Technological Forecasting and Social Change*, *158*, 120146. <https://doi.org/10.1016/j.techfore.2020.120146>
- Costa, J. P., Rei, L., Stopar, L., Fuat, F., Grobelnik, M., ... Wallace, J. (2021). Newsmesh: A new classifier designed to annotate health news with mesh headings. *Artificial Intelligence in Medicine*, *114*, 102053. Retrieved from <https://www.sciencedirect.com/science/article/pii/S09333365721000464>. <https://doi.org/10.1016/j.artmed.2021.102053>, PubMed: 33875160
- Deng, W., Huang, X., & Zhu, P. (2019). Facilitating technology transfer by patent knowledge graph. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*.

- Dessì, D., Osborne, F., Recupero, D. R., Buscaldi, D., & Motta, E. (2021). Generating knowledge graphs by employing natural language processing and machine learning techniques within the scholarly domain. *Future Generation Computer Systems*, 116, 253–264. <https://doi.org/10.1016/j.future.2020.10.026>
- Dörpinghaus, J., & Jacobs, M. (2020). Knowledge detection and discovery using semantic graph embeddings on large knowledge graphs generated on text mining results. In *2020 15th Conference on Computer Science and Information Systems (FedCSIS)* (pp. 169–178). <https://doi.org/10.15439/2020F36>
- Färber, M. (2019). The Microsoft Academic Knowledge Graph: A linked data source with 8 billion triples of scholarly data. In *International Semantic Web Conference* (pp. 113–129). https://doi.org/10.1007/978-3-030-30796-7_8
- Fathalla, S., Auer, S., & Lange, C. (2020). Towards the semantic formalization of science. In *Proceedings of 35th Annual ACM Symposium on Applied Computing* (pp. 2057–2059).
- Grimpe, C., & Hussinger, K. (2013). Formal and informal knowledge and technology transfer from academia to industry: Complementarity effects and innovation performance. *Industry and Innovation*, 20(8), 683–700. <https://doi.org/10.1080/13662716.2013.856620>
- Groth, P., Gibson, A., & Velterop, J. (2010). The anatomy of a nanopublication. *Information Services & Use*, 30(1–2), 51–56. <https://doi.org/10.3233/ISU-2010-0613>
- Hanieh, A. A., Abdelal, S., Krajnik, P., & Hasan, A. (2015). Industry-academia partnership for sustainable development in Palestine. *Procedia CIRP*, 26, 109–114. <https://doi.org/10.1016/j.procir.2014.07.184>
- Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G. D., ... Zimmermann, A. (2021). Knowledge graphs. *ACM Computing Surveys (CSUR)*, 54(4), 1–37. <https://doi.org/10.1145/3447772>
- Huang, M.-H., Yang, H.-W., & Chen, D.-Z. (2015). Industry-academia collaboration in fuel cells: A perspective from paper and patent analysis. *Scientometrics*, 105(2), 1301–1318. <https://doi.org/10.1007/s11192-015-1748-6>
- Jaradeh, M. Y., Auer, S., Prinz, M., Kovtun, V., Kismihók, G., & Stocker, M. (2019). Open research knowledge graph: Towards machine actionability in scholarly communication. *arXiv preprint arXiv:1901.10816*.
- Jose, V., Jagathy Raj, V. P., & George, S. K. (2021). Ontology-based information extraction framework for academic knowledge repository. In X.-S. Yang, S. Sherratt, N. Dey, & A. Joshi (Eds.), *Proceedings of Fifth International Congress on Information and Communication Technology* (pp. 73–80). Singapore: Springer Singapore. https://doi.org/10.1007/978-981-15-5859-7_6
- Knoth, P., & Zdrahal, Z. (2011). CORE: Connecting repositories in the open access domain. In *CERN Workshop on Innovations in Scholarly Communication (OAI7)*. Retrieved from <https://oro.open.ac.uk/32560/> (Poster Session ID: 53).
- Knoth, P., & Zdrahal, Z. (2012). CORE: Three access levels to underpin open access. *D-Lib Magazine*, 18(11/12), 1–13. <https://doi.org/10.1045/november2012-knoth>
- Kuhn, T., Chichester, C., Krauthammer, M., Queralt-Rosinach, N., Verborgh, R., ... Dumontier, M. (2016). Decentralized provenance-aware publishing with nanopublications. *PeerJ Computer Science*, 2, e78. <https://doi.org/10.7717/peerj-cs.78>
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago Press.
- La Bruzzo, S., Manghi, P., & Mannocci, A. (2019). OpenAIRE's DOIBoost - Boosting Crossref for Research. In P. Manghi, L. Candela, & G. Silvello (Eds.), *Digital libraries: Supporting open science* (pp. 133–143). Cham: Springer. https://doi.org/10.1007/978-3-030-11226-4_11
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174. <https://doi.org/10.2307/2529310>
- Larivière, V., Macaluso, B., Mongeon, P., Siler, K., & Sugimoto, C. R. (2018). Vanishing industries and the rising monopoly of universities in published research. *PLOS ONE*, 13, 1–10. <https://doi.org/10.1371/journal.pone.0202120>, PubMed: 30107002
- Ley, M. (2009). DBLP: Some lessons learned. *Proceedings of the VLDB Endowment*, 2(2), 1493–1500. <https://doi.org/10.14778/1687553.1687577>
- Löffler, F., Wesp, V., Babalou, S., Kahn, P., Lachmann, R., ... König-Ries, B. (2020). Scholarlensviz: A visualization framework for transparency in semantic user profiles. In K. Taylor, R. Gonçalves, F. Lecue, & J. Yan (Eds.), *Proceedings of the ISWC 2020 Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 19th International Semantic Web Conference (ISWC 2020)*, November 1–6.
- Lula, P., Dospinescu, O., Homocianu, D., & Sireteanu, N.-A. (2021). An advanced analysis of cloud computing concepts based on the computer science ontology. *Computers, Materials & Continua*, 66(3), 2425–2443. <https://doi.org/10.32604/cmc.2021.013771>
- Mannocci, A., Osborne, F., & Motta, E. (2019). The evolution of IJHCS and CHI: A quantitative analysis. *International Journal of Human-Computer Studies*, 131, 23–40. <https://doi.org/10.1016/j.ijhcs.2019.05.009>
- Marinakakis, Y. D. (2012). Forecasting technology diffusion with the Richards model. *Technological Forecasting and Social Change*, 79(1), 172–179. <https://doi.org/10.1016/j.techfore.2011.02.010>
- Michaudel, Q., Ishihara, Y., & Baran, P. S. (2015). Academia–industry symbiosis in organic chemistry. *Accounts of Chemical Research*, 48(3), 712–721. <https://doi.org/10.1021/ar500424a>, PubMed: 25702529
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems – Volume 2* (pp. 3111–3119). USA: Curran Associates Inc. Retrieved from <https://dl.acm.org/citation.cfm?id=2999792.2999959>
- Nayyeri, M., Cil, G. M., Vahdati, S., Osborne, F., Rahman, M., ... Lehmann, J. (2021). Trans4E: Link prediction on scholarly knowledge graphs. *Neurocomputing*, 461, 530–542. <https://doi.org/10.1016/j.neucom.2021.02.100>
- Nuzzolese, A. G., Gentile, A. L., Presutti, V., & Gangemi, A. (2016). Semantic web conference ontology—A refactoring solution. In *European Semantic Web Conference* (pp. 84–87). https://doi.org/10.1007/978-3-319-47602-5_18
- Osborne, F., & Motta, E. (2015). Klink-2: Integrating multiple web sources to generate semantic topic networks. In M. Arenas et al. (Eds.), *The Semantic Web – ISWC 2015*. Lecture Notes in Computer Science, vol. 9366. Cham: Springer. https://doi.org/10.1007/978-3-319-25007-6_24
- Osborne, F., Muccini, H., Lago, P., & Motta, E. (2019). Reducing the effort for systematic reviews in software engineering. *Data Science*, 2(1–2), 311–340. <https://doi.org/10.3233/DS-190019>
- Osborne, F., Salatino, A., Birukou, A., & Motta, E. (2016). Automatic classification of Springer Nature proceedings with Smart Topic Miner. In P. Groth et al. (Eds.), *The Semantic Web – ISWC 2016* (pp. 383–399). Cham: Springer. https://doi.org/10.1007/978-3-319-46547-0_33
- Peroni, S., & Shotton, D. (2018). The SPAR Ontologies. In D. Vrandečić et al. (Eds.), *The Semantic Web – ISWC 2018*. Lecture Notes in Computer Science, vol. 11137. Cham: Springer. https://doi.org/10.1007/978-3-030-00668-6_8

- Peroni, S., & Shotton, D. (2020). Opencitations, an infrastructure organization for open scholarship. *Quantitative Science Studies*, 1(1), 428–444. https://doi.org/10.1162/qss_a_00023
- Powell, W. W., & Snellman, K. (2004). The knowledge economy. *Annual Review of Sociology*, 30(1), 199–220. <https://doi.org/10.1146/annurev.soc.29.010202.100037>
- Ramadhan, M. H., Malik, V. I., & Sjafrizal, T. (2018). Artificial neural network approach for technology life cycle construction on patent data. In *2018 5th International Conference on Industrial Engineering and Applications (ICIEA)* (pp. 499–503).
- Rossanez, A., dos Reis, J. C., & da Silva Torres, R. (2020). Representing scientific literature evolution via temporal knowledge. <http://ceur-ws.org/Vol-2821/paper5.pdf>
- Saier, T., & Färber, M. (2020). unarXive: A large scholarly data set with publications' full-text, annotated in-text citations, and links to metadata. *Scientometrics*, 125(3), 3085–3108. <https://doi.org/10.1007/s11192-020-03382-z>
- Salatino, A. A., Osborne, F., Birukou, A., & Motta, E. (2019a). Improving editorial workflow and metadata quality at Springer Nature. In C. Ghidini et al. (Eds.), *The Semantic Web – ISWC 2019* (pp. 507–525). Cham: Springer. https://doi.org/10.1007/978-3-030-30796-7_31
- Salatino, A. A., Osborne, F., Thanapalasingam, T., & Motta, E. (2019b). The CSO classifier: Ontology-driven detection of research topics in scholarly articles. In A. Doucet, A. Isaac, K. Golub, T. Aalberg, & A. Jatowt (Eds.), *Digital libraries for open knowledge* (pp. 296–311). Cham: Springer. https://doi.org/10.1007/978-3-030-30760-8_26
- Salatino, A. A., Thanapalasingam, T., & Mannocci, A. (2019c). *angelsalatino/cso-classifier: CSO Classifier v2.3.2*. Zenodo. <https://doi.org/10.5281/zenodo.2660819>
- Salatino, A., Osborne, F., & Motta, E. (2020a). Researchflow: Understanding the knowledge flow between academia and industry. In *Knowledge Engineering and Knowledge Management 22nd International Conference, EKAW 2020*. https://doi.org/10.1007/978-3-030-61244-3_16
- Salatino, A. A., Thanapalasingam, T., Mannocci, A., Birukou, A., Osborne, F., & Motta, E. (2020b). The computer science ontology: A comprehensive automatically-generated taxonomy of research areas. *Data Intelligence*, 2(3), 379–416. https://doi.org/10.1162/dint_a_00055
- Salatino, A. A., Thanapalasingam, T., Mannocci, A., Osborne, F., & Motta, E. (2018a). Classifying research papers with the computer science ontology. In *ISWC (p&d/industry/bluesky). CEUR Workshop Proceedings* (Vol. 2180).
- Salatino, A. A., Thanapalasingam, T., Mannocci, A., Osborne, F., & Motta, E. (2018b). The computer science ontology: A large-scale taxonomy of research areas. In D. Vrandečić, et al. (Eds.), *The Semantic Web – ISWC 2018* (pp. 187–205). Cham: Springer. https://doi.org/10.1007/978-3-030-00668-6_12
- Sarica, S., Luo, J., & Wood, K. L. (2019). Technology knowledge graph based on patent data. *arXiv:1906.00411* [cs.LG].
- Satopaa, V., Albrecht, J., Irwin, D., & Raghavan, B. (2011). Finding a “kneedle” in a haystack: Detecting knee points in system behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops* (pp. 166–171). <https://doi.org/10.1109/ICD-CSW.2011.20>
- Schneider, J., Ciccarese, P., Clark, T., & Boyce, R. D. (2014). Using the micropublications ontology and the open annotation data model to represent evidence within a drug-drug interaction knowledge base. In *CEUR Workshop Proceedings*.
- Schwartz, D. L., & Sichelman, T. (2019). Data sources on patents, copyrights, trademarks, and other intellectual property. In *Research handbook on the economics of intellectual property law*. Chichester: Edward Elgar. <https://doi.org/10.4337/9781789903997.00040>
- Shotton, D. (2009). Semantic publishing: The coming revolution in scientific journal publishing. *Learned Publishing*, 22(2), 85–94. <https://doi.org/10.1087/2009202>
- Sinha, A., Shen, Z., Song, Y., Ma, H., Eide, D., Hsu, B.-J., & Wang, K. (2015). An overview of Microsoft Academic Service (MAS) and applications. In *Proceedings of the 24th International Conference on World Wide Web* (pp. 243–246).
- Stilgoe, J. (2020). *Who's driving innovation? New Technologies and the Collaborative State*. Cham: Palgrave Macmillan. <https://doi.org/10.1007/978-3-030-32320-2>
- Thanapalasingam, T., Osborne, F., Birukou, A., & Motta, E. (2018). Ontology-based recommendation of editorial products. In D. Vrandečić et al. (Eds.), *The Semantic Web – ISWC 2018* (pp. 341–358). Cham: Springer.
- Vergoulis, T., Chatzopoulos, S., Dalamagas, T., & Tryfonopoulos, C. (2020). Veto: Expert set expansion in academia. In M. Hall, T. Merčun, T. Risse, & F. Duchateau (Eds.), *Digital libraries for open knowledge* (pp. 48–61). Cham: Springer. https://doi.org/10.1007/978-3-030-54956-5_4
- Visser, M., van Eck, N. J., & Waltman, L. (2021). Large-scale comparison of bibliographic data sources: Scopus, Web of Science, Dimensions, Crossref, and Microsoft Academic. *Quantitative Science Studies*, 2(1), 20–41. https://doi.org/10.1162/qss_a_00112
- Wang, K., Shen, Z., Huang, C., Wu, C.-H., Dong, Y., & Kanakia, A. (2020). Microsoft Academic Graph: When experts are not enough. *Quantitative Science Studies*, 1(1), 396–413. https://doi.org/10.1162/qss_a_00021
- Wang, R., Yan, Y., Wang, J., Jia, Y., Zhang, Y., Zhang, W., & Wang, X. (2018). AceKG: A large-scale knowledge graph for academic data mining. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (pp. 1487–1490). New York: Association for Computing Machinery. <https://doi.org/10.1145/3269206.3269252>
- Weinstein, L. B., Kellar, G. M., & Hall, D. C. (2016). Comparing topic importance perceptions of industry and business school faculty: Is the tail wagging the dog? *Academy of Educational Leadership Journal*, 20(2), 62.
- Wolstencroft, K., Haines, R., Fellows, D., Williams, A., Withers, D., . . . Goble, C. (2013). The Taverna workflow suite: Designing and executing workflows of web services on the desktop, web or in the cloud. *Nucleic Acids Research*, 41(W1), W557–W561. <https://doi.org/10.1093/nar/gkt328>, PubMed: 23640334
- Zang, X., & Niu, Y. (2011). The forecast model of patents granted in colleges based on genetic neural network. In *2011 International Conference on Electrical and Control Engineering* (pp. 5090–5093).
- Zhang, X., Chandrasegaran, S., & Ma, K.-L. (2021). Conceptscope: Organizing and visualizing knowledge in documents based on domain ontology. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1–13).
- Zhang, Y., Zhang, F., Yao, P., & Tang, J. (2018). Name disambiguation in AMiner: Clustering, maintenance, and human in the loop. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 1002–1011).

APPENDIX

We report in this appendix several exemplary SPARQL queries on AIDA. The aim is to show the flexibility of AIDA and the complexity of the queries that can be formulated. We also hope that these examples will offer a good starting point for users that intend to reuse AIDA. All the following queries can be run on the AIDA SPARQL endpoint, available at <https://w3id.org/aida/sparql>.

The following performs a describe query for the paper with id 2040986908.

```
DESCRIBE <http://aida.kmi.open.ac.uk/resource/2040986908>
```

The following query returns all papers written by authors from the industrial sector *computing and it* associated with the topic *robotics*:

```
PREFIX aida-ont:<http://aida.kmi.open.ac.uk/ontology#>
PREFIX aida:<http://aida.kmi.open.ac.uk/resource/>
PREFIX aidaDB: <http://aida.kmi.open.ac.uk/resource/DBpedia/>
PREFIX cso: <http://cso.kmi.open.ac.uk/topics/>

SELECT ?paperId
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?paperId aida-ont:hasIndustrialSector aida:computing_and_it .
    ?paperId aida-ont:hasTopic cso:robotics .
}
LIMIT 20
```

The following query counts how many papers have been written by authors from an industrial affiliation.

```
PREFIX aida:<http://aida.kmi.open.ac.uk/ontology#>
SELECT (COUNT(?sub) as ?count)
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?sub aida:hasAffiliationType "industry"
}
```

The next query counts how many authors are affiliated with *The Open University*.

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
PREFIX schema:<http://schema.org/>
SELECT (COUNT(DISTINCT(?sub)) as ?count)
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?sub schema:memberOf ?aff .
    ?aff foaf:name "the_open_university"
}
```

The following query returns the industrial sectors of all the papers having *Semantic Web* as a topic.

```
PREFIX aida:<http://aida.kmi.open.ac.uk/ontology#>
PREFIX cso: <http://cso.kmi.open.ac.uk/topics/>
SELECT DISTINCT ?ind
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?sub aida:hasTopic cso:semantic_web .
    ?sub aida:hasIndustrialSector ?ind
}
```


The following query returns the papers associated with the topic *Semantic Web* and written in collaboration by authors from industry and academia, where those from academia are more than 80%.

```
PREFIX aida:<http://aida.kmi.open.ac.uk/ontology#>
PREFIX cso: <http://cso.kmi.open.ac.u/topics/>
PREFIX schema: <http://schema.org/>
SELECT ?paper ?ind (count(?author) as ?nauthor)
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?paper aida:hasTopic cso:semantic_web .
    ?paper aida:hasIndustrialSector ?ind .
    ?paper aida:hasPercentageOfAcademia ?x .
    ?paper schema:creator ?author .
    FILTER (?x>80)
}
ORDER BY ?paper
```

The following query returns the number of publications in a topic (in this case *Neural Networks*) during the last 5 years. It can be used to analyze the trend of this topic in time.

```
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX prism: <http://prismstandard.org/namespaces/basic/2.0/>
PREFIX aida:<http://aida.kmi.open.ac.uk/ontology#>
PREFIX cso: <http://cso.kmi.open.ac.uk/topics/>

SELECT ?year (count(?paper) as ?n_publications)
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?paper aida:hasTopic cso:neural_networks .
    ?paper prism:publicationDate ?year .
    FILTER(xsd:integer(?year)>=2016 && xsd:integer(?year)<=2020)
} GROUP BY ?year
ORDER BY DESC(?year)
```

The following query returns the topic distribution of a given affiliation (in this case The Open University). It can be used to characterize an organization according to its relevant topics.

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
PREFIX aida:<http://aida.kmi.open.ac.uk/ontology#>
PREFIX schema: <http://schema.org/>
SELECT ?topic(count(distinct(?paper)) as ?count)
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?paper schema:creator ?author .
    ?author schema:memberOf ?aff .
    ?aff foaf:name "the_open_university" .
    ?paper aida:hasTopic ?topic .
} GROUP BY ?topic
ORDER BY DESC(?count)
```

This query ranks affiliations according to their number of publications in a given topic (in this case *Semantic Web*):

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
PREFIX aida:<http://aida.kmi.open.ac.uk/ontology#>
PREFIX cso: <http://cso.kmi.open.ac.uk/topics/>
PREFIX schema: <http://schema.org/>
SELECT ?aff ?aff_name (count(distinct(?paper)) as ?count)
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?paper aida:hasTopic cso:semantic_web .
    ?paper schema:creator ?author .
    ?author schema:memberOf ?aff .
    ?aff foaf:name ?aff_name
} GROUP BY ?aff aff_name
ORDER BY DESC(?count)
LIMIT 100
```

This query returns the academic affiliations that collaborates most (in terms of publication number) with industrial organizations:

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
PREFIX aida:<http://aida.kmi.open.ac.uk/ontology#>
SELECT ?aff ?name (COUNT(?paper) as ?n_collaborations)
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    ?paper aida:hasAffiliationType "collaborative" .
    ?paper aida:hasAffiliation ?aff .
    ?aff aida:hasGridType "education" .
    ?aff foaf:name ?name .
} GROUP BY ?aff ?name
ORDER BY DESC(?n_collaborations)
```

The following query returns the DBpedia concepts associated to a given paper (id: 2300368847 in this case) using the mapping between CSO and DBpedia.

```
PREFIX aida: <http://aida.kmi.open.ac.uk/ontology#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX aida: <http://aida.kmi.open.ac.uk/resource/>
SELECT *
FROM <http://aida.kmi.open.ac.uk/resource>
WHERE {
    aida:2300368847 aida:hasTopic ?topic .
    ?topic owl:sameAs ?obj .
    FILTER(regex(str(?obj), "dbpedia" ) )
}
```