

# Intelligence Catalog-guided Tracking of the Evolution of (machine) Intelligence: Preliminary results

Dagmar Monett<sup>1,2</sup>, Niklas Lampe<sup>1</sup>, Max Ehrlicher-Schmidt<sup>1</sup>, and Nikolaj Bewer<sup>3</sup>

<sup>1</sup> Computer Science Dept., Berlin School of Economics and Law, Germany  
dagmar.monett-diaz@hwr-berlin.de

{s.lampe18, s.ehrlicherschmid18}@stud.hwr-berlin.de

<sup>2</sup> AGISI.org

<sup>3</sup> Universität Potsdam, Germany  
bewer@posteo.de

**Abstract.** This paper<sup>4</sup> presents some preliminary results after having analyzed the metadata of 146,585 and the full text of 9,879 relevant AI-related papers. Other datasets collected from different sources were prepared for further analysis, too. The idea is to investigate the use of capabilities that denote intelligence in the AI scientific discourse with the goal to track the evolution of the narrative around intelligence and its manifestations in humans and machines over time. Intelligence capabilities and related properties that are used when defining intelligence were extracted from previous work in this area in the form of a catalog and spanning over different fields, AI included. Although still an elusive, difficult to define concept, analyzing and understanding the discourse around intelligence may shape both the way we use it and how intelligent artifacts are and will continue to be developed in the future.

**Keywords:** Intelligence · Evolution · Human intelligence · Machine intelligence · NLP · Intelligence catalog.

## 1 Introduction

To the best of our knowledge, there is no known study that tracks the historical development of the terminology used in intelligence research, nor how it has affected both the evolution of machine or artificial intelligence (AI) and the language used by the machine intelligence community to define it.

The reasons are many and have been accentuated along the years: the absence of a common language around AI is a missing piece for a wide-ranging understanding of what intelligence means; there is no broadly accepted definition of machine intelligence; the field is multidisciplinary and almost each researcher defines intelligence through the lens of their discipline; and the exacerbated media AI hype from the last few years attempts against more understanding, fuels division (not only) inside the AI community, and deviates investments to the areas where only rapid gains can be attained, among other reasons.

<sup>4</sup> Copyright ©2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

But that situation must change and that is the main objective of the research project introduced in this paper. The research is based on an exhaustive analysis of more than 1,6 million of scientific papers. The central goal driving the research lies in discovering of how the vocabulary on intelligence, in particular, and on AI, in general, have shaped the discourse of the scientific community.

This would have several positive implications not only for the AI community, but also for those stakeholders outside it. For example, journalists and PR people, politicians, regulators, AI novices, and researchers from other fields, to name just some of the stakeholders that have been using more and more the AI vocabulary in their work, would be aware of the preferred scientific terminology by the AI community, would be able to study its history in a more supporting way, and would be exposed to a common language that facilitates understanding and dialogue across disciplines.

## 1.1 Related Work

Although there has been a recent and growing interest in characterizing, structuring, and tracking the evolution of AI in particular and of intelligent systems in general, there has been very little research on finding a common language around both intelligence and AI, and no research at all, to our best knowledge, on tracking the evolution of definitions of intelligence in the scientific literature. This is why the research presented here may be the first of its kind.

Already in [9] it was discussed how different works (e.g. [1,3,4,5]) have tried to analyze trends and the historical development in a broader sense of AI as a field. The research project presented in this paper uses the results obtained in [9], especially the intelligence vocabulary or catalog that the authors created, as well as it is inspired by the work done by other authors, as introduced above. The idea now is to study how the concepts or properties that could denote intelligence or intelligence capabilities have been used in the scientific literature in order to track the evolution of the concept of intelligence along the years.

## 1.2 Research Goal and Sub-goals

The central research goal of the project is to study the shifting focus of intelligence research. This goal is subdivided into the following sub-goals:

- **Sub-goal 1:** To study how the vocabulary on intelligence has been created and used by the research community along the years. This sub-goal can be accomplished by analyzing how the concepts from the intelligence catalog that was created in [9] have been used in different scientific publications along the years, as well as those attributes used to define intelligence from [11]. The fulfillment of this sub-goal presupposes the existence of a corpus of scientific papers that should be collected in advance (see Section 3 for more).
- **Sub-goal 2:** To delineate how the terminology on intelligence has been published in the scientific literature. After the most important concepts are determined in the first sub-goal, then the common trends of their evolution are to be identified.

- **Sub-goal 3:** To analyze how the language of intelligence affects a common understanding around both intelligence and AI. The trends in intelligence research along the years from the second sub-goal will contribute to a focused discussion on the role that the language of intelligence has played in defining and understanding both intelligence and AI.
- **Sub-goal 4:** To depict the historical development of the concept of (machine) intelligence. All results achieved in the previous sub-goals will be graphically synthesized, documented, and prepared for dissemination.

This paper contributes to most of the research sub-goals mentioned above.

## 2 Methodology and research phases

The research project consists of the following phases:

1. Data collection and preparation: Gradual, automated download of metadata to create a new corpus of scientific literature on intelligence and AI:
  - (a) All papers published in the Computer Science category “artificial intelligence” (cs.ai) of the arXiv e-prints.
  - (b) All papers published in other Computer Science categories of the arXiv e-prints that might be closely related.
  - (c) Other AI-related papers and scientific literature (depending on availability), as studied in [1,3,4,5].
  - (d) Other sources like leading conferences and workshops on AI, for instance.
2. Data pre-processing and analysis using machine learning techniques like natural language processing (NLP), as well as statistical analysis. The intelligence catalog available at <https://tinyurl.com/intelligenceCtlg>, as well as other results from [9], should be especially considered. This phase implements the research sub-goals 1 to 3.
3. Interpretation and explanation of the obtained results. This phase complements and documents the research sub-goals 1 to 3.
4. Final discussion, output, and summarization of the obtained results; dissemination of results: This phase implements the last research sub-goal.

It is expected that the phases involving the collection of the data and its pre-processing might likely consume a huge amount of time compared to other phases, as it is usually the case in related work.

## 3 Intelligence Capabilities under Study and Sources

Two main sources for tracking the use of intelligence capabilities in the scientific literature were considered for analysis.

The first source includes attributes that contributors to two symposia on intelligence used in their definitions. These were compiled by Sternberg and Detterman in [11]. Both symposia, celebrated in 1921 and 1986, respectively, are widely known in the

field of Psychology and other related fields. The definitions of human intelligence from the 1921 symposium focused on the “prediction of behavior.” Their contributors made emphasis on attributes like *adaption*, higher-level components (e.g. *problem solving*, *decision making*, *abstract reasoning*), and the *ability to learn*, among others. The definitions of intelligence (mostly human) from the 1986 symposium focused on the “understanding of behavior.” Their contributors used more attributes in general, some of them with a similar frequency or even higher than the ones from the 1921 symposium, like the speed of *mental processing*, *knowledge*, and *metacognition*, among others.

The second literature source includes words extracted from a corpus of definitions and experts’ comments on defining intelligence that conform the *Intelligence Catalog* or vocabulary presented in [9] and referred to in Section 2 above. This time, the definitions of intelligence and the corresponding experts’ comments made emphasis on the “computation of intelligence.” They were collected in the survey presented and discussed in [7,8] which had contributions not only from psychologists but also from computer scientists, engineers, philosophers, and academics and practitioners from others fields.

The vocabulary used to define intelligence in the first literature source might be biased, however: it considers contributions by only 39 psychologists in total that come only from US and Europe. Yet, the vocabulary from the second literature source is much broader and may counter the bias: a total of 567 individuals coming from more than 57 countries and from more than 15 different fields participated. Furthermore, more than 343 definitions of intelligence, both of human and machine (or artificial) intelligence, were examined.

### 3.1 Datasets

Table 1 summarizes the most important information about the datasets that were considered for analysis. Most of them comprise scientific papers published in the proceedings of leading AI conferences and around 50 years of documents from the AAI’s AITopics database on research, people, and applications of AI.<sup>5</sup> Also papers uploaded to the open-access repository arXiv have been taken into account.

### 3.2 Data collection

Collecting, converting, and preparing the data made it to be the most consuming tasks of the whole project. In particular, the data was collected in the following way:

- The datasets *D01*, *D02*, and *D03* were provided by Fernando Martínez-Plumed, the leading author in [5], in a personal communication. *D01* and *D03* contain the metadata of papers that were published in the proceedings of AAI and IJCAI, two leading conferences in AI worldwide, respectively. *D02* contains the metadata of more than one century of research as documented in the AAI’s AITopics database.
- The dataset *D04* contains the metadata of all papers published in the arXiv Computer Science categories *cs.ai* (Artificial Intelligence), *cs.ml* (Machine Learning), *cs.cl* (Computational Linguistics), *cs.ne* (Neural and Evolutionary Computing), and *cs.cv* (Computer Vision) until February 2018.

<sup>5</sup> See <https://aitopics.org/> for more.

Dataset	Source	No. of papers	Time period	Related work	Relevant data
<i>D01</i>	AAAI	4,691	1997-2017	[5]	title, keywords
<i>D02</i>	AITopics	113,558	1903-2018	[5]	title, keywords, abstract
<i>D03</i>	IJCAI	4,113	1997-2017	[5]	title, keywords
	arXiv>cs>ai, arXiv>cs>cl,				
<i>D04</i>	arXiv>cs>cv, arXiv>cs>ne, arXiv>cs>ml	41,000	1993-2018	[10]	title, abstract
<i>D05</i>	arXiv>cs>ai	16,625	1993-2018	[3]	title, abstract
<i>D06</i>	arXiv>cs	167,998	1993-2018	[2]	title, abstract
<i>D07</i>	arXiv	1,480,122	1993-2018	[2]	title
<i>D08</i>	ECAI	1,908	2000-2018	This paper	title, keywords, abstract
<i>D09</i>	ICML	5,690	1988-2019	This paper	title
<i>D10</i>	IJCAI	9,879	1997-2020	This paper	whole text

Table 1: Datasets, information, and metadata of more than 1.6 million papers that were collected.

- The dataset *D05* contains the metadata of papers published in the arXiv category cs.ai, available at [https://github.com/karenhao/techreview\\_arxiv\\_scrape](https://github.com/karenhao/techreview_arxiv_scrape). It might be a subset of *D04*; thus, we considered repetitions of single papers for deletion.
- The dataset *D06* contains the metadata of all arXiv papers published in all Computer Science categories (40 in total) through the end of 2018; thus, the datasets *D04* and *D05* might be subsets of it. It should be noted that papers uploaded to arXiv can be submitted to more than one category. Thus, it was important to check for duplicates in the dataset *D06* since it is formed by putting together 40 different categories (or files with metadata) from Computer Science.
- The dataset *D07* contains only the titles of all arXiv papers published in all categories (175 in total) through the end of 2018. It helps in identifying how some concepts from the intelligence catalog have been used in other fields (e.g. in Astrophysics, Quantitative Biology, Economics, and so on), at least in their titles.
- The dataset *D08* contains the metadata of all papers published in the ECAI proceedings from 2000 to 2016, as it can be found in <http://frontiersinai.com/>.
- The dataset *D09* was constructed after downloading the metadata of the ICML papers from dblp (Digital Bibliography and Library Project). This information can be found at [urlhttps://dblp.org/db/conf/icml/index](https://dblp.org/db/conf/icml/index).
- Finally, the dataset *D10* was constructed after downloaded the IJCAI papers from the website of the conference (see <https://www.ijcai.org/Proceedings/>). This dataset contains the whole text of the papers published in the IJCAI proceedings until the year 2020 (all as PDF files).

Most of these datasets have been already considered in some detail in the first phase of the research project; see the next sections for more.

## 4 Data Pre-processing and Analysis

Python was used for processing and analyzing the data.<sup>6</sup> Among the libraries used are `tz` and `PyMuPDF`, which are used to read the text from the PDF files (in the case of the IJCAI papers), and `Matplotlib` and `NumPy`, which are used to generate different diagrams. `Pandas` is used for reading the Excel files, data cleansing, data formatting and analysis, for it provides several useful operations for one- and multidimensional arrays.

Figure 1 illustrates the major phases and data flow that constitute our pipeline when processing the datasets.

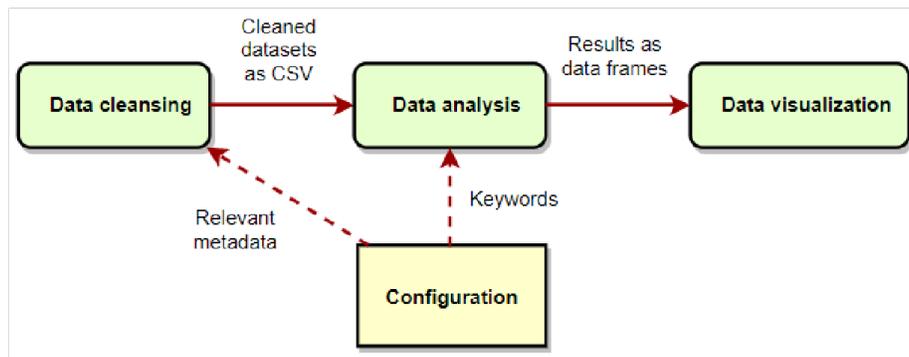


Fig. 1: Components of the analysis tool.

The configurator reads the configuration files, one with the column names of the relevant metadata, one with the keywords or intelligence capabilities to be searched for.<sup>7</sup> Then, the configurator matches that information with the data contained in the datasets and prepares it in order to be processed by the other components of the analysis tool. The datasets are then cleansed and irrelevant metadata and content are deleted. Other common NLP pre-processing tasks are performed as well, like lower-casing the keywords, deleting emoticons and other especial characters, and converting the data to CSV data files for later processing, among others tasks. The remaining content is formatted to a convenient representation. The analysis then starts by considering the year of publication of the different scientific papers and the frequency of their publication. Moreover, the occurrences of keywords are counted thereby avoiding occasional repetitions from the same entry. Percentages relative to the number of papers per year

<sup>6</sup> Both the source code of the Python scripts and the data that the authors of this paper collected are available upon request. They will be made available for others to download and work with.

<sup>7</sup> In what follows, we refer to *keywords* to denote the intelligence capabilities from the intelligence catalog presented in [9].

are calculated and the data prepared for the last step. Finally, different diagrams are produced, ready for visualization.

The corresponding analysis tool is much simpler for the case of the IJCAI papers which content is available. It consists of only two components: the PdfMiner and the analyzer. The former extracts the text from the PDF files and calculates the frequency of the keywords that are of interest. The latter reads the processed data, generates suitable diagrams, and saves the results in a specified folder. Also convenient NLP tasks of the pre-processing are performed here. It is worth mentioning that reading the PDF files with the appropriate functions from the PyMuPDF library is much faster (about 0.001 seconds per PDF file) than doing it with other libraries (like textract, Apache Tika or PyPDF2).

## 5 Preliminary Results and Discussion

More than 100 diagrams were produced for visualization and analysis. Some of them will be commented in what follows, mainly in a descriptive way.

For all the considered datasets, the number of papers published per year has increased dramatically in the last decade (more exactly: since 2010), almost showing an exponential growth.

A first analysis of the frequency which with the intelligence capabilities and related attributes from the intelligence catalog appear in the datasets *D01* to *D10* helped to distinguish three different stages:

1. AITopics: Long-term, historical development.
2. AAAI, IJCAI, ECAI and arXiv>cs>ai: Recent past.
3. ICML: Machine Learning in focus.

Figure 2 depicts the historical development in the use of some intelligence-related words in papers from AITopics, allowing for a broader, long-term analysis of some trends along more than 60 years. Notice how the word *learn* had its lowest use during the two AI winters (1974-1980 and 1987-1993) and started to gain momentum from the 90's on.

Figure 3 takes a closer look at some of the terminology often used to refer to the internal processing of intelligent artifacts. In this case, the whole texts (and not only the metadata) of AI papers published in the IJCAI proceedings are considered. The darker the color, the more frequent such words were used in a certain year.

Figure 4 shows some intelligence capabilities from the Intelligence Catalog and the frequency of their use in the text of IJCAI papers. Notice how there are words whose use has decreased with time (e.g. *act*, *communicate*, *perceive*, *reason*, *think*), some faster than others. They might be gradually losing their importance in the scientific discourse and, thus, in the vocabulary around machine intelligence. Once considered the paramount of intelligence in machines, as in Winston's well-known definition, "*Artificial Intelligence is . . . the study of the computations that make it possible to perceive, reason, and act*" [13], those words might be succumbing to the severity of the "AI effect," as McCorduck refers to in her book [6].

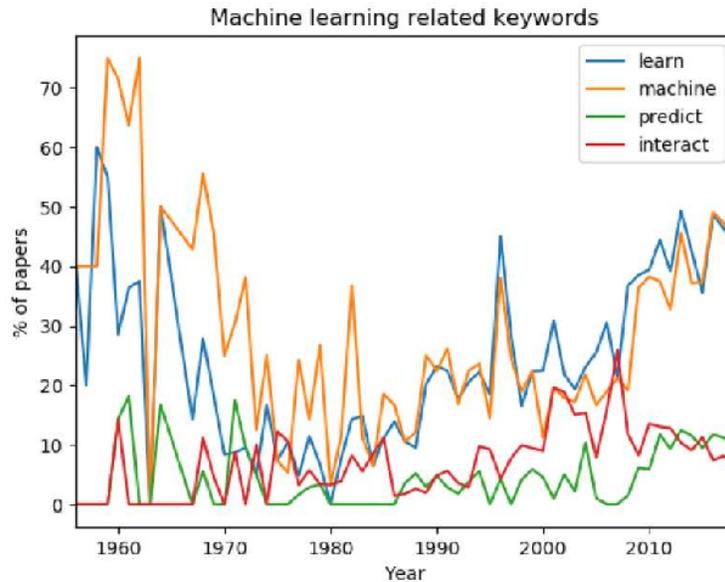


Fig. 2: Proportion of AITopics papers’ metadata containing intelligence-related words.

Other words, however, are getting more importance as time pass (e.g. *adapt*, *learn*, *predict*) and in so doing, they are documenting the increasing momentum that the AI field has been experiencing in the last decades. It is not surprising that some of them are even considered to be essential capabilities of machine intelligence, like in the widely accepted definition of intelligence by Wang: “*The essence of intelligence is the principle of adapting to the environment while working with insufficient knowledge and resources. Accordingly, an intelligent system should rely on finite processing capacity, work in real time, open to unexpected tasks, and learn from experience*” [12].

## 6 Limitations

The arXiv papers go back only to 1993 and the collected metadata from AAAI and IJCAI goes back only to 1997. The AI field is much older, though.

The ECAI proceedings are a very good source of information and the conference takes place since 1974, but only a few proceedings are available online (from the year 2000 onward, in a two-year frequency).

Furthermore, there are many other AI conferences and workshops that take place every year, some of them with a long tradition inside the AI field, that are not considered in this research project. For example, NeurIPS electronic proceedings are available online for the time period 1987-2018, but this is a conference with focus on specific areas of Machine Learning that does not cover other important AI sub-fields. Nevertheless, it will be considered in further work. Data has been also collected and is available for similar conferences, but the time periods are very limited or the proceedings are not

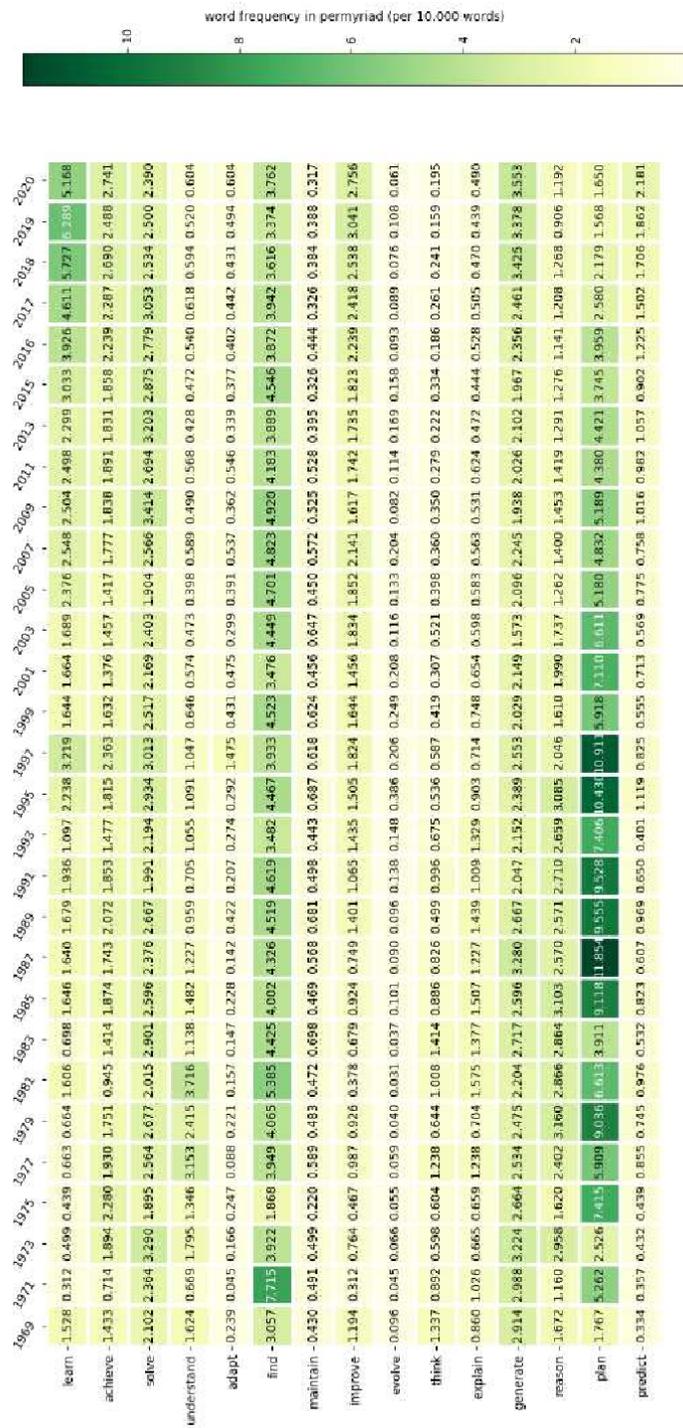


Fig. 3: Intelligence capabilities used when referring to the internal processing and their use in the text of IJCAI papers.

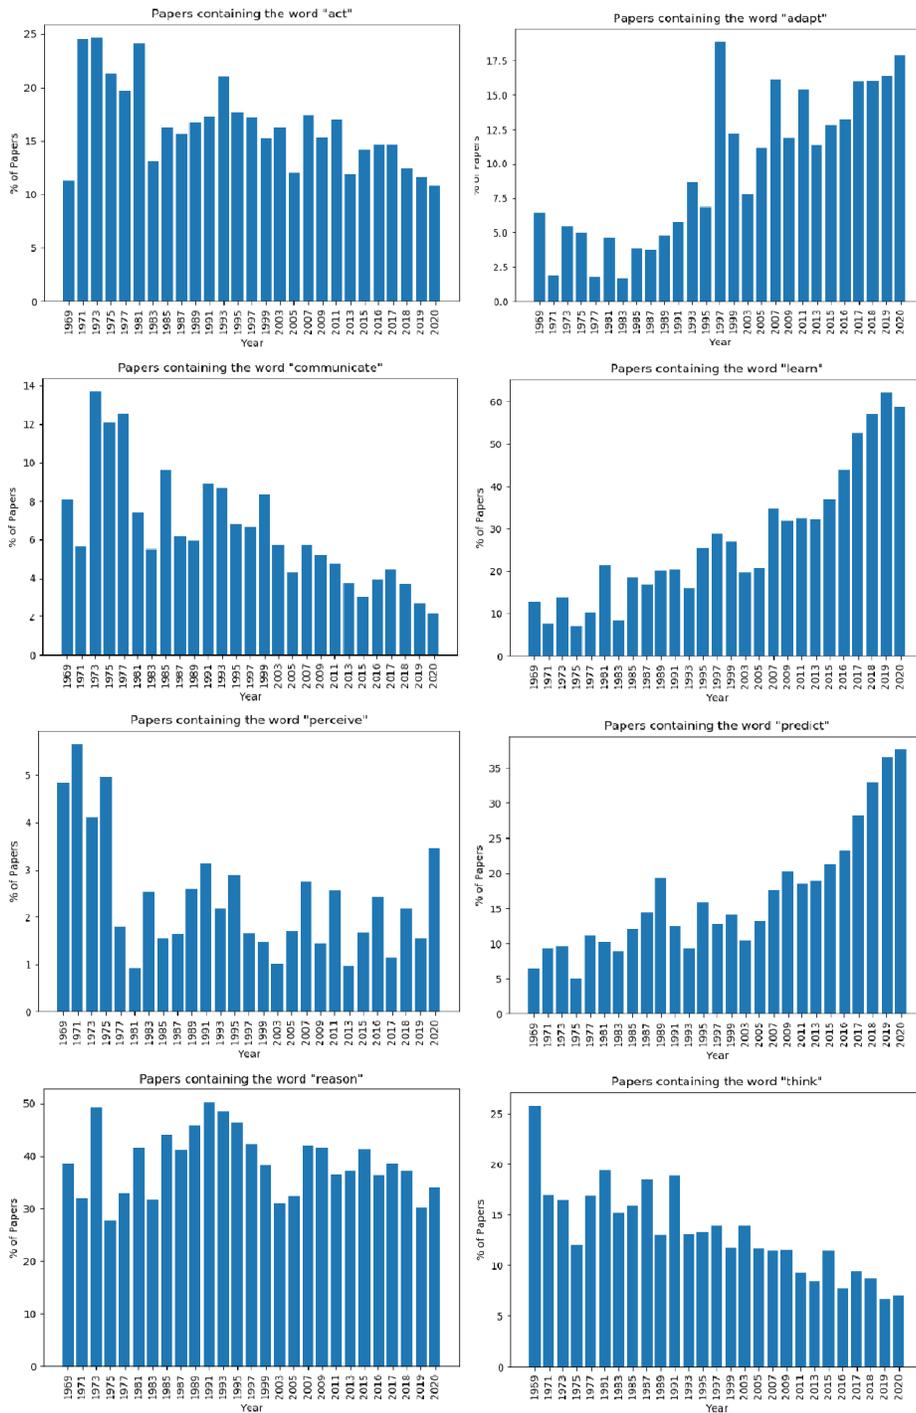


Fig. 4: Intelligence capabilities used in the texts of IJCAI papers.

available for download. The code for scraping the data might be available, too. Some of these conferences have a narrower focus, though.

Another possible limitation of this work and the ongoing bigger research project that includes it, as it was rightly pointed out by one of the anonymous reviewers,<sup>8</sup> is that we don't take into account how the English language nor how some words' meanings, especially in the field and sub-fields of AI, have evolved over time. However, we are of the opinion that our work would contribute to reaching that goal in a better way, especially when considering a semantic analysis (topic of current work), because it would be easier to detect in which contexts which intelligent capabilities and narratives have been used by the scientific community.

## 7 Conclusions

Although preliminary, the data collection, preparation, and analysis, together with the first results presented in this paper, set the basis for a much broader analysis of how the terminology is used to define intelligence has evolved over time. More work in this regard is needed, however. For example, semantic aspects that allow for distinguishing specific contexts on one side, and the relative importance of the words used to define intelligence on the other side, would help in framing a more accurate use of the terminology on intelligence over time. This is the topic of ongoing work. A similar analysis was already initiated in [9] for the definitions of human and artificial intelligence provided in [7,8]. A fundamental adaptation of those methods, which are already available, is not needed but their extension for using the new data introduced in this paper.

Finally, there is an extensive body of work on intelligence outside the AI community (for example in Psychology and Neuroscience, to name two of the most prominent fields in that research area) that must also be considered in future work. Although the vocabulary on intelligence used in these fields does not suffer from the same lack of consensus the AI community cannot put aside, acknowledging their lessons learned and the good practices would benefit both the AI field and the discourse on intelligence inside and outside it.

In order to accomplish these new goals, it would first be needed to harvest and process the new data, i.e. to collect, analyse, and pre-process, for example, already published papers and their metadata in the fields of Psychology, Neuroscience, and the like. These tasks should not be underestimated; they are actually the most time consuming tasks in any machine learning project.

## References

1. Elsevier: Artificial intelligence: How knowledge is created, transferred, and used. Trends in China, Europe, and the United States. <https://bit.ly/2Lajwaa> (2018)
2. Geiger, R.S.: ArXiv archive: A Tidy and Complete Archive of Metadata for Papers on arxiv.org. Zenodo. <http://doi.org/10.5281/zenodo.1463242> (2019)
3. Hao, K.: We analyzed 16,625 papers to figure out where AI is headed next. Intelligent Machines, MIT Technology Review. <http://bit.ly/2TfLfK1> (2019)

---

<sup>8</sup> We thank to all of them for their valuable comments and suggestions for improvement.

4. Klinger, J., Mateos-Garcia, J., Stathoulopoulos, K.: Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology. arXiv:1808.06355 (2018)
5. Martínez-Plumed, F., Loe, B.S., Flach, P.A., Ó hÉigeartaigh, S., Vold, K., Hernández-Orallo, J.: The Facets of Artificial Intelligence: A Framework to Track the Evolution of AI. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018. pp. 5180–5187. Stockholm, Sweden (2018)
6. McCorduck, P.: *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. A. K. Peters, Ltd., Natick, MA, second edn. (2004)
7. Monett, D., Hoge, L., Lewis, C.W.P.: Cognitive Biases Undermine Consensus on Definitions of Intelligence and Limit Understanding. In: Furbach, U., Hölldobler, S., Ragni, M., Rzepka, R., Schon, C., Vallverdu, J., Włodarczyk, A. (eds.) Joint Proceedings of the IJCAI-2019 Workshops on Linguistic and Cognitive Approaches to Dialog Agents and on Bridging the Gap Between Human and Automated Reasoning. pp. 51–58. CEUR-WS, Macau, China (2019)
8. Monett, D., Lewis, C.W.P.: Getting clarity by defining Artificial Intelligence—A Survey. In: Müller, V.C. (ed.) *Philosophy and Theory of Artificial Intelligence 2017*, vol. SAPERE 44, pp. 212–214. Springer, Berlin (2018)
9. Monett, D., Winkler, C.: Using ai to understand intelligence: The Search of a Catalog of Intelligence Capabilities. In: Alam, M., Basile, V., Dell’Orletta, F., Nissim, M., Novielli, N. (eds.) *Proceedings of the 3rd Workshop on Natural Language for Artificial Intelligence, NL4AI 2019*. pp. 1–15. CEUR-WS, Rende, Italy (2019)
10. Shah, N.: ArXiv data from 24,000+ papers. Papers published between 1992 and 2017. kaggle. <https://www.kaggle.com/neelshah18/arxivdataset> (2018)
11. Sternberg, R.J., Detterman, D.K.: *What is Intelligence? Contemporary Viewpoints on its Nature and Definition*. Ablex Publishing Corporation (1986)
12. Wang, P.: What do you mean by “AI”? In: Wang, P., Goertzel, B., Franklin, S. (eds.) *Artificial General Intelligence 2008, Proceedings of the First AGI Conference, Frontiers in Artificial Intelligence and Applications*. vol. 171, pp. 362–373. IOS Press Amsterdam, The Netherlands (2008)
13. Winston, P.H.: *Artificial Intelligence*. Addison-Wesley Publishing Company, third edn. (1992)