

**DIRECTORATE FOR EDUCATION AND SKILLS
EDUCATION POLICY COMMITTEE****Future of Education and Skills 2030: Curriculum Analysis****Preparing humanity for change and artificial intelligence: Learning to learn as a safeguard against volatility, uncertainty, complexity and ambiguity****8th Informal Working Group (IWG) Meeting
29-31 October 2018
OECD Conference Centre, Paris, France**

This draft paper was written by Ruben LAUKKONEN, Hannah BIDDEL and Regan GALLAGHER from the University of Queensland, Australia. This paper describes the challenges posed by the VUCA world, and how humans and artificial intelligence respond to it. For ACTION participants are invited to SEND their comments before 5 November 2018.

Miho TAGUMA, Senior Analyst, Miho.Taguma@oecd.org
Florence GABRIEL, Analyst, Florence.Gabriel@oecd.org

JT03437073

Table of contents

Abstract	3
1. Introduction	4
1. PART 1: VUCA	6
1.1. The Predictive Brain	6
1.2. Humans	7
1.3. Machines	9
1.4. Summary of Part 1	12
2. PART 2: Agency	13
2.1. Humans	13
2.2. Machines	15
2.3. Summary of Part 2	16
3. Discussion	17
4. Concluding Remarks	19
References	20

Figures

Figure 1. The exponential progress of computing power	4
Figure 2. Perceptual illusions: the power of the predictive brain	6
Figure 3. Decision making by human and machine intelligence.....	10
Figure 4. Reinforcement learning.....	15

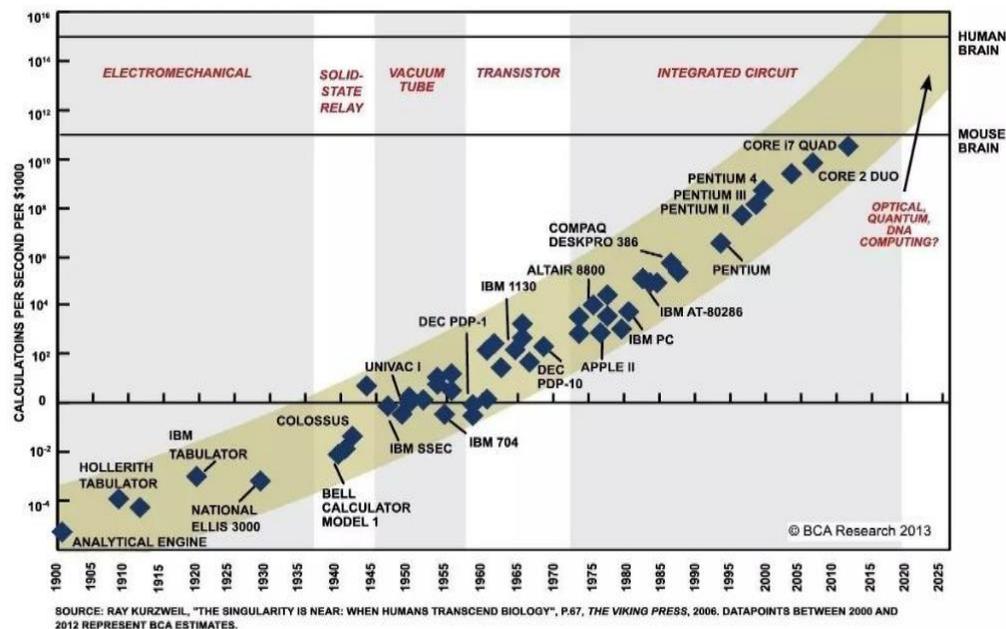
Abstract

In this paper, we describe the challenges posed by VUCA: Volatility, Uncertainty, Complexity, and Ambiguity. We begin by providing a brief review of human behaviour and decision making in VUCA situations. In doing so, we hope to illustrate how human minds and brains typically respond in unpredictable situations, and then based on this review, provide some recommendations about how to prepare for an increasingly dynamic world. We then compare human and machine abilities. The comparison to “intelligent” machines is valuable because it highlights the specific human capacities that are not easily replicated in computers. The challenge of building artificial intelligence (AI) also reveals important insights into the human skills and competencies that are likely to be important in the future. The distinguished fellow at the Institute for the Future, and best-selling author, Bob Johansen (2013^[1]), said that, “The ultimate dilemma is to take the VUCA world and change it from a threatening thing, which it certainly is, into a world that is not only threatening but also laden with opportunity.” To this end, we provide actionable methods for individuals and society to become robust and adaptable to many future outcomes and possible opportunities.

1. Introduction

Technology is progressing exponentially fast. Gordon Moore (1975^[2]) observed that computing power (defined by the number of transistors per integrated circuit) was doubling approximately every two years, a pattern that remains consistent to this day (known as “Moore’s Law”). If the trend continues, the processing capacity of machines *will* surpass the capacity of the human brain, and continue well beyond; each year making more progress than it did the year before (this intimidating process is illustrated in Figure 1). Exponential progress is notoriously difficult for humans—who are adapted to thinking linearly—to appreciate (Wagenaar and Sagaria, 1975^[3]). Technology and society are heavily interconnected, so faster changes in technology often translate to faster changes in society. Moreover, change, more often than not, translates to greater unpredictability in the form of VUCA. To the human mind that is seeking to make sense of the world, this is a potentially troubling state of affairs (as we will soon see). With growing automation, machines also have the potential to take away the pursuits that bring meaning and purpose to the lives of millions of people—their careers, their passions, and their way of life. While it is clear that there are potential costs to technological progress, there are also many benefits. Technology can transform society for the better by improving decision-making, healthcare, transport, security, and education. Artificial intelligence can liberate humans from many tasks that would ordinarily be demanded of them, potentially returning us to a society that is “time rich” instead of “time poor”. The key then, is in making the most of the benefits while being well prepared for the hazards.

Figure 1. The exponential progress of computing power



This figure shows the exponential progress of computing power from 1900 to 2013, with projections into 2025, where artificial intelligence is projected to parallel the processing capacity of the human brain

Source: Figure adapted from (Kurzweil, 2005^[4])

The acronym VUCA has been used often in the context of leadership theories (Pasmore and O'Shea, 2010^[5]), and the US Army originally used the term to refer to the increasingly VUCA-esque world following the Cold War (Bennis and Nanus, 1985^[6]). What are the different aspects of VUCA? Volatility is the liability of something to change rapidly and unpredictably. Stock markets, for example, are considered volatile because of how quickly they change and therefore how notoriously challenging they are to predict. Uncertainty relates to the quality of information one has—or the degree to which the outcome of an event is knowable in advance. Complexity increases when there is a greater number of relevant variables or interrelationships; the more variables, the more complex the situation. For instance, managing a classroom with five children is much less complicated than managing 100, and doing so is simpler if each student speaks the same language. Ambiguity occurs when an event, situation, or context is unclear, either because information is missing, inconsistent, contradictory, or obscured in some way. For humans—and indeed machines—each of these components of VUCA ultimately make for a less predictable world.

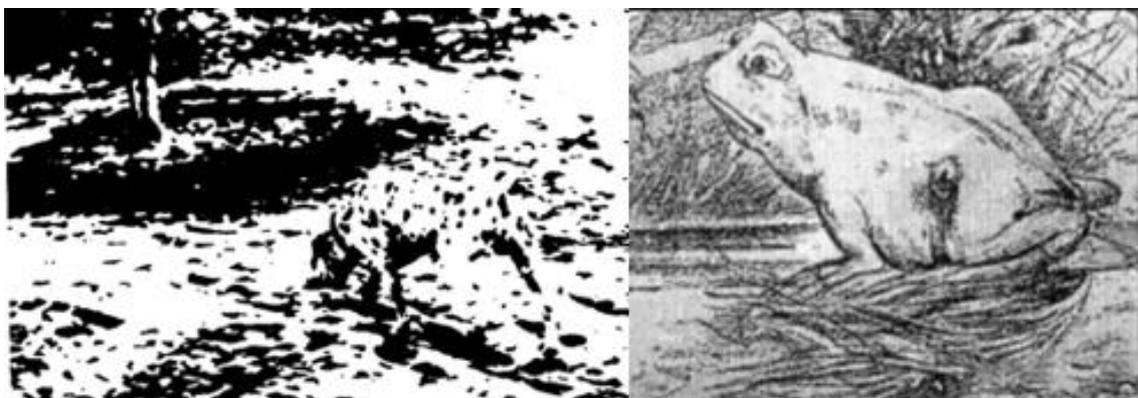
The following paper is divided into two parts (1: VUCA and 2: Agency), followed by a general discussion. Part 1 begins with a succinct introduction to a theory of the brain known as predictive processing. The aim here is to provide scaffolding for understanding how humans respond in VUCA situations. We then briefly review the psychology, neurochemistry and neuro-physiology literature to evaluate the negative and positive consequences of VUCA for human well-being. We then consider how people can learn to respond to VUCA effectively. In the final section of Part 1, we provide a broad overview regarding the current state of artificial intelligence (AI)—including specific and general AI—and then discuss how autonomous machines perform in VUCA situations. In Part 2, we turn our attention to Agency, and consider the possibility that a sense of agency might serve a valuable function for humans in uncertain situations. We also consider whether the processes that underlie autonomy in machines resemble primitive forms of agency, and whether AI is prepared for morally ambiguous situations. In the general discussion, we synthesise the review and propose that a malleable attribute of humans—their ability to learn new skills and apply meta-learning principles—will be particularly valuable for ensuring that humans can adapt to change in a VUCA world. We suggest that one of the key goals of education ought to be to prepare students to be adaptable learners equipped with meta-learning skills (Maudsley, 1980^[7]), so that when inevitable changes occur, people and communities are robust enough to adopt new skills and practices, and effectively transfer learning across situations. Meta-learning capacities may also help provide a sense of agency and self-efficacy in an otherwise chaotic world—it may not be possible to know *what* to expect, but it is perhaps possible to learn *how* to respond when the unexpected occurs.

1. PART 1: VUCA

1.1. The Predictive Brain

To understand how brains and minds respond to VUCA situations, it is useful to have a basic idea of how the brain works. A sense for human psychology is also valuable for later sections discussing AI, because the mind sciences and machine intelligence research are highly complementary (insights in one field often lead to insights in the other). Although there exists much debate on specific questions about how the mind and brain achieve certain feats, there are basic processes that are reasonably well agreed upon among scientists, and this short section focuses on these latter cases. There is no doubt, for example, that the way we perceive and experience the world is different from the way it actually is. That is to say, we do not see the world objectively; we see it based on a particular kind of interpretation that is constrained by our biology, developmental history, culture, and evolution. As a biological constraint, we have certain kinds of eyes that can only process certain kinds of information. Relative to the mantis shrimp for example—which has eyes that can move independently and can perceive ultraviolet light (Marshall and Oberwinkler, 1999^[8])—human vision is highly constrained (and our other senses are similarly limited). Our minds further constrain and influence our perceptions, as evidenced by visual illusions:

Figure 2. Perceptual illusions: the power of the predictive brain



To most people, the figure on the left initially looks like a series of dots lacking any particular meaning, and the figure on the right looks like a frog.

Notice any first impressions of the images, and then how they can change with a simple update to expectations (in cognitive science these are called *predictions*). The figure on the left can suddenly appear as a dog sniffing the ground, if one is looking for it. And the apparent frog in the other figure, if you tilt your head to the right, can suddenly reveal a ‘hidden’ horse. There is of course no true interpretation of the image—the visual information does not inherently contain a dog or a frog or a horse—but the predictions that we make drive what we perceive. We often face similarly ambiguous situations at night, where the mind needs to make sense of a confusing interaction of shadows and movement, for example, in order to distinguish whether there is a bat or a squirrel in a tree, or if it is just the wind. It is easy to appreciate how our predictions shape our perceptions in these

ambiguous environments, but perhaps more difficult to admit that all perception is essentially a product of our predictions. In some cases, our perceptions are such good predictors of the sensory environment that they seem very real, and so we have a lot of confidence in them. For example, there is little ambiguity between cars and trees, and so their consistency and obviousness make it hard to believe that the perception is in fact a prediction, just like the Dalmatian dog in Figure 2. The predictive processing view of the brain is influenced by a long list of prominent theories and researchers (possibly going back to Immanuel Kant (Swanson, 2016^[9]), but is most famously polished in the work of neuroscientist Karl Friston (2003^[10]), and the philosophers Andy Clarke (2013^[11]) and Jakob Hohwy (2013^[12]). We briefly provide a more detailed summary of the predictive processing view below.

Why does the brain make predictions? The simple reason is that predictions are an efficient way to navigate the world. Without a prediction process, the brain would need to gradually put together the world as if it were seeing it for the first time (every time!), carefully making sure that every bit of sensory information (for example the current taste of your own saliva) was being processed and perceived. Doing so would provide a more accurate representation of the sensory information, but it would be much slower than assuming what is “out there” and only updating those assumptions if there is an inconsistency (i.e., something surprising). By making a prediction about what ought to be experienced—and only changing those predictions when an error is made—one can make some sense of the situation and therefore act quickly (a valuable resource for the survival of any creature). There is also much redundant information that can be discarded from perception. For example, in reading this sentence there is little use in feeling the chair you are sitting on, or hearing the background hum of the air conditioner. The brain opts to ‘predict them out’ until they become relevant again (if you start feeling cold or hot then the presence or absence of a jacket is suddenly worth predicting about).

It could be said that the brain creates a virtual version or “mental model” of each moment based on its best interpretation of sensory information (this has also been called a ‘controlled hallucination’, (Ananthaswamy, 2017^[13]) and what is likely to be relevant in any given moment (Clark, 2013^[11]). When our mental model is inconsistent with sensory input, the brain encodes a ‘prediction error’, which corresponds to surprise or uncertainty (a state that the brain aims to remedy (Friston, 2009^[14]; Friston, 2010^[15]). The brain deals with surprises by either changing its predictions so that they better correspond to the input, or by acting in a way that makes sensory input more consistent with the predictions (Friston, 2013^[16]). Thus, in an uncertain situation—which the brain seeks to make predictable—one can either change the internal environment (i.e., the prediction), or the external environment (i.e., by changing behaviour) so that it matches the internal. The goal in both cases is to minimise the difference between one’s beliefs or predictions and the world as perceived by the senses. This view of the brain suggests that people try to make unpredictable circumstances predictable; so what are the consequences for human psychology as the world becomes increasingly volatile, uncertain, complex, and ambiguous?

1.2. Humans

Psychologists and neuroscientists have long studied the healthy and unhealthy ways to respond to VUCA. Each of the constructs—volatility, uncertainty, complexity, and ambiguity—have unique features but are also highly interrelated, and therefore often discussed under the umbrella of uncertainty. Below, we provide a brief review of the literature on how humans deal with unpredictable situations.

1.2.1. Uncertainty in the brain

In general, the human brain responds to uncertainty with enhanced arousal in order to dedicate more energy and resources to finding a resolution (Hirsh, Mar and Peterson, 2012^[17]). This boost of energy co-occurs with changes in neuromodulators such as dopamine, noradrenaline and acetylcholine (Marshall et al., 2016^[18]). Acetylcholine is believed to relate to expected uncertainty, or situations where predictions are unreliable. Acetylcholine has also been linked to increased effort to learn about an uncertain environment, suggesting that unhelpful top-down models are discarded and more focus is given to the current environment via sensory data (Marshall et al., 2016^[18]; Yu and Dayan, 2005^[19]). There is also evidence that depletions in dopamine impair adaptive responding to errors in predictable environments, and so changes in dopamine likely serve an important guiding function in response to uncertainty (Bestmann et al., 2014^[20]). Noradrenaline (norepinephrine)—another neurotransmitter—has also been linked to heightened arousal and alertness. In the short term, this arousal leads to heightened vigilance and a more focused attention that allows for rapid responding (Servan-Schreiber, Prinz and Cohen, 1990^[21]) and better detection of changes in the environment (Marshall et al., 2016^[18]).

While short-term arousal can be an adaptive response to uncertain and volatile environments, longer periods of arousal or stress can have negative consequences (Hirsh, Mar and Peterson, 2012^[17]; Peters, McEwen and Friston, 2017^[22]). For example, chronic stress due to uncertainty can lead to baseline noradrenaline becoming ‘tonic’ (remaining perpetually high). When the baseline firing rates are chronically high, the usual role that noradrenaline plays in goal orientation and learning from errors is hindered, so the brain struggles to predict goal orientated interpretation and action, and task performance declines (Aston-Jones and Cohen, 2005^[23]). It can become harder to focus on the task at hand, as alternative frameworks, activities and stimuli begin capturing attention (Eysenck et al., 2007^[24]; Yu and Dayan, 2005^[19]). If it is not possible to relieve the arousal of uncertainty, prolonged exposure to the stress response and its high energetic cost can drive the organism to overload (Hirsh, Mar and Peterson, 2012^[17]; McEwen, 1998^[25]). This overload can result in high blood pressure, cognitive dysfunction, depressed mood, accelerated disease progression, and permanent changes in brain architecture, minimising the organism’s capacity to deal with uncertainty in the future (Peters and McEwen, 2015^[26]; Peters, McEwen and Friston, 2017^[22]). The high energetic cost of uncertainty may explain why humans tend to experience sustained uncertainty as aversive (Peters et al., 2004^[27]). These findings also highlight why an increasingly unpredictable future poses significant challenges for humans; uncertainty can lead to both high functioning and problem solving on the one hand, or toward stress, fatigue and burnout on the other.

1.2.2. Responding to VUCA

From the perspective of the predictive brain, when an environment is predictable, then previous decisions are a good guide for future behaviour. In uncertain or volatile environments, previous decisions, and the models underlying those decisions, may no longer be effective. A healthy approach in these circumstances is to adjust the weighting of previous assumptions and allow future decisions to be guided by feedback from the environment. However, the way people respond to VUCA depends on a number of factors, including individual differences, state differences (i.e. stress, boredom, fatigue), and environmental context (i.e. perceived threat).

People who are less anxious about change (volatility), unpredictability (uncertainty), information overload (complexity), and conflicting mental models (ambiguity) function

better in changing environments than those who show high anxiety (Behrens et al., 2007_[28]). The same person can show differences in their tolerance of VUCA across contexts, and different people adapt to situational demands in different ways. Some people find VUCA a challenge to be invited, while others might perceive uncertainty as threatening. For example, the uncertainty about who a stranger is and why they are approaching you is likely to be experienced very differently if you are walking through a public park in the middle of the day versus walking through the same park in the middle of the night. Whether you are alone or in a group might also impact your experience. In the same way, an increasingly VUCA world can be perceived as either challenging or threatening to different people in different contexts. One solution to overcoming both threats and challenges is to have the skills, safeguards and tools to adapt to new environments.

Humans can be extraordinarily adaptive to change. We are intrinsically motivated by uncertainty, and driven to learn new and important skills to turn uncertain contexts into predictable ones. When people’s arousal responses are properly attuned to the uncertainty in the world, they are better at learning (De Berker, 2016_[29]). The same arousal can be exciting or anxiety-provoking, depending on the context and the individual. Indeed, openness to continually exploring, learning and growing—learning how to learn—may be one key to future success in a world approaching widespread automation and career displacement, which can be a direct threat to the livelihoods of many people. The challenge for the future of education, training, and work, will therefore be to leverage the naturally curious and exploratory nature of people, and develop social safeguards for those who are more directly or negatively impacted by the threatening aspects of VUCA. In later sections, we will address how individuals with a sense of control, or agency, fare better in VUCA situations; how agency can be enhanced; and how individuals with agency tend to show higher levels of adaptability and lower stress in unpredictable environments. Before we discuss the skills needed for future adaptability, we will first discuss the role played by technology in creating and navigating a VUCA world.

1.3. Machines

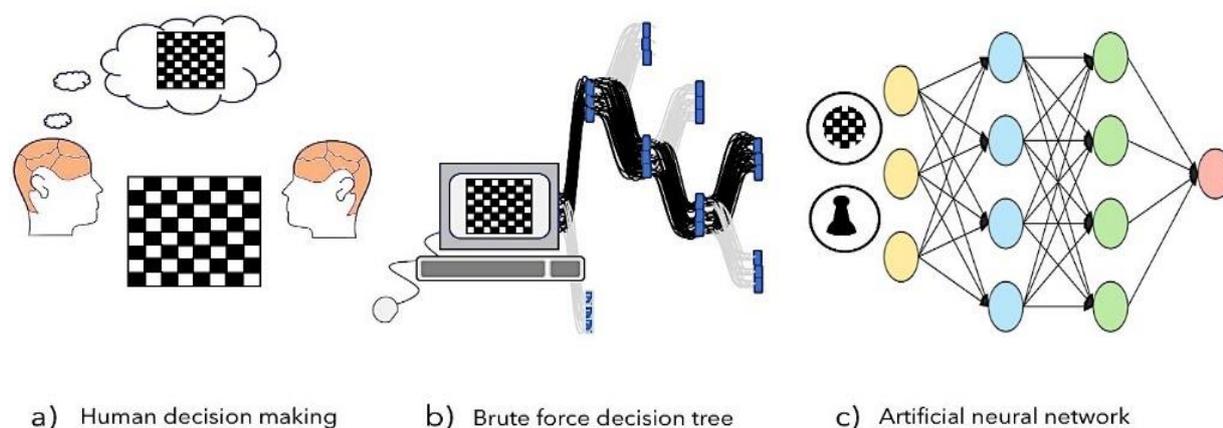
Humans are increasingly relying on smart tools to help us accomplish simple and complex tasks. Our devices are powered by technology that can learn from us, learn about us, assist us reaching our goals, shape our habits, and augment our natural problem-solving capabilities to help us make better decisions about our world. This technology is broadly defined as machine or artificial intelligence (AI) and will likely remain a cornerstone of human decision making into the foreseeable future. AI technology can also help humans overcome many of our own shortcomings, but how do volatility, uncertainty, complexity, and ambiguity impact the ability of AI to effectively solve problems? How does technology deal with VUCA in general?

There are two broad spheres of AI: the “general” AI of popular culture, which refers to the smart, adaptable machines that have the ability to approach and understand a variety of situations with human-level capability (i.e. running, driving, cooking, fighting, climbing, etc.; all the general-purpose tasks humans can or could perform proficiently); and “specific” AI, which can perform a small set of narrowly-defined tasks with human or superhuman capability. The general form of AI is the goal for research and development in machine intelligence, but it so far remains outside the short-term grasp of science and engineering (Lake et al., 2017_[30]). Already, however, the specific form of AI is permeating many aspects of our day-to-day lives; we have smart algorithms curating our search engine

results, robots defeating our world champions in chess and Go, intelligent optimisation and tailoring of our news and social media feeds, virtual machines mapping and navigating our paths using GPS, and a score of present and emerging applications within the domains of research, industry, and consumer products.

The success of machine learning and AI more broadly, has resulted from the development of computer programs (or algorithms) that can learn from complex data. There are many approaches to machine learning, but some of the most effective are the artificial neural networks (ANNs; see Figure 3), which are directly inspired by, and conceptually imitate, some of the functions of the human brain (Turing, 1950^[31]; Newell and Simon, 1961^[32]; LeCun, Bengio and Hinton, 2015^[33]). Smart algorithms are intelligent in the sense that they continuously improve their own performance (in prediction accuracy or decision-making, for example) without explicit programming. As a result, a software engineer need not know precisely how an intelligent algorithm arrives at a solution. Rather, the learning process can be used to find intelligent solutions to complex problems.

Figure 3. Decision making by human and machine intelligence.



a) Humans have developed rule-based mental models that allow us to compete, strategise, and outsmart each other; b) computers first developed the ability to defeat humans at games of strategy by sheer processing power, which involved calculating and accounting for as many outcomes as possible; c) brain-inspired neural networks now develop rule based strategies and learn by playing humans and other AI.

Of the challenges that VUCA presents, complexity is an important domain within which the specific form of AI can be very accomplished relative to humans (Palmer and Chakravarty, 2014^[34]). The amount and complexity of data generated in today's connected world is far beyond the processing capacity of humans, or even most standard data-processing software. An intelligent algorithm can be used to sort through and classify information consistent with a pre-specified goal (supervised learning) or to cluster information according to the underlying structure of the data the algorithm receives (unsupervised learning). Depending on the goal, either of these approaches can be exceedingly useful for synthesising vast amounts of information ("Big Data") and refining it into a form better suited to human intuitions (e.g. in the form of visual relationships, maps, suggestions and recommendations, and searchable databases).

AI can also be well-equipped to deal with a relatively high degree of uncertainty. Some of the early advances in computer intelligence came not from learning per se, but from the ability to calculate many potential solutions in a short time (known as a 'brute force' approach (Campbell, Hoane and Hsu, 2002^[35]). In situations of uncertainty, computers can generate a large and comprehensive contingency plan by simulating many possible futures (a decision tree), with enough accuracy and precision to compete with humans. While not typically considered 'intelligence', because it relies on static retrieval of stored information (Hsu, Campbell and Hoane, 1995^[36]), it was this brute force approach that led to the first computer, IBM's Deep Blue, unseating the world chess champion Garry Kasparov in a series of games in 1997. The speed and processing power of computational processing means that machine intelligence can be well-equipped to deal with complex information, and many possible and uncertain futures. However, resolving uncertainty and complexity is not without cost. Typically, this cost comes in the form of a trade-off between speed and accuracy, so AI can perform poorly when tasked with resolving a high number of relevant variables, if required to do so in real-time (Lake et al., 2017^[30]).

While the challenges of complexity and uncertainty can, to some degree, be mitigated by prioritising processing power or processing speed according to situational demands, volatility and ambiguity can represent a far more delicate challenge to the competency of both general and specific AI, and for quite different reasons. Specific AI, being adept at completing specific tasks, can lack the flexibility to deal with volatility in task demands or contexts (e.g. generalisation to another, even closely related, task). Unexpected inputs or demands on an algorithm can cause functioning to cease entirely (Mnih et al., 2015^[37]), and ambiguity can lead to an intelligent algorithm (even a general AI) to performing its function even when it is inappropriate to do so (Wang, Ashfaq and Fu, (n.d.)^[38]). To function in volatile and ambiguous domains, AI requires a more general approach to problem solving: the ability to learn how to learn.

A central barrier to developing AI that can think flexibly in volatile and ambiguous situations is the ability of an algorithm to generate 'mental' models of the world that are sufficiently general, and yet abandon them in circumstances where it is more prudent to do so (Lake et al., 2017^[30]). Humans, by contrast, can be strongly guided by our mental models of physics (the effect of gravity, the persistence of unseen items, causal relationships between objects) or psychology (what another person might think, intend, or not know), but can easily discard these intuitions if they do not serve the present purpose (e.g. learning a new skill). Specific AI lacks this flexibility. An AI with superhuman capability in one context (getting a high-score in a game of strategy), can take hundreds or even thousands of hours to learn to adapt its behaviour to arbitrary changes in goal (getting the lowest score) or context (changes in the shape or colour of a target) (Collins and Frank, 2013^[39]). Learned models of the world can interfere with AI adapting to novel situations with

different rules, even situations that would pose no barrier for the predictive processing capacities of a child (Eliasmith et al., 2012^[40]; Rougier et al., 2005^[41]).

Rapid technological progress has resulted in AI capable of dealing effectively with complexity and uncertainty. The human capability of adapting to novel situations is currently outside the scope of most intelligent algorithms. At present, context changes and situational ambiguity still provide a distinct challenge for moving from specific AI to more general AI. However, changes in technology are occurring at an ever-increasing pace. In the near future, machines will gain more autonomy and greater adaptability, and the technological limitations of today might be short-lived.

1.4. Summary of Part 1

Clearly, VUCA situations challenge both human and machine intelligence. We have so far suggested that people can respond in healthy and unhealthy ways to VUCA situations (e.g., perceiving change as a challenge, versus change as a threat). We explored how humans are potentially highly adaptive and in many instances, and can thrive despite unpredictable conditions. According to the predictive processing account of the mind, humans can reduce uncertainty through their actions in the world, through updating their beliefs, and by their ability to discard beliefs when they are unhelpful. In other words, they can deal effectively with uncertainty by being adaptable learners. Unpredictability can also lead people to feel anxious, stressed, or threatened when rigid beliefs are challenged. Therefore, the question is how best to prepare humans to deal with a future that is increasingly unpredictable, and to minimise the threatening aspects of rapid and disruptive changes.

We also explored the current state of research in artificial intelligence, and pointed to some crucial differences in human and machine intelligence. When placed in an unusual circumstance—such as a new country, new school or new workplace—humans can effectively learn the structure of the environment and replace old structures or beliefs that are no longer relevant. On the other hand, in many VUCA situations, machines are yet to demonstrate efficacy in principle, particularly in situations characterised by ambiguity and volatility (Lake et al., 2017^[30]). AI can be exceedingly efficient in completing specific tasks, and respond effectively to complexity and some features of uncertainty, but can be unable to deal with changes in goals and contexts of a task (Collins and Frank, 2013^[39]). Put simply, humans possess the capacities to deal with VUCA but sometimes fail to do so productively, whereas, in some cases machines lack the capacities altogether. In Part 2, we discuss how agency may be a valuable tool in the reduction of uncertainty, and another domain where differences between machines and humans are likely to be revealing.

2. PART 2: Agency

2.1. Humans

Western society is predicated on the notions of individual liberty and personal responsibility. Both concepts rest on the idea of agency—that the individual can make their own choices and act accordingly. A strong feeling of agency has been associated with higher life satisfaction, self-efficacy and self-control, a higher sense of meaning, and lower stress (Baumeister and Brewer, 2012^[42]). The attribution of agency to our own actions, and the actions of others, also facilitates social cohesion and responsibility (Baumeister, Crescioni and Alquist, 2011^[43]). Beliefs about one’s own and others free will (the capacity to make free choices and act accordingly) have also been shown to significantly affect people’s moral attitudes and behaviour, as well as their sense of agency (Vohs and Schooler, 2008^[44]; Shariff et al., 2014^[45]; Lynn et al., 2014^[46]). As we shall see, the sense of agency seems to arise from our capacity to act in line with our belief systems, goals, and values. These structures not only help us interpret uncertain stimuli, but seem to also guide our behaviour, decisions, and actions in an uncertain world.

One theory supposes that if the arising thoughts and intentions to act (predicted outcome) precede and are consistent with the action (actual outcome) then a sense of agency occurs, even if our action did not in fact cause the outcome. People can easily be manipulated to experience a false sense of agency. For example, simply having participants think about an upcoming action on a pointer (similar to a ouiji board) increased their sense of agency over that action, even though they were not in fact responsible (Wegner & Wheatley, 2009). From pressing the lights at a pedestrian crossing that is in fact on a timer, to sticking pins through a voodoo doll—our prior beliefs about our agency, seem to affect our sense of control over the outcome regardless of their veracity. The sense of agency is not necessarily illusory, however. Rather, it is highly flexible and therefore liable to making mistakes. Our beliefs about agency, our actions, and sensory feedback work together to adaptively inform our perception of agency. Most agency researchers now believe that both predictive and postdictive processes contribute to the experience of agency in action. Synofzik, Vosgerau, and Voss (2013^[47]) argue that internal predictions based on representations of motor commands, or “virtual versions” are matched to sensorimotor output (akin to predictive processing). Where the expected state meets the actual state, we feel a sense of agency for the action, and if it does not, then no sense of agency occurs (Frith, 2012^[48]). For example, choosing to pick up a water bottle, and then perceiving one’s hand moving to pick up the water-bottle would result in a sense of agency. Moore and Haggard (2008^[49]) suggest that when sensory data are reliable, sensory feedback will be weighted heavily in informing the sense of agency. When sensory data are weak, estimates will rely more on prior information (i.e. beliefs, knowledge, and contextual information). Like all of our perceptions, the sense of agency is based in a self-updating and flexible processing system and is therefore liable to errors in the service of maintaining both structure and adaptability. The price of errors may be worth the payoff. Because of the flexibility of our processing, we are able to have a pragmatic sense of agency even where a more literal agency does not exist. This sense of agency complements goal-directed action, constraining the number of possible affordances of interpretation and action, and making it clear that our own actions are capable of creating change in the world.

2.1.1. Beliefs, Agency, and Uncertainty

It is undeniable that most of the events in our lives are far more unpredictable and less certain than we like to admit. Just imagine what your teenage self would have predicted about your career, your beliefs, or any other facet of your current situation. Humans have clearly developed sophisticated methods for dealing with the inherently uncertain world in which we live. According to the Meaning Maintenance Model (Heine, 2006_[50]), having structures of value, such as belief systems, goals and identities can imbue meaning to otherwise ambiguous stimuli, directing our perception and action and reducing experiences of uncertainty.

How do beliefs reduce uncertainty? According to the Entropy Model of Uncertainty (Hirsh, Mar and Peterson, 2012_[17]), both goals and belief structures help to reduce uncertainty because they reduce the number of possible interpretations and actions. In line with the prospect that belief systems act as useful constraints, fundamentally religious people are more effective at regulating uncertainty, and even show a reduced neural response to uncertainty and error (Inzlicht et al., 2009_[51]; Kossowska et al., 2016_[52]). This protective function of beliefs in the face of uncertainty may partially explain the lure of conspiracy theories, and also why we tend to cling to our treasured beliefs (Nyhan and Reifler, 2010_[53]). As described earlier, our models of the world result in ongoing predictions, and these predictions are examined against and updated according to new (discrepant) information. The brain typically responds to uncertainty by either changing the predictions to better correspond to the input (updating the mental model), or by ‘sculpting and selecting’ sensory inputs to be more consistent with the predictions (changing behaviour to adapt to the prevailing environment) (Clark, 2013_[11]). Without mental flexibility—the willingness to change one’s models, we can keep operating with outdated models—by picking and choosing the information that suits them—and can become out of touch with the dangers and opportunities presented in a changing world. In some instances, people respond to uncertainty with rigid and repetitive ritual frameworks in order to regain a sense of control (Lang et al., 2015_[54]) and decrease anxiety (Brooks et al., 2016_[55]). Extreme manifestations of this can be seen in OCD, which has previously been associated with intolerance of uncertainty (Tolin et al., 2003_[56]). While a-priori structures (including cognitive and behavioural frameworks) are necessary to facilitate perception and action (Heine, 2006_[50]), by holding on to these structures too tightly we risk becoming rigid and dogmatic, unable to genuinely consider conflicting information or alternative perspectives.

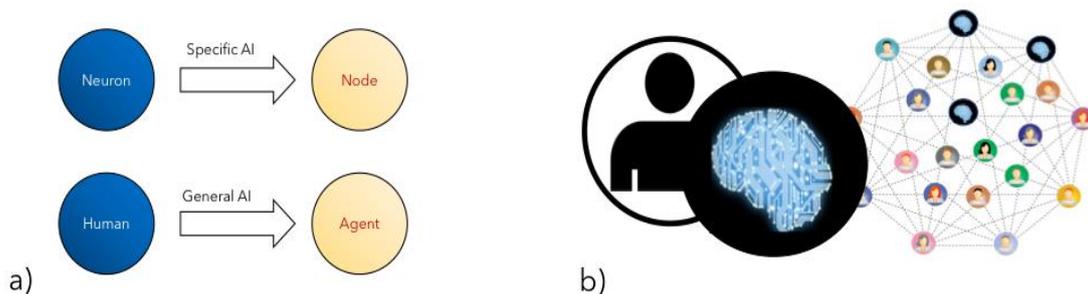
Individuals, groups and societies will always experience tension between the desire for stability and the inevitability of change. As we have seen, we meet our ever-changing world with models and beliefs—the ways of seeing things that have worked for us in the past. These models typically work well enough, and allow us to function with relative ease in an incredibly complex environment. However, we have also seen that there can be psychological resistance to information and ideas that challenge or contradict our models. Avoiding uncertainty is not a viable long-term strategy. In order to learn from uncertainty, we first need to tolerate its existence. In facing uncertainty, psychological flexibility and openness to learning may be key (Kasdan and Rottenberg, 2010_[57]). Fortunately, in the case of humans, our hardware is flexible (Draganski et al., 2004_[58]), and therefore adjusting to optimally meet uncertainty is challenging but not impossible. Can machines similarly use structures and agency to minimise uncertainty, or are there unique barriers in artificial intelligence?

2.2. Machines

The increasing integration of smart technology at many levels of society will require careful consideration of AI's capacity for dealing with VUCA situations. So far, though, smart technology has been considered primarily, or perhaps entirely, a tool for the benefit of human goals. In the future, the generalisability and capability of artificial intelligence will raise questions about the status of autonomous technologies as agents in their own right.

A modern approach in machine learning takes inspiration from the naturally adaptable way that humans approach and learn to deal with different situations, in particular, by trying to maximise the value of a future reward. This process, called reinforcement learning (Mnih et al., 2015^[37]), provides a value structure for the program to orient toward. Rather than account for as many alternatives as possible, like the brute force approach, the ANN (or agent) will attempt a single solution, initially at random, and use the accrual of reward as a measure of the effectiveness of the approach. This process is particularly effective at learning how to play games and solve problems that have small goals that comprise a larger objective. With time, the 'agent' can produce behaviours that solve complex problems because the algorithm continually selects for solutions that result in greater rewards. The reinforcement learning approach, inspired as it is by human learning, might provide programs with the learning-to-learn ability that would allow AI to function as autonomous agents in real-world contexts (illustrated in Figure 4).

Figure 4. Reinforcement learning



a) The functional equivalence of neurons and nodes, in specific domains, will be followed by the equivalence of human and AI capability in general domains; b) At some point, Human and AI competence will see functional equivalence (the singularity) and autonomous agents will be integrated into the community.

The challenge for a general AI becoming an autonomous agent is not only about gaining the ability to navigate a VUCA world with human-level competency. There are also practical and ethical challenges. Consider a situation in which a driverless car is faced with a specific goal: transport its inhabitants safely from location A to location B. Imagine now that the car is approaching a busy pedestrian crossing and the brakes have failed. The autonomous car then has a decision to make: save you—the passenger—from injury, at the expense of multiple lives, or swerve into a wall to avoid the pedestrians and potentially injure you. The 'hard-coded' value structure for the AI (keep the passenger safe) certainly will help guide behaviour, but even well-intended solutions can be disastrous when inflexibly applied. More practically, though, would you buy a car that chooses to kill you? Can you justify the car that kills pedestrians? Even if autonomous systems make 'correct' decisions, and function as intended, there are practical barriers to granting autonomy to AI technology in a volatile and ambiguous world.

There are also practical considerations for not granting agency to AI technology. At present, we have AI systems that guide and facilitate human decision-making, but this intermediate approach is unlikely to continue indefinitely. At some point, it might become unethical to allow humans to make autonomous decisions because the precision, accuracy, and response time of intelligent machines will be far superior to human operators. Moreover, the point at which human and AI competency is equivalent (the singularity) will be brief, and the period in which AI surpasses human competency will be indeterminately long.

2.3. Summary of Part 2

Humans possess a sense of agency, which is imperfect, and error-prone, but guides decisions in a productive way, often effectively mitigating uncertainty. For people, the question is ‘how to make the most of values, beliefs, and goals as functions of agency’, whereas machines do not yet possess the capabilities that would allow the same sort of autonomy (although some mechanisms such as reinforcement learning appear to come close). Difficulties emerge particularly when situations are ethically ambiguous. AI is good at following rules, but many situations in life are embedded in specific contexts that change the rules. In such circumstances, the complex interactions between values, beliefs, and goals can provide axioms—and agency—for decision-making in thorny moral scenarios. For now, the ability to approach uncertainty, develop new beliefs and values, and act according to those beliefs in a productive way according to malleable goals—accompanied by a distinct sense that one *has* agency and efficacy in doing so—remains unique to humans.

3. Discussion

“We’re in an unprecedented situation in history in the sense that nobody knows what the basics about how the world will look like in 20 or 30 years. Not just the basics of geopolitics but what the job market would look like, what kind of skills people will need, what family structures will look like, what gender relations will look like. This means that for the first time in history we have no idea what to teach-in-schools.” Yuval Noah Harari (de Freytas-Tamura, 2018_[59])

Is it possible to prepare for the unknown? If a traveller embarks to a new destination with an unfamiliar culture, an unknown terrain, and they do not know what language is spoken, it is not useful for them to try and learn Norwegian, or to hike through snowy mountains. They cannot, by definition, prepare with specific skills, since there is no compass for determining what specific skills will be valuable. They can, on the other hand, prepare themselves in another way. They can equip themselves with the states of mind, the strategies, and the self-efficacy, required to learn the new skills, learn the structure of the new culture, and learn the new language, so that they are ready to adapt when they arrive at their destination.

Our literature review indicates that, with regard to the human mind, VUCA territories are potentially hazardous; but not necessarily so. Uncertainty, when registered as threatening, can cause a cascade of sympathetic responses associated with stress and anxiety (de Berker et al., 2016_[60]). The brain, as it fails to predict the environment, registers this failure as problematic. Prolonged exposure to uncertainty related stress can result in severe states of confusion, distraction, depressed moods, and may even accelerate disease progression (Eysenck et al., 2007_[24]; Peters and McEwen, 2015_[26]). A low tolerance of uncertainty is also associated with worry, anxiety disorders, phobias, and depression (Buhn and Dugas, 2002_[61]; Dugas et al., 1998_[62]; Freeston et al., 1994_[63]; McEvoy and Mahoney, 2011_[64]). However, these are not inevitable reactions. In the best-case scenarios, the arousal caused by uncertainty can be a strong motivator that increases sensory vigilance, attention, and can improve task performance (Servan-Schreiber, Printz and Cohen, 1990_[21]; Proulx and Heine, 2009_[65]). Individuals who are tolerant of uncertainty may see VUCA situations as opportunities for learning—as exciting (perhaps even awe-inspiring) events that can with time and effort, be understood.

As discussed in PART 2: Agency, a sense of agency can be a powerful way to mitigate the harmful effects of VUCA. If one has values, beliefs, and goals that help to constrain the myriad possible actions and options—and has the sense that they are free to act according to them—then one may be less likely to experience uncertainty as threatening (Inzlicht and Tullett, 2010_[66]; Kossowska et al., 2016_[67]; Proulx, Inzlicht and Harmon-Jones, 2012_[68]). How might society and education systems provide a genuine sense of agency to people dealing with an increasingly dynamic and changing world? If one loses their source of employment to automation, for example, how can they be encouraged to see this uncertainty as an opportunity rather than a threat? To answer these questions, we need to consider what opportunities are genuinely afforded by change and VUCA. Perhaps the clearest opportunity presented is that of learning. Almost any situation where genuine learning can occur, there is necessarily uncertainty. For example, walking into a new classroom and being presented with new material inherently exposes a prospective student to uncertainty. If the material in the classroom is seen as non-threatening, and one has a sense of agency and self-efficacy in learning the material, then a positive learning experience can occur (Zimmerman, 2000_[69]). By analogy, VUCA situations that occur naturally in the world can also be opportunities for learning—because where there is uncertainty there is also something new to understand. And since any one skill that is

learned may have a short half-life as an employment prospect, then fostering the ability to learn itself may be the only viable solution.

Learning to learn is also valuable because it remains a key challenge for artificial intelligence. In many specific domains, machines are already outperforming humans. However, if you place a machine that is an expert in Chess, and ask it to play Tetris, it may never learn the rules without further human tinkering. On the other hand, if you teach a child to play Chess it might not make them better at playing Tetris, but they *will* be able to learn to play Tetris (or indeed basketball!) without too much difficulty. This ability to learn the structure of one environment, and then put aside those structures (if they are irrelevant) in order to learn new structures, is a valuable advantage. It is not only an effective way to make humans robust to change, but it is also one of the capacities that is least likely to be computerised. Machines also must adapt to change, and if humans are more efficient in this regard, then the niche for human beings does not need to be a specific skill, instead, it may be the ability to pick up new skills itself.

What is meta-learning, and can it be improved? According to Maudsley (1979^[70]), meta-learning is “...the process by which learners become aware of and increasingly in control of habits of perception, inquiry, learning, and growth that they have internalised.” John Biggs (Biggs, 1985^[71]) used the term to refer to a state of “...being aware of and taking control of one’s own learning.” To put it simply, meta-learning is knowing how best to learn. To improve meta-learning one can learn the principles and strategies that make learning more efficient. Retrieval practice, spacing, interleaving, feedback, elaboration, and testing, can all make an enormous impact on the quality and efficiency of learning that occurs (Roediger III and Butler, 2011^[72]; Karpicke and Blunt, 2011^[73]). Students can also be taught to monitor their progress and to respond positively in the face of failure (Dweck et al., 1978^[74]; Boekaerts and Corno, 2005^[75]). Interventions that involve teaching metacognitive strategies and self-regulated learning also show substantial improvements in academic outcomes, and these strategies can be used throughout one’s life (Dignath and Büttner, 2008^[76]). The ability to transfer knowledge from one domain to another—where such transfer is relevant—appears also to be a learnable skill (Brown and Kane, 1988^[77]). Clearly, there are many effective strategies for making humans better learners, and if an individual is equipped with these meta-learning capacities, VUCA may not appear so threatening.

How can society be structured to incentivise lifelong learners? If nearly every individual will need to change professions and update their skills during their adult careers, then many more teachers, mentors, and instructors, will be needed. Individuals can be prepared during primary and secondary education to be effective meta-learners, but in order to ‘upskill’ out in the world, they also require someone to teach them. If, for example, a major technological advancement replaces an entire sector in the workforce with machines, then there will need to be willing and incentivised instructors to accommodate their transition into new industries. If individuals from a broad range of careers with unique expertise are encouraged to be mentors and to hire apprentices — and have, the means to support those apprentices—then switching careers may be far less daunting. Teaching and instructing itself is also one of the least likely jobs to be automated according to the 2013 report from Oxford University (Frey and Osborne, 2013^[78]). Therefore, when one door closes (e.g., due to automation), then new employment is also created for those who are willing to teach the transitioning workforce. Here, social policies that incentivise instructors are valuable. For example, by providing a stipend to employers and experts to teach and to hire apprentices, all parties benefit. The employers have the opportunity to teach, and the unemployed can fluently upskill and transition into a new sector.

4. Concluding Remarks

The world was once like an ocean on a calm day—the movements were still difficult to predict but the waves of change were slower and less dramatic. Today, technology, globalisation, and population growth contribute to increasingly stormy seas. It is good fortune then, that humans are endowed with a certain plasticity that makes them adaptable to change through learning. In a VUCA world, this may be their greatest asset. While uncertainty has its perils for humans, it is the machines that are not yet sufficiently advanced to deal with the ambiguity and the volatility of modern life. When equipped with meta-learning strategies—and access to teachers and mentors—humans may very well be in a position to thrive, as lifelong learners, in the same circumstances where machines blunder. Training meta-learning strategies can also provide individuals with a tool—and a valuable sense of agency—to embrace uncertainty as opportunity for learning.

References

- Ananthaswamy, A. (2017), “That's a termite colony between your ears”, *New Scientist*, Vol. 233/3112, pp. 42-43. [13]
- Aston-Jones, G. and J. Cohen (2005), “An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance”, *Annual Review of Neuroscience*, Vol. 28, pp. 403-450. [23]
- Baumeister, R. and L. Brewer (2012), “Believing versus disbelieving in free will: Correlates and consequences”, *Social and Personality Psychology*, Vol. 6/10, pp. 736-745. [42]
- Baumeister, R., A. Crescioni and J. Alquist (2011), “Free will as advanced action control for human social life and culture”, *Neuroethics*, Vol. 4/1, pp. 1-11. [43]
- Behrens, T. et al. (2007), “Learning the value of information in an uncertain world”, *Nature neuroscience*, Vol. 10/9, p. 1214. [28]
- Bennis, W. and B. Nanus (1985), *Leaders: The Strategies for Taking Charge*, Harper & Row. [6]
- Bestmann, S. et al. (2014), “The role of dopamine in motor flexibility”, *Journal of cognitive neuroscience*, Vol. 27/2, pp. 365-376. [20]
- Biggs, J. (1985), “The role of metalearning in study processes”, *British Journal of Educational Psychology*, Vol. 55, pp. 185-212. [71]
- Boekaerts, M. and L. Corno (2005), “Self-regulation in the classroom: A perspective on assessment and intervention”, *Applied Psychology*, Vol. 54/2, pp. 199-231. [75]
- Brooks, A. (2014), “Get excited: Reappraising pre-performance anxiety as excitement”, *Journal of Experimental Psychology: General*, Vol. 143/3, p. 1144. [84]
- Brooks, A. et al. (2016), “Don't stop believing: Rituals improve performance by decreasing anxiety”, *Organizational Behavior and Human Decision Processes*, Vol. 137, pp. 71-85. [55]
- Brown, A. and M. Kane (1988), “Preschool children can learn to transfer: Learning to learn and learning from example”, *Cognitive Psychology*, Vol. 20/4, pp. 493-523. [77]
- Buhn, K. and M. Dugas (2002), “The intolerance of uncertainty scale: Psychometric properties of the English version.”, *Behaviour research and therapy*, Vol. 40/8, pp. 931-945. [61]
- Campbell, M., J. Hoane and F. Hsu (2002), “Deep Blue”, *Artificial intelligence*, Vol. 134/1-2, pp. 57-83. [35]
- Clark, A. (2013), “Predictive brains, situated agents, and the future of cognitive science”, *Behavioral and brain sciences*, Vol. 36/3, pp. 181-204. [11]
- Collins, A. and M. Frank (2013), “Cognitive control over learning: Creating, clustering, and generalizing task-set structure”, *Psychological review*, Vol. 120/1, p. 190. [39]

- De Berker, A. (2016), “Acute stress selectively impairs learning to act”, *Scientific Reports*, Vol. 6/29816. [29]
- de Berker, A. et al. (2016), “Computations of uncertainty mediate acute stress responses in humans”, *Nature communications*, Vol. 7, p. 10996. [60]
- de Freytas-Tamura, K. (2018), *What’s Next for Humanity: Automation, New Morality and a ‘Global Useless Class’*. [59]
- Dignath, C. and G. Büttner (2008), “Components of fostering self-regulated learning among students. A meta-analysis on intervention studies at primary and secondary school level”, *Metacognition and learning*, Vol. 3/3, pp. 231-264. [76]
- Draganski, B. et al. (2004), “Neuroplasticity: changes in grey matter induced by training”, *Nature*, Vol. 427/6972, p. 311. [58]
- Dugas, M. et al. (1998), “Generalized anxiety disorder: A preliminary test of a conceptual model”, *Behaviour research and therapy*, Vol. 36/2, pp. 215-226. [62]
- Dweck, C. et al. (1978), “Sex differences in learned helplessness: II. The contingencies of evaluative feedback in the classroom and III. An experimental analysis”, *Developmental psychology*, Vol. 14/3, p. 268. [74]
- Eliasmith, C. et al. (2012), “A large-scale model of the functioning brain”, *Science*, Vol. 338/6111, pp. 1202-1205. [40]
- Eysenck, M. et al. (2007), “Anxiety and cognitive performance: attentional control theory”, *Emotion*, Vol. 7/2, p. 336. [24]
- Freeston, M. et al. (1994), “Why do people worry?”, *Personality and individual differences*, Vol. 17/6, pp. 791-802. [63]
- Frey, C. and M. Osborne (2013), *The Future of Employment: How Susceptible are jobs to Computerisation*, http://sep4u.gr/wp-content/uploads/The_Future_of_Employment_ox_2013.pdf. [78]
- Friston, K. (2013), “Learning and inference in the brain”, Vol. 16/9, pp. 1325-1352. [16]
- Friston, K. (2010), “The free-energy principle: a unified brain theory?”, *Nature Reviews Neuroscience*, Vol. 11/2, p. 127. [15]
- Friston, K. (2009), “The free-energy principle: a rough guide to the brain?”, *Trends in cognitive sciences*, Vol. 13/7, pp. 293-301. [14]
- Friston, K. (2005), “A Theory of Cortical Responses”, *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, Vol. 360/1456, pp. 815-836. [79]
- Friston, K. (2003), “Learning and inference in the brain”, *Neural Networks*, Vol. 16/9, pp. 1325-1352. [10]

- Frith, C. (2012), “Explaining delusions of control: The comparator model 20 years on”, *Consciousness and cognition*, Vol. 21/1, pp. 52-54. [48]
- Greene, B. (2006), *The Universe on a String*, [83]
<https://www.nytimes.com/2006/10/20/opinion/20greenhed.html>.
- Heine, S. (2006), “The meaning maintenance model: On the coherence of social motivations”, *Personality and Social Psychology Review*, Vol. 10/2, pp. 88-110. [50]
- Hirsh, J., R. Mar and J. Peterson (2013), “Personal narratives as the highest level of cognitive integration”, *Behavioral and Brain Sciences*, Vol. 36/3, pp. 216-217. [82]
- Hirsh, J., R. Mar and J. Peterson (2012), “Psychological entropy: A framework for understanding uncertainty-related anxiety”, *Psychological review*, Vol. 119/2, pp. 304-320. [17]
- Hohwy, J. (2013), *The predictive mind*, Oxford University Press. [12]
- Hsu, F., M. Campbell and A. Hoane (1995), *Deep Blue system overview*, ACM. [36]
- Inzlicht, M. et al. (2009), “Neural markers of religious conviction”, *Psychological Science*, Vol. 20/3, pp. 385-392. [51]
- Inzlicht, M. and A. Tullett (2010), “Reflecting on God: Religious primes can reduce neurophysiological response to errors”, *Psychological Science*, Vol. 21/8, pp. 1184-1190. [66]
- Johansen, B. (2013), “Navigating the VUCA world”, *Research-Technology Management*, Vol. 56/1, pp. 10-15. [1]
- Karpicke, J. and J. Blunt (2011), “Retrieval practice produces more learning than elaborative studying with concept mapping”, *Science*, p. 1199327. [73]
- Kasdan, T. and J. Rottenberg (2010), “Psychological flexibility as a fundamental aspect of health”, *Clinical psychology review*, Vol. 30/7, pp. 865-878. [57]
- Kossowska, M. et al. (2016), “Self-image threat decreases stereotyping: The role of motivation toward closure”, *Motivation and emotion*, Vol. 40/6, pp. 830-841. [67]
- Kossowska, M. et al. (2016), “Anxiolytic function of fundamentalist beliefs: neurocognitive evidence”, *Personality and Individual Differences*, Vol. 101, pp. 390-395. [52]
- Kurzweil, R. (2005), *The Singularity is Near*, New York: Viking Books. [4]
- Lake, B. et al. (2017), “Building Machines That Learn and Think Like People”, *Behavioral and Brain Sciences*, Vol. 40. [30]
- Lang, M. et al. (2015), “Effects of anxiety on spontaneous ritualized behavior”, *Current Biology*, Vol. 25/14, pp. 1892-1897. [54]
- LeCun, Y., Y. Bengio and G. Hinton (2015), “Deep Learning”, *Nature*, Vol. 521, pp. 436-444. [33]

- Lynn, M. et al. (2014), “Priming determinist beliefs diminishes implicit (but not explicit) components of self-agency”, *Frontiers in psychology*, Vol. 5, p. 1483. [46]
- Marshall, J. and J. Oberwinkler (1999), “Ultraviolet vision: The colourful world of the mantis shrimp.”, *Nature*, Vol. 401/6756, p. 873. [8]
- Marshall, L. et al. (2016), “Pharmacological fingerprints of contextual uncertainty”, *PLoS Biology*, Vol. 14/11, p. e1002575. [18]
- Maudsley, D. (1980), “A theory of meta-learning and principles of facilitation: An organismic perspective”, *Unpublished Masters Thesis*. [7]
- Maudsley, D. (1979), *A theory of meta-learning and principles of facilitation: An organismic perspective (Unpublished Master’s thesis)*., University of Toronto, Canada.. [70]
- McEvoy, P. and A. Mahoney (2011), “Achieving certainty about the structure of intolerance of uncertainty in a treatment-seeking sample with anxiety and depression”, *Journal of Anxiety Disorders*, Vol. 25/1, pp. 112-122. [64]
- McEwen, B. (1998), “Stress, adaptation, and disease: Allostasis and allostatic load”, *Annals of the New York academy of sciences*, Vol. 840/1, pp. 33-44. [25]
- McEwen, B. and E. Stellar (1993), “Stress and the individual: mechanisms leading to disease”, *Archives of Internal Medicine*, Vol. 153/18, pp. 2093-2101. [80]
- Mnih, V. et al. (2015), “Human-level control through deep reinforcement learning”, *Nature*, Vol. 518/7540, p. 529. [37]
- Moore, G. (1975), “Progress In Digital Integrated Electronics”, *Technical Digest*, pp. 11-13. [2]
- Moore, J. and P. Haggard (2008), “Awareness of action: Inference and prediction”, *Consciousness and cognition*, Vol. 17/1, pp. 136-144. [49]
- Newell, A. and H. Simon (1961), *GPS, a program that simulates human thought*. [32]
- Nyhan, B. and J. Reifler (2010), “When corrections fail: The persistence of political misperceptions”, *Political Behavior*, Vol. 32/2, pp. 303-330. [53]
- Palmer, J. and A. Chakravarty (2014), *Supervised machine learning*, John Wiley & Sons Inc. [34]
- Pasmore, B. and T. O’Shea (2010), “Leadership agility: A business imperative for a VUCA world”, *People and Strategy*, Vol. 33/4, p. 32. [5]
- Peters, A. and B. McEwen (2015), “Stress habituation, body shape and cardiovascular mortality”, *Neuroscience & Biobehavioral Reviews*, Vol. 56, pp. 139-150. [26]
- Peters, A., B. McEwen and K. Friston (2017), “Uncertainty and stress: Why it causes diseases and how it is mastered by the brain”, *Progress in neurobiology*, Vol. 156, pp. 164-188. [22]

- Peters, A. et al. (2004), “The selfish brain: competition for energy resources”, *Neuroscience & Biobehavioral Reviews*, Vol. 28/2, pp. 143-180. [27]
- Proulx, T. and S. Heine (2009), “Connections from Kafka: Exposure to meaning threats improves implicit learning of an artificial grammar”, *Psychological Science*, Vol. 20/9, pp. 1125-1131. [65]
- Proulx, T., M. Inzlicht and E. Harmon-Jones (2012), “Understanding all inconsistency compensation as a palliative response to violated expectations”, *Trends in cognitive sciences*, Vol. 16/5, pp. 285-291. [68]
- Roediger III, H. and A. Butler (2011), “The critical role of retrieval practice in long-term retention”, *Trends in cognitive sciences*, Vol. 15/1, pp. 20-27. [72]
- Rougier, N. et al. (2005), “refrontal cortex and flexible cognitive control: Rules without symbols”, *roceedings of the National Academy of Sciences*, Vol. 102/20, pp. 7338-7343. [41]
- Servan-Schreiber, D., H. Printz and J. Cohen (1990), “A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior”, *Science*, Vol. 249/4971, pp. 892-895. [21]
- Shariff, A. et al. (2014), “Free will and punishment: A mechanistic view of human nature reduces retribution”, *Psychological Science*, Vol. 25/8, pp. 1564-1570. [45]
- Swanson, L. (2016), “The predictive processing paradigm has roots in Kant”, *Frontiers in systems neuroscience*, Vol. 10, p. 79. [9]
- Synofzik, M., G. Vosgerau and M. Voss (2013), “The experience of agency: an interplay between prediction and postdiction”, *Frontiers in psychology*, Vol. 4, p. 127. [47]
- Tolin, D. et al. (2003), “Intolerance of uncertainty in obsessive-compulsive disorder”, *Journal of anxiety disorders*, Vol. 17/2, pp. 233-242. [56]
- Turing, A. (1950), “COMPUTING MACHINERY AND INTELLIGENCE”, *Mind*. [31]
- Vohs, K. and J. Schooler (2008), “The value of believing in free will: Encouraging a belief in determinism increases cheating”, *Psychological Science*, Vol. 19/1, pp. 49-54. [44]
- Wagenaar, W. and S. Sagaria (1975), “Misperception of exponential growth”, *Perception & Psychophysics*, Vol. 18/6, pp. 416-422. [3]
- Wang, X., R. Ashfaq and A. Fu((n.d.)), “Fuzziness based sample categorization for classifier performance improvement”, *ournal of Intelligent & Fuzzy Systems*, Vol. 29/3, pp. 1185-1196. [38]
- Wegner, D. and T. Wheatley (1999), “Apparent mental causation: Sources of the experience of will”, *American Psychologist*, Vol. 54/7, p. 480. [81]
- Yu, A. and P. Dayan (2005), “Uncertainty, neuromodulation, and attention”, *Neuron*, Vol. 46/4, pp. 681-692. [19]

Zimmerman, B. (2000), “Self-efficacy: An essential motive to learn”, *Contemporary educational psychology*, Vol. 25/1, pp. 82-91. [69]