

Artificial Intelligence Seen Through the Lens of Bateson's Ecology of Mind

Dai Griffiths*

Research Institute for Innovation & Technology in Education (UNIR iTED)
Universidad Internacional de La Rioja (UNIR). 26006, Logroño, La Rioja (Spain)

Received 1 April 2021 | Accepted 20 June 2021 | Published 3 August 2021



ABSTRACT

Gregory Bateson developed a number of ideas which are relevant to artificial intelligence, and in particular to the ascription of qualities such as mind, consciousness, spirituality and the sacred. Relevant sections of Bateson's key works are discussed, and his intellectual framework for an ecology of mind is summarized, and in particular his concepts of mind, learning, and the sacred. These are then applied to discuss whether artificial intelligence applications can be considered to possess 'mind'. It is concluded that symbolic artificial intelligence falls short of Bateson's criteria for mind, as do neural networks, although approach more closely. Nor are computers based on the rules of formal logic able to engage with the sacred, which is paradoxical in nature. However, artificial intelligence applications can form part of an ecology of mind and can be involved in the experience of the sacred. Bateson's writing remains a fertile source of ideas relevant to an understanding of the nature and capabilities of artificial intelligence.

KEYWORDS

Bateson, Mind, Artificial Neural Networks, Symbolic Artificial Intelligence, Sacred.

DOI: 10.9781/ijimai.2021.08.004

I. INTRODUCTION

ACCORDING to Marvin Minsky, "we are on the threshold of an era that will be strongly influenced, and quite possibly dominated, by intelligent problem-solving machines" [1]. These words were written over sixty years ago, but only today are we stepping over the threshold. Since Minsky's paper Artificial Intelligence (AI) has demonstrated its huge potential, but the nature and degree of its influence on human beings, and the ways in which it may dominate our society remain uncertain. Indeed, it remains the case that, as Minsky commented, "there is no generally accepted theory of intelligence" [1], and consequently there is no consensus on the nature of artificial intelligence. Given the rapidly increasing power of AI, a clarification of these open questions is ever more urgent.

This special issue seeks to contribute to an understanding of these matters by discussing how the capabilities of AI could be interwoven with the human phenomena of spirituality and analogue thinking. This raises further questions about human biology and thought, the nature of machine thinking, the digital and the analogue, and of the nature of spirituality. To make progress in understanding the relationship between these complex topics, it is necessary to view them not as separate phenomena, each determined according to their own rules, but rather to establish an overarching theory within which they can all be conceptualized. Such theories are few and far between. The work of Gregory Bateson (1904-1980) encompasses all these aspects and

provides a theoretical position from which an interwoven discussion of AI, spirituality and analogue thinking becomes possible. This paper builds on Bateson's ideas to reflect on whether an AI can be said to have a mind, and on the possible relationship of such a mind with the spiritual, or as Bateson would have termed it, the sacred. It is hoped that a reexamination of Bateson's ideas, which are now unfamiliar to many, may offer a valuable perspective from which to view the complex and deep issues raised by AI. The intention is to be of interest to both readers familiar with AI who know nothing of Bateson, and readers familiar with Bateson who know nothing of AI. Consequently, no prior knowledge of either area is assumed.

The radical and interdisciplinary nature of Bateson's ideas has led them to remain outside the mainstream. He is perhaps best known for his assertion "In fact, what we mean by information, the elementary unit of information, is a difference which makes a difference" [2] p.459. This remains widely cited, including by Floridi, [3] p.85, perhaps the most prominent current theorist of information (who also points out that Bateson's definition was anticipated by Donald MacCrimmon MacKay [3] p.85). However, this aphorism is just one of a set of interrelated ideas which together defined Bateson's concept of an ecology of mind. Bateson was well aware of computers and their potential power but could not have foreseen the developments which have taken place since his death.

Bateson's thought was extraordinarily wide ranging, unusual, and spread across publications on many themes during a long career. The reader should be aware that this paper offers an individual reading of Bateson, based on a particular route through his writing undertaken for a specific purpose. It does not aspire to giving a complete overview of his work, and most obviously leaves to one side his contributions

* Corresponding author.

E-mail address: david.griffiths@unir.net

to anthropology and psychology. If the interest of the reader is piqued by this discussion, there is no substitute for engaging with the books in which his ideas were set out [2], [4], [5], [6]. A number of valuable studies of Bateson's work are also available, including David Lipset's biographical study [7] and the synthesis and interpretation provided by Peter Harries-Jones [8], [9] and Noel Charlton [10].

II. METHODOLOGY

This paper considers the ways in which Bateson's concepts of an ecology of mind, and of the sacred, can shed light on the nature of AI. Within this context, the tasks undertaken by this paper are to:

- outline Bateson's position on dualism and information, which underlies his understanding of mind
- summarize and clarify Bateson's concepts of mind, learning and the sacred
- apply these concepts as criteria for the presence of mind and the sacred in both symbolic AI (which was well established when Bateson was writing), and in neural networks (which are prominent today).

Bateson develops and discusses his concepts of mind and the sacred in four books: *Steps to an Ecology of Mind* [2], *Mind and Nature, a Necessary Unity* [4], *Angels Fear, Towards an Epistemology of the Sacred* [5] (written in collaboration with his daughter Mary Catherine Bateson and published posthumously), and *A Sacred Unity: Further Steps to an Ecology of Mind* [6] (a posthumous selection of unpublished writings). The relevant ideas are not presented in unitary manner, but rather are spread throughout these books. A first task was therefore to read the texts, and to take notes of the key formulations of Bateson's theoretical position, and of those instances where he discusses AI. This was followed by the identification of representative claims made for AI, as set out in the literature. This reading and note taking may be characterized as a 'snowball' approach, as conducted, for example, by Hepplestone et al. [11]. Papers and books were selected based on their title and abstract (where available), and additional sources were identified from references within the articles. The direction of exploration was from the present to the past, gradually uncovering the antecedents of the current discourse. The two bodies of notes were raw material for an analysis of the relationship between the two sets of ideas, evolved in the successive drafts of the paper over a period of three months.

This paper explores the degree to which the application of Bateson's ideas may facilitate an understanding AI. It is not a critique of Bateson's thinking, nor an attempt to update his concepts in the light of recent developments, valuable though such contributions would be. Consequently, Bateson's terminology is given precedence. Thus, the paper discusses 'mind' and 'mental processes' rather than 'intelligence', and 'the sacred' rather than 'the spiritual'. This is a pragmatic response to the task in hand and avoids falling into an examination of competing definitions which would take the place of a substantive discussion.

III. BATESON'S VIEW OF MIND

A. Differences and Distinctions

Bateson was adamantly opposed to explanations which depended on the identification of a dualism. In particular he rejected Descartes' distinction between "mind or thinking substance" and "extended substance or body" (i.e. the characteristics of physical bodies). [12] p.208-209, which he saw as a strategy for avoiding the problem. In his view dualism is

a device for removing one half of the problem for explanation from that other half which could be more easily explained. Once separated, mental phenomena could be ignored. This act of subtraction, of course, left the half that could be explained as excessively materialistic, while the other half became totally supernatural. ... The materialist superstition is the belief (not usually stated) that *quantity* (a purely material notion) *can determine pattern*. On the other side, the antimaterialist claims *the power of the mind over matter*. [5] p.59

Leaving on one side the responsibility of Descartes for the prevalent dualist view of mind and matter (see [13] for a sympathetic view of Descartes thinking) it remains deeply entrenched in our current thinking about mind. Indeed, it is entangled with ideas of the spirit or soul, and of what it means to be alive. These ideas appear to be of such value to people that they are willing to live with, or even embrace, the contradictions which they generate. Bateson set out to overcome the "formidable barrier" [5] p.12 of Cartesian dualism, and to formulate a system which could accommodate his conviction that "mind and nature form a necessary unity, in which there is no mind separate from body, and no god separate from his creation" [5] p.12. Bateson's proposal of the ecology of mind takes as its starting point this rejection of dualism and is rigorous in following through its implications.

In identifying what Bateson meant by *mind* we must start with his conception of the animate and the inanimate. It should be noted that this is not a dualist explanation, but rather a distinction made in identifying the phenomena to be explained.

...we will use Jung's term *Pleroma* as a name for that unliving world described by physics which in itself contains and makes no distinctions, though we must, of course, make distinctions in our description of it.

In contrast, we will use *Creatura* for that world of explanation in which the very phenomena to be described are among themselves governed and determined by difference, distinction, and information. [5] p. 18.

Within *Pleroma*, interactions take place through the transfer of energy, for example in the friction of a meteorite entering the atmosphere, and its subsequent impact crater. This energy is conserved, so the size of the impact crater will be proportional to the velocity of the meteorite and the resistance of the site. In *Creatura* the situation is entirely different, as there is no relationship between the energy involved in distinction making and the scale of the consequences for the organism. The thunder of a passing truck may lead to less activity than the quiet rattle of a rattlesnake. Indeed, Bateson points out that there can be an inverse energy relationship, as, for example, when an amoeba moves more because it has been deprived of nutrients and is searching for food [2] p.490.

There is a near-infinite number of discontinuities in the environment, and all of these could be inputs into *Creatura*. This is problematic, as no organism can regulate its relationship with its environment by adsorbing an infinite variety of inputs. In the cybernetic literature this principle has the status of a law, Ashby's law of requisite variety, which states that a regulator's capacity cannot exceed its capacity as a channel for variety [14]. The variety in the *Pleroma* to which the organism is exposed is attenuated by the act of making distinctions. This attenuation is achieved by selecting from the infinite range of differences in the environment those which are relevant to the ongoing life of the organism. The selection is carried out in part by limits of the organism's sense organs, and in part through the operation of its nervous system.

B. The Nature of Information

The energy involved in distinction making is often minuscule, initiated by the tiny impact of photons on the retina, or sound waves on the ear drum. The energy required for the difference to have any consequences is provided by the organism itself, both in activating the neural pathways, and in subsequent muscular activity, is obtained from within, from its metabolism.

Bateson describes this decoupling of energetic cause and effect in terms of transformation or coding within a circuit. He often illustrated this with the example of a blind person with a stick. Interaction with the environment creates transforms that are transmitted up the stick as vibrations, and then further transformed into neural activity. Bateson warns that “What is transmitted on a neuron is not an impulse, it is news of a difference” [2] p.490. In other words, the news of difference does not carry with it its own interpretation, which is dependent on the system through which it is traveling. It is this “news of difference” which constitutes information. Although Bateson did not discuss the ontological implications of his position, it implies the existence of a ‘real world’, but one which can only be apprehended at one remove, and which is constructed by the subject. “The mind contains no things, no pigs, no people, no midwife toads, or what have you, only ideas (i.e., news of difference), information about “things” in quotes, always in quotes.” [4] p.132.

C. The Nature of Mental Processes

As indicated above, Bateson was attempting the very challenging task of defining mind in terms which avoided proposing mind and matter as different substances [5] p.16. Within this context, and given the concepts outlined above, what constitutes a mental process in Bateson’s thinking? What is a mental process composed of? Where is it located? How can it be identified? Fortunately, Bateson was very explicit about the criteria for the existence of a mind¹:

1. A mind is an aggregate of interacting parts or components
2. The interaction between parts of mind is triggered by difference, and difference is a nonsubstantial phenomenon not located in space or time, difference is related to negentropy and entropy rather than to energy.
3. Mental process requires collateral energy
4. Mental process requires circular (or more complex) chains of determination
5. In mental process, the effects of difference are to be regarded as transforms (i.e. coded versions) of events which preceded them. The rules of such transformation must be comparatively stable (i.e., more stable than the content) but are themselves subject to transformation
6. The description and classification of these processes of transformation disclose a hierarchy of logical types immanent in the phenomena. [4] p.92 (italics in the original)

Circular causation is required to sustain a mental process, as without it there would be only an isolated event. It should be noted that “a change in any part of the circle can be regarded as cause for change at a later time in any variable anywhere in the circle” [5] p.60. Bateson gives the simple example of a thermostat, in which a rise in ambient temperature can be seen as causing a change in the switch of the thermostat, or the thermostat can be seen as controlling the temperature of the room. Bateson (in common with Hofstadter [15]) ascribes consciousness to recursive circular causation, defining it as “A reflexive aspect of mental process that occurs in some but not all minds, in which the knower is aware of some fraction of his knowledge or the thinker of some fraction of his thought” [5] p.207. It should be noted, however, that Bateson’s criteria do not include consciousness, nor do they specify that mind should be contained within a single organism [5] p.210.

In accordance with the six criteria for mental processes above, Bateson was clear that mental processes are digital in nature. This is because mental processes require coded transforms of difference. These in turn require distinction making which turns any analogue value into a digital one through a distinction between the two sides of a threshold. He observed that in animals “the central nervous system and DNA are in large degree (perhaps totally) digital, but the remainder of the physiology is analogic” [4] p.180.

¹ Bateson offered a set of four criteria in an earlier paper of 1969 [2] p.490, which evolved into the definitive set of six discussed here, published in 1979 [4] p. 92 and repeated in a slightly simplified form in 1987 [5] p.18-19.

D. The Ecology of Mind

Bateson’s criteria for mental processes are straightforward, and easy to accept, at least for those sympathetic to his non-dualist starting position. Nevertheless, the criteria have implications which are not immediately obvious, and are, indeed, startling. Bateson argues that

...any ongoing ensemble of events and objects which has the appropriate complexity of causal circuits and the appropriate energy relations will surely show mental characteristics. It will *compare*, that is, be responsive to difference (in addition to being affected by the ordinary physical “causes” such as impact or force). It will “process information. [2] p.315 (italics in the original).

Accordingly, Bateson includes within the category of *mental process* “a number of phenomena which most people do not think of as processes of thought” [5] p.16, including embryology, evolution, and “all those lesser exchanges of information and injunction that occur inside organisms and between organisms, and that, in the aggregate, we call *life*.” [5] p.17. This, he implies, is the logical consequence of rejecting a dualist view of mind and matter. To understand these mental processes, he proposed the concept of an *ecology of mind*, which is ecological in the sense that it concerns the interrelations and dependencies between mental systems of all sorts and their environments. In his view, in explaining the behavior of a human being or other organism, “this “system” will usually *not* have the same limits as the “self” – as this term is commonly (and variously) understood.” He gives the example of felling a tree with an axe, in which each stroke is modified according to the shape of the cut face of the tree. He sets out the mental process as

(differences in tree)-(differences in retina)-(differences in brain)-(differences in muscles)-(differences in movement of axe)-(differences in tree), etc. What is transmitted around the circuit is transforms of differences. And as noted above, a difference which makes a difference is an idea or unit of information. [2] p.317.

E. A Tenuous Tradition Building on the Ecology of Mind

Bateson was a unique figure, but he was not entirely alone in his view that the mind was not contained in the brain, and there is a tenuous thread of related work leading to the present which should be briefly discussed here to give context for our discussion.

In their highly influential book *The Embodied Mind* (1991), Varela Thompson and Rosch write that by embodied they mean:

...first, cognition depends upon the kinds of experience that come from having a body with various sensorimotor capacities, and second, that these individual sensorimotor capacities are themselves embedded in a more encompassing biological, psychological, and cultural context. [16] p.173.

This view clearly has a great deal in common with Bateson’s ecology of mind. Both Varela and Thompson were aware of Bateson’s thinking [17] although they do not cite him in their book

Edwin Hutchins analyzed the processes of navigation by teams in ships, and came to the conclusion that “The central computations of the navigation tasks are accomplished by the propagation of information across representations and representational media.” [18] p.218. He recognized the influence of Bateson in his focus on mapping propagation of information beyond the limits of an individual, writing “I take the fundamentals of an architecture of cognition and a sense of a unit of analysis from Gregory Bateson”. [18] p.291.

In 1998 Andy Clark and David Chalmers wrote an influential paper called ‘The Extended Mind’ [19], which built on Varela’s ideas. In recent years there has been an upsurge of interest in ‘4e cognition’ which brings together Gibson, Varela, Hutchins and Chalmers “under one heading and conceives of them as coherently opposed to the internalist, brain-centered views of cognitivism” [20] p.4. A related area in which there has been an active discussion in recent years has been that of cognition in plants, or as it is perhaps more correctly

termed, plant gnosophysiology [21], which is explicitly linked with extended cognition [22].

The fact that this tenuous thread of thinking around extended cognition now seems to be gathering some degree of prominence makes it timely to reassess the relevance of Bateson's, particularly as his contribution has been so widely forgotten or ignored.

IV. BATESON AND SYMBOLIC ARTIFICIAL INTELLIGENCE

A. Symbolic Artificial Intelligence

The development of artificial intelligence technology is a convoluted story, with many interconnecting strands, which this paper will not attempt to describe. The reader who would like to explore this history is referred to 'Mind as Machine', Margaret Boden's two volume history of the field [23]. For present purposes, it is sufficient to distinguish between systems based on formal logic (which were available during Bateson's lifetime), and those making use of neural networks (which had been conceived of but were not developed in a practical form).

Bateson was central figure in the establishment of cybernetics in the 1940s and a prominent participant in the seminal Macy conferences [24]. As such he would have been well aware of the ongoing development of artificial intelligence, and knew a number of the leading figures, including John von Neuman, Warren McCulloch and Norbert Wiener [25]. Indeed, artificial intelligence was established as a distinct field in a split from cybernetics at the Dartmouth Summer Research Project in 1956, partly as a result of personality clashes with Norbert Wiener, who was perceived to be the public face of cybernetics [26] p.78.

In 1957 Herbert Simon and Allen Newell made the bold claim that machines making use of heuristic problem-solving methods would, within ten years, be the world's chess champion, prove an important new mathematical theorem, write music accepted by critics as having considerable aesthetic value, and that most theories in psychology would take the form of computer programs [27]. Looking back on their work in 1976 they noted that it was based on the hypothesis that "A physical symbol system has the necessary and sufficient means for general intelligent action." [28]. The approach taken by Simon and Newell became dominant and was famously characterized by Haugeland in 1985 as Good Old Fashioned Artificial Intelligence, or GOFAI, as it became known. He identified the essential claims of GOFAI theories to be:

1. our ability to deal with things intelligently is due to our capacity to think about them reasonably (including sub-conscious thinking); and
2. our capacity to think about things reasonably amounts to a faculty for internal "automatic" symbol manipulation. [29] p.82

In line with Haugeland's second claim, this tradition has been known as 'symbolic AI', which is the term we will use here. It has also been referred to as 'cognitivist', since "according to this view cognition occurs by taking in information provided by the environment, forming this into representations which can then be processed to provide logical responses by way of activity." [30] p.106. This tradition continues to be of significance in the present day, often in combination with newer machine learning methods as proposed, for example, by Gary Marcus in his 2020 paper *The Next Decade in AI* [31].

In 1972, when Bateson published *Steps to an Ecology of Mind*, symbolic AI researchers were still confident of the success of their paradigm. However, this was also the year that Hubert Dreyfus published a book arguing that a system that could use natural language and recognize complex patterns would need a body, and that such robots would need to be entirely different from present digital computers [32]. p.216. The furious response to Dreyfus' book suggested that his admittedly

sharp critique of symbolic AI had touched a nerve. By the time that *Mind and Nature* was published in 1979 and it had become clear that the trajectory of research established at Dartmouth had encountered major problems. Bateson did not participate in the heated discussion around AI, but the ambitions and shortcomings of AI research formed an important part of the backdrop to his thinking, which was highly relevant to the debate.

B. Symbolic AI and Bateson's Criteria for Mind

Bateson's inquiry into mind was an unusual one within the cybernetics and AI communities. One of the few who discussed computers in terms of mind was Simon, who is reported by McCorduck [33] p.151 as saying that through his work in artificial intelligence he had arrived at "a notion that a mind was something which took some program inputs and data and had some processes which operated on the data and produced output." In other words, an understanding of mind did not inform Simon's understanding of AI, rather it was the emerging capabilities of AI which were taken as a basis for understanding the nature of mind. Similarly, the book *The thinking computer: mind inside matter* by Bertram Raphael, published in 1976 [34], not only proclaims in its title the dualism that Bateson abjured, it also consider avoids any discussion of what might constitute a mind. Raphael simply claims that if a computer successfully models the processes followed by a human in carrying out a task, then "we can view the flow chart of its program as a plausible guide to the logic of the inner workings of the mind." [34] p.300.

Bateson's criteria for mind (see section III.C, above) offer a perspective from which to view the claims of symbolic AI. It is clear that 1, 2, & 3 are fulfilled, since computers running AI applications are (1) an aggregate of interacting components, (2) triggered by difference and (3) employ collateral energy. Criteria 4, 5, & 6 need more consideration.

Criterion 4 "requires circular (or more complex) chains of determination" [4] p.92. There are certainly plenty of loops in computer code, for example a sub-routine may be called frequently in the execution of a program, always returning to the starting point so that the program can proceed. Nevertheless, circular chains of determination contravene the formal logic applied in symbol manipulation and are treated as a bug or a malware attack by both conventional computers and by symbolic AI. In other words, the system is designed to prevent a program changing its own functioning or the computational environment in which it operates, in response to its own operation. In biology, on the other hand, "many regularities are part of – contribute to – their own determination." [5] p.161.

Criterion 5 is fulfilled, in that computer circuits do indeed involve coded transforms of events which preceded them. However, the rules which govern these transformations are not, in classical computer architectures, subject to transformation without the intervention of a programmer.

Criterion 6 is that "The description and classification of these processes of transformation disclose a hierarchy of logical types immanent in the phenomena." In this regard Simon's *The Sciences of the Artificial* demonstrates that the symbolic AI community was well aware of hierarchy, which Simon discusses in terms that Bateson might well have approved of:

...complexity frequently takes the form of hierarchy and ... hierarchic systems have some common properties that are independent of their specific content. [35] p. 87.

Still more Batesonian is Simon's argument that hierarchies can take the form of different levels of description, commenting that the genetic description of a single cell may therefore take a quite different form from the genetic description that assembles cells into a multi celled organism. [35] p.115.

In practical computing applications, however, these structures were limited to branching classification trees, determined by the programmer. The contrast with Bateson's approach can be seen clearly in the question posed by Raphael "If we wish to insert knowledge into a computer what kinds of concepts must we represent?" [34] p.48. This conception of knowledge as something independent of the knower, which can be 'injected', is far removed from Bateson's view that the meaning of information is dependent on its context.

A deeper problem implied by criterion 6 is that it refers not only to a hierarchy, but to a "hierarchy of logical types immanent in the phenomena". Bateson's understanding of logical types was based on Whitehead and Russell's *Principia Mathematica*. He summarized the principles as being:

no class can, in formal logical or mathematical discourse, be a member of itself; that a class of classes cannot be one of the classes which are its members; that a name is not the thing named... a class cannot be one of those items that are correctly classified as its non members... if these simple rules of formal discourse are contravened, paradox will be generated... [2] p.280.

In Bateson's view, the information flow in an organism is logically typed, but not neatly separated in the way that a programmer might define a set of classes. Rather, the organism generates an enormous and tangled network of messages [4] p.109, within which complex relationships of logical types emerge, although the observer may find these hard to identify. I understand this to be an implication of Bateson's use of the word 'immanent' in his criterion.

C. Symbolic AI in an Ecology of Mind

We have seen in the previous section that computers running symbolic AI applications are lacking two of Bateson's criteria for mind: circular chains of causation, and complex interactions between logical types. However, this does not mean that they do not constitute parts of in an ecology of mind.

Bateson argues that

...in no system which shows mental characteristics can any part have unilateral control over the whole. In other words, *the mental characteristics of the system are immanent, not in some part, but in the system as a whole.* [2] p.316.

The question then arises of what constitutes 'the system' in a symbolic AI application. As we have seen in section III.D, viewed within an ecology of mind, a system does not usually have the same limits as the self, but rather is constituted by the limits of the flow of information within an ecology of mind. Consequently, the questions "Can a computer think?" and "Is the mind in the brain?" should be answered in the negative [2] p.316 (with the possible exception of processes which monitor the internal states of the computer or brain). More precisely:

... it would be incorrect to say that the main business of the computer – the transformation of input differences into output differences – is "a mental process". The computer is only an arc of a larger circuit which always includes a man and an environment from which information is received and upon which efferent messages from the computer have effect. This total system, or ensemble, may legitimately be said to show mental characteristics. It operates by trial and error and has creative character.

Similarly, we may say that "mind" is immanent in those circuits of the brain which are complete within the brain. Or that mind is immanent in circuits which are complete within the system, brain *plus* body. Or, finally, that mind is immanent in the larger system – man *plus* environment. [2] p.317.

Discussing the blind person with a stick, Bateson asks where that person's self begins. For Bateson the answer was *at the tip of the stick*, because the any other location would "draw a delimiting line across this pathway is to cut off a part of the systemic circuit which

determines the blind man's locomotion." [2] p.318. On this basis, we make the same mistake if we draw a delimiting line between a human and artificial intelligence.

We can now provide a response to the question "Does a symbolic AI application have a mind?". From the perspective of the ecology of mind the answer is "No, but that is not the right question to ask". It is more valuable to ask "What is the structure of the ecology of mind generated when a person interacts with an AI?" The answer will involve mapping information flows (i.e. coded transforms of distinctions), while paying attention to their logical types and to recursive causation.

V. BATESON'S IDEAS AND AI IN THE 21ST CENTURY

A. Neural Networks and Deep Learning

In 1985 Haugeland was able to write that "AI more or less ignores learning" [29] p.11. This is by no means the case today, when technologies related to 'machine learning' using neural networks, often adopting a connectivist approach, have become established as the focus for most AI research, while symbolic AI retains importance as an established method.

The history of neural networks is usually traced back to a paper of 1943 by McCulloch and Pitts which provided a Logical Calculus for Nervous Activity. Both the authors knew Bateson and coincided with him at the seminal Macy conferences [24]. We may therefore assume that Bateson would have known of their approach to AI and would not have ignored it in developing his thinking on the ecology of mind. However, given the low level of achievements of neural networks at the time, there was no need to pay particular attention to this line of work in his writing. Bateson's criteria for mind make no reference to the structure of information flows, nor the material through which they pass. Consequently, the discussion of mind in relation to symbolic AI can, in principle, also be applied to neural networks.

Since Bateson was writing, however, neural networks have become vastly more powerful, and have demonstrated astonishing capabilities. Machine learning methods examine the relationship between inputs and outputs in a set of data, with the following requirements set out in [36] p.22.

1. Data (a set of historical examples).
2. A set of functions that the algorithm will search through to find the best match with the data.
3. Some measure of fitness that can be used to evaluate how well each candidate function matches the data.

The measure of fitness is used iteratively to adjust the functions to arrive at the best available fit with the data. The term 'deep learning' has been coined to refer to machine learning methods which deploy algorithms in layers, each of which is optimized by the network. That such systems are powerful is not in doubt. They "control operations as diverse as labeling images, recognizing speech, translating texts, playing strategy games, predicting protein folds, detecting new exoplanets, analyzing fMRI data, and driving automobiles autonomously" [37] and there is no limit in sight to what may be achieved in future. Be this as it may, the aspect of deep learning which most concerns us here is its recursive structure, in which the results of information processing change the way in which this processing is carried out. This requires a reconsideration of Bateson's criteria 4 & 5, both of which were partially fulfilled by symbolic AI.

1. Bateson's Criteria for Mind Applied to Deep Learning

As regards criterion 4, there is no doubt that deep learning applications have a greater capability to change their own functioning

than do symbolic AI applications. However, the reach of these changes is tightly constrained to specified algorithms in a fixed architecture.

The situation for the closely related criterion 5 is similar. In common with symbolic AI, deep learning fulfills the requirement that effects of difference are transforms of events which preceded them. As we noted above, however, the problem arises with Bateson's additional requirement that the rules which govern that transformation should be subject to transformation. To some extent, the implementation in machine learning of recursive change of the functions used in transformation addresses this requirement. Moreover, the use of 'evolutionary' or 'genetic' algorithms, which has been developing in AI since the 1960's, has become an established technique (see De Jong [38] for a historical overview), and this enables more substantial changes to be made to the rules. However, as a biologist, when Bateson wrote about changing rules he was surely thinking of evolution, and, as we have already noted, he viewed embryology and evolution as mental processes [5] p.16. From the perspective of biological evolution, the evolution that takes place in deep learning is rather superficial. This is because the applications which determine the changes in the algorithms are themselves excluded from evolutionary change, and the same is true for the hardware on which the system runs. Zaadnoordijk, Besold, and Cusack [39] point out that machine learning has been based on adult learning, but that a study of specific processes in the cognitive development infants might produce valuable techniques. Such a focus on developmental change could, perhaps, move towards a more complete fulfillment of criterion 5.

Thus, the development of deep learning takes a step towards fulfillment of criteria 5 & 6, but still requires the presence of a human to meet the requirements.

2. Levels of Learning

Bateson makes a perceptive comment in this context: "The question is not "Can machines learn?" but what level or order of learning does a given machine achieve?" [2] p.284. He is precise about what he means by levels of learning, and his definition is closely related to the logical types introduced in section IV.B (see *The Logical Categories of Learning and Communication* [2] p.279-308 for a full discussion of the relationship summarized below).

Zero learning is defined as that in which information is simply stored and reproduced at the appropriate time, for example learning the time of an appointment. Bateson comments that "many very simple mechanical devices show at least the phenomenon of zero learning" [2] p.284, adding that "a very high (but finite) order of complexity may characterize adaptive behavior based on nothing higher than zero learning" [2] p.284. This complexity is possible because information of many different logical types may be managed within a finite and constrained architecture, even though the AI application that is doing the learning is constrained to zero learning.

Level I learning involves "the class of phenomena which are appropriately described as *changes* in zero learning (as "motion" describes change of position)" [2] p.287. In level 1 learning the entity gives different responses at different times. In an organism, this may be, for example, a result of habituation or reinforcement. In the machine, level 1 learning is absent in symbolic AI, but is clearly present in deep learning applications.

Level II learning, put simply, is 'learning to learn', for example one might learn to perform better at rote learning tasks. This involves changes in the process of learning, and recognition of new contexts which require different responses. Bateson terms this "changes in the manner in which the stream of action and experience is segmented or punctuated into contexts together with changes in the use of context markers" [2] p.293. He illustrates this with the example of 'reversal learning' experiments in which the subject is taught that $X = R1$, and

that $Y = R2$. Once this has been learned the relationship is reversed. Level II is manifested in an improvement of the subject in recognizing the reversal and adapting to it.

Our discussion suggests that deep learning is pushing towards achievement of Level II but has not achieved it. There is no equivalent in deep learning to the developmental changes of children, or when an entirely new set of capabilities is acquired through learning mathematics or a musical instrument from a position of complete ignorance. The developmental approach recommended by Zaadnoordijk, Besold, and Cusack [39], cited above, indicates a possible route forward for deep learning in this respect.

Bateson also discusses a rather more elusive *level III learning*, which he describes as "likely to be difficult and rare even in human beings", involving "profound reorganization of character" [2] p.301. However, this is out of scope for a discussion of current AI.

VI. THE SACRED

A. The Sacred and the Spiritual

This special issue is concerned with AI and spirituality, but Bateson largely avoided the word *spirituality* in his writing. This is perhaps because the meaning of the word is so tightly bound up with the dualism spiritual-material. An exception to this rule is that he defines *sacrament* as the "outward and visible sign of an inward and spiritual grace" [4] p.230. In this case he was perhaps describing religious practices as an anthropologist, using the terminology of the field. In any event, he never invokes the spiritual as part of an explanation for the phenomena which he discusses. On the other hand, the sacred is a frequent theme in his writing. The meaning which he attached to the sacred was elusive, for reasons which will become clear in our discussion below, but two characteristics can be made clear as a starting point. Firstly, Bateson did not see the sacred as in any way opposed to a scientific understanding of the world, and he was convinced that "there are strong arguments for the necessity of the sacred, and that these arguments have their base in an epistemology rooted in improved science and in the obvious." [5] p.11. Secondly, he situated the sacred (as Mary Catherine Bateson puts it) in "the integrated fabric of mental processes that envelops all our lives" [5] p.200, and consequently the concept of the sacred is an integral part of the ideas which we have been discussing in this paper.

B. Paradox and Causation

Bateson frequently referred to Epimenides paradox.

...the ancient paradox of Epimenides - "Epimenides was a Cretan who said, 'Cretans always lie'" - was built upon classification and metaclassification. I have presented the paradox here in the form of a quotation within a quotation, and this is precisely how the paradox is generated. The larger quotation becomes a classifier for the smaller, until the smaller quotation takes over and reclassifies the larger, to create contradiction. [4] p.116-117.

The same oscillation can be seen in physical systems, for example in a buzzer circuit:

- If contact is made at A, then the magnet is activated.
- If the magnet is activated, then contact at A is broken.
- If contact at A is broken, then the magnet is inactivated.
- If magnet is inactivated, then contact is made. [4] p.59.

The question arises as to why the Cretan liar is paradoxical, whereas a buzzer is an unproblematic piece of everyday equipment. The answer, argues Bateson, lies in a confusion about the meaning of the word 'if', which can either refer to causal relationships (if an induction magnet is activated, then a nail will be attracted to it) or

to logical relationships (if all men are mortal, and Socrates is a man, then Socrates is mortal). Thus, the sequence in the buzzer circuit makes perfect sense when seen as a causal description, but none whatsoever as a sequence of logical propositions. The difference is the result of the inevitable inclusion of time in causal relationships, so that the description of the buzzer circuit is a set of sequential steps, each of which supplants the previous one. In contrast, the steps of the syllogism showing that Socrates is mortal are valid simultaneously, and permanently. Bateson argues that this has major implications for computers.

The “if ... then ...” of logic contains no time. But in the computer, cause and effect are used to *simulate* the “if ... then ...” of logic, and all sequences of cause and effect necessarily involve time. [2] p.281.

The code which is run on a computer is an abstract logical structure which stands outside of time. However, when the code is instantiated in digital circuits, it operates as sequences of cause and effect which exist in time. The result is, as Norbert Wiener pointed out, that a computer would encounter the Cretan liar not as a paradox, but rather as an oscillation YES . . . NO. . . YES . . . NO . . . until it runs out of energy. [4] p.117. This, if it were permitted by the programmers of the computer, would be experienced by the user as a malfunction of the computer.

Among other consequences, computers as we as we are familiar with them, with their ground rules of logic and respect for logical types, appear to be precluded from the possibility of being conscious in Bateson’s sense of reflexive and recursive mental processes, a perspective which is explored at length by Hofstadter in his discussion of ‘strange loops’ [15].

More generally, Bateson argues that biological systems, including brains, are networks of causal links. Furthermore, every circuit of causation in biology, in physiology and neural processes, and in ecological and cultural systems, “conceals or proposes those paradoxes and confusions that accompany errors and distortions in logical typing.” [4] p.109. In practice, it is exceedingly difficult to decide that aspects of an organism’s activity are in a meta-level relationship to others [4] p.117. It is not possible to render such process in a set of logical links, without violating the rules of logical types established by Whitehead and Russell to exclude paradox [40].

C. Paradox and the Sacred

A metaphor is a kind of syllogism, but one which is not held together by the logical links of Socrates’ mortality, discussed above. Bateson gives the example, [2] p.205, of

- Grass dies
- Men die
- Men are grass.

He argues that this is the way that biological homology is best understood, as for example, “a formal similarity that suggests a relationship, like that between a human hand and the wing of a bat” [5] p.192. Such formal similarities emerge not from logical connections, but from a vast network of causal relationships full of circularity and contradictions in logical typing, with an associated lack of clarity of what is causing what, and at what level of logical type. He also proposes that this is how poets think, and, we might add, other artists. In this he echoes his contemporary Arthur Koestler, whose concept of bisociation [41] analyzed all creativity in terms of bringing together intersecting planes of associations, with concomitant violation of logical categories. Bateson knew Koestler from the Macy conferences, but strongly rejected some of his ideas [5] p.57-58.

In our normal waking states, we make internal or external reports of our perceptions, in a state which Bateson refers to as *prose consciousness*, and which he associated with the left hemisphere

of the brain. In this state we are quite able to label the thing that we perceive as a symbol, for example a stop sign on the road. We can even label it as a metaphor, and parse that metaphor into its components. But we also have other states, where the identification of hand and bat, or wine and blood, is not labeled with a logical hierarchy, but experienced as an identity. This mode of thinking, familiar from dreams, and also present in (for example) aesthetic experiences, trances of various sorts, religious experience and the intensity of love. In these states the difference between the logical types of the map and the territory is dissolved, and we return to the “innocence of communication by means of pure mood-signs” [2] p.183. This state is the ‘inward and spiritual grace’ of which the sacrament is an ‘outward and visible sign’. From this perspective, the sacrament is more than a metaphor, but is rather seen as the thing itself, leading, for example, soldiers to sacrifice themselves to save a flag, and for martyrdom to be embraced to defend the idea of the transubstantiation of the host.

In our interior life, and in our relations with our environment, human beings participate in both logical and causal circuits, in prose consciousness and the transcendent, and in the rational and the emotional. Indeed, paradox is central to the most widely recognized sacrament in western society, the mass, in which the bread and wine are both themselves and the body and blood of Christ. How is this to be understood. Is the bread transformed into the body of Christ during the mass, through which we can experience union with Christ? Or is it a symbol for the body of Christ, whose contemplation can lead to religious insight?

Bateson suggests that “the richest use of the word “sacred” is that use which will say that what matters is the *combination* of the two... any fracturing between them is, shall we say, anti-sacred” [6] p.267. Any attempt to analyze a specific example of this cohabitation between the different visions requires, instead of a unified experience, the alternating view of that experience from the two different standpoints, dissolving the phenomena which we hope to analyze, an alternation reminiscent of the computer’s response to the Cretan liar.

Such a combined experience of opposites involves paradox not only in operating with the conflicting premises of two mutually incompatible types of interaction, but also in considering the nature of the resulting combined entity. It was because of this that Bateson stated that “To be conscious of the nature of the sacred or of the nature of beauty is the folly of reductionism” [4] p.214.

It follows from this position that

noncommunication of certain sorts is needed if we are to maintain the “sacred.” Communication is undesirable, not because of fear, but because communication would somehow alter the nature of the ideas [5] p.80.

It is this which leads Bateson to be elusive in his descriptions of the sacred. He suggests that this is part of wider phenomenon, whereby there may be processes in all living systems such that “if news or information of these processes reaches other parts of the system, the working together of the whole will be paralyzed or disrupted” [5] p.81.

The conceptual framework outlined above places the sacred outside the domain of AI as we know it. Any engagement with the sacred requires and engagement with and tolerance of paradox. AI applications that are currently conceived of, running as they do on von Neumann architectures, are unable to encounter paradox. Consequently, they cannot, in themselves, engage in the mental tight-rope walk involved in the merging these perspectives which Bateson sees as being the core of the sacred. Thus, an AI built on current design principles is systemically unable to experience, or even to represent, an important aspect of the human mind. This implies a constraint on the ability of AI to interact with a human being in a way which would enable it to substitute for a human caregiver or teacher.

In our discussion of mind in AI we saw that although AI does not in itself fulfill the criteria for mind, it can be a significant element within a wider ecology of mind. Something similar may apply to the sacred. There is no reason why the extraordinary logical structures generated by computers should not be a powerful component of the sacred. Perhaps the increasing power of simulations will give rise to new opportunities for experience of the sacred, as one pole of a combined experience. Readers who find the possible association of simulations, including those involving sex or violence, with the sacred, should bear in mind that the Latin root of the word, *sacer*, referred not only to the extremes of holy and pure, but also those of the unholy and impure [6] p.267.

VII. CONCLUDING REMARKS

The application of Bateson's ideas to AI is not intended to constitute a solution to the difficult questions which surround AI. Nor is it suggested that the insight obtained supplants other work carried out since Bateson's death. It is, however, proposed that there are valuable characteristics in Bateson's thought which can inform the current debate on AI.

Firstly, Bateson's work is based on strong foundations. His analysis starts with an explicit statement of the nature of information, but, in contrast, much of the literature of AI is silent on this. Similarly, Bateson is rigorous in his rejection of dualism, following through the implications for the nature of mind. A lack of clarity on these issues may or may not be a problem in the practical tasks of building AI applications, but a reading of Bateson suggests that this lack is a barrier to conceptualizing the phenomena generated by those applications. Whether or not one agrees with Bateson's views, the admirable clarity of his position provides an example which could usefully inform current attempts to improve our understanding of what AI is, and how humans interact with it.

Secondly, the explanations offered by Bateson are functional, and he ascribes the properties of things to their structure. There is therefore no obstacle in principle to AI achieving human mental abilities. The constraints on AI which we have identified in this article are related to the structure of computers as we know them, and as we can presently conceive of them. There is every reason to suppose that Bateson would have agreed with Chalmers when he argued that a neural description of the brain, translated into a combinatorial-state automata, would have experiences indistinguishable from the brain [42] p.321. Neural networks have moved some distance in this direction with increasingly sophisticated models of the behavior of neurons, see for example [43], and further progress is surely to be expected. In this context Bateson's ideas can make a valuable contribution by focusing attention on the levels of learning which are exhibited in machine learning, and on the scope of adaptive change which is required if AI is to become equivalent to its organic counterpart.

Thirdly, as Denning and Tedre pithily put it, in deep learning applications "All there is inside is an inscrutable, complex mass of connections." [44] p.173. This aspect of deep learning is intriguing, because it moves AI in the direction of Bateson's description of the equally inscrutable tangled network of messages in organic brains, within which complex relations of logical types are imminent. However, the rather rigid layering of the algorithms which run deep learning applications would seem to militate against the development of recursion in the mass of connections in deep learning applications. Leaving to one side the complex architectural issues which arise, Bateson's ideas suggest that it would be interesting to explore the results of loosening the prohibition of recursion in the networks of connections in machine learning, and indeed encouraging it.

Lastly, one of the most challenging aspects of Bateson's ecology of mind is the idea that mind does not end at the physical limits of an organism or machine, but rather at the limits of the information flows which constitute the mind. However, the alternative to this view is equally difficult to assimilate, i.e. that a mind is constituted of something other than information flows, by a mental stuff which is present in brains, and perhaps in AI, but which we have yet to detect. Some nodes in the ecology of mind are clearly more powerful than others. When I interact with a dog, I am aware that I have mental capabilities which the dog does not have (although it doubtless has some important capabilities, for example relating to smell, which I lack). The same is true of my interactions with the computer on which I am typing this text. As a result, the search for, and the potential deification of, a discrete super-mind is misleading from Bateson's perspective. Whatever is developed in the future will participate in an ecology of mind with all the organisms and AIs with which it is in contact. Indeed, from Bateson's perspective, it is hard to see how the AI could be useful or effective without that network of information flows, within and between components of the ecology.

Nevertheless, just as humans are peak predators in the ecology of energy, they are also peak nodes in the information flows of an ecology of mind. The singularity, popularized by Kurzweil [45], suggests that once AI surpasses human capabilities, it will accelerate exponentially past us, and become superhuman. There is understandable concern regarding what such a superhuman entity might choose to do to its progenitors. But Bateson's writings suggest that we ask another kind of question, one which should not wait until the postulated singularity arrives. Our immediate concern should be "what is the impact of increasing AI capabilities on the ecology of mind, and how does this change the niche of human beings within that ecology".

REFERENCES

- [1] M. Minsky, "Steps Toward Artificial Intelligence," *Proceedings of the Institute of Radio Engineers*, vol. 49, no. 1, pp. 8–30, 1961, doi: 10.1109/JRPROC.1961.287775.
- [2] G. Bateson, *Steps to an Ecology of Mind*, 2nd ed. Chicago, IL: Chicago University Press, 2000, first published 1972.
- [3] L. Floridi, *The Philosophy of Information*. Oxford, UK: Oxford University Press, 2011.
- [4] G. Bateson, *Mind and Nature: A Necessary Unity*. New York, NY: E. P. Dutton, 1980, first published 1979.
- [5] G. Bateson and M. C. Bateson, *Angels Fear: Towards an Epistemology of the Sacred*, 2nd ed. Cresskill, NJ: Hampton Press, 2005, first published 1987.
- [6] G. Bateson, *A Sacred Unity: Further Steps to an Ecology of Mind*. New York, NY: Bessie/HarperCollins, 1991.
- [7] D. Lipset, *Gregory Bateson: The Legacy of a Scientist*. Eaglewood Cliffs, NJ: Prentice Hall, 1980.
- [8] P. Harries-Jones, *A Recursive Vision: Ecological Understanding and Gregory Bateson*. Toronto, Ontario: University of Toronto Press, 1995.
- [9] P. Harries-Jones, *Upside-Down Gods: Gregory Bateson's World of Difference*. New York, NY: Fordham University Press, 2016.
- [10] N. G. Charlton, *Understanding Gregory Bateson: Mind, Beauty, and the Sacred Earth*. Albany, NY: State University of New York Press, 2008.
- [11] S. Hepplestone, G. Holden, B. Irwin, H. J. Parkin, and L. Thorpe, "Using technology to encourage student engagement with feedback: A literature review," *Research in Learning Technology*, vol. 19, no. 2, pp. 117–127, 2011, doi: 10.1080/21567069.2011.586677.
- [12] R. Descartes, *The Philosophical Writings of Descartes*. Cambridge, UK: Cambridge University Press, 1984.
- [13] G. Baker and K. Morris, *Descartes's Dualism*. London, UK: Routledge, 2005.
- [14] R. Ashby, "Requisite variety and its implications for the control of complex systems," *Cybernetica*, vol. 1, no. 2, pp. 83–99, 1958.
- [15] D. Hofstadter, *Godel, Escher, Bach: An Eternal Golden Braid*. New York, NY: Basic Books.

- [16] F. J. Varela, E. Thompson, and E. Rosch, *The Embodied Mind*. Cambridge, MA: MIT Press, 1993.
- [17] E. Thompson, "Life and Mind: From Autopoiesis to Neurophenomenology," *Phenomenology and the Cognitive Sciences* vol. 3, no. 4, pp. 381-398, 2004.
- [18] E. Hutchins, *Cognition in the Wild*. Cambridge, MA: MIT Press, 1995.
- [19] A. Clark and D. Chalmers, "The extended mind," *Analysis*, vol. 58, no. 1, pp. 7-19, 1998.
- [20] A. Newen, L. De Bruin, and S. Gallagher (Eds.), *The Oxford Handbook of 4e Cognition*. Oxford, UK: Oxford University Press, 2018.
- [21] D. Michmizos and Z. Hilioti, "A roadmap towards a functional paradigm for learning & memory in plants," *The Journal of Plant Physiology*, vol. 232, pp. 209-215, 2019.
- [22] A. G. Parise, M. Gagliano, and G. M. Souza, "Extended cognition in plants: is it possible?," *Plant SignalLing and Behavior*, vol. 15, no. 2, pp. 1-10, 2020.
- [23] M. A. Boden, *Mind as Machine: a History of Cognitive Science, volumes 1 & 2*. Oxford: Oxford University Press, 2006.
- [24] C. Pias, *Cybernetics - The Macy Conferences 1946-1953: the complete transactions*. Berlin: diaphenes, 2003.
- [25] S. P. Heims, "Gregory Bateson and the Mathematicians: From Interdisciplinary Interaction to Societal Functions," *Journal of the History of the Behavioral Sciences*, vol. 13, no. 1977, pp. 141-159, 1977.
- [26] N. J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge, UK: Cambridge University Press, 2010.
- [27] H. A. Simon and A. Newell, "Heuristic Problem Solving: The Next Advance In Operations Research," *Operations Research*, vol. 6, no. 1, 1958.
- [28] A. Newel and H. A. Simon, "Computer Science as Empirical Inquiry: Symbols and Search," *Communications of the Association for Computing Machinery*, vol. 19, no. 3, 1976.
- [29] J. Haugeland, *Artificial Intelligence: The Very Idea*. Cambridge, MA: The MIT Press, 1985.
- [30] J. Mingers, "Embodying Information Systems: The Contribution of Phenomenology," *Information and Organization*, vol. 11, no. 2, pp. 103-128, 2001.
- [31] G. Marcus, "The next decade in AI: Four steps towards robust artificial intelligence," *arXiv*, no. February abs/2002.06177., 2020.
- [32] H. L. Dreyfus, *What Computers Can't Do: A Critique of Artificial Reason*. Harper & Row, New York, 1972.
- [33] P. McCorduck, *Machines Who Think*. Natick, MA: A K Peters, Ltd, 2004.
- [34] B. Raphael, *The thinking computer: mind inside matter*. San Francisco, CA: W. H. Freeman, 1976.
- [35] H. A. Simon, *The Sciences of the Artificial*, Cambridge, MA: MIT Press, 1969.
- [36] J. D. Kelleher, *Deep Learning*. Cambridge, MA: MIT Press, 2019.
- [37] C. Buckner, "Deep learning: A philosophical introduction," *Philosophy Compass*, vol. 14, no. 10, pp. 1-19, 2019.
- [38] K. A. De Jong, *Evolutionary Computation: A Unified Approach*. Cambridge, MA: MIT Press, 2006.
- [39] L. Zaadnoordijk, T. R. Besold, and R. Cusack, "The Next Big Thing(s) in Unsupervised Machine Learning: Five Lessons from Infant Learning," *arXiv*, 2020.
- [40] A. N. Whitehead and B. Russell, *Principia Mathematica*. Cambridge: Cambridge University Press, 1910.
- [41] A. Koestler, *The Act of Creation*. London: Pan books, 1977.
- [42] D. J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory*. Oxford, UK: Oxford University Press, 1996.
- [43] A. Jeyasothy, S. Sundaram, and N. Sundararajan, "SEFRON: A New Spiking Neuron Model with Time-Varying Synaptic Efficacy Function for Pattern Classification," *IEEE transactions on neural networks and learning systems*, vol. 30, no. 4, pp. 1231-1240, 2019.
- [44] P. J. Denning and M. Tedre, *Computational Thinking*. Cambridge, MA: MIT Press, 2019.
- [45] R. Kurzweil, *The Singularity is Near*. London, UK: Viking, 2005.

Dai Griffiths

Dai Griffiths, also known as David, has a background in the arts, and holds a PhD from Universitat Pompeu Fabra in the area of ICT. He spent the first part of his career working as a teacher in primary, secondary and higher education, as well as in interpersonal skills training in industry, before becoming fascinated by the potential of computers in education. For the past twenty-five years he has worked in the development of educational applications, and as an educational technology researcher. He has published extensively, and details are available at <https://orcid.org/0000-0002-6863-2456>. In this work he became deeply engaged with the tradition of cybernetics, and he remains active in this field. He was Professor of Educational Cybernetics at the University of Bolton. At the University of Bolton he was a member of the Institute for Educational Cybernetics, and of Cetus. He then took on a role in the Department of Education of Bolton University, leading the Department's PhD and Doctor of Education programs. Dai Griffiths is currently a Senior Researcher at the Research Institute for Innovation & Technology in Education (UNIR iTED), Universidad Internacional de La Rioja (UNIR). He is an Associate of Cetus LLP.