

Mapping Intelligence: Requirements and Possibilities

Sankalp Bhatnagar^{1,2}, Anna Alexandrova^{1,4}, Shahar Avin³, Stephen Cave¹,
Lucy Cheke^{1,11}, Matthew Crosby^{1,7}, Jan Feyereis^{8,9}, Marta Halina^{1,4},
Bao Sheng Loe⁵, Seán Ó hÉigeartaigh^{1,3}, Fernando Martínez-Plumed⁶,
Huw Price^{1,3}, Henry Shevlin¹, Adrian Weller^{1,12}, Alan Winfield^{1,10}, and
José Hernández-Orallo^{1,6,13}

¹ Leverhulme Centre for the Future of Intelligence, UK.

² The New School, New York.

³ Centre for the study of Existential Risk, University of Cambridge, UK.

⁴ Department of History and Philosophy of Science, University of Cambridge, UK.

⁵ Psychometrics Centre, University of Cambridge, UK.

⁶ Universitat Politècnica de València, Spain.

⁷ Imperial College, UK.

⁸ AI Roadmap Institute, Czech Republic.

⁹ GoodAI, Czech Republic.

¹⁰ Bristol Robotics Laboratory, UWE Bristol, UK.

¹¹ Department of Psychology, University of Cambridge, UK.

¹² Alan Turing Institute, UK.

¹³ Corresponding Author: jorallo@dsic.upv.es.

Abstract. New types of artificial intelligence (AI), from cognitive assistants to social robots, are challenging meaningful comparison with other kinds of intelligence. How can such intelligent systems be catalogued, evaluated, and contrasted, with representations and projections that offer meaningful insights? To catalyse the research in AI and the future of cognition, we present the motivation, requirements and possibilities for an atlas of intelligence: an integrated framework and collaborative open repository for collecting and exhibiting information of all kinds of intelligence, including humans, non-human animals, AI systems, hybrids and collectives thereof. After presenting this initiative, we review related efforts and present the requirements of such a framework. We survey existing visualisations and representations, and discuss which criteria of inclusion should be used to configure an atlas of intelligence.

1 Introduction

Despite significant AI progress, its pace and direction are largely unassessed and hard to extrapolate. The main reason for this is that we lack the tools to properly evaluate, compare and classify AI systems, and thus determine the future of the field. The comparison of AI systems with human and non-human intelligence is typically performed in an informal and subjective way, often leading to contradicting assessments, especially in hindsight (Hayles, 1996; Brooks, 1997; Pfeifer,

2001; Shah et al., 2016). The comparison between non-human animals and AI ranges from setting the goal of designing artificial agents with the behaviour of an earwig (Kirsh, 1991) to the intelligence of a rat (Cadman, 2014; Shead, 2017), without further specification of what these capabilities or dimensions for comparison should be. The comparison with humans is not much more precise. For instance, two decades ago it was cognitive functions related to perception and action that seemed unattainable – “the gardeners, receptionists, and cooks are secure in the decades to come” said Steven Pinker in 1994. Now, these are the functions that look easier to be automated (Frey & Osborne 2017) – “if a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future” (Ng, 2016). Today, it is higher-level cognition (causal reasoning, compositionality, theory of mind, meta-cognition, etc.) that seems more out of reach (Marcus, 2018).

The assessment is especially difficult as academia and industry in AI are rushing to achieve breakthroughs for specific problems, which often require massive data, computation power, embedded heuristics, strong bias, etc., undermining generality, autonomy and efficiency. For instance, AI can now play most video games (Hessel et al., 2017) and board games (Silver et al., 2017) better than humans, but the immediate training data and computational power that are needed are – as for today – orders of magnitude higher than those used by a human. As a result, it is difficult for policy makers to assess what AI systems will be able to do in the near future, and how the field may get there. There is no common framework to determine which kinds of AI systems are even desirable.

This contrasts with empirical science, where measurements, comparisons, representations and taxonomies are widespread. These characterisations can be theory-driven, such that a prior conceptual framework is used to categorise system features, or can be data-driven, which is increasingly important in many scientific disciplines (Marx, 2013; Landhuis, 2017; Einav & Levin, 2014). Conceptual progress partly relies on finding and testing hypotheses through the computational analysis of large amounts of shared data (Gewin, 2002), using open data science tools (Lowndes et al., 2017). In AI, we would like to analyse the state and progress of artificial systems based on data-grounded investigations. Research priorities and safety concerns depend on this analysis. We need to assess whether new AI systems and techniques are simply an incremental improvement for a narrow collection of applications or a real breakthrough representing a more general cognitive ability, which can be established in relation to comparable abilities in humans and other animals.

This wider view of AI, in the context of all kinds of intelligence, dates back to Sloman’s “space of possible minds” (Sloman, 1984). Figure 1 compares (a) a figurative plot (Shanahan, 2016), covering a wide range of systems (also see (Yampolskiy, 2014, Fig. 3b), (Arsiwalla et al., 2017, Fig. 3c), and (Solé, 2017, Fig. 3d)), with (b) a plot depicting precise experimental results for several ape species on a battery of tests (Herrmann et al., 2007). The figure illustrates a visible trade-off between completeness and empirical grounding. What we need is to leverage the best of both worlds: a data-based representation of very different

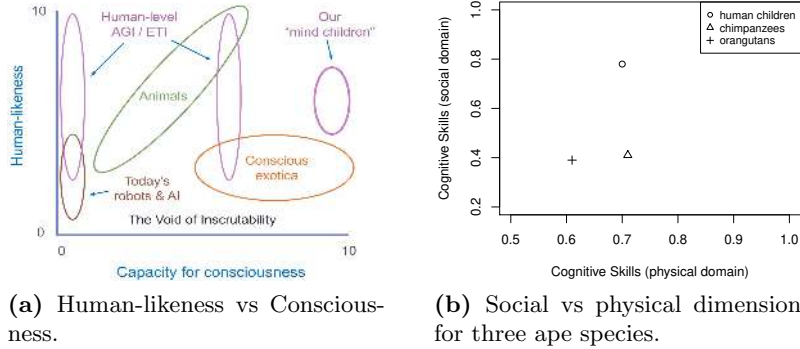


Fig. 1: Different kinds of minds represented according to several dimensions. Left: Figurative human-likeness vs consciousness (from Shanahan 2016). Right: Two dimensions of cognitive skills (social vs physical domain) according to the results of a test battery on three different groups of apes (adapted from Herrmann et al. 2007).

cognitive systems, including humans, non-human animals, AI systems, hybrids and collectives, where actual measurements can be aggregated and combined.

This requires a novel platform, an ‘atlas of intelligence’, that integrates an extensive inventory of cognitive systems, a behavioural test catalogue (with test batteries that could be aggregated into dimensions) and an experimentation repository (results from measurements). The platform would be populated collectively, facilitating cross-comparison and reproducibility (Aarts et al., 2015; Vanschoren et al., 2015). The atlas would represent a new cartographical endeavour for a better understanding of the geography of the space of intelligence.

This paper explores the motivations, the requirements and the possibilities of such an atlas. Section 2 explores in more depth why the atlas is needed in terms of four lists of items specifying the motivations, applications, dimension manipulations and entities to be covered in the atlas. Section 3 focuses on the idea of an atlas as a set of maps, and configures a partial specification in terms of the maps we would like it to have. This section includes a collection of maps and graphical representations, some of them already proposed in the literature (but most without real data) and some desired representations. We close the paper with a discussion about future work. Finally, appendix A gives a short overview of similar initiatives in other areas, and how these relate to the atlas.

2 Motivations, applications, dimensions and kinds

This section presents a series of lists of items covering the motivations and applications of the atlas (why, and for what, an atlas is needed), and the potential dimensions and kinds of systems to be included (what the contents should be). The lists are not meant to be exhaustive and free from overlaps (some ideas are represented by several items with different perspectives), but rather to serve as initial items for discussion and refinement.

Motivations

The motivations are meant to highlight the needs for an atlas of intelligence. We identify them following scientific, technological and societal needs that are recognised at present or in the near future. Most of them focus on better understanding, representing and cataloguing what we know about different kinds of intelligence. Still, we do not exclude the needs for anticipation, so we also cover those motivations that are related to having better predictions about the existing and future changes of human and artificial intelligence.

- **Milestones and Pathways:** Unlike most non-human biological cognition, human cognition is changing: the average IQ in many countries is increasing (the *Flynn effect*), our memory (Sparrow et al., 2011) is changing due to the *Google effect* (digital amnesia), navigation abilities (McKinlay, 2016; Milner, 2016) atrophied because satnavs, cognitive rewards mechanisms are changing because of gamification, etc. This is a process that is accelerated by technology, and will be magnified by the use of cognitive assistants and cognitive prosthetics, especially for the elderly. AI itself and human-machine hybrids (either as individual cyborgs or as mixed collectives) are progressing in directions that we are not able to compare with the past or extrapolate, in order to understand where all this is leading, and the associated opportunities and risks (research priorities and safety concerns).
- **Laypeople Understanding:** In those cases where comparisons can be made by looking at a set of traits, it is usually too complex for non-experts to understand what the key differences are between two cognitive systems, especially when one is natural and the other is artificial. Visual representations are appropriate, as humans are good at understanding geographical analogies (e.g., the 1948 book “the map that came to life” helped children understand the countryside where a trajectory and a story were accompanied by maps).
- **Crossover Measuring:** Data-driven comparison is usually based on measurement instruments, reporting a series of measured values that can be represented. But we do not have many test batteries that can be applied across species or even AI systems. The generalisation of representations where different natural species and AI technologies are put together would encourage the adoption and definition of more universal tests having better measurement invariance across different entities.
- **Behavioral Taxonomies:** If we go beyond life in our comparisons, especially if they are based on similarity, the dominant genotypic approach cannot be used broadly. Taxonomies and models must be mostly informed by behavioural analyses, in contrast to phenotypic, ethological, genotypic or neurological approaches (Cattell & Coulter, 1966; Miller, 1967). But we contemplate cladistic principles (using hierarchies or dendrograms) and we consider morphological or functional similarities as far as they affect behaviour.
- **Testing New Intelligence:** The progress in AI suggests that task-oriented evaluation (i.e., the performance of an AI system for a particular task) may be insufficient. Other ways of characterising and measuring AI are needed. While some capabilities and tests can be inherited or extended from psychometrics

or animal cognition, there may be some other capabilities or skills that are completely new, especially when we analyse the cognitive profile of human-machine hybrids or collectives.

- **Critical Perspective:** There is an urgent need for better understanding the way the intelligence landscape is changing, for both humans and AI systems, in areas such as automation, education and ethics. It is hard to regulate or incentivise some actions not knowing how they affect the intelligence landscape.
- **Beyond Anthropocentrism:** While it is generally accepted that intelligence is the product of evolution, it is still hard to recognise intelligence in other species or in AI systems, and compare it without using humans as a yardstick.
- **Grand Goals:** While interdisciplinarity in the study of intelligence has increased, there are still many attributes and behaviours that are not properly mapped between disciplines, and there is no wide recognition of a shared space. The geographical analogy of an intelligence landscape as an opportunity for exploration and discovery can help inspire the next generation of researchers in areas such as comparative cognition, psychology, philosophy and artificial intelligence, and, most especially, in multidisciplinary domains.
- **Replicability and Reuse:** New research procedures and visualisations for the analysis of cognitive systems are difficult to apply to other systems or in other contexts. This limitation is more blatant when we see similar ideas, representations or experimental protocols appear in different disciplines.
- **Data-driven and Hypothesis-driven:** When cognition is analysed in one species or a particular AI technology, there is a lack of a sufficiently wide sample to infer and reject hypotheses. The recent trend of a more collaborative data science approach should encourage initiatives where data from different disciplines can be put together to test hypotheses about cognition.

Some of the motivations above have deep roots in cognitive science, comparative psychology, philosophy and AI (Macphail, 1987; Thagard, 2009; Gentner, 2010), but others are more specific to some particular areas or emphasise the need of better representations and comparisons.

Applications

Moving from what is needed to the things an atlas will make possible leads us to the identification of new possibilities and transformations. The criteria for inclusion are such that the list covers potential applications for scientists, philosophers, educators, policy-makers and the general public, directly using the platform or as an indirect result of its use:

- **(Re-)Education:** Traditionally, children and adults used animals as models of different personalities and capabilities, interacting with them regularly. Today, in urban societies with less contact with animals, it is becoming easier to portray and transmit some concepts using robots as models, as cinema and advertising (especially when targeting children) have already understood. An atlas covering animals and robots could be used in museums, schools and universities as a way of articulating over this intelligence landscape.

- **Effective Navigation:** An atlas, with different representations, would help us locate where we are (humans and AI), the trajectories taken in the past years and the destinations we are heading to, helping to visualise whether some targets or trajectories can take us to dangerous areas.
- **Ethical Assessments:** Visual representations make some ethical dilemmas about moral agency and patiency more explicit, as we can see whether the way we look at animals and artificial agents is different from the way measured traits put them on some representations. This will make some ethical issues more conspicuous (animal, robot or human suffering, uncanny valleys, etc.).
- **Consequences:** Not only the locations but also the distributions and densities would help us analyse (especially in advance) the population of creatures affected by research, law, environment, technology, etc., in a critical way. In other words, the maps could also be used to represent the areas and entities (and how many) would be affected by a phenomenon.
- **De/Re-Centre Humans:** Humans, as a species, groups and individuals could be located at different locations depending on the representation, making more explicit that there is a Copernican revolution in the way intelligence is seen today, sustained in the progress of comparative cognition, evolutionary psychology and, increasingly, artificial intelligence.
- **Metaphors and Narratives:** An atlas would build upon the perception we have about animal behaviour. This would help us better understand and locate where we are in AI, in a more meaningful way than just saying AI is at the level of the rat. Instead we would like to align the cognitive profile of a rat with the cognitive profile of a particular AI system, and see the differences in a less monolithic way.
- **Archival Exploration:** An atlas of intelligence would also help to see a history of intelligence, where we would go from extinct animals and past computer/AI systems to the present day, seeing the directions their evolution has taken according to different dimensions.
- **Morgan’s Canon:** C. Lloyd Morgan stated: “In no case is an animal activity to be interpreted in terms of higher psychological processes if it can be fairly interpreted in terms of processes which stand lower in the scale of psychological evolution and development” (Morgan 1903). An atlas would help interpret, extend or overhaul this canon for artificial systems, hybrids or collectives.
- **Unification:** An atlas would require and hence would encourage the definition of more general tests and metrics, embracing natural and artificial systems, and would aim at more unified theories of cognition, going beyond human psychology and evolution to consider every possible cognitive system, especially looking at those places in the maps where there are gaps, whether it is possible to have entities there and how they would be interpreted.

Some of the applications clearly derive from the motivations (e.g., beyond anthropocentrism and de/re-centre humans) but others represent possibilities that perhaps were not even recognised as a necessity, such as the use of the atlas for archival exploration, which may lead to unforeseen purposes. Concerning the needs and possibilities introduced above, represented by the motivations and

applications, we add the *dimensions* and the *kinds of systems* we want to cover, which specify the atlas in general terms.

Dimensions

We are aware of the lack of consensus about the most relevant attributes for the analysis of cognition – not to mention general theories. Because of this disagreement, we want the atlas to be able to integrate different perspectives and attributes of the interest. Consequently, rather than enumerating the specific dimensions of representation that could be used, which could ultimately be created and refined by the users, we clarify how these dimensions operate in general terms.

- **Observation-Based:** the dimensions of representation should be agnostic to particular hypotheses, so that the users could do their theories from the values observed. Of course, there are always some underlying assumptions (and the influence of underlying theories) whenever an observation or measurement is made, but this should be as explicit as possible.
- **Multiple Interface:** the atlas should allow users to project or aggregate the data and derive some maps and other representations from these transformations, as usual in other visualisation frameworks.
- **Interactive Querying:** the atlas could be interrogated through queries, including filters and joins across different data sources, in an interactive way, as in tools of analytical processing.
- **Creative and Constructive:** the atlas should allow users to combine elements, creating new features (and hence new spaces) and creating new entities, such as populations or individuals, combining their cognitive entities under some specified models.
- **Populational / Theoretical:** the elements to represent could correspond to actual populations or subgroups but also to theoretical elements and groups.
- **Bottom-Up / Top-Down:** the dimensions could correspond to basic psychological mechanisms or to more abstract, integrated skills. The atlas should allow users to aggregate and disaggregate these dimensions.
- **Transversal Connections:** the atlas would allow users to combine behavioural traits (skills, functions, capabilities) with non-behavioural traits (physical traits, computational effort, evolutionary traits, etc.).
- **Topographical/Geographical Visualisation:** the atlas should combine as many elements of visualisation and representation (colours, contours, textures) as may be found useful to show the information in insightful ways.

Despite the intended flexibility, some of these dimensional operations give a more precise account at the specification level for the atlas on how data, hypotheses and visualisations must be connected. For instance, the multiple interface, the interactive querying and the topographical and geographical representations very much resemble some common retrieval and representational systems powered by data visualisation tools. On the other hand, the other dimensional characteristics are more aligned with the management of conceptual ontologies and taxonomies.

Kinds of systems

Finally, regarding the kinds of cognitive systems to be represented, we want to cover all possible ranges, according to several criteria: integration, nature, time, distribution and existence. This comprehensive view of cognitive systems would ultimately allow us to put very different types of entities into comparison.

- **General and Narrow:** specific systems aiming at a single task or species in a narrow environment could be covered, as well as those systems that are flexible in a broad range of environments.
- **Individual and Collective:** individual entities could be located as well as collectives (along with their components).
- **Biological and Artificial:** living beings, including plants and animals, and artificial systems, including autonomous agents, robots, corporations, etc.
- **Hybrid (Extended/Enhanced Minds):** humans improved by technology, either internally (enhanced, as cyborgs or through nootropics) or externally (extended by assistants), as well as AI systems using human computation.
- **Novel and Old:** covering current living beings and AI systems, but also extinct species and AI systems of the past.
- **Distributed and Centralised:** systems that are identified by a single body, but also natural and artificial swarms as well as distributed intelligence, including societies.
- **Alien and Fictional:** even for speculation or theorisation, the atlas could also show some imaginary entities.

Apart from the scientific questions needed to build such a platform, its success depends on the engagement of the (research) community and other stakeholders. It is crucial then to identify whether the needs, dimensions and elements represented are well aligned with the potential users and contributors. Consequently, we conducted a preliminary survey to get feedback from researchers and other potential users in many different areas, using the items described in this section. We targeted different communities: artificial intelligence, animal cognition, psychology, philosophy, design and some others. The results of the questionnaire were positive in general. This was not taken as a justification or validation of the categories presented here but, more on the contrary, as a way of recognising omissions, duplications or desiderata nobody is asking for. We focused especially on the open comments from some respondents who were more critical¹⁴.

We considered the previous motivations, intended applications, dimensions to consider, and entities to cover to be a sufficient reason for starting the construction of an atlas, with the necessary caution about potential pitfalls and the need of selecting pieces of the atlas that could be chosen as more low-hanging fruits of the whole project. The previous lists are preliminary, and the priorities for selecting which categories are most important to start with — e.g., prototypes or first cornerstones of the project — are still subject to debate.

Next, we refine this first conception of the atlas by considering existing representations and maps.

¹⁴A detailed analysis of the questionnaire can be found in (Bhatnagar et al., 2017).

3 Collections of maps: representational possibilities

As an atlas is a set of maps, in this section we collect and recreate some of the maps that have been proposed in the past, most of them at a figurative level, and discuss representations that we would like to include in the future. Figure 1 contained examples of a classical multidimensional representation (although two dimensions are especially fitted for paper and screens). The axes represent dimensions of interest and the points represent the entities (the cognitive systems) we want to compare. We will see many others of these, being different because of the dimensions that are chosen or the elements that are represented. In other cases, the representations detach from this multidimensional view but still remain meaningful in geographical or topological terms.

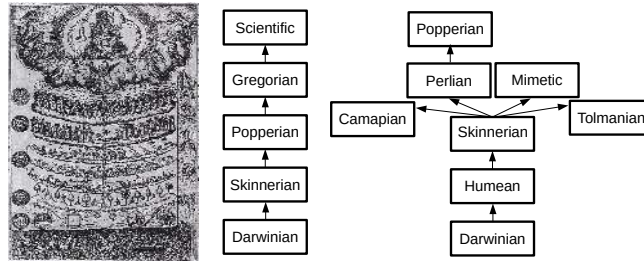



Fig. 2: *Left:* Scala naturae, as depicted in the 16th century (de Valadés, 1579). *Middle:* a representation of Dennett’s Tower of Generate and Test, which depicts creatures according to when and how they adapt (Dennett, 1995), *Right:* Godfrey-Smith’s refinement of the bottom part of Dennett’s tower (the part corresponding to cognitive evolution) in the form of a tree (Godfrey-Smith, 2015, Fig. 2).


We start with the oldest and simplest representations, those inspired in the *scala naturae*, which are monolithic, or at most, arboreal (see Figure 2), where membership to a species is replaced by other criteria for classification. At some point the categorical representations (monolithic or hierarchical) led to more quantitative and multidimensional representations, as we see in Figure 3.

Moravec was not the first one to compare animals and computers according to several dimensions, but some of his plots had an important effect on the narratives about how far AI had come in the 1990s. For instance, Figure 3a compares computational power (speed and storage capacity) for a wide range of entities.


Some other representations have tried to compare animals and artificial systems for other dimensions. For instance, computational efficiency can be replaced by an estimation of energy consumption (Winfield, 2014), which is a physical property that can be used as a dimension alongside some other more behavioural traits. One common representation is based on Venn diagrams, where the sizes and locations are completely arbitrary, and the only purpose is to show a diversity and inclusions/overlaps between sets, such as Figure 3b from Yampolskiy (2014). Some other plots are more speculative, especially when the goal is to represent consciousness, such as the one from Arsiwalla et al. (2017) in Figure 3c.




(a) Comparison of hardware between several living and inanimate objects (Moravec, 1998, Fig. 1)




(b) Figurative space of minds (Yampolskiy, 2014, Fig. 1)



(c) “Morphospace of consciousness” from (Arsiwalla et al., 2017, Fig. 3)



(d) “Biological computational morphospace” (Solé & Macia, 2011, Fig. 7)



(e) Cognitive space of human-robot interactions (Solé, 2017, Fig. 2)



(f) Illustration of Hans Moravec’s “landscape of human competence” (Tegmark, 2017, Fig. 2.2)

Fig. 3: A collection of figurative maps of intelligence.

Other comparisons are at a much more physical (or implementational) level, such as the one from Solé (2017), representing the “morphospace” in terms of “embedding”, “diversity” and “parallelism”, shown in Figure 3d, or represent some aspects of human-robot interaction, again figurative (Figure 3e). An interesting twist is given when the space represents the tasks or abilities (without any clear criterion for proximity), but the Z-dimension (height) is represented

by time (or progress in AI). According to this, we have a figurative plot like Tegmark’s representation (Fig. 3f) of Moravec’s landscape (Tegmark, 2017).

A more thoughtful analysis of dimensions may lead to more than three elements, whose representation (if all of them are quantitative) is more cumbersome. Star (cobweb) plots are a practical option here, although they can get too messy if too many individuals are shown. Also, trajectories are more difficult to represent in these plots. Figure 4 shows how four dimensions are used to compare the intelligence of several organisms.

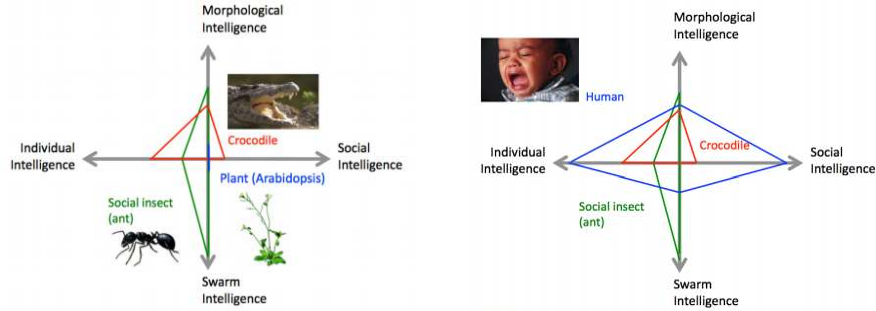


Fig. 4: Comparison of different systems on a space of four dimensions, using star plots (Winfield, 2017, Figs. 2 and 3).






Fig. 5: Collective diversity in terms of psychometric profiles for two figurative groups of five agents, shown with circles. For each plot, the x dimension represents IQ score and the y dimension represents social sensitivity. The mean point is shown with the cross and the maximum and minimum envelopes are represented with a triangle and a square respectively (Hernández-Orallo, 2017, Fig. 15.3).

Following the comments of some of the respondents of the questionnaire, we are also interested in representations of ‘collective intelligence’, even if figurative. For instance, Fig. 5 represents a profile of members of a team and tries to derive aggregate values (minima, maxima and means) for the group.


So far, all the previous representations were figurative, in the sense that there was no measured data or observations from which the maps were represented, but just some general knowledge and intuitions of these magnitudes. In what follows, we include some representations that are using real data. For instance, the easiest way of comparing two systems or species is by comparing their results




(a) Comparison between average humans and AI on CIFAR (Eckerley et al., 2017)




(b) Comparison between crows and monkeys for working memory (Balakhonov & Rose, 2017)




(c) Learning curves by frames seen on the ALE benchmark (Hessel et al., 2017, Fig. 1)



(d) Progress (average performance) by year on ALE benchmark (Eckerley et al., 2017)



(e) Comparing inductive reasoning and numeric ability over time (Schaie, 1996)



(f) Comparing inductive reasoning and numeric ability directly (Schaie, 1996)

Fig. 6: A collection of existing empirically-grounded maps.

for the same task, as in Figure 6a and 6b. But we can also compare abstract or aggregated traits or skills, as we showed in Figure 1. A representation that is becoming very common in AI is to show the results normalised by human performance (Figures 6c and 6d), even in cases where many tasks are aggregated.

While these representations are common and useful, they do not fit the geographical representation of the atlas well. In other words, these plots are not meant to compare AI systems and humans. They are just meant to compare AI systems, where human data is just used to make the results for several games somewhat commensurate when aggregated. This means that the space is anthropocentric, where humans would always be at 100% – a Ptolemaic model. Indeed, for both plots one of the dimensions does not apply to humans. For instance, in Figure 6c, human accuracy is achieved with a number of frames that is at most in the small millions, and also in Figure 6d we cannot properly compare the year humans were introduced with the year a ML technique was introduced.




Fig. 7: Different cognitive systems according to the perception by people (from a survey, where the dimensions have been reduced to two dimensions by PCA). (Gray et al., 2007; Wegner & Gray, 2017)

Trajectories can also be compared over time, as shown in Figure 6e. Here, time is applied to the same entity, so we see how the entity (a population in this case) changes with time. But a trajectory is better seen when the dimensions of the plot are not time – time is not usually represented in a static map. Instead, one can see how an individual or group moves in a space of dimensions chronologically (learning episodes, cognitive decline or enhancement, etc.), illustrated in Figure 6f.

Actual data can also be obtained and processed from subjective perception. For instance, Gray et al. (2007) extract two principal components: agency and experience (what we could also refer to as ‘patience’) in order to quantify how much mind people ascribe to different kinds of cognitive systems, from robots to dead people, as illustrated in Figure 7.

After all these graphical representations, the question is how these can help us configure a set of relevant maps we would like the atlas to have. First, we can look at the elements: many are multidimensional and it is just the dimensions and the elements portrayed which make them really distinctive. This is an advantage, as many of these plots could be generated with a standard tool and interface if we had the data and we could choose the dimensions and elements. An interactive interface could be used as in other exploration tools (e.g., analytical processing or visualisation tools). Second, it is appropriate to look at the

purpose of each of these representations and see whether they correspond to the needs and applications we identified in previous sections. For instance, Figures 1, 2, 3b, 3c, 3d, 3e, 4, 6b, 6c, 6e, and 6f are mostly explanatory or differential in purpose, while Figures 3a, 3d, 3f, 5, 6a and 6d seem to have a more forecasting intention. Some have a broader coverage of kinds of intelligence (Figures 3a, 3b, 3c, 4 and especially 7) and others are more specific.

4 Conclusions

This paper has presented the first steps of an atlas of intelligence, which at this stage must focus on the elicitation of needs (in terms of motivations and requirements) and possibilities (applications, representations and kinds of entities covered). After this analysis, we now have a much better account of how wide the initiative is. The next steps should focus on recognising the applications that might have more impact and are more feasible in the short term. This assessment would allow us to establish the specification of the atlas in a progressive way, so that an essential part of it can be designed and enriched over time. For such an ambitious approach, it is important to think big, as we have done here, while starting small, and grow incrementally.

Apart from the instrumental purpose of this paper as a first step in the development of an atlas of intelligence, this work (independently of how far the atlas develops in the future) brings attention to methodological issues (and related philosophical and theoretical) issues in all disciplines related to intelligence and cognition. Scientists in these disciplines usually see themselves as explorers, but exploration involves much more than discovering and inventing. Scientists also need (to be) cartographers, curators and taxonomists in order to structure, facilitate and disseminate what is known, and assess their unknowns, prioritise their goals and see their progress in perspective. In the same way Linnaeus changed the way living beings were described, catalogued and named, motivating new lines of research, this initiative will help to establish the parameters and the instruments to properly handle and understand the space of existing and future cognitive systems, and exploit its research possibilities.

Acknowledgements The initiative was supported by the Leverhulme Trust via the Leverhulme Centre for the Future of Intelligence. J. H-Orallo and F. M-Plumed were supported by EU (FEDER) and the Spanish MINECO under grant TIN 2015-69175-C4-1-R and by GVA under grant PROMETEOII/ 2015/013 and by the Air Force Office of Scientific Research under award number FA9550-17-1-0287. J. H-Orallo also 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 also at the CFI. F. M-Plumed was also supported by INCIBE (*Ayudas para la excelencia de los equipos de investigación avanzada en ciberseguridad*). A. Weller acknowledges support from the David MacKay Newton research fellowship at Darwin College, the Alan Turing Institute under EPSRC grant EP/N510129/1 & TU/B/000074, and the Leverhulme Trust via the CFI.

References

- Aarts, A., Anderson, J., Anderson, C., Attridge, P., Attwood, A., & Fedor, A. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251), 1–8.
- Allen Institute for Brain Science. (2016). *Allen brain observatory*. Retrieved from <http://observatory.brain-map.org/visualcoding>
- Arsiwalla, X. D., Moulin-Frier, C., Herreros, I., Sanchez-Fibla, M., & Verschure, P. (2017). The morphospace of consciousness. *arXiv preprint arXiv:1705.11190*.
- Balakhonov, D., & Rose, J. (2017). Crows rival monkeys in cognitive capacity. *Scientific reports*, *7*(1), 8809.
- Bhatnagar, S., Alexandrova, A., Avin, S., Cave, S., Cheke, L., Crosby, M., ... Hernández-Orallo, J. (2017). *A First Survey on an Atlas of Intelligence*. Retrieved from <http://www.dsic.upv.es/~flip/papers/Bhatnagar18.SurveyAtlas.pdf>
- Boero, F., & Bernardi, G. (2014). Phenotypic vs genotypic approaches to biodiversity, from conflict to alliance. *Marine genomics*, *17*, 63–64.
- Brooks, R. A. (1997). From earwigs to humans. *Robotics and Autonomous Systems*, *20*(2-4), 291–304.
- Cadman, E. (2014). AI not just a game for DeepMind’s Demis Hassabis. *Financial Times*. Retrieved from <https://www.ft.com/content/1c9d5410-8739-11e3-9c5c-00144feab7de>
- Castelvecchi, D. (2016). Tech giants open virtual worlds to bevy of ai programs. *Nature*, *540*(7633), 323–324.
- Cattell, R. B., & Coulter, M. A. (1966). Principles of behavioural taxonomy and the mathematical basis of the taxonome computer program. *British Journal of Mathematical and Statistical Psychology*, *19*(2), 237–269.
- Dennett, D. C. (1995). Darwin’s dangerous idea. *The Sciences*, *35*(3), 34–40.
- de Valadés, D. (1579). *Rhetorica christiana*.
- Eckersley, P., Nasser, Y., et al. (2017). *EFF AI Progress Measurement Project*. Retrieved 2017-01-10, from <https://www.eff.org/es/ai/metrics>
- Einav, L., & Levin, J. (2014). Economics in the age of big data. *Science*, *346*(6210), 1243089.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: how susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280.
- Gentner, D. (2010). Psychology in cognitive science: 1978–2038. *Topics in Cognitive Science*, *2*(3), 328–344.
- Gewin, V. (2002). Taxonomy: All living things, online. *Nature*, *418*(6896), 362–363.
- Godfrey-Smith, P. (2015). Towers and trees in cognitive evolution. In B. Huebner (Ed.), *The philosophy of daniel dennett* (chap. 8.1). Oxford University Press.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, *315*(5812), 619–619.

- Hayles, N. K. (1996). Narratives of artificial life. *Future Natural: Nature, Science, Culture*. London: Routledge, 146–64.
- Hernández-Orallo, J., Baroni, M., Bieger, J., Chmait, N., Dowe, D. L., Hofmann, K., ... Thórisson, K. R. (2017). A new AI evaluation cosmos: Ready to play the game? *AI Magazine*, 38(3), 66–69. Retrieved from <https://www.aaai.org/ojs/index.php/aimagazine/article/view/2748>
- Hernández-Orallo, J. (2017). *The measure of all minds: evaluating natural and artificial intelligence*. Cambridge University Press.
- Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B., & Tomasello, M. (2007). Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science*, 317(5843), 1360–1366.
- Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., ... Silver, D. (2017). Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*.
- Kirsh, D. (1991). Today the earwig, tomorrow man? *Artificial intelligence*, 47(1-3), 161–184.
- Landhuis, E. (2017). Neuroscience: Big brain, big data. *Nature*, 541(7638), 559–561.
- Leslie, M. (2005). Calling all taxonomists. *Science*, 307(5712), 1021–1022.
- Lowndes, J. S. S., Best, B. D., Scarborough, C., Afflerbach, J. C., Frazier, M. R., OHara, C. C., ... Halpern, B. S. (2017). Our path to better science in less time using open data science tools. *Nature Ecology & Evolution*, 1, 0160.
- Macphail, E. M. (1987). The comparative psychology of intelligence. *Behavioral and Brain Sciences*, 10(4), 645–656.
- Marcus, G. (2018). Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*.
- Marx, V. (2013). Biology: The big challenges of big data. *Nature*, 498(7453), 255–260.
- McKinlay, R. (2016). Use or lose our navigation skills: automatic wayfinding is eroding natural abilities, warns roger mckinlay. *Nature*, 531(7596), 573–576.
- Midford, P. E. (2004). Ontologies for behavior. *Bioinformatics*, 20(18), 3700–3701.
- Miller, R. (1967). Task taxonomy: Science or technology? *Ergonomics*, 10(2), 167–176.
- Milner, G. (2016). *Pinpoint: How gps is changing technology, culture, and our minds*. WW Norton & Company.
- Moravec, H. (1998). When will computer hardware match the human brain. *Journal of evolution and technology*, 1(1), 10.
- National Research Council. (2011). *Toward precision medicine: building a knowledge network for biomedical research and a new taxonomy of disease*. National Academies Press.
- Ng, A. (2016). What artificial intelligence can and can't do right now. *Harvard Business Review*, 9.
- Parr, C. S., Wilson, N., Leary, P., Schulz, K. S., Lans, K., Walley, L., ... R. J. Corrigan, J. (2014). *The encyclopedia of life v2: Providing global access to knowledge about life on earth* (Vol. 2). Retrieved from <http://www.eol.org/>

- Pfeifer, R. (2001). Embodied artificial intelligence 10 years back, 10 years forward. In *Informatics* (pp. 294–310).
- Pinker, S. (1994). *The language instinct: How the mind creates language*. William Morrow and CompanyK.
- Roskov, Y., Abucay, L., Orrell, T., Nicolson, D., Bailly, N., Kirk, P., ... Penev, L. (2018). *Species 2000 ITIS catalogue of life*. Aeon. Retrieved from <http://catalogueoflife.org/col>
- Schaie, K. W. (1996). Intellectual development in adulthood. In J. E. Birren & K. W. Schaie (Eds.), *Handbook of the psychology of aging* (Vol. 4, pp. 266–286). Academic Press Inc.
- Shah, H., Warwick, K., Vallverdú, J., & Wu, D. (2016). Can machines talk? comparison of eliza with modern dialogue systems. *Computers in Human Behavior*, *58*, 278–295.
- Shanahan, M. (2016). Conscious exotica. from algorithms to aliens, could humans ever understand minds that are radically unlike our own? *Aeon*. Retrieved from <https://aeon.co/essays/beyond-humans-what-other-kinds-of-minds-might-be-out-there>
- Shead, S. (2017). Facebook’s AI boss: In terms of general intelligence, we’re not even close to a rat. *Business Insider*. Retrieved from <http://uk.businessinsider.com/facebooks-ai-boss-in-terms-of-general-intelligence-were-not-even-close-to-a-rat-2017-10>
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., ... others (2017). Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.
- Slovan, A. (1984). *The structure and space of possible minds*. School of Cognitive Sciences, University of Sussex.
- Solé, R. (2017). Rise of the humanbot. *arXiv preprint arXiv:1705.05935*.
- Solé, R. V., & Macia, J. (2011). Synthetic biocomputation: the possible and the actual. In *Ecal* (pp. 29–36).
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *science*, 1207745.
- Stuart, S., Wilson, E., McNeely, J., Mittermeier, R., & Rodríguez, J. (2010). The barometer of life. *Science*, *328*(5975), 177–177.
- Tegmark, M. (2017). *Life 3.0: Being human in the age of artificial intelligence*. Knopf.
- Thagard, P. (2009). Why cognitive science needs philosophy and vice versa. *Topics in Cognitive Science*, *1*(2), 237–254.
- Vanschoren, J., Bischl, B., Hutter, F., Sebag, M., Keggl, B., Schmid, M., ... Lawrence, N. (2015). Towards a data science collaboratory. *Lecture Notes in Computer Science (IDA 2015)*, 9385.
- Vanschoren, J., van Rijn, J. N., Bischl, B., & Torgo, L. (2013). Openml: Networked science in machine learning. *SIGKDD Explorations*, *15*(2), 49–60. Retrieved from <http://doi.acm.org/10.1145/2641190.2641198> doi: 10.1145/2641190.2641198
- Wegner, D. M., & Gray, K. (2017). *The mind club: Who thinks, what feels, and why it matters*. Penguin.

- Winfield, A. F. (2014). Estimating the energy cost of (artificial) evolution. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, & H. Lipson (Eds.), *Proc. 14th int. conf. on the synthesis and simulation of living systems (alife)* (pp. 872–875). MIT Press.
- Winfield, A. F. (2017). How intelligent is your intelligent robot? *arXiv preprint arXiv:1712.08878*.
- Yampolskiy, R. V. (2014). The universe of minds. *arXiv preprint arXiv:1410.0369*.

A Appendix: Why is an atlas needed? Similar initiatives

While identifying the need for an atlas, we look at how it fits in cognitive science as a whole and also whether there are initiatives in other fields that could be inspirational.

Regarding cognitive science, it is true that its goal is to cover all possible cognitive systems, understand their behaviour and mechanisms, and establish meaningful comparisons. However, the field has not yet been able to portray a systematic representation covering both natural and artificial systems. But if we do not find this systematic representation in cognitive science, do we find it in related subdisciplines? The answer is that some similar initiatives in other disciplines do exist¹⁵:

- **Life forms:** Examples are Wikispecies (Leslie, 2005), the All Species Foundation (Gewin, 2002), the Catalogue of Life and the Encyclopedia of Life (Roskov et al., 2018; Hayles, 1996; Parr et al., 2014; Stuart et al., 2010).
- **Neuroscience:** the Cognitive Atlas¹⁶ and related repositories for neuroscience¹⁷ include an ontology of human cognitive functions and related tasks, and the pathologies affected. The Allen brain observatory¹⁸ (Allen Institute for Brain Science, 2016) is a more visually-oriented platform that maps perception and cognition to parts of the human brain (National Research Council, 2011).
- **Psychometrics:** There are several initiatives bringing together test batteries and repositories: the mental Measurement yearbook¹⁹, and with a more open character, the International Personality Item Pool²⁰ and the International Cognitive Ability Resource²¹.
- **Machine learning and data science research:** Kaggle²², OpenML²³ (Vanschoren et al., 2013) and many other platforms (e.g., `gitxiv.com`) provide benchmarks for ML. OpenML also includes experimental results that can be compared, aggregated and represented with powerful analytical packages.

¹⁵Some of these initiatives are in genomics and brain imaging (Midford, 2004; Boero & Bernardi, 2014).

¹⁶<http://www.cognitiveatlas.org>

¹⁷<https://poldracklab.stanford.edu/>

¹⁸<http://observatory.brain-map.org/visualcoding/>

¹⁹<http://buros.org/mental-measurements-yearbook>

²⁰<http://ipop.ori.org>

²¹<http://icar-project.com>

²²<http://www.kaggle.com>

²³<http://www.openml.org>

- **Artificial intelligence:** there are many collections of benchmarks and associated results, such as ALE²⁴, OpenAI universe/gym²⁵, Microsoft Malmo²⁶, Facebook’s CommAI-env²⁷, DeepMind Lab²⁸ (see Hernández-Orallo et al., 2017 for a summary) and meta-views, such as a recent EFF analysis²⁹ and the AI index report³⁰. This is a sign that AI is looking in this direction (Castelvecchi, 2016; Hernández-Orallo, 2017). The tasks are rarely arranged into abilities and the data usually compares specialised AI systems against average humans.

A partially overlapping initiative is the AI Roadmap Institute³¹, which encourages, compares and studies various AI and general AI roadmaps. It focuses on the future and on AI primarily, with representations that are usually flowcharts and pathway comparisons. Besides identifying where the field of AI stands as a whole, it also aims to identify dead-ends and open research problems on the path to the development of general AI systems.

The data and conceptual framing of the above projects can be used to inform an atlas of intelligence. Still, no repositories or taxonomies exist focusing mostly on behaviour, encompassing natural and artificial systems, as we are undertaking. Of course, the fact that something does not exist yet is not a sufficient reason that it should. The need for an atlas has to be supported by a series of motivations and applications, which we do in section 2.

²⁴<http://www.arcadelalearningenvironment.org/>

²⁵<https://gym.openai.com/>

²⁶<https://www.microsoft.com/en-us/research/project/project-malmo/>

²⁷<https://research.fb.com/projects/commai/>

²⁸<https://deepmind.com/blog/open-sourcing-deepmind-lab/>

²⁹<http://www.eff.org/ai/metrics>

³⁰<http://aiindex.org>

³¹<http://www.roadmapinstitute.org>