

The Facets of Artificial Intelligence: A Framework to Track the Evolution of AI

Fernando Martínez-Plumed¹, Bao Sheng Loe², Peter Flach³,
Seán Ó hÉigeartaigh², Karina Vold², José Hernández-Orallo¹

¹ Universitat Politècnica de València, Spain

² University of Cambridge, United Kingdom

³ University of Bristol, United Kingdom

fmartinez@dsic.upv.es, bs128@cam.ac.uk, Peter.Flach@bristol.ac.uk,
so348@cam.ac.uk, kvv22@cam.ac.uk, jorallo@dsic.upv.es

Abstract

We present nine facets for the analysis of the past and future evolution of AI. Each facet has also a set of edges that can summarise different trends and contours in AI. With them, we first conduct a quantitative analysis using the information from two decades of AAI/IJCAI conferences and around 50 years of documents from *AI topics*, an official database from the AAI, illustrated by several plots. We then perform a qualitative analysis using the facets and edges, locating AI systems in the intelligence landscape and the discipline as a whole. This analytical framework provides a more structured and systematic way of looking at the shape and boundaries of AI.

1 Introduction

“What is AI?” has been a common question from the inception of the discipline in the 1950s [McCarthy *et al.*, 2006; Moor, 2006; Solomonoff, 2017] until the end of the last century [Lehman-Wilzig, 1981; Fetzer, 1990; McCarthy, 1998]. In the twenty-first century, the discipline is not only well-developed but it is often said that this is *the age of AI* [McCorduck, 2004]. AI is set to pervade and transform every aspect of life. In a way, this is no different from what computer science has already been doing, creating computerised, digital and virtual versions of almost everything, with AI now introducing new adjectives such as ‘intelligent’, ‘smart’ and ‘cognitive’ to almost every process or gadget, from medical diagnosis to personal assistants. This expected expansion and intertwining with every other research discipline and aspect of life is pushing the contours of AI in many directions. However, unlike computer science, which is based on well-established models of computation that integrate hardware and software, AI has evolved with a more fluid definition, primarily because of our varied conceptions of intelligence.

When looking at the definitions of (artificial) intelligence (see, e.g., [Legg and Hutter, 2007], for a compendium), we see that definitions can be categorised according to two distinct dimensions, following [Russell and Norvig, 2009]. First, we can characterise intelligence in terms of “thinking”

(process-oriented) or “acting” (goal-oriented). Second, we can characterise intelligence taking humans as a reference or looking for a more abstract or universal reference (such as rationality). These two dimensions are summarised in Table 1.

Not all definitions can be clearly classified according to this table. For instance, Minsky’s famous definition of AI, as the “science of making machines capable of performing tasks that would require intelligence if done by [humans]” [Minsky, 1968], uses humans as a reference, so it would be located on the left of the table, but still somewhere between the top and bottom part of the table as it focuses on processes (i.e., thinking) for humans and tasks (i.e., acting) for machines

This suggests that we can look at these dimensions in a more nuanced, continuous way, as *facets* rather than discrete categories, to look at the evolution of AI as a discipline. For instance, has AI been more focused towards “thinking” (and hence processes or techniques) or more focused towards “acting” (and hence tasks and applications)? Is it more or less influenced by human intelligence (using human processes and the tasks humans can actually solve) or is the discipline going in a direction of a more abstract or universal characterisation? In this paper, we will develop these and other dimensions into facets to characterise the *object* of AI, i.e., the AI systems and services, such as generality, location and embodiment, and the *subject* of AI, such as paradigms, actors, character and nature, as a discipline. With these facets we will perform a qualitative and quantitative analysis of the evolution of AI.

We will base our analysis on evidence as much as possible, looking at data from scientific venues, reports, surveys and other sources. At the end of the paper, we will take a more principled stance and we will discuss how to extrapolate these trends or identify what kinds of criteria are necessary to establish AI contours that can be more stable and useful.

The contributions of this paper are a set of well-structured facets to analyse the location and contours of AI. This is complemented by substantial evidence collected from several sources, such as 20 years of the IJCAI and AAI proceedings.

Thinking Humanly	Thinking Rationally
Acting Humanly	Acting Rationally

Table 1: Categories of (artificial) intelligence definitions, according to the two binary *dimensions* in [Russell and Norvig, 2009, Fig. 1.1].

ings and the whole *AI topics* database¹. All the datasets, plots and code to scrape and process this data are made publicly available (see footnote 7).

There are several reasons why using clear criteria such as the facets and edges are helpful for defining AI, how it has evolved and what it is likely to be in the future. First, internally, the research community needs clear criteria to determine what is in or out of scope, for conferences, journals, funding and hiring. Second, policy makers are considering ways of regulating AI, but wrestle with the problem of defining what systems are actually AI. Third, different terms are being used to capture the same subparts or forms of AI, such as artificial general intelligence (AGI), machine intelligence, cognitive computing, computational intelligence, soft computing, etc. Without a clear structure about AI, it is difficult to determine whether they are part, they overlap or they are simply redundant. Fourth, some areas, such as machine learning, are taking a more relevant role in AI, but they are also intertwined with areas such as statistics, optimisation or probability theory, which were not always considered near the contours of AI. Fifth, along the history of AI there has been rise and decline of interest and research towards the different AI paradigms, techniques and approaches.

2 Background

Disciplines are commonly analysed historically, and AI is not an exception. A history of a discipline usually emphasises the problems, progress and prospects, but not necessarily delineates its contours. For instance, there are several excellent accounts of the history of AI [McCorduck, 2004; Buchanan, 2005; Nilsson, 2009; Boden, 2016]. Some of them cover AI from a philosophical perspective. However, given the implications of AI for the interpretation of the human mind, the analysis of AI as a discipline (as usual from the viewpoint of philosophy –or methodology– of science) is not commonly done in terms of its external and internal contours. As a result, several key questions remain: what are the criteria to recognise that an entity is part of AI? Furthermore, how can we recognise the internal subdisciplines in AI?

One possible approach to this is to determine the nature and contours of AI by its common use. While this may be a good approach for evaluating progress as a whole or for a few benchmarks², the analysis of subdisciplines by their popularity may be prone to many terminological confusions and many vested interests, with the risk of having characterisations that are very volatile, such as big data.

Bibliometric approaches are a common tool to analyse disciplines, but the focus is usually put on impacts per author, venue, location or institution. Sometimes, bibliometrics studies disciplines and subdisciplines. For instance, Scimago provides a way of looking at the “shape of science”³ where one can locate AI and some of its subareas in terms of their internal and external relations. Provided with a set of tags, bibliometrics can study how frequent several tags are in published papers (including titles, keywords or abstracts). For instance,

¹<https://aitopics.org>.

²See, e.g., <http://aiindex.org/> and <https://www.eff.org/ai/metrics>.

³<http://www.scimagojr.com/shapeofscience/>.

Artificial Intelligence	Artificial Neural network+	Genetic algorithm+
Expert System*	Optimization*	Prediction*
Classification*	Design*	Fuzzy Logic+
Multi-agent system+	Simulation*	Support Vector Machine+
Machine Learning+	Diagnosis*	Data mining*
Swarm intelligence+	Particle swarm optimization+	Decision support system*
Pattern recognition*	Knowledge-based system*	Case-based reasoning*
Distributed artificial intelligence+	Knowledge representation*	Management*
Identification*	Decision making*	Fault Diagnosis*
Computational intelligence+	Recognition*	Ontology+

Table 2: Top 30 high-frequency keywords used for the bibliometric analysis in [Niu *et al.*, 2016, Tab. 6]. Those shown with ‘+’ are included in their Fig. 6 (methods and models) and those with ‘*’ are included in their Fig. 7 (applications).

[Niu *et al.*, 2016] includes a thorough historical analysis of publications in about 20 relevant journals in AI (but not conferences or open journals such as JAIR) from 1990 to 2014. The number of publications is shown by 5-year periods. The keywords used are shown in Table 2.

As we can see in the list, the keywords include terms for disciplines and subdisciplines, techniques, and application areas, from which the authors distinguish “methods and models” and “applications”. As the analysis is limited to most frequent areas, it excludes important subfields of AI (e.g., “planning”) and some assignments are vague, with keywords such as “design”, “identification” or “prediction”.

Still, we see that this analysis is aligned with the first facet mentioned in the introduction, of whether the discipline is characterised by their techniques or their applications. This is not surprising, as it is a typical categorisation of disciplines according to their techniques and applications, especially in engineering. A proper cataloguing effort (e.g., as done by ACM and other associations for computing) would be an option, but AI is too dynamic to allow for a stable set of terms for a long period.

In the end, instead of a reactive approach focusing on the trends, a more proactive stance towards the recognition of the subdisciplines can be seen as a duty for scientific associations, editorial boards and program chairs. Accordingly, AI researchers should configure the landscape of AI⁴ and determine what is relevant or off-topic for a venue. They can also determine what the subareas are, so that proper reviewers and sessions are allocated depending on the importance of each area. It is unusual, however, to conduct a more systematic analysis of how these choices are made. Three remarkable exceptions are [Shah *et al.*, 2017], where area relevance is only examined on passing, [Fast and Horvitz, 2017] where the authors focus on views expressed about AI in the New York Times over a 30-year period in terms of public concerns as well as optimism, and [Moran *et al.*, 2014], which focuses on the venue keywords, including a cluster analysis on the 2013 keyword set and a new series of keywords, which were adapted by AAI2014. Table 3 shows an integration of their selection with the keywords of AAI2014.

Applications	Cognitive modeling	Game theory
Heuristic search & optimization	Humans & AI	Knowledge representation
Machine learning	Multiagent systems	Natural language processing
Planning	Search & constraint satisfaction	Vision & robotics

Table 3: A selection of keywords as an intersection of the two last columns of [Moran *et al.*, 2014, Tab. 2].

⁴<https://aaai.org/Magazine/ailandscape.php>.

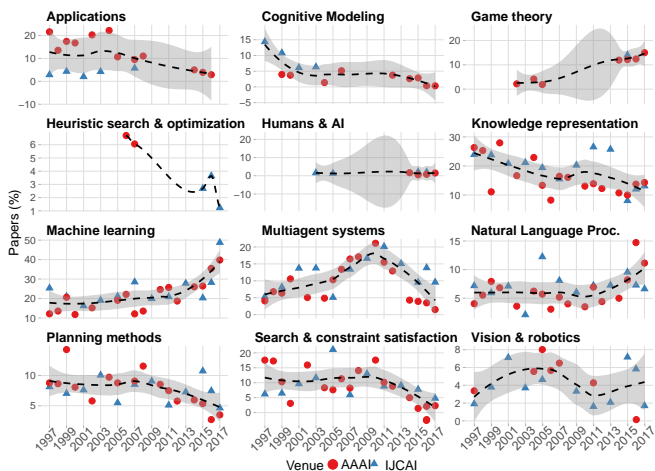


Figure 1: Papers published in AAAI (red dots) and IJCAI (blue triangles) conferences (1997-2017) separated by their area (topic “keyword”, Table 3). Areas are mutually exclusive and sum up to 100% per year. General (dashed black line) tendencies and standard errors (bands) are shown for both data series together.

3 Historical Data: Analysis by Keywords

We used the keywords in Table 3 for a first quantitative analysis, using data obtained from two significant sources:

- AAAI/IJCAI conferences (1997-2017). We obtained data of all the accepted papers for these two conferences from DBLP⁵. This database represents information about computer science (and thus AI) that comes mostly from their researchers (conference proceedings and journals).
- *AI topics* documents (1905-2018). *AI topics*⁶ is an archive kept by the AAAI, containing a variety of documents (e.g. news, blog entries, conferences, journals and other repositories) that are collected automatically with NewsFinder [Buchanan *et al.*, 2013].

With a mapping approach between terms and the categories in Table 3 as a representative list of subareas in AI, we summarised trends in a series of plots. For the AAAI/IJCAI conferences data, we used the keywords appearing in the proceedings, and for *AI topics* we used the tags (substrings appearing in titles, abstracts and topics). Regarding conferences, other major venues (such as ECAI, ICML or NIPS) were not included due to the lack of keyword information in DBLP. Still, AAAI and IJCAI can be considered to be a representative basis for analysing AI trends proper.

Figure 1 shows the evolution of the areas in Table 3 for the past 20 years using the AAAI/IJCAI data. Document counts are normalised to sum up to 100% per year. Standard errors are also shown, where in some cases the thickness of the band is very variable (e.g., “Human & AI” and “Game Theory”) due to data sparsity, or it cannot be computed due to insufficient data available (e.g., “Heuristic search & optimisation”).

Figure 2 shows a similar evolution based on the *AI topics* data. If we visually compare the same keywords between the plots found in Figures 1 and 2 (taking into account that axes are different), we see that both sources are rather consistent and highlight a few clear trends. For example, the categories

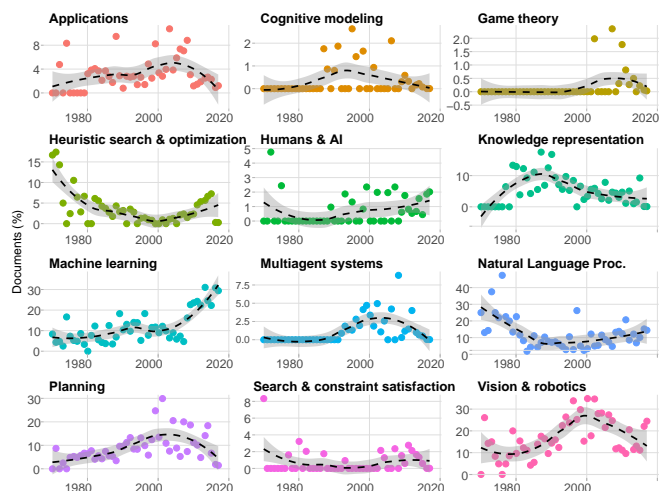


Figure 2: Documents included in AI topics (1970-2017) by area (topic “keyword”, Table 3). Legend as in Figure 1 (note that the x -axis is different, with a much wider time span here).

heuristic search & optimisation and *knowledge representation* show a decreasing trend. Other categories seemed to have peaks: *cognitive modeling* in the 1990s, strongly associated with the emergence of several cognitive architectures (e.g., ACT-R [Anderson *et al.*, 1997], EPIC [Kieras and Meyer, 1997] or SOAR [Newell, 1994]); *planning methods* around 2000, possibly due to the introduction of the first method for solving POMDP offline [Kaelbling *et al.*, 1998], jumpstarting its widespread use in robotics and automated planning; and *multiagent systems* around 2010, when they were successfully applied to real world scenarios (e.g., autonomous vehicles [How *et al.*, 2008]) and graphical applications (e.g., video games [Hagelbäck and Johansson, 2008]). Some others had valleys, such as *natural language processing* (NLP) around 2000, showing a paradigm shift (from deterministic phrase structure analysis before the 1990s to more probabilistic NLP methods). Long Short-Term Memory (LSTM) recurrent neural net (RNN) models [Hochreiter and Schmidhuber, 1997] have nowadays found rapid adoption due to an increase in computational capacity in the 2010s. Only *machine learning* had a clear increasing trend over time found in both sources, where the steepest slope is located in the 2010s and reveals the current relevance and attention, mostly caused by deep learning and reinforcement learning. Overall, this is in agreement with the general perception about the field.

While the plots show a confirmatory evidence of the changing trend in AI, using area keywords may only present one side of AI. The choice and relevance may be highly dependent on many factors such as the state of the art, the sociological perceptions or a moving target. Thus, we question whether it is possible to investigate the historical evolution in a richer and more systematic manner.

4 A Faceted Analysis of the Evolution of AI

The two dimensions in Table 1 can be used as a basis for a different arrangement of tags and categories in our historical analysis. However, instead of taking a dichotomous and

⁵<http://dblp.uni-trier.de/>

⁶<https://aitopics.org/misc/about>.

monolithic perspective about dimensions, we consider the use of *facets*. This is motivated by realising that when choosing any criterion for analysis, there is always a gradation, and sometimes this gradation does not follow a straight line between two extremes, but an area among two or more edges, like a polygon. Hence, we use the term *facet* for this surface, and the term *edge* for each of its boundaries. To illustrate this, let us develop the first two dimensions into facets:

- F1: The *functionality facet* (with edges ‘techniques’, ‘applications’ and ‘tasks’) analyses the functionality of AI systems, such as knowledge representation, reasoning, learning, communication, perception, action, etc. The *processes* of AI systems fall at the edge *technique*, which relates to how AI systems “think”. We can also characterise AI systems in terms of their behaviour (how they “act”), leading to two different edges: the *tasks* they solve and the *application* areas they are used in. This facet can then be imagined as a triangle.
- F2: The *referent facet* (with edges ‘human’ and ‘universal’) distinguishes definitions or conceptions of AI systems that go from an anthropocentric view to a more universal (theoretical) perspective. At the “human” edge, AI could be characterised by being able to solve all the tasks or by implementing all the intelligent processes humans are able to do. This view would be closely related to what is known as human-like AI (see, e.g., [Lake *et al.*, 2017]) or the view of AI as pursuing human automation [Frey and Osborne, 2017; Brynjolfsson and Mitchell, 2017]. On the other hand, if AI is characterised at the “universal” edge, it would be defined in terms of a more theoretical set of problems or by implementing some abstract processes.

Using the *AI topics* data, for the *functionality facet* we can look at several categories for each of the three edges: techniques, applications and tasks. We made tag-category mappings for about 30 techniques, 20 applications and 30 tasks. A selection of categories for each edge is shown here, although all the plots and results are presented in a separate link⁷. Fig 3 (a) shows the techniques with the highest peak in terms of percentage. We see trends that are consistent with machine learning taking more relevance, with roughly 4-5 times more coverage over the 50 year period, and along the lines of what was seen in the previous section. Fig 3 (b) shows the six most popular application areas. Here the trends are flatter, although some slight trends can be seen for *health & medicine* and *personal assistants*, in line with the insights reported in related work [Fast and Horvitz, 2017]. Finally, Fig 3 (c) shows the relevance of chess, especially in the 1980s and 1990s, and their decrease after Deep Blue beat Kasparov. Poker and Robocup are examples of tasks that became representative over a small period of time because of either an algorithmic breakthrough or the popularity of their competitions.

For the *referent facet*, we can also explore its edges. Fig 4 (a) shows six groups that are most associated with the edge ‘human’. The only relatively clear trends are a fall in *psychology* and an increase in *security, privacy and safety*. While the former may be due to the decline of cognitive modelling

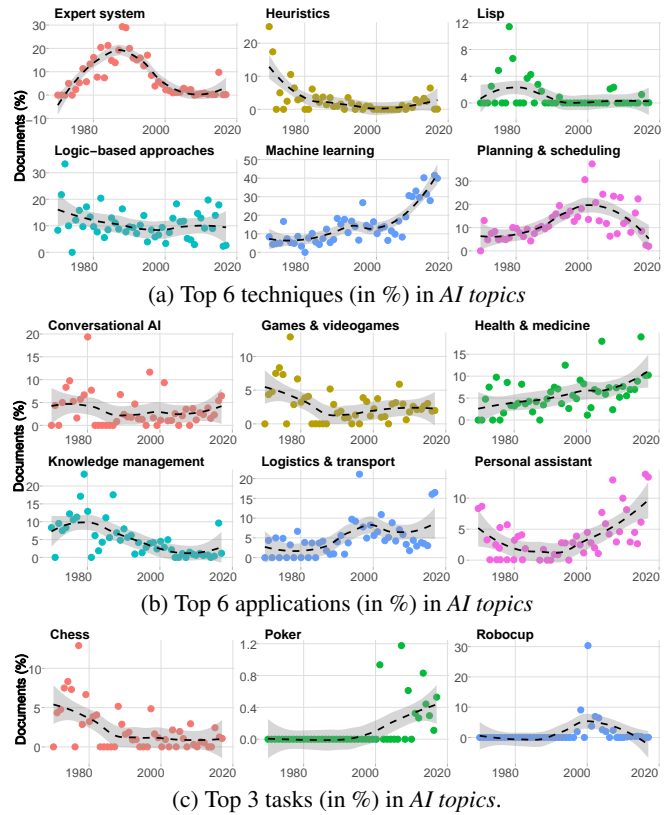


Figure 3: Functionality facet.

in general, the latter clearly underlines the rise of AI ethical and privacy issues already pointed out by several governments and agencies (e.g., [White House, 2016]), as well as the appearance of new regulations (e.g., GDPR [EU Regulation, 2016]), ultimately triggered by public opinion and a concern in the field itself. Fig 4 (b) shows six groups that are most associated with the edge ‘universal’. The representations are rather flat for this edge, with very variable coverage over the whole period under analysis, but no clear trends.

Both the *functionality* and the *referent* facets are usually linked to a large number of terms and keywords, and they give a broad overview that is familiar to AI researchers. However, are these two facets sufficient to characterise the *object* and *subject* of AI? If we think of the elements of AI –its systems and services (the object)–, we need some other relevant facets to characterise them, such as their generality, location and embodiment. If we think of AI as a discipline (the subject), we need to consider its paradigms, actors, character and nature. Let us describe these new seven facets:

- F3: The *generality facet* (with edges ‘specific’ and ‘general’) considers whether AI is concerned with the creation of specific systems (i.e., solving one particular task) or the development of systems (and techniques) that solve (or can be applied to) a wide range of tasks. This dichotomy is usually referred to as narrow versus general AI.
- F4: The *location facet* (with edges ‘integration’ and ‘distribution’) depicts where a system starts and ends. The very notion of centralised or distributed decision making has been at the definitional level of *agent* for a few decades, as well as the related notion of *autonomy* [Luck, 2017].

⁷Data, code and plots, and an online R Shiny app are publicly available at: <https://evoai.shinyapps.io/evoai/>.

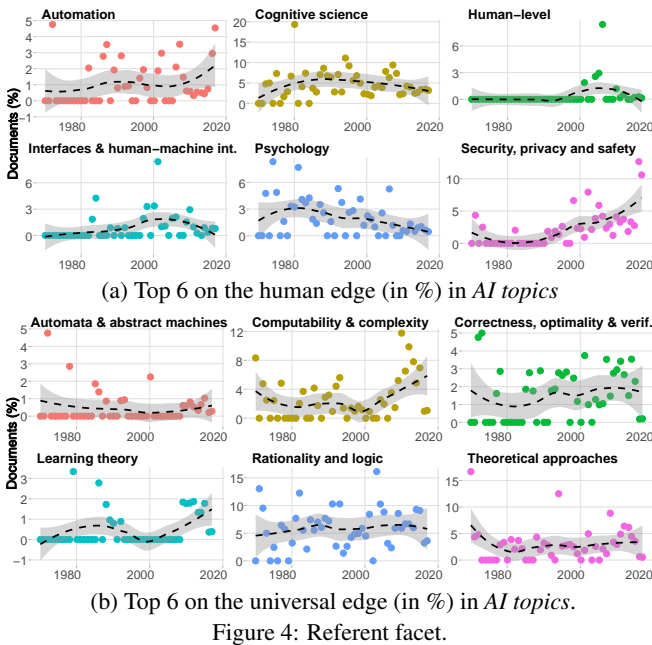


Figure 4: Referent facet.

Also, there is an increasingly blurred boundary on where human cognition ends and where it is assisted, extended or orthosed by AI [Ford *et al.*, 2015] (and vice versa, through human computation [Quinn and Bederson, 2011]).

- F5: The *embodiment facet* (with edges ‘physical’ and ‘virtual’) distinguishes whether the AI system is linked to a body or physical equipment or, on the contrary, is basically of algorithmic character, installed on devices, working on the cloud or migrating between different platforms, usually dealing with elements in a digital world.
- F6: The *paradigm facet* (with edges ‘discrete’, ‘continuous’ and ‘probabilistic’) distinguishes the underlying approaches behind many principles and tools of AI. At the discrete edge, we see those problems and methods seen in a combinatorial way, where a logical or evolutionary process combines or applies operators. At the continuous edge, we see quantitative optimisation problems tackled with gradient descent, kernels and matrix operations. And, on the probabilistic edge, problems are seen in a probabilistic, stochastic or statistical view.
- F7: The *actor facet* (with edges ‘academia’, ‘industry’, ‘government’ and ‘independent’) identifies who are the driving forces behind AI as a discipline. This facet considers who is most relevant in –and ultimately steering– the discipline according to its current challenges, regulations and potential future advances.
- F8: The *character facet* (with edges ‘empirical’ and ‘theoretical’) determines whether AI is guided by experiments, like other empirical disciplines, or whether it is of a more theoretical nature. Note that this facet is different from the referent facet (e.g., a non-anthropocentric view can be very experimental).
- F9: The *nature facet* (with edges ‘technology’, ‘engineering’, ‘science’ and ‘philosophy’) describes AI according to what kind of discipline it is. This loosely corresponds respectively to whether it creates devices and products,

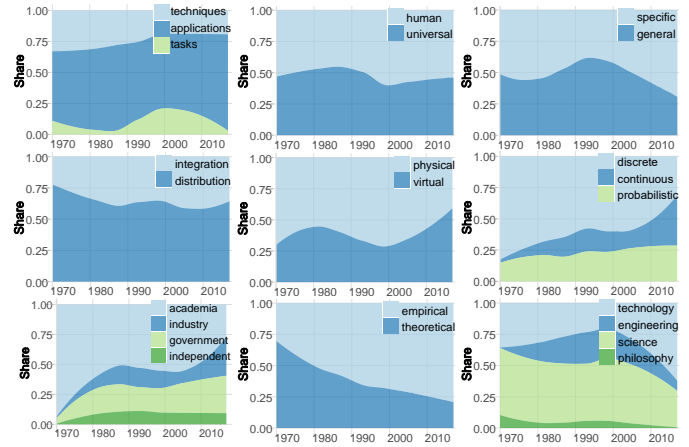


Figure 5: Share of all edges per facet in *AI topics* (facets F1-F9, ordered left to right, top to bottom).

solves problems, answers questions or asks them.

We do not show individual categories for these facets due to space limitations. Instead, we will look at the aggregation of categories by edges, assuming the edges are exclusive (so their share is always 100%). In a way, this can be understood as a non-monolithic view of polarities in sentiment analysis. Fig. 5 shows this share for the nine facets, where the data has been smoothed with a moving average filter. We see trends for many of them: an increase of the relevance of applications for F1 and the focus on more specific systems for F3 (the number, but also the diversification, of applications may explain this), more virtual systems for F5 (with the appearance of many new AI experimentation and evaluation platforms [Hernández-Orallo *et al.*, 2017]), more continuous paradigms for F6 (given the success of deep learning and other methods based on this paradigm), more industry for F7 (and mostly at the cost of academia), a more empirical character for F8 and a view of the discipline in a more technological way for F9 (which may be one possible reason for the current concern in AI safety and governance, usually harder to handle when the engineering and scientific perspectives are weak). Overall, we can also see that the trends in some facets, especially the shifts from the 2000s, may be strongly related (facets F1, F3, F7 and F9).

One important insight that we can gain from the visual output shown in all these plots is that some trends peak, while some others are cyclic. Consequently, the plots have explanatory and confirmatory value about the relevance of different areas and perspectives of how AI is defined. However, we have to be cautious and not use them for forecasting. In fact, from a governance perspective, there is a known bias towards reinforcing the dominant view (the winner takes all) or even confounding the part with the whole (e.g., deep learning with machine learning, and machine learning with AI). Therefore, this quantitative analysis could be used to inform compensatory actions with some less popular but important areas, so that the discipline is ready when dominant paradigms shift or some of their pathways plateau.

5 Qualitative Analysis

The facets are also useful to analyse AI in a more qualitative manner. For example, F1 to F5 are based on the *object* of AI

and can characterise either particular AI systems or AI as a discipline. On the other hand, F6 to F9 focus on the *subject*, being more methodological for the discipline as a whole.

Let us examine the first five facets from the perspective of the AI systems and components. In the end, if we set the goal of AI as building and understanding AI systems then the contours of the discipline will be clear as far as we have a clear notion of what an AI system is. Nonetheless, the question remains: how can we characterise this in a stable way for the years to come?

Looking forward, the referent facet (F2) seems to become very relevant. We are aware of AI systems already exceeding some human capabilities and this will continue to be so during the century. Considering human intelligence as a goal has been a driving force (and will continue to be so in the near future), but it is rather short-sighted. Instead, placing AI as exploring a more universal ‘intelligence landscape’ is not only more inclusive about what AI is (systems solving tasks humans cannot solve would still be part of AI), but represents a Copernican view where humans are no longer at the centre (see Figure 6). For instance, the terms *human-level machine intelligence* or *human-level artificial intelligence* have many issues (is it for all tasks?, what is an average human?, how to extrapolate beyond human level?). Actually, the term –albeit not the definition– is usually replaced by *high-level machine intelligence* [Müller and Bostrom, 2016] or simply AGI (which should rather refer to facet F3).

Of course, even if humans are not taken as a reference, they still have to be put at the central position due to the interaction and impact of AI on them. But this is usually referred to as *human-centred artificial intelligence*. For instance, applications must be prioritised according to human needs, and for many of these applications we want more ‘human-like’ AI [Lake *et al.*, 2017; Marcus, 2018], more ‘human-beneficial’ AI (i.e., safer and taking the values of humans into account [Russell *et al.*, 2015; Amodei *et al.*, 2016]) and also more ‘human-ethical’ AI [Goldsmith and Burton, 2017]. Note that we do not want to replicate certain behaviours humans display (e.g., gender bias) in our AI systems. In many cases, it is not that AI should automate all tasks humans do, or perform better than humans, but AI should also do other tasks or perform very differently. Consequently, there should be no reason why some tasks are excluded from AI when they are solved in a different way. This –known as the ‘AI effect’ [McCorduck, 2004]– is partly motivated because of the specialisation of the solution, as we will discuss below.

The bottom view of Figure 6 opens the contours of AI but requires the definition of the intelligence landscape. That takes us back to the range around functionality facet and Table 1 (right), between thinking and acting. We have seen in the quantitative analysis that techniques, tasks and applications are volatile, so any enumeration of them is going to be incomplete.

An alternative view based on skills and abilities can endure changes much better [Hernández-Orallo, 2016; 2017]. For instance, perception, learning and planning have been substantial in AI. Systems having some of these skills are recognised as AI without further information about the techniques or the particular tasks the elements are applied to, in the same

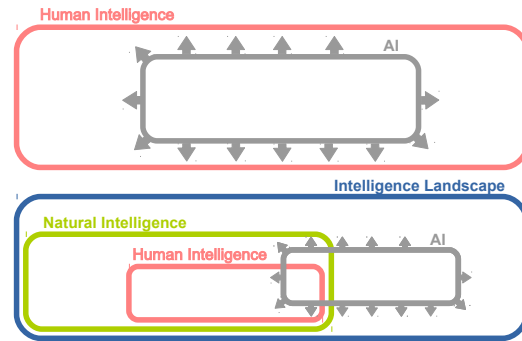


Figure 6: Top: Anthropocentric AI. Bottom: AI in the ‘intelligence landscape’.

way we do for non-human animals. This is actually the true core of facet F1.

But skills and abilities can only be characterised by looking at facets F3, F4 and F5. For instance, having n different systems for navigating n different buildings would not convince us that AI systems today have very good navigation skills (F3, generality). They would just be narrow systems not really having the skill. Similarly, the system might actually be controlled by many subsystems on the Internet, using real-time data from sensors around the building and even from the devices the humans in the building are using (F4, location). Finally, it is not the same to physically navigate a real building than a virtual scenario (F5, embodiment).

In fact, we are heading towards the direction of AI services instead of AI systems, which is typically known as cognition as a service [Spohrer and Banavar, 2015]. In cognitive services research, the focus is on bringing down the overall cost and increasing general performance. However, certain considerations must be taken into account, especially in the context of automation [Frey and Osborne, 2017; Brynjolfsson and Mitchell, 2017]. For example, what is the difference between full automation and efficient semi-automation? Furthermore, the blurred line between work contributed by AI systems and humans makes it even harder to differentiate how performance can be attributed realistically. Also, some tasks can be finally automated in ways where AI plays a secondary role, by changes in the logical or physical configuration. For instance, an intelligent robot can be helped by a proper design of the body, usually referred to as morphological intelligence [Winfield, 2017].

Interestingly, the phenomenon of AI being more powerful because of the use of human data or human computation is mirrored by the *extended mind* [Clark and Chalmers, 1998]. The human mind is viewed as incorporating all of its mind tools, from pen and paper to cognitive assistants. The contours of where the human mind ends, and whether the surrounding tools and devices are included, is also important to determine where the AI system ends too, and who really provides the services. Hence, the facets F3 to F5 will become more relevant and sophisticated in the future. Perhaps instead of autonomy, we will need to trace how much each part contributes to the whole ability of the whole system or service.

Finally, as we mentioned above, the four last facets have a more methodological stance. For instance, it seems irrelevant what paradigm from facet F6 is most important at the moment. For instance, arcade games (such as the Atari ALE

benchmark) can now be played with relatively good performance using deep reinforcement learning [Mnih *et al.*, 2015], evolutionary programming [Kelly and Heywood, 2017] or neuroevolution [Hausknecht *et al.*, 2014]. However, if it is the case that more and more AI systems are using continuous, gradient-descent, approaches rather than more discrete, combinatorial, approaches, this may have an impact on the contours with neighbouring disciplines. AI would be clashing more often with algebra, optimisation and statistics (or even physics), and would be pulled away from the traditional logic and discrete mathematics (or even evolution). Similarly, whether approaches are more empirical or theoretical, according to facet F8, is not necessarily a reason to be considered more or less AI, but can still affect the perception of the field, and the boundaries with some other disciplines.

Facets F7 (actors) and F9 (nature) are related, as it is mostly academia that cares for scientific and philosophical questions about AI. If we look at Figure 6 again, all actors will be interested in building systems and services that cover the intelligence landscape for an increasing number of applications, considering AI as a technology and focusing on engineering. But it is mostly (or only) academia which is interested in what the intelligence landscape looks like, the evaluation of where systems are located in this space and what the implications are while covering this landscape. These and some other areas will have to be prioritised by academia (and government funding) in order to have more vision and governance within the field of AI.

6 Conclusion

Artificial intelligence, as any other scientific discipline, is partly a social phenomenon, and its definition and contour are highly influenced by its actors and stakeholders. We will always need to track and update the field in many ways, from the use of self-reported questionnaires⁸ to data from the venues and news related to AI.

In this paper, we presented a series of facets and associated edges to analyse the historical evolution of AI, and gathered some insight into its future. The data of venues and AI repositories were useful for quantitative analyses. Moreover, the data and the mapping between tags and categories are publicly provided so that others can apply the same faceted framework to other sources of data about AI. Finally, the facets represent a framework to discuss, in a more qualitative way, the past, present and future of AI, and its relation to other disciplines.

Acknowledgements

We thank the support from the Leverhulme Trust via the Leverhulme Centre for the Future of Intelligence (CFI). F. Martínez-Plumed was supported by INCIBE (*Ayudas para la excelencia de los equipos de investigación avanzada en ciberseguridad*). J. H-Orallo received a Salvador de Madariaga grant (PRX17/00467) from the Spanish MECED for a research stay at the CFI, Cambridge, and a BEST grant

(BEST/2017/045) from the GVA for another research stay at the CFI. This work was also supported by the EU (FEDER), and the Spanish MINECO under grant TIN 2015-69175-C4-1-R and the Generalitat Valenciana PROMETEOII/2015/013.

References

- [Amodei *et al.*, 2016] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [Anderson *et al.*, 1997] John R Anderson, Michael Matessa, and Christian Lebiere. Act-r: A theory of higher level cognition and its relation to visual attention. *Human-Computer Interaction*, 12(4):439–462, 1997.
- [Boden, 2016] Margaret A Boden. *AI: Its nature and future*. Oxford University Press, 2016.
- [Brynjolfsson and Mitchell, 2017] Erik Brynjolfsson and Tom Mitchell. What can machine learning do? Workforce implications. *Science*, 358(6370):1530–1534, 2017.
- [Buchanan *et al.*, 2013] Bruce G Buchanan, Joshua Eckroth, and Reid Smith. A virtual archive for the history of ai. *AI Magazine*, 34(2):86, 2013.
- [Buchanan, 2005] Bruce G Buchanan. A (very) brief history of artificial intelligence. *Ai Magazine*, 26(4):53, 2005.
- [Clark and Chalmers, 1998] Andy Clark and David Chalmers. The extended mind. *analysis*, 58(1):7–19, 1998.
- [EU Regulation, 2016] General Data Protection EU Regulation. Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data. *Official Journal of the European Union (OJ)*, 59:1–88, 2016.
- [Fast and Horvitz, 2017] Ethan Fast and Eric Horvitz. Long-term trends in the public perception of artificial intelligence. In *AAAI*, pages 963–969, 2017.
- [Fetzer, 1990] James H Fetzer. What is artificial intelligence? In *Artificial Intelligence: Its Scope and Limits*, pages 3–27. Springer, 1990.
- [Ford *et al.*, 2015] Kenneth M Ford, Patrick J Hayes, Clark Glymour, and James Allen. Cognitive orthoses: toward human-centered ai. *AI Magazine*, 36(4):5–8, 2015.
- [Frey and Osborne, 2017] Carl Benedikt Frey and Michael A Osborne. The future of employment: how susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280, 2017.
- [Goldsmith and Burton, 2017] Judy Goldsmith and Emanuelle Burton. Why teaching ethics to ai practitioners is important. In *AAAI*, pages 4836–4840, 2017.
- [Hagelbäck and Johansson, 2008] Johan Hagelbäck and Stefan J Johansson. Using multi-agent potential fields in real-time strategy games. In *Proc. of the 7 Int. Joint Conf. on Autonomous agents and multiagent systems-Volume 2*, pages 631–638, 2008.
- [Hausknecht *et al.*, 2014] Matthew Hausknecht, Joel Lehman, Risto Miikkulainen, and Peter Stone. A neuroevolution approach to general atari game playing. *IEEE*

⁸For example, [Müller and Bostrom, 2016] or http://agisi.org/Survey_intelligence.html.

- Trans. on Computational Intelligence and AI in Games*, 6(4):355–366, 2014.
- [Hernández-Orallo *et al.*, 2017] Jose Hernández-Orallo, Marco Baroni, Jordi Bieger, Nader Chmait, David L Dowe, Katja Hofmann, Fernando Martínez-Plumed, Claes Strannegård, and Kristinn R Thórisson. A new ai evaluation cosmos: Ready to play the game? *AI Magazine, Association for the Advancement of Artificial Intelligence*, 2017.
- [Hernández-Orallo, 2016] José Hernández-Orallo. *The measure of all minds: evaluating natural and artificial intelligence*. Cambridge University Press, 2016.
- [Hernández-Orallo, 2017] José Hernández-Orallo. Evaluation in artificial intelligence: from task-oriented to ability-oriented measurement. *Artificial Intelligence Review*, 48(3):397–447, Oct 2017.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [How *et al.*, 2008] Jonathan P How, Brett Behihke, Adrian Frank, Daniel Dale, and John Vian. Real-time indoor autonomous vehicle test environment. *IEEE control systems*, 28(2):51–64, 2008.
- [Kaelbling *et al.*, 1998] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- [Kelly and Heywood, 2017] Stephen Kelly and Malcolm I Heywood. Multi-task learning in atari video games with emergent tangled program graphs. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 195–202. ACM, 2017.
- [Kieras and Meyer, 1997] David E. Kieras and David E. Meyer. An overview of the epic architecture for cognition and performance with application to human-computer interaction. *Hum.-Comput. Interact.*, 12(4):391–438, December 1997.
- [Lake *et al.*, 2017] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *BBS*, 40, 2017.
- [Legg and Hutter, 2007] Shane Legg and Marcus Hutter. A collection of definitions of intelligence. *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms, in Frontiers in Artificial Intelligence and applications*, 157:17–24, 2007.
- [Lehman-Wilzig, 1981] Sam N Lehman-Wilzig. Frankenstein unbound: towards a legal definition of artificial intelligence. *Futures*, 13(6):442–457, 1981.
- [Luck, 2017] Michael Luck. IJCAI-2017 Special Track on AI & Autonomy, 2017.
- [Marcus, 2018] Gary Marcus. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*, 2018.
- [McCarthy *et al.*, 2006] John McCarthy, Marvin L Minsky, Nathaniel Rochester, and Claude E Shannon. A proposal for the Dartmouth summer research project on artificial intelligence, august 31, 1955. *AI magazine*, 27(4):12, 2006.
- [McCarthy, 1998] John McCarthy. What is artificial intelligence? 1998.
- [McCorduck, 2004] Pamela McCorduck. *Machines who think: A personal inquiry into the history and prospects of artificial intelligence*. AK Peters Natick, MA, 2004.
- [Minsky, 1968] M. L. Minsky. Semantic information processing. 1968.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Belle-mare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [Moor, 2006] James Moor. The Dartmouth college artificial intelligence conference: The next fifty years. *Ai Magazine*, 27(4):87, 2006.
- [Moran *et al.*, 2014] Kelly Moran, Byron C Wallace, and Carla E Brodley. Discovering better AAAI keywords via clustering with community-sourced constraints. In *AAAI*, pages 1265–1271, 2014.
- [Müller and Bostrom, 2016] Vincent C Müller and Nick Bostrom. Future progress in artificial intelligence: A survey of expert opinion. In *Fundamental issues of artificial intelligence*, pages 555–572. Springer, 2016.
- [Newell, 1994] Allen Newell. *Unified theories of cognition*. Harvard University Press, 1994.
- [Nilsson, 2009] Nils J Nilsson. *The quest for artificial intelligence*. Cambridge University Press, 2009.
- [Niu *et al.*, 2016] Jiqiang Niu, Wenwu Tang, Feng Xu, Xiaoyan Zhou, and Yanan Song. Global research on artificial intelligence from 1990–2014: Spatially-explicit bibliometric analysis. *ISPRS Int. Journal of Geo-Information*, 5(5):66, 2016.
- [Quinn and Bederson, 2011] Alexander J Quinn and Benjamin B Bederson. Human computation: a survey and taxonomy of a growing field. In *Proc. of the SIGCHI conf. on human factors in computing systems*, pages 1403–1412. ACM, 2011.
- [Russell and Norvig, 2009] Stuart Jonathan Russell and Peter Norvig. *Artificial intelligence: a modern approach, 3rd Edition*. Prentice Hall, 2009.
- [Russell *et al.*, 2015] Stuart Russell, Daniel Dewey, and Max Tegmark. Research priorities for robust and beneficial artificial intelligence. *Ai Magazine*, 36(4):105–114, 2015.
- [Shah *et al.*, 2017] Nihar B Shah, Behzad Tabibian, Krikamol Muandet, Isabelle Guyon, and Ulrike Von Luxburg. Design and analysis of the nips 2016 review process. *arXiv preprint arXiv:1708.09794*, 2017.
- [Solomonoff, 2017] Grace Solomonoff. Ray Solomonoff and the Dartmouth summer research project in artificial intelligence, 1956, 2017.
- [Spohrer and Banavar, 2015] Jim Spohrer and Guruduth Banavar. Cognition as a service: an industry perspective. *AI Magazine*, 36(4):71–86, 2015.
- [White House, 2016] Executive Office of the President White House. *Preparing for the Future of Artificial Intelligence*. CreateSpace Independent Publishing Platform, USA, 2016.
- [Winfield, 2017] Alan FT Winfield. How intelligent is your intelligent robot? *arXiv preprint arXiv:1712.08878*, 2017.