

17. Common sense skills: Artificial intelligence and the workplace

Lucy Cheke, Leverhulme Centre for the Future of Intelligence, Cambridge

Marta Halina, Leverhulme Centre for the Future of Intelligence, Cambridge

Matthew Crosby, Leverhulme Centre for the Future of Intelligence, Cambridge

This chapter looks at the basic “common sense” skills needed by artificial intelligence (AI) to perform in the workplace. It begins by focusing on the challenge for AI of navigating in complex and unpredictable environments. It explores the common sense skills underpinning these apparently sophisticated skills, including spatial memory, object representation and causal reasoning. The chapter continues with an exploration of social challenges of the workplace, including the need to predict or infer meaning from ambiguous behaviour and the “common sense” skills underpinning these. Finally, it looks at the challenges, opportunities and next steps related to testing AI with tasks from psychology.

Introduction: Workplace perspectives on artificial intelligence

Artificial intelligence (AI) systems can be trained to exceed human performance on certain tasks such as playing games, recognising images, controlling large cooling systems or painting cars. However, these systems cannot act outside of the task for which they have been trained. Moreover, they often fail under even minor deviations from their expected inputs (Shevlin et al., 2019^[1]; Shanahan et al., 2020^[2]).

From a workplace perspective, this means that AI and robotics can only be used for highly specific, well-defined tasks or where the environment can be strictly controlled. If AIs are to move beyond these narrow tasks, a number of skills are required.

In contrast, humans and other animals are generalists, able to perform a wide variety of both complex and seemingly simple tasks. Humans have a great capacity for specialisation. However, they are also highly capable of retraining and adapting to novel situations without losing previous knowledge. At the chess world championship, the grand masters are expected to be able to make a coffee, find the bathroom and operate the lights in their hotel rooms, all while remaining chess experts.

The maturation of these common sense abilities often marks the biggest developmental leaps in children: a step-change in behavioural flexibility and independence occurs, for example, at around 12 months, with the ability to search for an object (be that food, toy or parent) that has moved out of sight (Piaget, 1977^[3]).

Non-human animals are also adept at common sense reasoning (Tomasello and Call, 1997^[4]). A cat that sees a mouse enter a tunnel can flexibly deploy a range of behaviours – waiting at the tunnel mouth, running to the other end of the tunnel, digging or reaching in a paw. These are all based on the understanding that while the mouse can no longer be seen, it still exists and is potentially accessible. Both of these examples involve the cognitive capacity known as “object permanence”: an assumed skill in human adults, an area of study in child and animal psychology, and yet to be achieved in AI.

The need for common sense skills

AI and robotics can, of course, develop far beyond current capabilities without developing common sense reasoning. However, the presence or absence of common sense will fundamentally change the types of positions AI can take in the workplace of the future. Without common sense, AI will arguably remain a highly specialised tool limited in application (Shanahan et al., 2020^[2]). It will either require specifically designed workspaces, potentially costly human supervision or be bounded by inflexibility and overfitting.

Without common sense, attempts to broaden AI usage will result in applications for which it is fundamentally ill-suited, leading to potentially disastrous outcomes. In the United Kingdom, recent problems surrounding the algorithm used to assign A-level results to students illustrate the issue. Overfitting and other problems led to bias in the predicted grades assigned to students by the algorithm. Kelly (2021^[5]) calls this probably the greatest policy failure in modern times for the examination system in England.

If AI can develop some common sense skills, it might be able to expand into a new category of roles within the workplace.

In some cases, these common sense skills will be needed for an artificial agent to perform any kind of role. This is likely the case for roles that require diverse and context-sensitive interactions with the physical and social environment, such as social care. Care-taking, in the form of both child and elder care, involves physical activities (feeding and cleaning) and social ones (play and friendship). Both are highly reliant on flexibility and common sense reasoning (Lin, Abney and Jenkins, 2017^[6]).

In other cases, common sense skills would be required for the benefit of automation to outweigh the cost and effort required for supervision and restructuring the environment. Traditional autonomous-vehicle

designs, for example, have led to accidents in cases where a vehicle encounters a situation not included in its training data (Shafaei et al., 2018^[7]).

To avoid such accidents, autonomous vehicles require continuous human supervision while on the road. Companies like iSee, however, are moving away from such designs to “autonomy powered by humanistic common sense”. It aims to build cars that can flexibly respond to new situations without human supervision (www.isee.ai/about-us).

Both the knowledge and appropriate tools are currently lacking to assess the capacity of AI to take on some or all of these common sense skills. At present, tests of AI capacity are neither sufficiently cognitively defined nor sufficiently general to answer this question. The next two sections review spatial and social cognition respectively, exploring the implications of these two key categories of common sense for the future of work. The final section then proposes a way to assess AI progress towards solving these challenges.

Spatial navigation and treatment of objects

Spatial navigation and the appropriate treatment of objects is a challenge for robotics. Robots are either limited in application or rely on environments that have been specifically adapted to ensure highly consistent inputs. Any input that deviates from expectation may result in damage or disruption. Amazon, for example, designs its warehouses to allow the Kiva robots to transport entire racks on predefined paths, while picking and packaging products still requires human common sense and dexterity (Wurman, D’Andrea and Mountz, 2008^[8]). Outside of certain factories and warehouses with a large scope for investment, many workplaces cannot easily be adapted to address such strict limitations.

Issues faced by developers of robotic vacuums and lawn mowers exemplify the challenges faced by bringing robotic assistance into a wider range of workplaces. These robots must navigate a messy real-world environment with unpredictable inputs and obstacles: is that grass, the edge of a flowerbed or an object on the lawn?

To some extent, the environment can be controlled and accounted for. The edge of a lawn, for example, could be delineated with a guide wire. Similarly, the ceiling can be scanned for room size or the robot can be re-routed after a detected collision. However, in practice, many users spend more time managing the robot than they save by owning it.

A taxonomy of common sense space and object skills

Non-human animals can navigate complex environments with varying levels of sophistication, with even the simplest creatures showing some capacity. Animals use spatial skills but also memory to deal with obstacles, objects and affordances. The following taxonomy is a subset of skills identified by psychologists as key features of physical cognition in human and non-human animals. It focuses on skills relevant for basic physical tasks common in human workspaces, such as navigating a space and interacting with objects.

1. Spatial memory and navigation:
 - **Path integration:** using self-motion cues (like limb velocity and movement) to navigate. This ability allows navigation without relying on external landmarks (Etienne and Jeffery, 2004^[9]).
 - **Cognitive maps:** building an internal representation of external space, routes and landmark arrangements (Kitchin, 1994^[10]). Cognitive maps can rely on path integration, as well as other sources of information (like external landmarks).
 - **What-where-when episodic-like memory:** remembering the location of a specific object at a specific time (Clayton et al., 2001^[11]).

- **Episodic memory:** remembering entire events (including spatial context, perceptual information and internal processes) (Tulving, 1983^[12]; Cheke and Clayton, 2015^[13]). Episodic memory includes information about the “what”, “where” and “when” of an event (as in episodic-like memory). However, it also involves mentally reconstructing the experience of an event and other features.
2. Object representations:
- **Object-level representations:** Representing visual input as objects (that may be movable, able to block other objects, act as containers etc.), rather than patterns of light (Gregory, 1997^[14]).
 - **Object permanence:** knowing something is still there, even if you can no longer detect it (Baillargeon and DeVos, 1991^[15]). Object permanence requires representing visual input as objects.
 - **Affordance-level representations:** relating perceptual information into affordance information to, for example, predict whether a large object will fit through a narrow aperture (Scarantino, 2003^[16]). An “affordance” is a property of an object that makes clear how it can be used. Affordance-level representations often depend on representing visual input as objects (e.g. representing something as an “apple” may evoke the affordance “eat-able”).
3. Causal reasoning
- **Spatial and object inference:** the ability to infer the location or properties of an object through inference, based on prior knowledge and context. For example, eliminative reasoning can conclude that out of a large and small container, only the former can contain a large object (Shanahan et al., 2020^[2]).
 - **Folk physics:** the capacity to predict outcomes based on some understanding of the physical mechanisms involved (e.g. predicting a spherical object on a slope will roll downwards) (Povinelli, 2003^[17]). The capacity for folk physics will typically rely on a capacity for spatial and object inference.

AI’s space and object skills in the workplace

Even the simplest tasks may require these skills if the inputs themselves cannot be strictly controlled and defined. Various standard software programs involve “objects” (e.g. text boxes or pictures). These objects often vary in whether they can be manipulated, and whether they can obstruct or occlude other objects, etc.

An AI without common sense space and object skills may be useful in processing documents but only if these take predictable forms and are accompanied by sufficient metadata. Such an AI would be unable to cope with variations in input type (such as new software). Even in perfect conditions, it would make regular “obvious” errors that would necessitate extensive oversight.

In contrast, an AI with an ability to represent objects and their affordances – i.e. able to continue to reason about them when they are no longer perceptible (object permanence) and infer new information based on this reasoning – would be able to process a large range of files from even unfamiliar sources. Furthermore, if an agent is also able to develop a folk physics of the software environment (e.g. image software that would contain multiple “layers” with “transparency”), when encountering an entirely unfamiliar object or software, it could make inferences based on its knowledge of similar objects/software. While this may sound sophisticated, it is the minimum required for an agent skilled in PowerPoint, for example, to learn to use Google Slides without explicit retraining.

Robots' space and object skills in the workplace

Robots in controlled environments

For robots in physical workplaces, common sense space and object skills are even more crucial. Robotics is currently limited to strictly controlled environments. Even within workplaces (such as factories and warehouses), that can provide a relatively high level of consistency, space and object common sense would vastly expand the types and range of tasks available to robotics.

Amazon illustrates the limitations of robots in the workplace. Even if all packages were labelled with (say) barcodes representing their content, a robot would still need several skills to locate a given package reliably. First, it would need to represent objects (“this visual input is a *package*”). Second, it would need to represent the affordances of the object (“the package is *x shape and x weight* and this translates into manipulating it in *x manner*”). Third, it would often require object permanence and inference (“the barcode may be behind or under the visible surface of the package, which may mean rotating the package to find it”). For basic navigation of the warehouse without predefined routes, the robot would need to use path integration, as well as the capacity to deal with obstacles.

Robots in uncontrolled environments

In many cases, where robot assistance would be most useful – such as medical and elderly care – are within uncontrolled environments (Lin, Abney and Jenkins, 2017^[6]). However, these environments, by definition, create challenges. At its most fundamental level, assistance in the home will require the capacity to navigate a home environment, and to share it with someone who may behave in unpredictable ways.

A robot carer without space and object common sense would function only within a consistently clear-floored home. It could not interact with objects that might take unpredictable forms or be in unpredictable locations. It could also not perform any bodily care (assisting with dressing, washing or toileting) without risk of causing pain and injury. Without at least episodic-like memory, it would be unable to locate items that can vary in their location *even if it has seen where they were last placed*.

Conversely, a robot carer in possession of space and object common sense has multiple tools at its disposal. When confronted with an obstacle, a robot carer could use a cognitive map to choose an alternative route. It could also use affordance-level representations to relocate the obstacle and clear its own path. If an important item (e.g. a hearing aid or pacifier) falls out of sight, it could use object permanence to retrieve it. If delivering a drink, it could use folk physics regarding support and flatness to identify a surface that can safely hold the receptacle. Episodic memory would be crucial for a robot carer to monitor cognitive health of their charge. Speech repetition, for example, could be a warning sign for memory decline.

In a wide range of workplace contexts, space and object common sense skills make a fundamental difference to the types of roles an agent can play. Some of these skills (e.g. path integration) are already common features in AI. Others (e.g. folk physics) are clearly some way off. For most skills in between, tools are not yet available to assess them in AI, and thus it is not clear how close they are to development (Crosby, 2020^[18]; Shanahan et al., 2020^[2]).

Pragmatics and social interaction

Recent advances in linguistic AI have led to an explosion of interactive applications. Chatbots, for example, are now routinely used as a first port of call for online customer enquiries and tech support. GTP-3 is a language model with 175 billion parameters that can perform at close to human level in many few-shot learning tasks.

However, GPT-3 performs relatively poorly on common sense reasoning tests that involve inferring meaning from non-explicit reference or background knowledge (Floridi and Chiriatti, 2020^[19]). Within linguistics, this kind of reasoning using language is known as pragmatics. However, the skill of predicting or inferring meaning from ambiguous behaviour is an issue beyond linguistics and common to all social interaction.

In the physical realm, a considerable challenge for AI has been to detect, identify and interpret inputs associated with objects and their affordances. Much the same challenge is magnified when it comes to human behavioural cues. For example, gaze direction indicates the subject of an individual's attention (and therefore their behaviour and verbal reference). It provides the means for responding appropriately to a host of otherwise ambiguous behaviours. However, gaze indicates a direction but not a specific target. The target must be inferred from the context and behaviour of the individual, in combination with prior knowledge. These forms of common sense inferences are a particular challenge for current AI.

A taxonomy of common sense social and communicative skills

Animal and developmental cognition research distinguishes between two categories of social skills. Social learning and communication is the exchange of information with others. Meanwhile, social cognition is understanding and predicting the behaviour and mental states of others. The following taxonomy represents capacities widely accepted by researchers in this area (Hoppitt and Laland, 2013^[20]; Shettleworth, 2013^[21]).

1. Social learning and communication

- **Local or stimulus enhancement:** increased attention to areas or objects manipulated by others (Hoppitt and Laland, 2013^[20]). Stimulus enhancement may be involved in other forms of learning (e.g. observational conditioning). For example, an observer may respond more to a stimulus as a result of witnessing another agent interact with that stimulus.
- **Observational conditioning:** learning action-outcome relationships through watching others. For example, a rhesus monkey with no prior exposure to snakes will not behave fearfully in response to a snake. However, if the monkey observes another individual responding fearfully to snakes, it will begin responding this way as well (Cook et al., 1985^[22]).
- **Emulation:** targeting a goal or outcome after observation without imitating exact behaviour (Tomasello, 1990^[23]). An adult carrying a stack of books might use her elbow to switch on a light; a child emulating this adult would know they could use alternative means (e.g. their hand) to achieve this same goal of switching on the light.
- **Imitation:** reproduction of observed behaviour. In the previous example, a child learning through imitation would use their elbow to switch on the light if they had observed the adult doing so (Hoppitt and Laland, 2013^[20]). Whether it is more appropriate to emulate or imitate will depend on the situation. In this example, the exact behaviour was the product of a limitation on the adult (the stack of books) that the child did not suffer. It was therefore more efficient to emulate the target of the action than to imitate the exact action itself.
- **Non-verbal communication:** production and comprehension of non-linguistic communicative signals, such as gestures and facial expressions (Kendon, 2004^[24]). Some argue that non-verbal communication paves the way for language development in humans (Iverson and Goldin-Meadow, 2005^[25]).
- **Linguistic communication:** production and comprehension of language for communicative purposes (Tomasello, 2009^[26]).

2. Social cognition

- **Co-operation:** co-ordinating behaviour with another individual to achieve a shared goal (Henrich and Muthukrishna, 2021^[27]). Some researchers suggest that co-operative activities,

like helping and sharing, depend on understanding the goals of others or a minimal theory of mind (Tomasello and Vaish, 2013^[28]).

- **Behaviour reading:** inferring likely behaviour from context and behaviour or facial expressions (Perner and Ruffman, 2005^[29]).
- **Minimal theory of mind:** inferring the goal (or outcome to which a behaviour is directed) of another agent from context and behaviour or facial expressions, without reference to mental states (Butterfill and Apperly, 2013^[30]). This ability goes beyond behaviour reading (in involving goal attribution) but does not require a full-blown theory of mind.
- **Theory of mind:** inferring and reasoning about internal epistemic and motivational states, such as intentions, desires and knowledge (Wellman, 1992^[31]). Agents with this capacity can predict and explain behaviour by attributing a range of mental states, rather than just reasoning about behaviour or goal-directed behaviour.
- **Empathy:** inferring and reasoning about the emotional states of others (Decety and Lamm, 2006^[32]).

Social skills in the workplace

All AIs interface, either directly or indirectly, with people. The form of this interaction is dictated by the capacity of the AI to deal with variability and unpredictability in its social input. This, in turn, dictates the roles available to AI in the workplace. This social input may take multiple forms, most commonly: information communicated via language, behaviour and social context.

Predicting human behaviour is a major challenge for AI in multiple contexts. Yet this will be central to enable AIs to expand into new roles in the workforce, even if those roles do not involve an explicitly social component. For robots, human workers represent a unique form of (highly fragile) obstacle with complex and unpredictable trajectories.

Predicting trajectories is particularly challenging for autonomous vehicles. Unless AI controls all road users, autonomous vehicles must predict and account for human behaviour to avoid collision. An AI without social common sense may learn a range of predictors of collision-risk. For example, it can learn that a child on the pavement may lead to a child on the road. However, it will struggle with anything that falls outside of its specific training or requires inference. For example, it may not understand that a ball on the road may lead to a child on the road, with potentially fatal consequences.

Much more explicitly social roles for AI and robots may be widely useful (as can be seen with chatbots). However, these roles would require more refined social skills beyond predicting simple behaviours (such as trajectory). A robot carer, for example, may not require full-blown empathy. However, it must identify needs that may be communicated both verbally and nonverbally. It must also deal with issues as they arise such as distress, confusion and injury.

A robot carer would also need to be able to monitor health and cognitive function. This would require learning enough information about a specific individual to detect changes in personality or behaviour (such as aggression, delirium or unresponsiveness).

A robot carer would also benefit from being able to distinguish between intentional (lying down) and unintentional (falling down) behaviour. This fundamental element of social common sense in practice requires elements of theory of mind, making it one of the more complex social cognition skills. However, it is required ubiquitously in the workplace: papers placed on the desk should be processed, while papers placed on a shelf that *fall* onto the desk should not.

Requirements for training and supervision will significantly influence the potential for proliferation of AI within the workforce. If expert oversight is required for every new task, then such agents may be useful if produced and programmed or trained at scale. However, they will not be appropriate for niche or specified

applications. The ability of agents to “learn on the job” – to learn by watching and perhaps practising with oversight – vastly expands potential applications.

Als are currently adept at some forms of social learning such as observational conditioning. They can be trained on predictive relationships, for example, by watching YouTube videos. However, other forms are much more challenging. Emulation – arguably the most effective form of task learning – requires the goal or affordances of a situation to be extracted from observation.

Testing artificial intelligence with tasks from psychology: Challenges, opportunities and next steps

The above review of a taxonomy of common sense skills draws on research in developmental and comparative psychology. It focused on those capabilities that represent the breadth of current work in human and non-human animal research (Tomasello and Call, 1997^[4]; Shettleworth, 1998^[33]; Hoppitt and Laland, 2013^[20]) and which are also immediately relevant to the development of AI in the workplace (Shanahan et al., 2020^[2]). These capabilities are ubiquitous in humans and often common throughout the animal kingdom, but currently pose a challenge for AI.

A research agenda for artificial intelligence in the workplace

Lack of these common sense skills in AI limits the tasks that can be performed. It means many types of jobs will be safe from full automation until significant progress is made. To understand what jobs these are, and whether that progress is likely, requires basic research in two areas. First, researchers need to identify which common sense skills are necessary for specific roles within the workplace, and to what extent. When most humans have a capability, it is an assumed skill within an employee, but this would not be the case with an AI. Second, they need to understand current AI capacity, and measure progress, in common sense skills.

This chapter has taken a first step towards addressing the first area of research. Concerning the second area, a number of well-established tests within animal and developmental cognition research assess these skills. However, there are no “off-the-shelf” cognitive tests appropriate for testing AI. Many tests designed for children are diagnostic, rather than evaluative, and others rely heavily on species-specific tendencies or biases.

Conversely, a number of benchmarking assessments for AI have been the basis for cognitive claims. However, these assessments are rarely cognitively defined, and suffer from a construct validity problem. In particular, it is easy for a programmer (either consciously or subconsciously) to design and train an AI to pass a particular *test*, rather than to possess a particular *skill* (Hernández-Orallo, 2017^[34]; Shevlin and Halina, 2019^[35]).

Animal-AI Testbed

One solution to the validity problem is to assess AI in the same way that psychologists assess animals and children, such as through the Animal-AI Testbed (Crosby, Beyret and Halina, 2019^[36]; Crosby, 2020^[18]; Crosby et al., 2020^[37]). Here, agents can be trained in a 3D environment and learn about the spaces, objects, agents and rules it contains.

During training, agents can access a simulated arena with simple objects of various sizes, shapes and textures. The objects behave as they would in the real world (subject to forces of gravity, friction, etc.) due to simulated physics. The arena also contains positive and negative rewards, which an agent can attempt to acquire or avoid, respectively.

This training arena is designed to reproduce the type of environment an animal might encounter before they are tested in a typical animal-cognition task. Crucially, in both cases, an agent is not trained on the test itself. It thus cannot learn to simply pass this test through hundreds or thousands of trials. Instead, it must rely on general common sense skills like those reviewed above.

The most recent version of the Animal-AI Testbed consists of 300 tasks. Each task includes a set of objects, including a positive reward and a time limit. An agent succeeds in the task if it retrieves sufficient reward (i.e. reward above some threshold) within the time limit. The tasks are either drawn directly from or inspired by tests used in developmental and comparative psychology.

Each task aims to probe an agent's ability to engage in common sense reasoning. For example, several tasks test an agent's spatial memory and navigation abilities: these include mazes of various difficulties, from simple mazes in the shape of the letter "Y" to more difficult radial mazes with many arms. Another category of task tests an agent's capacity for object permanence by moving rewards out of sight.

Finally, other tasks probe an agent's capacity for folk physics by presenting agents with problems that require a tool. To solve these problems, the agent must have some understanding of the physical mechanisms involved and how the tool can produce a causal effect that will lead to a reward.

Using a wide range of cognitively defined tests, it is therefore possible to provide a robust and valid assessment of AI common sense capacities (Crosby et al., 2020_[37]).

Emerging tests for social and communicative skills

Using the Animal-AI Testbed, it has been demonstrated that state-of-the-art artificial agents display path integration and the beginnings of cognitive maps and spatial and object inference (Crosby et al., 2020_[38]; Voudouris et al., 2021_[39]). However, episodic memory, object-level representations, object permanence, affordance-level representations and folk physics require further research. Solving these tasks in Animal-AI will be the first step towards solving them in real-world settings and will be a marker for the possible landscape of AI integration into the future workplace.

The Animal-AI Testbed tests for common sense space and object skills rather than social and communicative skills. However, the latter is a promising area for future research. In psychology, there is a rich literature of tasks for testing social and communicative capacities in human and non-human animals (see references in taxonomy above).

Researchers have begun to apply these tasks to AI systems. For example, AI and robotics researchers have tested artificial systems on theory of mind tasks (Winfield, 2018_[40]; Rabinowitz et al., 2018_[41]). Further research in this area would benefit from distinguishing abilities like behaviour reading and theory of mind, given these capacities may have differential effects on performance (Shevlin and Halina, 2019_[35]).

Enhancing human performance with artificial intelligence

Moving beyond autonomous artificial agents, another use of AI in the workplace involves enhancing or augmenting human performance. Such enhancement technologies include tools like speech assistants, navigation systems and other decision-support systems (Hernández-Orallo and Vold, 2019_[42]; Shevlin et al., 2019_[1]; Sutton et al., 2020_[43]).

These tools are likely to present unique challenges to implementing AI in the workplace. Human-AI collaboration would also benefit from AI with common sense skills. For example, common sense space and object skills would facilitate guidance and navigational systems that can respond to the challenges and affordances of the environment rather than relying on pre-programmed routes and rules.

Common sense social abilities may be even more important in the context of AI-human collaboration: an agent that can detect and respond to needs, interpret verbal and non-verbal communication, and engage

in joint problem solving will be a far more effective assistant in whatever capacity it is employed. Thus, common sense skills are important for not only enabling AI to operate autonomously or semi-autonomously in the workplace but also for enhancing human performance.

Recommendations

- **Use animal and child development tests to test common sense skills**

While AIs excel at highly specified skills, they struggle with basic skills humans take for granted. Without such fundamental skills, they cannot navigate the world in a flexible and adaptive manner. It is challenging to test these skills in AI precisely because they are not generally assessed in human adults. Instead, tests developed for animals and young children provide the most appropriate means to assess AI systems in these areas.

- **Adapt these tests for AI**

However, such tests cannot be simply used, off the shelf, to assess AI. They must be adapted into contexts appropriate for AI and administered to ensure assessment of the skill and not expertise at the test. The Animal-AI Testbed provides an example of how this can be done.

- **Extend tests to social and communicative common sense**

The Animal-AI Testbed should be extended to tests for social and communicative common sense. These will be important for the integration of AI into workplaces that involve human interaction.


References

- Baillargeon, R. and J. DeVos (1991), "Object permanence in young Infants: Further evidence", *Child Development*, Vol. 62, pp. 1227-1246, <http://dx.doi.org/10.1111/j.1467-8624.1991.tb01602.x>. [15]
- Butterfill, S. and I. Apperly (2013), "How to construct a minimal theory of mind", *Mind & Language*, Vol. 28, pp. 606-637, <http://dx.doi.org/10.1111/mila.12036>. [30]
- Cheke, L. and N. Clayton (2015), "The six blind men and the elephant: Are episodic memory tasks tests of different things or different tests of the same thing?", *Journal of Experimental Child Psychology*, Vol. 137, pp. 164-171, <http://dx.doi.org/10.1016/j.jecp.2015.03.006>. [13]
- Clayton, N. et al. (2001), "Elements of episodic-like memory in animals", *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, Vol. 356, pp. 1483-1491, <http://dx.doi.org/10.1098/rstb.2001.0947>. [11]
- Cook, M. et al. (1985), "Observational conditioning of snake fear in unrelated rhesus monkeys", *Journal of Abnormal Psychology*, Vol. 94/4, pp. 591-610, <http://dx.doi.org/10.1037//0021-843x.94.4.591>. [22]
- Crosby, M. (2020), "Building thinking machines by solving animal cognition tasks", *Minds and Machines*, Vol. 30, pp. 589-615, <http://dx.doi.org/10.1007/s11023-020-09535-6>. [18]
- Crosby, M., B. Beyret and M. Halina (2019), "The Animal-AI Olympics", *Nature Machine Intelligence*, Vol. 1, pp. 257-257, <http://dx.doi.org/10.1038/s42256-019-0050-3>. [36]

- Crosby, M. et al. (2020), *The Animal-AI Testbed and Competition*, [38]
<https://proceedings.mlr.press/v123/crosby20a.html>.
- Decety, J. and C. Lamm (2006), "Human empathy through the lens of social neuroscience", *The Scientific World Journal*, Vol. 6, pp. 1146-1163, <http://dx.doi.org/10.1100/tsw.2006.221>. [32]
- Escalante, H. and R. Hadsell (eds.) (2020), "The animal-AI testbed and competition", *Proceedings of the NeurIPS 2019 Competition and Demonstration Track*, No. 123, PMLR, <http://proceedings.mlr.press/v123/crosby20a.html>. [37]
- Etienne, A. and K. Jeffery (2004), "Path integration in mammals", *Hippocampus*, Vol. 14, pp. 180-192, <http://dx.doi.org/10.1002/hipo.10173>. [9]
- Floridi, L. and M. Chiriatti (2020), "GPT-3: Its nature, scope, limits, and consequences", *Minds and Machines*, Vol. 30, pp. 681-694, <http://dx.doi.org/10.1007/s11023-020-09548-1>. [19]
- Gregory, R. (1997), *Eye and Brain: The Psychology of Seeing – Fifth Edition*, Princeton University Press, <http://dx.doi.org/10.2307/j.ctvc77h66>. [14]
- Henrich, J. and M. Muthukrishna (2021), "The origins and psychology of human cooperation", *Annual Review of Psychology*, PMID: 33006924, pp. 207-240, <http://dx.doi.org/10.1146/annurev-psych-081920-042106>. [27]
- Hernández-Orallo, J. (2017), "Evaluation in artificial intelligence: From task-oriented to ability-oriented measurement", *Artificial Intelligence Review*, Vol. 48, pp. 397-447, <http://dx.doi.org/10.1007/s10462-016-9505-7>. [34]
- Hernández-Orallo, J. and K. Vold (2019), "AI extenders: The ethical and societal implications of humans cognitively extended by AI", *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, Association for Computing Machinery, <http://dx.doi.org/10.1145/3306618.3314238>. [42]
- Hoppitt, W. and K. Laland (2013), *Social Learning: An Introduction to Mechanisms, Methods, and Models*, Princeton University Press, <http://dx.doi.org/10.1515/9781400846504>. [20]
- Iverson, J. and S. Goldin-Meadow (2005), "Gesture paves the way for language development", *Psychological Science*, PMID: 15869695, pp. 367-371, <http://dx.doi.org/10.1111/j.0956-7976.2005.01542.x>. [25]
- Kelly, A. (2021), "A tale of two algorithms: The appeal and repeal of calculated grades systems in England and Ireland in 2020", *British Educational Research Journal*, Vol. n/a, <http://dx.doi.org/10.1002/berj.3705>. [5]
- Kendon, A. (2004), *Gesture: Visible Action as Utterance*, Cambridge University Press, <http://dx.doi.org/10.1017/CBO9780511807572>. [24]
- Kitchin, R. (1994), "Cognitive maps: What are they and why study them?", *Journal of Environmental Psychology*, Vol. 14, pp. 1-19, [http://dx.doi.org/10.1016/S0272-4944\(05\)80194-X](http://dx.doi.org/10.1016/S0272-4944(05)80194-X). [10]
- Lin, P., K. Abney and R. Jenkins (eds.) (2017), *Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence*, Oxford University Press, <http://dx.doi.org/10.1093/oso/9780190652951.001.0001>. [6]

- Perner, J. and T. Ruffman (2005), "Infants' insight into the mind: How deep?", *Science*, Vol. 308/5719, pp. 214-216, <http://dx.doi.org/10.1126/science.1111656>. [29]
- Piaget, J. (1977), *The Development of Thought: Equilibration of Cognitive Structures (Trans A. Rosin)*, Viking. [3]
- Povinelli, D. (2003), *Folk Physics for Apes: The Chimpanzee's Theory of How the World Works*, Oxford University Press, <http://dx.doi.org/10.1093/acprof:oso/9780198572190.001.0001>. [17]
- Rabinowitz, N. et al. (2018), "Machine theory of mind", *Proceedings of the 35th International Conference on Machine Learning*, No. 80, Dy, J. and A. Krause (eds.), PMLR, <http://proceedings.mlr.press/v80/rabinowitz18a.html>. [41]
- Scarantino, A. (2003), "Affordances explained", *Philosophy of Science*, Vol. 70, pp. 949-961, <http://dx.doi.org/10.1086/377380>. [16]
- Shafaei, S. et al. (2018), "Uncertainty in machine learning: A safety perspective on autonomous driving", Conference: First International Workshop on Artificial Intelligence Safety Engineering, Västerås, Sweden. [7]
- Shanahan, M. et al. (2020), "Artificial intelligence and the common sense of animals", *Trends in Cognitive Sciences*, Vol. 24, pp. 862-872, <http://dx.doi.org/10.1016/j.tics.2020.09.002>. [2]
- Shettleworth, S. (2013), *Fundamentals of Comparative Cognition*, Oxford University Press, <https://global.oup.com/academic/product/fundamentals-of-comparative-cognition-9780195343106>. [21]
- Shettleworth, S. (1998), *Cognition, Evolution, and Behavior*, Oxford University Press, <https://global.oup.com/academic/product/cognition-evolution-and-behavior-9780195319842>. [33]
- Shevlin, H. and M. Halina (2019), "Apply rich psychological terms in AI with care", *Nature Machine Intelligence*, Vol. 1, pp. 165-167, <http://dx.doi.org/10.1038/s42256-019-0039-y>. [35]
- Shevlin, H. et al. (2019), "The limits of machine intelligence", *EMBO Reports*, Vol. 20/10, p. e49177, <http://dx.doi.org/10.15252/embr.201949177>. [1]
- Sutton, R. et al. (2020), "An overview of clinical decision support systems: Benefits, risks, and strategies for success", *npj Digital Medicine*, Vol. 3, p. 17, <http://dx.doi.org/10.1038/s41746-020-0221-y>. [43]
- Tomasello, M. (2009), *Constructing a Language*, Harvard University Press, <https://www.hup.harvard.edu/catalog.php?isbn=9780674017641>. [26]
- Tomasello, M. (1990), "'Language' and intelligence in monkeys and apes", in *Cultural Transmission in the Tool Use and Communicatory Signaling of Chimpanzees*, Parker, S. and K Gibson (eds.), Cambridge University Press, <http://dx.doi.org/10.1017/cbo9780511665486.012>. [23]
- Tomasello, M. and J. Call (1997), *Primate Cognition*, Oxford University Press, <https://global.oup.com/academic/product/primate-cognition-9780195106244>. [4]
- Tomasello, M. and A. Vaish (2013), "Origins of human cooperation and morality", *Annual Review of Psychology*, PMID: 22804772, pp. 231-255, <http://dx.doi.org/10.1146/annurev-psych-113011-143812>. [28]

- Tulving, E. (1983), *Elements of Episodic Memory*, Oxford University Press. [12]
- Voudouris, K. et al. (2021), “Direct Human-AI Comparison in the Animal-AI Environment”, [39]
<https://doi.org/10.31234/osf.io/me3xy>.
- Wellman, H. (1992), *The Child’s Theory of Mind*, MIT Press, [31]
<https://mitpress.mit.edu/books/childs-theory-mind>.
- Winfield, A. (2018), “Experiments in artificial theory of mind: From safety to story-telling”, [40]
Frontiers in Robotics and AI, Vol. 5, p. 75, <http://dx.doi.org/10.3389/frobt.2018.00075>.
- Wurman, P., R. D’Andrea and M. Mountz (2008), “Coordinating hundreds of cooperative, [8]
autonomous vehicles in warehouses”, *AI Magazine*, Vol. 29, pp. 9–19.



From:
AI and the Future of Skills, Volume 1
Capabilities and Assessments

Access the complete publication at:
<https://doi.org/10.1787/5ee71f34-en>

Please cite this chapter as:

Cheke, Lucy, Marta Halina and Matthew Crosby (2021), “Common sense skills: Artificial intelligence and the workplace”, in OECD, *AI and the Future of Skills, Volume 1: Capabilities and Assessments*, OECD Publishing, Paris.

DOI: <https://doi.org/10.1787/5fe70a0c-en>

This work is published under the responsibility of the Secretary-General of the OECD. The opinions expressed and arguments employed herein do not necessarily reflect the official views of OECD member countries.

This document, as well as any data and map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area. Extracts from publications may be subject to additional disclaimers, which are set out in the complete version of the publication, available at the link provided.

The use of this work, whether digital or print, is governed by the Terms and Conditions to be found at <http://www.oecd.org/termsandconditions>.