DOES NATURALISED EPISTEMOLOGY HAVE SOMETHING TO DO WITH COGNITIVE PSYCHOLOGY AND, PERHAPS, ARTIFICIAL INTELLIGENCE?

Ephraim Stephen Essien
Dept. of Philosophy, Practical Ethics and Systematic Theology
School of Humanities
University of SOUTH AFRICA

estephenessien@gmail.com

Abstract. This study aims to demonstrate the link between Artificial Intelligence and Naturalised Epistemology. Artificial neural networks simulate natural neural networks, which is a bio-psychological process. The central thesis of this research is that Artificial Intelligence and Naturalised Epistemology share a nexus in Cognitive Psychology.

Keywords: artificial intelligence, naturalised epistemology, psychology, Quine, cognitive science

INTRODUCTION

There has been little or no research on the nexus between artificial intelligence and naturalised epistemology. There is research on artificial intelligence, epistemology, and philosophy generally construed. However, less research focused primarily on psychology as the basic link between artificial intelligence and naturalised epistemology. This study aims to bridge the gap in epistemology and the philosophy of artificial intelligence.

Epistemology, the theory of knowledge, is one of the central areas of philosophy. The questions addressed by epistemology have historically included what knowledge is, how we can or should achieve it, and how much, if anything, we can know. Naturalism is the view that the world contains only natural

phenomena and that the appropriate methods for acquiring knowledge of the world are those of sciences.

NATURALISED EPISTEMOLOGY

The term "naturalised epistemology" was introduced by W.V. Quine in his 1969 essay *Epistemology Naturalised*, in which he argues that epistemology should be regarded as continuous with, or even part of, natural science. Epistemological naturalists often contrast their approach with that taken by René Descartes. Descartes held that knowledge has a foundational structure. At the foundation are beliefs which we "clearly and distinctly perceive", and about which we are therefore completely certain. For Descartes, these include beliefs about the contents and operations of our own minds. Other beliefs must be inferred from these foundational beliefs in order for us to be justified in holding them. Until we can show, based on foundational beliefs, that there is a world outside our minds, and that proper scientific methods will reliably give us information about it, we can have no confidence in the results of sciences.

Advocates of a naturalised epistemology see the role of epistemology very differently. For them, philosophy does not come prior to science. The starting point of epistemology should not be our introspective awareness of our own conscious experience, but rather the conception of the larger world that we get from common sense and science. Most naturalists would also reject many other features of Descartes' epistemology, including the view that knowledge requires certainty, the view that all our knowledge must be inferred from foundational beliefs, and the view that it is possible to know substantive facts about the world a priori, that is, without needing experience to provide evidence of their truth (Brown and Cutis, 2010).

Of the three main epistemological issues, namely, the nature of knowledge, the means of acquiring it, and its extent, Quine's

naturalised epistemology focuses on the second, the issue of how knowledge is acquired. In a famous passage, Quine describes what he sees as the proper subject of naturalised epistemology:

> It studies a natural phenomenon, *viz.* a physical human subject. This human subject is accorded a certain experimentally controlled input—certain patterns of irradiation in assorted frequencies, for instance—and in the fullness of time the subject delivers as output a description of the three-dimensional external world and its history. The relation between the meagre input and the torrential output is a relation that we are prompted to study for somewhat the same reasons that always prompted epistemology; namely, in order to see how evidence relates to theory, and in what ways one's theory of nature transcends any available evidence. (Quine 1969: 82–83).

For Quine, then, naturalised epistemology is the empirical study of how human beings develop a theory of the natural world on the basis of their sensory inputs. Given this understanding of epistemology, it is clear why Quine thinks that "epistemology, or something like it, simply falls into place as a chapter of psychology". Cognitive psychology and/or cognitive science connect perfectly with what Quine describes here as the domain of epistemology in so far as the processes of thinking and knowing are concerned.

Like Descartes, Quine takes epistemology to be "concerned with the foundations of science" (Quine 1969, 69). Addressing the logical empiricist project of rational reconstruction, he says that the Cartesian quest for certainty is the remote motivation of epistemology, both on its conceptual and doctrinal side (Quine 1969, 74).

About the epistemological project so understood, Quine's chief observation is not some news: the Cartesian quest is "a lost cause" (*Ibidem*). Whether in the form Descartes himself practised, or in any subsequent form up to and including the logical empiricists, its work on both the conceptual and the doctrinal side is bound to fail: no strict translation of the notion of "body" in sensory terms is possible, and "the inferential steps between

sensory evidence and scientific doctrine must fall short of certainty" (Quine 1969, 74–75).

What *is* new in *Epistemology Naturalised* is what Quine recommends in the face of this result: "Why all this creative reconstruction, all this make-believe? The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world. Why not just see how this construction really proceeds? Why not settle for psychology?" (*Ibidem*, 75). "If all we hope for is a reconstruction that links science to experience in explicit ways short of translation, then it would seem more sensible to settle for psychology. Better to discover how science is in fact developed and learned than to fabricate a fictitious structure to a similar effect" (*Ibidem,*78). Epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science. But a conspicuous difference between old epistemology and the epistemological enterprise in this new psychological setting is that we can now make free use of empirical psychology. (*Ibidem,* 82–83). Ryesiew observes that even if it would offend strong anti-psychologists, it is not the suggestion that epistemologists make "free use" of empirical psychology that is so radical; it is the suggestion that psychology can and should *replace* epistemology (Ryesiew 2020).

Quine goes on to argue that scepticism is an offshoot of science. The basis for scepticism is the awareness of illusion, the discovery that we must not always believe our eyes. But in what sense are they illusions? In the sense that they seem to be material objects which they in fact are not. Illusions are illusions only relative to prior acceptance of genuine bodies with which to contrast them. The positing of bodies is already rudimentary physical science; and it is only after that stage that the sceptic's invidious distinctions make sense. Rudimentary physical science, that is, common sense about bodies, is thus needed as a springboard for scepticism (Quine 1975, 67). But if scepticism itself is born of science, we can appeal to science in answering its

doubts. For instance, we can look to natural selection, and find "some encouragement in Darwin" in quelling doubts about the reliability of induction: creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die out before reproducing more of their kind (Quine 1969b, 126). Quine appears to be recommending replacement naturalism and, consequently, the elimination of terms of epistemic appraisal in favour of descriptions of psychological processes.

As we can see from the foregoing, Quine developed epistemological naturalism and gave epistemology the task of identifying a substantial and constructive role for the sciences in epistemological theorising. One popular way to think about the continuity between sciences and epistemology is in terms of how normative questions about how we ought to form our beliefs cannot be answered independently of descriptive questions about how we form beliefs. Understood thus, the challenge for the naturalised epistemologist is to spell out in more detail the respective contribution by (traditional) epistemology and the sciences, and in particular the extent to which the latter is to replace or simply complement the former.

Naturalized epistemology encompasses a range of philosophical perspectives on knowledge that prioritize the application of scientific methods. By emphasizing empirical approaches to understanding how knowledge is acquired, it moves away from many traditional philosophical debates and instead focuses on the processes through which knowledge is formed.

For Quine, attempts to pursue the traditional project of finding the meanings and truths of science philosophically have failed on their own terms and failed to offer any advantage over the more direct methods of psychology. Since traditional philosophic analysis of knowledge fails, those wishing to study knowledge ought to employ natural scientific methods. Scientific study of knowledge differs from philosophic study by focusing on how humans acquire knowledge rather than speculative analysis of knowledge (Quine 2004, 292-300). In short, Quine

identifies psychology as an entailment of epistemology because of his emphasis on focusing on the natural mental processes of how we know. For him, there are psychological processes involved in the art of knowing, and psychology falls within the domain of the natural sciences.

Among the various forms of epistemological naturalism, three stand out: *replacement naturalism*, which argues for discarding traditional epistemology in favour of scientific methodologies; *cooperative naturalism*, which suggests that epistemology can be enriched by insights from cognitive science; and *substantive naturalism*, which equates knowledge claims with natural facts. Quine appears to align most closely with *replacement naturalism*.

However, much of epistemology as traditionally conceived seems to be left out of Quine's picture, and contemporary epistemological naturalists differ in how they think these topics should be addressed. First, one of the main concerns of epistemology has been to understand what knowledge is, in the sense of identifying necessary and sufficient conditions for knowing something. This seems to require an analysis of the concept of knowledge rather than an empirical investigation of the natural world. Some naturalists believe that epistemology should simply abandon conceptual analysis; some accept that conceptual analysis is a necessary and non-scientific part of epistemology, and conclude that only parts of epistemology can be naturalised; and some hold that conceptual analysis itself should become an experimental discipline.

A second aspect of traditional epistemology that Quine seems to neglect concerns the second epistemological question, that of how we do or should acquire knowledge. Many critics of Quine have noted that by focusing exclusively on the descriptive issue of "How we base a rich theory of the world on limited evidence?", he appears to neglect normative issues about how we ought to modify our beliefs in light of new evidence. Some moderate epistemological naturalists concede that such issues

cannot be regarded as part of science, while others have suggested that even normative issues can be naturalised.

A final issue that Quine pays little attention to relates to the third epistemological issue, that of how much knowledge, if any, we can have. Quine recommends treating the issue of the extent of our knowledge as internal to science. However, a main focus of traditional epistemology has been to address whether it is possible to convincingly refute radical scepticism, the idea that all or most of our beliefs could be seriously mistaken. To address this question by appealing to the results of science seems to beg the question. Can there be a naturalistic response to radical scepticism? Most contemporary naturalists would concede that they cannot refute scepticism, but would also hold that the only sceptical doubts worth taking seriously are those that arise from within science itself.

Artificial Intelligence

Artificial intelligence (AI) is commonly used to describe machines that replicate cognitive functions associated with the human mind, such as learning and problem-solving (Russell & Norvig, 2003). Essentially, AI involves the imitation or simulation of human intelligence. Given its connection to knowledge, AI naturally intersects with epistemology. Research in AI focuses on various challenges, including reasoning, knowledge representation, planning, learning, natural language processing, perception, and the ability to interact with and manipulate objects. The field of AI draws from multiple disciplines, including computer science, information engineering, mathematics, psychology, linguistics, and philosophy. Among these, reasoning, cognition, knowledge representation, learning, and perception are particularly relevant to epistemology.

These problems are natural epistemic processes in the human mind, which naturalised epistemology emphasises as the domain of cognitive psychology. In artificial intelligence, these processes

are digitalised and technologically enabled. There, they are carried out by machines. Have machines taken over the cognitive role of the human mind? Have robots replaced human beings? The machines haven't taken over. Not yet, at least. However, they are seeping their way into our lives, affecting how we live, work and entertain ourselves. From voice-activated personal assistants like Siri and Alexa to foundational technologies such as behavioural algorithms, predictive search, and autonomous self-driving vehicles, artificial intelligence is embedded in numerous applications today.

AI IN ACTION: CURRENT APPLICATIONS AND INNOVATIONS

Today's AI systems are advanced machine learning programs equipped with complex behavioural algorithms that adapt to human preferences. While highly efficient, these systems are not becoming intelligent in an existential sense; rather, they enhance their capabilities and effectiveness through vast datasets. Below are some of the most widely used examples of artificial intelligence today.

1. *Siri*

Everyone is familiar with Apple's personal assistant Siri. She is a friendly voice-activated computer that we interact with daily. She helps us find information, gives us directions, adds events to our calendars, helps us send messages, and so on. Siri is a pseudo-intelligent digital personal assistant. She uses machine-learning technology to get smarter and better able to predict and understand our natural-language questions and requests.

2. *Alexa*

Alexa aims to become the smart home's hub. When Amazon first introduced it, it took much of the world by storm. However, its usefulness and uncanny ability to decipher speech from anywhere in the room has made it a revolutionary product that can help us scour the web for information, shop, schedule appointments, set alarms and a million other things, but also help power our smart homes and those that might have limited mobility.

3. *Tesla*

Tesla is quite possibly one of the best cars ever made, not only for the fact that it has received so many accolades but because of its predictive capabilities, self-driving features and sheer technological "coolness". Anyone who is into technology and cars likes to own a Tesla, and these vehicles are only getting smarter and smarter thanks to their over-the-air updates.

4. *Cogito*

Originally co-founded by CEO Joshua Feast and Dr Sandy Pentland, Cogito is quite possibly one of the most powerful examples of behavioural adaptation to improve the emotional intelligence of customer support representatives that exists on the market today. The company is a fusion of machine learning and behavioural science to improve customer interaction for phone professionals. This applies to millions of voice calls that are occurring on a daily basis.

5. *Boxever*

Boxever, co-founded by CEO Dave O'Flanagan, is a company that leans heavily on machine learning to improve the customer's experience in the travel industry and deliver "micro-moments", or experiences that delight the customers along the way. It's through machine learning and the usage of AI that the company has dominated the field, helping its customers find new ways to engage their clients in their travel journeys.

6. *John Paul*

This highly esteemed luxury travel concierge company, helmed by David Amsellem, is another powerful example of AI in the predictive algorithms for existing-client interactions, able to understand their desires and needs on an acute level. The company provides concierge services for millions of customers through the world's largest companies such as VISA, Orange, or Air France and was recently acquired by Accor Hotels.

7. *Amazon.com*

Amazon's transactional AI has been in use for quite some time, enabling the company to generate massive revenue online. With its algorithms continuously improving each year, Amazon has become remarkably adept at predicting consumer interests based on online behaviour. While the company envisions a future where products are shipped before customers even realize they need them, this level of predictive commerce has yet to be fully realized—but it is certainly on the horizon.

8. *Netflix*

Netflix provides highly accurate predictive technology based on customer's reactions to films. It analyses billions of records to suggest films that you might like based on your previous reactions and choices of films. This tech is getting smarter and smarter by the year as the dataset grows. However, one drawback of this technology is that lesser-known films often go unnoticed, while big-name movies dominate and gain even more visibility on the platform.

9. *Pandora*

Pandora AI is one of the most revolutionary techs that exists out there today. They call it their musical DNA. Based on 400 musical characteristics, each song is first manually analysed by a team of professional musicians based on this criterion, and the system has an incredible track record for recommending songs that would otherwise go unnoticed but that people inherently love.

10. *Nest*

The Nest Learning Thermostat, acquired by Google in January 2014, uses behavioural algorithms to learn from your heating and cooling patterns, adjusting the temperature in your home or office based on your personal preferences. It can now be voice-controlled by Alexa and has expanded to include a range of other products, such as Nest cameras.

Jeremy Sutton (2021) posits that the term AI is typically used to describe both the "technology designed to perform activities that normally require human intelligence" and the multidisciplinary field of science concerned with understanding and developing that technology (Luxton, 2014). These technological feats, however, have not yet enabled any AI to pass *The Turing Test*. We can only wait to see if that will ever happen in the future.

Psychology and Artificial Intelligence

How is AI related to psychology? We may think that psychology deals with the most human side and that we do not need technological tools to get better day by day. The truth is that the field of psychology keeps growing every day and AI is part of it because it is the attempt to make a machine reach the level of "intelligence" of a human being (Lozano, 2020). In its early years, AI was created to emulate the functioning of the mind and the brain in computer software. More specifically, in the field of Deep Learning, Artificial Neural Networks try to imitate the neural networks of a living organism. Its goal is to get a combination of parameters which better fit a specific problem.

We emphasise that psychology and artificial intelligence (AI) are closely related disciplines. For instance, in the psychology degree, some basic AI concepts are always covered, normally in the courses about the Psychology of Thought. Also, in computer science when studying courses related to artificial intelligence, there are usually references to theories of human cognition. Neuroscience has an eminently integrative approach and, in general, research in cognitive science usually brings together experts from different areas, such as neuroscience, psychology, artificial intelligence, philosophy of mind, robotics, *etc.*

Certainly, both AI and psychology have a common axis: understanding the processes that give rise to intelligent behaviour. In the case of psychology, the study focuses on human

beings, and we talk about mental processes. In the case of AI, the study focuses on machines, and we talk about information processing.

In general, psychology focuses on three main axes of the person: cognition, emotion and behaviour. From the point of view of "weak AI", machines do not think, although they process information, they do not feel either, although they can identify emotions, and they have behaviour which is determined by the output of their algorithms.

We could say that psychology deals with biological organisms (usually the human species), while AI deals with artificial cognitive systems. But deep down, both types of minds face the same essential challenge: adapting to the environment and solving problems efficiently even in situations of uncertainty, ambiguity, and noise. This capacity is what we commonly associate with intelligent beings.

Based on this analogy between intelligent biological and artificial systems, research in both disciplines is interrelated:

- Knowledge of the human mind can contribute to the design of more intelligent artificial systems.
- The use of computational models can contribute to research on the functioning of the human mind.

In other words, the use of *bio-inspiration* in AI implies that the design of some artificial systems is based on the dynamics observed in the human cognitive system or other species. At the same time, hypotheses about how the human mind works can be tested, at least partially, using computational models based on Artificial Intelligence.

This parallelism between "artificial minds" and "natural minds" is not the only link between the two disciplines. For example, from the point of view of the possible practical application of Artificial Intelligence in the area of psychology, there are multiple possibilities, such as:

Intelligent Systems based on Machine Vision for:

- Emotions; facial expression training.
- Detection of situations of health risk.

Intelligent Systems based on Natural Language Understanding for:

- Early detection of psychological problems.
- Automatic detection of personality traits.

Intelligent Systems based on Voice Signal Processing for:

- Identification of mood and level of physiological activation.
- Detection of symptoms of anxiety and depression.

Intelligent Systems based on Pattern Recognition on sensor data for:

- Automatic detection of a person's behaviour.
- Automatic detection of falls, accidents or assaults.

This technology is applied in different areas of psychology like talent management, education, psychotherapy, neuropsychology or prevention in mental health.

Conclusion

There is a basic connection between naturalised epistemology, artificial intelligence and cognitive psychology. Cognitive psychology attempts to understand cognition's complexity through research, testing, and building models of how the human mind handles and processes complex information during attention, memory, and perception (Zivony, 2019).

AI and cognitive psychology share similar aims—to understand the nature of intelligent behaviour—with the former attempting to build such processes using advanced technology. And while computational modelling and AI have subtle differences, they are both valuable approaches for understanding the nature of intelligent thinking and providing insights into the growing field of cognitive psychology. Computational modelling

involves "programming computers to model or mimic aspects of human cognitive functioning" (Eysenck & Keane, 2015).

Artificial intelligence's underlying processes, on the other hand, typically bear no resemblance to the mechanisms used by the human brain. Rather than attempt to develop computational models that help us understand human intelligence, the AI designer's goal is to produce an outcome that appears intelligent. Such processes do not need to be functionally similar to those of a human.

However, there is one particular model that appears to bridge the gap between the two approaches. *Connectionism* was originally inspired by the network of neurons that exists within the brain. "Connectionist models typically consist of interconnected networks of simple units exhibiting learning" and model cognition with no explicit rules (Eysenck & Keane, 2015). While the brain may be described as a highly complex neural network, and connectionist models have successfully modelled specific human-like processes (such as face recognition), the jury is still out regarding whether such models explain human cognition. However, deep neural networks, inspired by cognitive psychology theories and methods, have had some success in explaining how children learn labels for objects and offer a great example of the benefits of combining knowledge and expertise from multiple disciplines (Ritter, Barrett, Santoro, & Botvinick, 2017).

REFERENCES

Cutis, B & Luper, S. 2010. "Naturalised Epistemology." Routledge Encyclopaedia of Philosophy. Retrieved 15 Feb. 2022, from: https://www.rep.routledge.com/articles/thematic/naturalized-epistemology/v-2

Eysenck, M. W., & Keane, M. T. 2015. *Cognitive psychology: A student's handbook*. Psychology Press.

Lozano, C.R. 2020. "The World of Artificial Intelligence and Psychology."

Retrieved 15 Feb. 202, from https://openexpoeurope.com/en/the-world-of-artificial-intelligence-and-psychology/

Luxton, D. D. 2014. "Artificial intelligence in psychological practice: Current and future applications and implications." *Professional Psychology: Research and Practice*, 45 (5), 332–339.

Quine, W.V.O. 2004. *"Epistemology Naturalised."* In E. Sosa, and J. Kim (ed.). *Epistemology: An Anthology. Malden,* MA: Blackwell Publishing, 292-300.

Quine, W.V.O. 1969. "Epistemology Naturalised." *Ontological Relativity and Other Essays*, 69-90. New York: Columbia University Press.

Quine, W.V.O. 1969b. "Natural Kinds." *Ontological Relativity and Other Essays*, 114-138. New York: Columbia University Press.

Quine, W.V.O. 1975. "The Nature of Natural Knowledge." In Samuel Guttenplan (ed.), *Mind and Language*, 67-81. Oxford: Oxford University Press.

Ritter, S., Barrett, D. G. T., Santoro, A., & Botvinick, M. M. 2017. "Cognitive psychology for deep neural networks: A shape bias case study." *Proceedings of the 34th International Conference on Machine Learning, PMLR* 70, 2940–2949.

Russell, S.J., P. Norvig. 2003. *Artificial Intelligence: A Modern Approach*, (2nd ed.), Upper Saddle River, New Jersey: Prentice Hall.

Rysiew, P. 2020. "Naturalism in Epistemology." In *Stanford Encyclopaedia of Philosophy*. Retrieved 15 Feb. 2022, from https://plato.stanford.edu/entries/epistemology-naturalized/#QuiEpiNat

Sutton, J. 2021. "Artificial Intelligence in Psychology: 5 Revolutionary Examples." Retrieved 15 Feb. 2022 from https://positivepsychology.com/artificial-intelligence-in-psychology/

Zivony, A. 2019. "What is cognitive psychology?" The British Academy. Retrieved 15 Feb. 2022 from https://www.thebritishacademy.ac.uk/blog/what-is-cognitive-psychology/