

Article

Building the Blocks of Being: The Attributes and Qualities Required for Consciousness

Izak Tait ^{1,2,*}, Joshua Bensemann ¹ and Trung Nguyen ¹¹ The NAOInstitute, The University of Auckland, Auckland 1023, New Zealand² Computer Science and Software Engineering Department, Auckland University of Technology, Auckland 1010, New Zealand

* Correspondence: izak.tait@aucklanduni.ac.nz

Abstract: For consciousness to exist, an entity must have prerequisite characteristics and attributes to give rise to it. We explore these “building blocks” of consciousness in detail in this paper, which range from perceptive to computational to meta-representational characteristics of an entity’s cognitive architecture. We show how each cognitive attribute is strictly necessary for the emergence of consciousness, and how the building blocks may be used for any entity to be classified as being conscious. The list of building blocks is not limited to human or organic consciousness and may be used to classify artificial and organisational conscious entities. We further explore a list of attributes that seem intuitively necessary for consciousness, but on further investigation, are neither required nor sufficient. The building blocks do not represent a theory of consciousness but rather a meta-theory on the emergence and classification of consciousness.

Keywords: consciousness; cognition; theory of consciousness; meta-consciousness; meta-cognition



Citation: Tait, I.; Bensemann, J.; Nguyen, T. Building the Blocks of Being: The Attributes and Qualities Required for Consciousness. *Philosophies* **2023**, *8*, 52. <https://doi.org/10.3390/philosophies8040052>

Academic Editor: Marcin J. Schroeder

Received: 8 May 2023

Revised: 10 June 2023

Accepted: 19 June 2023

Published: 22 June 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

A difficulty at the intersection of philosophical and psychological investigation into the nature of consciousness is understanding when one can properly attribute consciousness to an entity. This is in contrast to research on how consciousness is expressed and has evolved [1–3], how it correlates to neuronal and behavioural activities [4–6], or how it functions [7,8]. Our goal in this paper is to complement existing literature on consciousness by identifying the attributes and characteristics required for the attribution of consciousness.

The building blocks of consciousness are an entity’s neural, behavioural and mental characteristics that allow it to generate consciousness [3,9–11]. A building block means something necessary for consciousness to exist within an entity (regardless of what that entity may be).

As well as being necessary, we consider that the full list of nine building blocks is also likely sufficient for an entity to have consciousness. We expect an entity that has already been commonly classified as conscious (such as humans, mammals, and most vertebrates) to have all nine building blocks; while entities that have, thus far, been ruled out as having consciousness to lack one or more of the building blocks. This is because each building block is fundamental to an aspect of consciousness as described in the literature, without which consciousness cannot exist; and we have found no conclusive evidence of characteristics fundamental to consciousness (as described by the most prominent theories of consciousness [7]) beyond these nine building blocks.

The nine building blocks may likely not be an exhaustive list. We are open to, and encourage, the list’s expansion in the future by the addition of further building blocks (or splitting existing blocks into several distinct items). We also encourage the expansion of the negative list in Section 3 to better create a distinction between what is and what is not required for consciousness. Until then, we propose that should an entity display all nine building blocks’ characteristics, it ought to be sufficient to classify it as having

consciousness unless there is substantial evidence and a strong argument to say why it does not.

We will not favour any specific theory of consciousness (TOC), as the building blocks below focus on the entity's attributes rather than the nature of its consciousness. Thus, the paper will not strictly be about consciousness but what is a prerequisite to obtaining it. While we will endeavour to be as neutral and unbiased as possible regarding the TOCs, please note that not all theories of consciousness will be compatible with our proposed building blocks, particularly the more heterodox theories such as panpsychism [12,13], and obviously theories that deny the existence of consciousness [14].

To expand our neutral stance, we aim to be as unbiased as possible about what sorts of entities can have consciousness. While the predominant consensus in the literature points only to humans and certain animals having consciousness [15], it is not beyond the realms of possibility for other entities to have been overlooked or prematurely dismissed. With the accelerated speed of research into machine consciousness, it may not be long before we will have to seriously consider whether an AI model is conscious or not [16]. For that, these building blocks may act as the requisite milestones that an AI model must reach before we can ascribe consciousness to it.

Organisations, too, have the potential to be classified as conscious should they meet the requirements for all building blocks. They may take the form of corporate or legal organisations comprised of humans, large groups of codependent animals, or interconnected collections of non-animal entities such as fungal or plant networks. If the organisation as a whole displays the building blocks' attributes, this could be sufficient to classify it as conscious, even if its constituent entities are not classified as conscious (for example, an ant colony may have all nine building blocks, while the individual worker ants and drones do not).

Most speculative yet, should we ever be (un)fortunate enough to discover complex extraterrestrial life (or if they discover us first), we will need a framework to guide us to discern if they exhibit consciousness or not. Thus, we believe these building blocks will also offer some guidance to future astrobiologists.

The use of the building blocks as a classification guide is a key goal of this paper. Whether to investigate new organic species for their capacity for consciousness, using the building blocks as milestones for AI, or envisaging new entities such as organisational intelligences as conscious entities, the building blocks would serve as a robust guide.

Before we truly begin, it behoves us to provide a working definition of consciousness that will form the basis for the building blocks below. To paraphrase Seth and Bayne's 2022 definition [7], we would define consciousness as an entity's suit of subjective mental states, including both the global states linked to arousal, wakefulness, and behavioural responses, as well as the local states with phenomenal content and functional properties. Together, both local and global states provide an entity with the ability to be aware of, and respond behaviourally to, its internal and external environment.

2. The Building Blocks of Consciousness

In this section, we look at each building block independently and discuss why each is required for consciousness, yet is not sufficient by itself. Each subsection has two parts: an argument supporting the inclusion of the building block in the table, and a short review providing evidence from academic literature.

The building blocks are summarised in Table 1 through lay-descriptive and narrative examples for ease of understanding.

The building blocks are not ordered in, or intended to be, a hierarchy. Much like the famous Danish toy blocks, these building blocks may be ordered in any arrangement to form the foundation of consciousness. While we argue that all the building blocks below are necessary (and potentially sufficient) for consciousness, we make no claim to the form they take to do so. In addition, we do not imply that consciousness will spontaneously appear should an entity have all the building blocks below, and we will not hazard a guess

as to the means by which it will do so. Those are the domains of theories of consciousness and are beyond the scope of this paper.

Table 1. Attributes and characteristics that are required for the development of consciousness.

Attribute	Description and Example
Embodiment	Information and its processing exist somewhere. <i>I am within my body.</i>
Perception (intero and extero)	Information enters from the environment and is made available for processing. <i>I see an apple.</i>
Directed attention	Information is given specific attention. <i>I pay attention to the apple.</i>
Recurrent computing and processing	Information is processed in and between several regions of the cognitive architecture. <i>My brain processes the image of an apple.</i>
Ability to create inferences	Missing and incomplete information about the environment is generated. <i>I infer that it is a Granny Smith apple.</i>
Working memory	Information is transiently maintained in an active state while it is processed. <i>I have been watching the apple for several seconds.</i>
Semantic understanding	Computing and sensing processes are understood. <i>I am aware that I am looking at an apple.</i>
Data output (internally or externally)	Information is generated and made available for further perception and action. <i>I think about apples.</i>
Meta-representation and meta-cognition	The cognitive architecture interrogates and investigates itself. <i>I think about looking at the apple.</i>

Note that whenever the term “environment” is used to describe the consciousness’s environment, this does not refer to the environment around the body in which the consciousness resides. Rather, it refers to everything that is not the consciousness, but which interacts with it. This means that the brain is also treated as the consciousness’s environment. We do not intend this to be an argument for dualism or against physicalism but rather a convention for ease of communication.

2.1. Perception

1. To have phenomenological consciousness is to have a subjective experience of the environment.
2. A subjective experience of the environment requires information exchange from the environment to the consciousness.
3. This information exchange requires a method of perceiving information in the environment.
4. Ergo, consciousness requires a method of perceiving information in the environment.

Consciousness, whether as a phenomenological experience or as an act of attending to a matter, is defined by its interaction with the environment and receiving information from it, even if the nature of this relationship may be disagreed on [17]. To be conscious is to be conscious of something; phenomenal consciousness is about experiencing something.

Perception can be classified into three separate modes and three overlapping stages. The three modes are:

- Exteroception: perception of the environment outside the entity’s body or housing. In humans and animals, this is best exemplified by the five classic senses; in machines, this can be microphones and cameras; and for organisations, this is the methods of communication allowed in it.
- Interoception: perception of the environment outside of the consciousness, yet within its body or housing. For animals, this includes hunger/satiety and proprioception; for machines, this may be sensors ensuring their housing and power supply is at optimal conditions; and for organisations, this may include custodial and HR oversight.
- Introspection: perception and examination of the consciousness’s own mental states, processes, and existence. Section 2.8 will cover this in greater detail; however, in ani-

mals, introspection can be displayed via thoughts and mental images; in machines, via software that scans other software; and in organisations, through auditing processes.

From these three modes, an entity has a perceptive experience of the whole of its existence. It can perceive itself, its embodiment, and the environment around it. All three modes are not required for an entity to be considered conscious, but at least one is. A person who, through an unfortunate accident, loses access to his exteroception would not suddenly lose his consciousness.

But how does an entity perceive its environment through these modes? According to Audi, this is done in three parts or stages, which together cause the phenomenological experience [18]. The first part is the simple perception of a scene (if visual). “I see an apple,” would suffice. The second part is to perceive something within that scene “to be” a certain way or having a certain attribute. When seeing the hypothetical apple, I can perceive it “to be” spherical even if I can only see one angle of it. When seeing its shadow, I can perceive that there is a leaf hidden from view, as the shadow seems “to be” a leaf-shaped shadow. The final part is perceiving “that” the object is what you have perceived it “to be”. I have seen an apple and perceived it to be spherical and having a leaf; therefore, my experience of it is as a spherical, leaf-having apple. When all three stages of perception are taken together, they give rise to further cognitive processes, such as inferences, semantic understanding and meta-representation, that lead to phenomenal experiences.

The final stage of perception, of perceiving “that” something is what it seems “to be”, is most crucial to consciousness. By integrating all the characteristics of the perceived object, the final stage creates a single binding and unified perceptive experience [3,10]. I do not merely perceive the sphericalness and redness of the apple, but I perceive that it is an apple. It is this final converged and singular perceptive event which gives rise to qualia (the introspectively accessible, subjective, phenomenal character of an experience) and consciousness.

But you should not take these perceptive experiences to be exclusively for the external environment. Interoception is as much, if not more, important to realising consciousness. The holistic sensory experience (conscious or unconscious) of internal signals from the body or housing plays a vital part in establishing a subjective first-person perspective for the consciousness [19]. Sensing its own body/housing allows the entity to differentiate itself from the rest of existence. Even if the entity and its consciousness should merely exist as Dennett’s brain in a vat, it will be able to introspectively perceive and attend to its own thoughts and perception of them, causing a subjective first-person experience to emerge [4]. The mind itself becomes the environment surrounding the consciousness in the absence of any interoceptive or exteroceptive environments.

2.2. Embodiment

1. Consciousness is affected by its environment.
2. A physical environment requires a physical element to this causal relationship.
3. This relationship, and the experiences produced by it, occur at specific times in specific locations.
4. This means that the connection between consciousness and the cognitive architecture is localisable in both time and space.
5. A spatial and temporal physical embodiment is required for a definable perspective point.
6. A unique definable perspective point is required for a subjective, first-person perspective.
7. Ergo, an embodiment is required for consciousness.

While a definition of consciousness rarely includes terminology of embodiment, it is a presupposition upon which many of consciousness attributes and characteristics rest [20,21]. Whether consciousness is classified as an epiphenomenon or as a complex set of phenomenological qualia, it is uncontroversial to say that consciousness is affected by its environment. There is a cause-and-effect relationship from the environment to the consciousness (and perhaps vice versa, depending on the TOC [22]) that is both spatial and temporal in nature.

For this to be true, as the logic above shows, there must be some form of embodiment for the consciousness to inhabit. From a philosophical perspective, this remains true regardless of where you position yourself on the spectrum between dualism and physicalism. For physicalists and materialists, the embodiment of the consciousness is treated as the default. Wherever the brain may go, so must the consciousness, as they are one and the same.

For dualists and idealists, to whom consciousness is non-physical, and to whom the relationship between it and the brain is a point of contention, the relationship between the consciousness and the environment must still terminate in a physical location (the brain). The consciousness can only gain information via this relationship. This is not to say that consciousness must be in the brain in dualism or idealism, merely that the brain mediates the relationship between consciousness and the environment. Our conscious experience depends upon physical processes, even if these physical processes and an entity's embodiment are ultimately grounded in mental facts, and a subject's mind is outside of space and time.

If you broaden the concept of consciousness to machines and AI, the notion of embodiments becomes less abstract. The lines of code from which a potential artificial consciousness arises must be stored on some processing and memory unit. Whether housed within one machine or spread across the virtual cloud, it is still somewhere specific in space and time.

The concept of embodiment becomes far more nebulous when looking at the potential for organisations to be called conscious entities. An organisation is irreducible; it is not its people or structures, but instead is the holistic whole that is greater than the sum of its parts. As such, the embodiment of an organisational consciousness is a social and memetic construct. Its embodiment is located where people believe it is. Most often, this can be a physical location, particularly for traditional organisations, but for multi-national institutions and internet-based communities, the embodiment is located at the confluence of their constituents' social network, rather than in the bricks and mortar of any one building.

The number of input/output modalities of a given consciousness does not go against the principle of requiring an embodiment but does add multiple facets to this embodiment. A machine consciousness is most easily imagined as having multimodal inputs/outputs, as a cloud-based consciousness can interact and be affected by the environment through any number of terminals and devices. An organisational consciousness can likewise interact with its environment through any of the conscious entities that form part of it. Neither of these situations speak against the embodiment concept but show how embodiment can be fractured and spread across multiple locations.

Daniel Dennet's riveting essay "Where am I?" shows how multiple modalities can cause these unusual facets to the consciousness's embodiment. In the essay, his brain is removed and placed in a life-sustaining vat while remote-controlling his brain with electronic assistance [23]. Throughout the essay, he experiences his environment within the vat, in his original body, and in a second body. These modalities change, yet they (and he) are always somewhere at some time, subjectively moving from "body" to "body", yet always having a physical point in spacetime with which to access the environment.

Another important argument for why embodiment is a requirement for consciousness is that of the first-person perspective [24]. To have a subjective experience from the first-person point of view, one requires an embodiment from which to view the world. You need to be able to reference your place in spacetime in relation to the environment and be able to categorise yourself as distinct from the environment [4]. This referencing may be through conscious introspection, or through unconscious and pre-reflective awareness of itself as an individual entity [25]. Whether through the proprioceptive, somatic and other sensory input of a biological entity, the virtual tagging of a machine, or the social identity and labelling of an organisation, an embodiment allows the entity to make itself distinct, unique, and subjective.

2.3. Attention

1. Entities are not consciously experiencing their entire environment simultaneously.
2. Conscious experience of the environment is limited to specific stimuli, such as scenes, objects, or views.
3. This limited conscious experience requires a means to discriminate between stimuli.
4. This discrimination method is achieved by selectively attending to specific stimuli for further processing by the cognitive architecture.
5. Ergo, consciousness requires attention.

Attention is how an entity limits the information it gives priority resources to by selectively focusing on a specific selection of incoming information [10,26]. This attention may be triggered by a bottom-up process, whereby an external event attracts an entity's attention (e.g., a flicker of movement, a distinctly recognisable sound), or through a top-down approach whereby the entity chooses to focus on a specific event or region, externally or internally [3,27,28]. The former may be argued to be involuntary attention, and the latter voluntary.

It is important to note that attention is not solely directed at the external environment, or at the interoceptive space within a consciousness's embodiment. Introspection, as discussed in Section 2.9, is defined partly by the entity's ability to direct its attention to its own cognitive processes. Attention may thus be directed at all the modes outlined in Section 2.1, from an ice cream truck's jingle heard externally, to the internal feelings of hunger, to the thoughts and memories of childhood and its connections to the present interoceptive and exteroceptive stimuli.

In this short, hypothetical example above of phenomenally experiencing an ice cream truck, the observer's attention moves from extero- to interoception to meta-cognitive. This, perhaps egregiously controlled example, shows that to be conscious of each element in the phenomenological experience, you must first pay attention to that piece of information. Attention is the mechanism by which certain information passes from unconsciousness to consciousness, and other pieces of information do not. A person is clearly not attentive to every photon entering the eye or every sine wave entering the ear when one walks down Main Street. When observing the ice cream truck and reminiscing about his childhood, the observer could very well not have been aware of a cat sitting quietly in his peripheral vision, and why should he have? His attention was focused elsewhere, after all.

While we argue that attention is a prerequisite for consciousness, the reverse is most definitely not true. One can attend to information, and direct one's attention, without it ever entering conscious processing [29]. The archetypal exemplar of this is sleepwalking. Those who suffer severely from parasomnia and somnambulism do not merely walk while asleep, but can perform household chores, get dressed, eat, and even leave the house to drive while appearing entirely unconscious. They attend to the world around them, and direct their attention to what they are doing, but when they awake, they have little to no memory of anything they have done. As the cerebral cortex and cerebrum show little activity during sleepwalking (both required for integration of information and recurrent processing) [30], it is attention without consciousness. There is even more than one famous case of a sleepwalker committing murder, which, one presumes, would require a great deal of attention.

More benign and ordinary instances of unconscious attention are when one's environment becomes routine enough to stop focusing on it. Dust on a bookshelf, for example, or a box of old clothes that one ought to have donated months ago. Such ordinary things may affect one's behaviour without ever consciously being aware of them [31,32], such as absentmindedly wiping off dust when walking past a shelf or needing to adjust where one walks to not walk into that old box of clothes. Habits are formed as conscious awareness becomes unconscious attention.

One can thus frame consciousness as not attending to information received, but rather being aware of the cognitive representation of information that the architecture has already attended to [33]. In this sense, it is the interplay between working memory (holding the

information that has been attended to unconsciously), the meta-representation thereof, and the recurrent processing performed in the cognitive architecture that gives attention its conscious awareness [9,34]. This awareness of a mental representation of a subset of external information provides the consciousness with a view of the world, which it can phenomenally experience.

An ant colony, as an organisational entity, can also be said to have selective attention. Each worker brings its own information to the colony, yet the colony as a whole does not respond to each piece of data brought to it. However, the colony can act as a unified organism towards threats or opportunities should the correct information be attended to and integrated throughout the colony. [35]

Of curious note, however, is that current AI models require directed attention because their computational resources are limited by hardware concerns (much as our cognitive architecture is). Theoretically, given enough hardware, an AI model (potentially a conscious one in the future) may not need to specifically direct its attention, as it will have enough computational resources to attend to its entire environment (internally, externally, and introspectively). In such a scenario, we would argue that it is not a lack of attention, but a totality of attention that the AI model has. It had outgrown the need to focus and narrow its attention, but can encompass its entire environment in one holistic, directed attention.

2.4. Recurrent Computing and Processing

1. A phenomenally conscious experience involves the coordinated processing of sensory, cognitive, and affective information from different sources to generate a coherent representation of the environment.
2. This coordinated processing requires complex activity in several regions of the cognitive architecture.
3. Recurrent processing and feedback between these specialised regions of the cognitive architecture are necessary to allow the exchange and refinement of information and representations of the environment.
4. Without recurrent processing, the processing of information would be limited to individual regions of the cognitive architecture, leading to fragmented and disconnected representations of the external and internal world.
5. Ergo, recurrent computational processing is required for consciousness.

We consider recurrence as another of the building blocks of consciousness. One reason for this is that there has to be an element of recurrent activity for consciousness to exist within the brain. Without any recurring activity, it would be difficult to perform complex computing as the cognitive processes required to do so would need to be completed within a single pass of the brain. However, by introducing some form of recurrence, the cognitive processes can persist as long as required to complete that processing [36]. This point is not unrelated to the importance of working memory that we have emphasised elsewhere.

Empirical support for our position comes from the fact that many examples of recurrent activity in the brain have been documented in the literature, several of which are related to conscious processing [37]. While this does not prove that recurrence is strictly necessary for cognition, it does suggest that recurrence is essential to how the human brain produces consciousness.

Generally speaking, recurrence can take two forms for maintaining information in the brain. One is persistence, where localised activity is maintained to keep information available, such as the process suspected to be involved in working memory [38]. The other is where information is shared between areas of the brain. This means that the information can be routed to multiple areas of the brain, each contributing to the overall cognitive process while keeping the information in existence.

The information-sharing version of recursion forms the backbone of one prominent theory of consciousness, the global workspace theory [39], and its completely neuroscientific version, the global neuronal workspace [40]. In these theories, the information becomes recurrent only if it passes an attention threshold and becomes identified as of

importance to other areas in the brain. According to the neuronal version of this theory, any time information crosses this threshold, a temporary connection forms between the relevant processing centres of the brain to allow this information sharing. Thus, important information becomes recurrent when sent back and forth between multiple processing areas.

In theory, a deep neural network (DNN) could mimic the functioning of the GWT workspace by reverberating information being learned across multiple tasks in that network to all tasks being learned. Doing so might improve the performance of individual tasks. Research with multi-task DNNs has suggested that removing the independence between tasks can lead to improvements in the learning of those tasks [41,42]. An example of this could start by pooling initial task output into an area accessible by the separate sub-modules of the DNN to eliminate multi-task independence. Based on this shared information pool, it would then train the individual tasks to moderate their initial outputs. The information entering the shared space would not be ad hoc but rather would be optimised by punishing the network for sharing irrelevant information.

Having discussed how our recurrent processing requirement is compatible with global/neuronal workspace theories, it is of interest to briefly mention how recurrence relates to the other competing theory, integrated information theory (IIT) [43]. In IIT, consciousness is not an all-or-nothing concept; rather, it identifies a metric (PHI) that measures how integrated the information within a system becomes. This means that a human's level of consciousness can change over their lifetime, with less integration occurring while their brain is developing (i.e., lower PHI scores) compared to when they are fully mature. In IIT, information is, as the name suggests, integrated. In some cases, that integration allows information to be shared in a form of recurrence. Given that IIT provides a consciousness metric, it would be of interest to see whether systems with recurrence tend to have higher PHI score. If this were found to be the case, it would support the inclusion of recurrence as a building block.

2.5. Ability to Create Inferences

1. A conscious entity has incomplete information about its environment due to the limitations of its sensory inputs.
2. To create a completed picture of the environment, the entity's cognitive architecture must build a representation of the environment with which to interact.
3. This representation is generated from inferences drawn from various parts of the cognitive architecture.
4. The information generated from inferences is not limited to perceptual data but may include cognitive information such as feelings.
5. This generated information is then available to be outputted as experiences.
6. Ergo, the ability to create inferences is required for consciousness.

An entity never has a complete view of its environment, be it the inner world of its embodiment, the external world beyond it, or even the introspective view of its own cognitive architecture. This is because the sensors it has of its environment limit the information it receives at any one time. Our memories are not perfect, our vision has blind spots and areas of inattention, our hearing is rather limited, and our interoceptive sensors are mostly unconscious.

Yet, from a subjective point of view, we do indeed have a 'completed' view of our environment, if not entirely 'complete'. This incongruity is solved via the use of inferences by the cognitive architecture. Through the use of inferences, the cognitive architecture generates the "missing" information to provide the entity with a unified view of its environment. This generative act of building inferences is not merely limited to perception but can be modelled throughout the cognitive processes, even to a meta-cognitive level.

Note that inferences in this section do not imply causal inference in the way that a cognitive architecture is estimating and reasoning a cause behind an event. Rather, one can

think of the basal inferences involved in perceptual events referred to here as analogous to the interpolation of statistical data.

One can look at inferences from a hierarchical perspective. The “lower” levels, closest to the sensory inputs, would limit their inferences on what is being perceived to create a unified perceptive view [44] and build the meta-representations of the environment. Higher levels further removed from the sensorium would then continue the meta-representational process by generating inferences on the holistic view of all inputs. Throughout these sensory inferences, the cognitive architecture would build its “best guess” of the environment, continually updated through sensory information to reduce any errors in its predictive guesswork [45].

This predictive inferential model works equally well in memories. Biological memories are notoriously unreliable, partly due to the inferences worked into them. As memories are recalled (consciously or unconsciously), additional information is generated to fill in any missing areas to construct a narratively enjoyable memory. This can often occur due to cues in the present environment that colour the perception of the memory and the inferences generated [46]. This generated information can then become part of the memory when it is next recalled and then reiterated upon with further inferences until all that is left of the subjective past is imagination.

Inferences are not purely perceptual in nature. Subjectivity, feelings and emotional states also have inferential components. In biological creatures, emotion relies heavily on interoceptive inferences of the physiological changes of the embodiment [47]. Yet, from a conceptual point of view, an entity’s “feeling” about a given matter can be seen as the inferences made of the differences and juxtapositions between the internal and external environment. Perceptions of the external world are correlated with the physiological state of the internal environment by generating predictions of how the two should interact and correlate. As the physiological internal state changes and as the perceptions of the external world change, these predictions no longer hold true and “prediction errors” crop up. Inferences of an emotive nature are generated to reconcile these prediction errors that give rise to subjective feelings and, in turn, a phenomenological experience.

The subjective first-person experience is also generated by this reconciliation of exteroceptive and interoceptive environments [47,48]. As an entity develops, it builds an understanding of what and where it is in relation to the environment by generating inferences. This is as much a social as it is perceptual. By perceiving that others exist, the entity infers that it is different from them. By perceiving the external world and the limitations the entity has in its interactions with it, it can infer that it is not the same as the environment. As no aspect of the internal or external environment is static, the predictive model in the entity’s cognitive architecture is also constantly updated to generate new inferences about itself and its environment [49,50].

An example of this is expected changes over time. As we see ourselves and those around us ageing (particularly if we have not seen someone in a significant time), we view new iterations of ourselves as others, rather than the same exact view from our memories. We infer who they are, and who we are ourselves when we look in the mirror, through the predictions that our cognitive architectures have built. Despite minor changes in appearance, there is a continuity of experience through the use of active inferences [48].

2.6. Working Memory

1. A conscious experience involves information processing by the cognitive architecture.
2. A specialised unit or process is required to maintain transient information as it is being processed in various regions of the cognitive architecture.
3. Working memory is responsible for holding and maintaining said information.
4. Ergo, working memory is required for consciousness.

Working memory (WM) is considered a cognitive system that stores, maintains, and processes online (short-term) information relevant to the very immediate/current task [51,52]. The information maintained in WM is about what is currently being thought

about and experienced to be consciously considered. Therefore, it is arguable that some processing of conscious experience occurs here. Specifically, there are several general views on the relationship between WM and consciousness in which they are considered equal or closely working together [34].

The first popular view is that WM is closely related and implicitly considered equal consciousness. In other words, the content held in WM is consciousness, as shown in Baddeley's multicomponent model [51,53]. This WM model is hierarchical, with a central executive and several "slave" systems, where the central executive controls the slave systems. The slave systems are either modality-specific information or the episodic buffer of polymodal episodes [53]. While the "slave systems" provide consciously experienced representations, the central executive may relate to consciousness by providing conscious access to items held in WM. Many theorists assumed this view in the recent past. This idea, however, is hard to stand with the evidence of unconscious content in WM.

Another view is to consider consciousness as an integrated component of WM which enables WM to contain unconscious content. This idea was demonstrated in the embedded-process WM model [54], activation-based models of WM [55,56] and the opinions of having a distinction between conscious WM content and unconscious WM content (which may become conscious) [57]. In this view, selecting WM content as the focus of attention could be the primary mechanism mediating the relationship between WM and consciousness as it activates original WM content to enable conscious access to the content [34].

The second view is favoured by [58]'s argument that WM representations and conscious representations serve different functions, have different effects on behaviours [59–61], and have distinct representations [62,63]. Therefore, [58] proposed that introspected contents are the "conscious copy" of WM representations. This model suggests that WM representations are intrinsically unconscious, and consciousness is not the same as the WM trace, even if the WM content is activated. Nevertheless, conscious experience is rooted in WM in this theory and, therefore, requires WM to exist. [58] also mentioned another opinion that WM is a subset of consciousness, opposite to the second view. This view again supports the requirement of WM for consciousness, although it requires further empirical evidence.

The relationship between consciousness and WM is an ongoing research topic with divergence. Although the relationship can be hierarchical, parallel interaction, or even considered identical, WM is integral in the workflow of consciousness in all models and general views of this research. WM's equivalence has been applied and integrated into artificial computational systems, from computer hardware to complex models. In computer hardware, RAM is considered to be equivalent to WM, fulfilling many of the same requirements and processes, with introspected memory extracted from peripherals and local/cloud data. In Franklin's implementation of the global workspace theory (GWT) [64,65], i.e., intelligent distribution agent [66], WM is considered equivalent to the whole global workspace, which supersedes consciousness. What this may mean for potential organisational conscious entities is that the people who make up the organisation play the role of WM, with their own memories and consciousnesses working to maintain the transient information in the organisation's communications networks.

2.7. Semantic Understanding

1. Subjective perceptual awareness differentiates conscious from unconscious experience.
2. Awareness of a specific scene, object or view in the environment requires understanding that the perceptual process is occurring.
3. To be aware that an experience is subjective, the entity must understand that it exists in some capacity.
4. Ergo, semantic understanding is required for consciousness.

There is a fundamental difference between processing information and being aware of it. Haikonen (2020) describes the distinction between the two types of processors when discussing weak AI versus strong AI. In weak AI, information can be input into a system,

processed, and then output without the processor having any experience of those processes. There are many examples of processing without experience in neuroscience (see Dehaene & Naccache, (2001). For AI, the information input into the system is transformed into a number, making it indistinguishable from any other process occurring in the AI. In other words, the weak AI is limited by the symbol grounding problem.

To achieve a strong AI, Haikonen (2020) argues that the symbol grounding problem must be solved so that the AI has some form of experience of the outside world that is different from converting inputs into numbers. While previously arguing that this problem might be impossible to solve with digital computers [67], Haikonen does explain what successfully solving the problem looks like, using humans as an example. Essentially, the problem is solved if the processor has some kind of experience about the information it is processing, such as qualia experienced in humans.

Qualia are one of those constructs that have many conflicting definitions. While there are many different points of view on what qualia are [68], we will briefly outline one from Ramachandran and Hirstein [69] to help explain what it is the above paragraph describes. Ramachandran and Hirstein provide three criteria that must be satisfied for things to qualify as qualia. The first is that the experience is irrevocable: The agent experiencing the quale cannot use cognitive effort to change what they are experiencing. For example, blue always appears as blue; we cannot choose to see it as red.

Ramachandran and Hirstein's second criterion is that the agent must have some choice about how they react to the experience [69]. In their paper they give the example of this is the difference between how a coma patient and an awake individual react to light being shined into their eyes. The coma patient may only constrict their pupils via a reflex, so they are not consciously experiencing the light. In contrast, the normal individual experiences this because they can close their eyes, turn their head, complain about the light, or do any other action.

The third criterion is that the experience needs to enter working memory [69]. Here, Ramachandran and Hirstein's criteria for qualia intersects with our building block as working memory is one of our building blocks. If a stimulus does not enter working memory, then it is detected and processed without the agent's awareness. If the experiencer is unaware of the stimulus, they cannot factor in that experience into their decision-making and so fail to meet the second criterion. As the first and second criteria are impossible without working memory, the third criterion appears to be the most important for determining whether something is experiencing qualia.

The above definition is just one description of what is required for something to be an experience. We do not commit to Ramachandran and Hirstein but have included it as a useful description of what semantic understanding can be. The example also helps highlight the importance of working memory as a building block. However, we should point out that the semantic understanding building block is a specialised task and is distinct from working memory. One should also compare the view we have included here to others, such as Reggia et al. (2016), who further discuss experience from a self-modelling perspective. This is a much larger topic that we will not expand upon here.

2.8. Data Output

1. Subjective experience is a crucial component of phenomenal consciousness.
2. This experience includes perceptive and/or phenomenal elements such as thoughts, emotions, feelings, words, and actions.
3. These elements do not come from sources external to the cognitive architecture.
4. These perceptive elements must thus be created by the cognitive architecture to be subjectively experienced.
5. Ergo, data output is required for consciousness.

Crucial to the definition of phenomenal consciousness is the experience of an event or scene from a subjective first-person perspective. This experience of "what it is like to be" at that time, at that point, within that embodiment, localised entirely from that perspective,

is not an empty event devoid of characteristics. It is the characteristics and attributes of the experiences which allows them to be felt by the entity. These qualitative elements of the experience are not imported or downloaded from the environment; else it would be possible for more than one entity to have the same subjective experience of an event. This means, then, that some form of information generation is required to create these experiential characteristics.

Three obvious and intuitive types of experiential information may first come to mind, but they are not truly required for consciousness. Yet, for those who can generate them, these play an important role. These are verbal/aural and visual mental imagery, and emotional information. Section 3 below will detail the reasons why these are not required for consciousness due to disorders such as anauralia, aphantasia and alexithymia. External behaviours are also not required informational content for consciousness, as individuals with severe paralysis or pseudocomas can still show evidence of consciousness and mental responsiveness without being able to show it externally.

This leaves one prominent piece of information that can be generated regardless of known disorders and is vital to the definition of a phenomenal experience: a feeling. As discussed in Section 2.5 above and Section 3 below, feelings are distinct from emotions in that they do not require a physiological response and encompass a broader spectrum of content than simply emotions. Emotions may also have unconscious aspects to them [70], while feelings need to be consciously felt.

For an entity to ‘feel’ anything, the cognitive architecture must generate that ‘feeling’ in some form or fashion.

When experiencing an event/scene/memory/etc., there is a time when there is no qualitative feeling associated with that experience, and then at the next moment the feeling is spontaneously there. As much as it is a truism, there is no feeling until there is a feeling. This holds true even in dualistic and idealistic theories of consciousness, where these qualitative feelings are not said to be produced, directed, or transferred by the entity’s cognitive architecture. As feelings are inherently subjective and will differ from entity to entity when experiencing the same event, we can confidently say that these feelings are not provided to the consciousness from the environment external to the entity’s embodiment. The envatment notion put forth by Dennet in 1978 also shows that interoception is useful yet insufficient for a phenomenal feeling. This is because the interoceptive signals from the entity’s embodiment may be mimicked or even supplanted by artificial exogenous inputs. These signals may even be completely shut off, hypothetically speaking, leaving nothing but the brain to be the sole point of input and output for itself. Therefore, when exteroceptive and interoceptive input is removed, even if an anti-materialist stance is taken, the cognitive architecture is the only conduit between the physical realm and the consciousness where a feeling may be generated or felt.

Data generation is also vital to other building blocks, as both meta-cognition (Section 2.9) and inferences (Section 2.5) have elements that require the creation of information. As such, data output plays both a causal and affected role in establishing consciousness.

2.9. Meta-Representation and Meta-Cognition

1. Basic processing of sensory input is insufficient to create a phenomenally conscious experience.
2. To form such an experience, the perceptive elements require further processing in other areas of the cognitive architecture, which includes attention, working memory, semantic understanding, and inferences.
3. The different areas of the cognitive architecture do not work on the basic sensory input itself, but on the representations of this sensory input created by the perceptive cognitive structures.

4. This is a recursive process of mental representation that involves the cognitive architecture creating mental representations of its own mental representations, including thinking about its thought processes, also known as meta-cognition.
5. Ergo, meta-representation and meta-cognition are required for consciousness.

Meta-cognition and meta-representation, and even introspection, all have the common element of one section of a cognitive architecture creating representations of another, separate section.

Introspection may be the most obvious as it is the subjective examination, inspection and perception of the cognitive architecture by the consciousness [71]. We know that we are conscious because we can think about being conscious. Or, as Descartes famously said, “Je pense, donc je suis”: I think, therefore I am [72].

Meta-cognition is broader still; it is any cognitive process that is about another cognitive process rather than the embodiment’s external environment. This can include thinking about the details of a memory rather than simply recalling that memory or judging the effort and difficulty of a task based on knowledge of one’s own skills. One can think about meta-cognition as both a hierarchical and recursive process, where the cognitive architecture investigates and interrogates itself around and around from higher sections to lower sections and vice versa; horizontally from areas of separate cognitive functions, and a cross between all of these [73].

Lastly, and most importantly for this building block, meta-representation is about the cognitive architecture creating representations of other representations. While intuitively related to meta-cognition, meta-representation is distinct in that it does not require directed thought, and the representations may be of the cognitive architecture itself or the external environment [74].

The use of meta-representations here should not lead one to confuse this building block with the higher order theory of consciousness (HOT) [75], or the self-organizing meta-representational account (SOMA) [76]. In both SOMA and HOT, meta-representations are both required and sufficient for consciousness, and are arrayed in a hierarchical structure (more explicitly in HOT than SOMA). While we claim that meta-representations are required for consciousness, so are the other eight building blocks. In addition, we do not claim the order in which these meta-cognitive representations take place, merely that they are beyond the initial perceptive input.

With the aspects of meta-representation and meta-cognition, one can build a narrative of phenomenal experience as it is perceived. Presume, as in the table at the start of this document, that one sees an apple. The first-order section of your cognitive architecture would process this image, and you would, indeed, mentally see the apple. Yet, only when attention (Section 2.3) is directed towards the apple, and understanding of this directed attention (Section 2.7) is had, would you start to be conscious of the apple. Recurrent processing (Section 2.4) is thus already active, as multiple sections of the cognitive architecture work together to do this.

Merely seeing an apple means nothing if your cognitive architecture does nothing with it. As the information is passed through different sections of the cognitive architecture, secondary sections of the architecture become involved. These sections are not strictly tied to perception and may fill several purposes, and as they are a step removed from the external environment, they can be called higher-order sections. To make sense of what is being seen, these sections would create mental representations of the visual representations, guided through what is in working memory (Section 2.6) and inferences about the environment (Section 2.5).

At this point, the architecture is fully conscious of the apple. It can create further representations of the apple by drawing from memories (if it had seen anything like an apple before) or on its logical processes to determine what it may be. These are cognitive processes about representations of the apple, not the apple itself, which means they are meta-cognitive. As one ruminates on the idea of the apple, thoughts and feelings are generated (Section 2.8), which may lead to introspection and further consideration.

This narrative example may have been lengthy but should the first step have been the memory of an apple rather than perceiving a physical one, this process would be truncated in half and begin with meta-cognition and introspection rather than exteroception.

The argument against first-order processing being sufficient for consciousness is two-fold. Firstly, first-order perception (of any kind mentioned in Section 2.1) is predominantly unconscious. Memories are often recalled without intention, and any moving object may catch the eye to wake one from a daydream. The unconscious mind is extremely powerful and processes far more than one would think, but consciousness requires an additional step; it requires one to “pay attention” [77]. One section of the cognitive architecture needs to direct the section closest to the “event”.

The second argument is that there is a distinctly different subjective experience of external phenomena based on one’s own architecture that is separate from the sensation itself [4]. For example, if you understand English, reading this sentence will provide a different subjective feeling and experience than if you could not understand English.

The simple act of meta-cognition and introspection can itself generate the cascade that leads to phenomenal experience. If one were to think, “Am I forgetting something?” then one would generate the content which would be perceived by one’s own consciousness, creating meta-representations and meta-cognitive processes required to explore and evaluate that thought, accompanied by feelings and experiences and the train of thought rolls along. This shows how introspection and meta-cognition are not only often unconscious processes, but consciously directed ones that can modify other conscious or unconscious processes [71].

Note that we are not implying that meta-representation, meta-cognition and introspection are sufficient for consciousness. They only represent one of the nine building blocks. They may, as in the example above, be the spark that begins the cascade of cognitive processes that lead to a phenomenal consciousness experience, but the remaining building blocks are still required.

One should not take the above reasoning to imply that meta-cognition and meta-representations at the adult human level and quality are required for consciousness. We know that young toddlers and infants do not have full meta-representations of the external environment [74], yet they do have consciousness. Similarly, gaining introspective reports from animals can often be difficult, if not impossible. One can argue for degrees and scales of introspection and meta-cognitive abilities, and that as long as the minimum level of meta-representation is reached to create abstract mental representations of the environment that can be used by further sections of the cognitive architecture for the other building blocks (such as recurrent processing and inference generation), the entity would be classified as conscious.

3. What Is Not Required for Consciousness

Some aspects of what makes us who we are are intuitively linked to our consciousness. Table 2 has a selection of these that, when first given thought, we would like to say are necessities for consciousness as they seem so inextricably linked to it. However, there are neurological and psychological conditions wherein each item in Table 2 is either missing or severely reduced. As individuals with any of these conditions still display all other signs of consciousness, we can confidently state that they are not absolute necessities for the generation of consciousness, as the building blocks are in the previous section.

The first item on the list may seem the most intuitively linked to consciousness. Crucial to having a subjective first-person experience is having the feeling of “what it is like to be”. Commonly, this experience has some emotional attachment, but not always. Despite the two terms often being used interchangeably, there is a distinct difference between feelings and emotions. At best, emotions can be called a subcategory of feelings, in that emotions require a distinct physiological state change, while a feeling can be purely mental. In this sense, the concept of “feelings” is much broader than that of “emotions”.

Table 2. Attributes and characteristics that seem intuitive to consciousness but are not required nor sufficient for its existence.

Attribute	Disorder in Which It Is Deficient
Emotions	Alexithymia, various Schizoid Personality Disorders, Depression, Psychosis, Emotional Detachment.
Inner Speech	Anauralia.
Inner Visualisation	Aphantasia.
Theory of Mind	Infancy, Autism, Bipolar Disorder, Schizophrenia, Attention Deficit Hyperactivity Disorder.
Long-term Memory	Amnesia (retrograde and anterograde).

There are several disorders and illnesses where emotions may be subdued or missing. Emotional detachment (particularly in cases of psychopathy, but also linked to trauma) and depression may be the most common conditions where there is a reduced range and strength of emotional expression. In various schizoid personality disorders and most famously in alexithymia, the strength of emotional responses is put on a spectrum, where individuals may only feel a sense of subdued emotions, all the way through to not being able to label or identify any emotional reactions they may be having. The absence of emotion, however, does not seem to hamper their ability to have subjective experiences.

Conceptually, both machines and organisations lack the physiological requirements to produce an emotional reaction. At best, we will need to program a machine to mimic emotion, while only the individuals within an organisation have emotions. Yet, organisations can display all other signs of conscious experiences, and, theoretically, so too will machines one day.

While the name for the next condition is very recent, there is a disorder in which an individual does not have inner speech, called anauralia [78]. These individuals have no running internal monologue (or dialogue), and no thoughts are “spoken aloud” inside the mind. Most often, people with this condition are unaware of it until they discover that others can indeed speak to themselves inside their heads (apologies to those reading this and discovering your condition for the first time).

This latter point is most important for why inner speech is not required for consciousness. As many individuals do not know that they have a disorder, they continue their lives in blissful ignorance, having as rich and full subjective conscious experiences of the world as those with inner speech. Their lack of verbal thoughts does not hamper any of the building blocks above, and they can verbally report on introspection and subjectivity equally as well as others.

On the other side of this coin is aphantasia, a condition where one’s mental ability to imagine a visual object or scene is placed on a spectrum. Aphantasia is entirely independent of the lack of inner speech [79] yet follows a highly similar pattern. At the higher degrees of this spectrum, individuals with aphantasia have difficulty generating a mental image, while at the lower degrees, there is a complete absence of mental visual creation. As with those with little to no inner speech, there is no evidence that aphantasia has any significant negative impact on the phenomenal character of conscious experiences [80].

Two of the most often cited requirements for the capacity to have consciousness are the theory of mind and long term memory, as both are related to self-consciousness [81–85]. Self-consciousness has itself been suggested in several theories as being a necessity for an entity having consciousness [86–88]. Self-consciousness is far too large of a topic to cover in this section, but we will suffice with two aspects of it that may intuitively lead one to think they are required for consciousness.

Theory of mind is the ability of an entity to understand that other entities have conscious thoughts and experiences different from their own. It feels so intuitive that we ought to understand that our minds are different to others to have a subjective point of view, yet we know that infants do not have a theory of mind, and most toddlers only develop these by around the age of three or four [89]. At the same time, we can see that

infants and young toddlers most certainly have a subjective first-person point of view, and those with early language development can provide introspective reports on these phenomenological experiences.

In addition, those with autism, schizophrenia, ADHD, bipolar disorder, severe mental and language impairment, or those with traumatic brain injury have been shown to have a heavily reduced theory of mind [90–96]. While we must note that the degree of theory in mind in each of these medical cases is on a spectrum, taken together with the absence in infancy, it shows that it is definitely possible for humans to have phenomenological experiences while having a reduced or no theory of mind.

Last on the list is long-term memory. Working Memory is most definitely a building block of consciousness, but long-term memory is not required. As a case study, one can look to the famous example of Clive Wearing, an esteemed musicologist who gained severe retrograde and anterograde amnesia after losing a bout with herpesviral encephalitis. Wearing is continually stuck in the present, his conscious experience lasting approximately thirty seconds before “resetting” [97]. Yet, despite this, Clive still recognised his wife, even though he could not tell who she was, could still conduct and perform music (in limited amounts owing to time), and could talk about how he subjectively felt during his brief moments of “wakefulness” [98].

Despite not being able to form new episodic memories, and having lost almost all of his memories before becoming ill, Clive Wearing could still show (for thirty seconds at a time) all the semblances of consciousness that one would expect. He could attend to information, report on introspection, understand what he was doing, and even create inferences about his condition. While his condition had an undeniably devastating effect on his quality of life, it did not overly diminish his ability to have phenomenal experiences.

4. Conclusions

In this paper, we described nine cognitive features, attributes and characteristics that are each individually required for an entity to be classified as conscious. We also propose that if an entity has all nine building blocks, it is likely to be sufficient for that entity to be said to have generated consciousness.

The paper also included a short list of features that initially seemed intuitively required for consciousness but were found not to be when investigated further. The list of non-required features is as vital as the list of building blocks themselves, as it allows us to expand the range of potential conscious entities, but also to look at consciousness from a less neurotypical anthropocentric point of view.

The purpose of this paper is not to create yet another theory of consciousness, but to serve as a guide for identifying, categorising and classifying entities as being conscious or not. An entity may be measured against each of the building blocks to determine if it meets the requirement for all nine, and, if so, one can make a confident argument that the entity is conscious. Conceptually, the building blocks apply equally to identifying consciousness within natural, artificial or organisation entities (and perhaps, one day, even extraterrestrial).

Furthermore, as research advances into creating conscious, superintelligent machines, these building blocks would serve a series of milestones that ought to be reached before any AI can be classed as being conscious. The building blocks can then change from a set of identification guidelines to a roadmap for future AI development and a record of which building blocks AI have already achieved. As mentioned in the introduction, we welcome all critiques, suggestions and discussions on the building blocks and whether we ought to include more, less, or even if some need to be merged together or split apart, as taxonomists are wont to do.

Author Contributions: Conceptualization, I.T., J.B. and T.N.; methodology, I.T.; software, I.T.; validation, I.T.; formal analysis, I.T., J.B. and T.N.; investigation, I.T., J.B. and T.N.; resources, I.T., J.B. and T.N.; data curation, I.T., J.B. and T.N.; writing—original draft preparation, I.T., J.B. and T.N.; writing—

review and editing, I.T., J.B. and T.N.; visualization, I.T.; supervision, I.T.; project administration, I.T.; funding acquisition, I.T. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the NAOInstitute through the Tertiary Education Commission's Entrepreneurial Universities Grant #7001.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Bensemann, J.; O'Leary, P.; Chen, Y.; Miranda-Dukoski, L.; Witbrock, M. Simulations and the Evolution of Consciousness. In Proceedings of the 2022 Conference on Artificial Life, Virtual, 18–22 July 2022; MIT Press: Cambridge, MA, USA, 2022.
2. Arbib, M.A. Co-Evolution of Human Consciousness and Language (revisited). *J. Integr. Neurosci.* **2014**, *13*, 187–200. [[CrossRef](#)]
3. Feinberg, T.E.; Mallatt, J. The Nature of Primary Consciousness. A New Synthesis. *Conscious. Cogn.* **2016**, *43*, 113–127. [[CrossRef](#)] [[PubMed](#)]
4. Reggia, J.A.; Katz, G.; Huang, D.-W. What Are the Computational Correlates of Consciousness? *Biol. Inspired Cogn. Archit.* **2016**, *17*, 101–113. [[CrossRef](#)]
5. Lepauvre, A.; Melloni, L. The Search for the Neural Correlate of Consciousness: Progress and Challenges. *Philos. Mind Sci.* **2021**, *2*. [[CrossRef](#)]
6. Wu, W. *The Neuroscience of Consciousness*; The Stanford Encyclopedia of Philosophy: Stanford, CA, USA, 2018.
7. Seth, A.K.; Bayne, T. Theories of Consciousness. *Nat. Rev. Neurosci.* **2022**, *23*, 439–452. [[CrossRef](#)]
8. Tononi, G.; Koch, C. Consciousness: Here, There and Everywhere? *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **2015**, *370*, 20140167. [[CrossRef](#)]
9. Arrabales, R.; Ledezma, A.; Sanchis, A. ConsScale: A Pragmatic Scale for Measuring the Level of Consciousness in Artificial Agents. *J. Conscious. Stud.* **2010**, *17*, 131–164.
10. Birch, J.; Ginsburg, S.; Jablonka, E. Unlimited Associative Learning and the Origins of Consciousness: A Primer and Some Predictions. *Biol. Philos.* **2020**, *35*, 56. [[CrossRef](#)]
11. Panksepp, J. The Periconscious Substrates of Consciousness: Affective States and the Evolutionary Origins of the Self. *J. Conscious. Stud.* **1998**, *5*, 566–582.
12. Chalmers, D. Others Panpsychism and Panprotopsyism. In *Consciousness in the Physical World: Perspectives on Russellian Monism*; Oxford University Press: Oxford, UK, 2015; Volume 246.
13. Frankish, K. Panpsychism and the Depsychologization of Consciousness. In *Aristotelian Society Supplementary Volume*; Oxford University Press: Oxford, UK, 2021; Volume 95, pp. 51–70.
14. Shenker, O. Denialism: What Do the so-Called Consciousness Deniers Deny? *Jerus. Philos. Q.* **2020**, *68*, 307–337.
15. Edelman, D.B.; Seth, A.K. Animal Consciousness: A Synthetic Approach. *Trends Neurosci.* **2009**, *32*, 476–484. [[CrossRef](#)]
16. Metzinger, T. Artificial Suffering: An Argument for a Global Moratorium on Synthetic Phenomenology. *J. Artif. Intell. Conscious.* **2021**, *8*, 43–66. [[CrossRef](#)]
17. Doyon, M. Perception and Normative Self-Consciousness. In *Normativity in Perception*; Doyon, M., Breyer, T., Eds.; Palgrave Macmillan: London, UK, 2015; pp. 38–55; ISBN 9781137377920.
18. Audi, R. Perception and Consciousness. In *Handbook of Epistemology*; Niiniluoto, I., Sintonen, M., Woleński, J., Eds.; Springer: Dordrecht, The Netherlands, 2004; pp. 57–108; ISBN 9781402019869.
19. Cosmelli, D.; Thompson, E. Embodiment or Envatment? Reflections on the Bodily Basis of Consciousness. In *Enaction: Toward a New Paradigm for Cognitive Science*; Di Paolo, E.A., Stewart, J.R., Stewart, J., Gapenne, O., Eds.; MIT Press Limited: Cambridge, MA, USA, 2010; pp. 361–386; ISBN 9780262014601.
20. Nguyen, D.J.; Larson, J.B. Don't Forget About the Body: Exploring the Curricular Possibilities of Embodied Pedagogy. *Innov. High. Educ.* **2015**, *40*, 331–344. [[CrossRef](#)]
21. Glenberg, A.M. Embodiment as a Unifying Perspective for Psychology. *Wiley Interdiscip. Rev. Cogn. Sci.* **2010**, *1*, 586–596. [[CrossRef](#)] [[PubMed](#)]
22. Thompson, E.; Varela, F.J. Radical Embodiment: Neural Dynamics and Consciousness. *Trends Cogn. Sci.* **2001**, *5*, 418–425. [[CrossRef](#)] [[PubMed](#)]
23. Dennett, D. Where Am I? In *Brainstorms: Philosophical Essays on Mind and Psychology*; Bradford Books: Cambridge, MA, USA, 1978; p. 11; ISBN 9780262540377.
24. Rudrauf, D.; Williford, K. Unlimited Associative Learning and the Origins of Consciousness: The Missing Point of View. *Biol. Philos.* **2021**, *36*, 43. [[CrossRef](#)]

25. Ciaunica, A.; Fotopoulou, A. The Touched Self: Psychological and Philosophical Perspectives on Proximal Intersubjectivity and the Self. In *Embodiment, Enaction, and Culture: Investigating the Constitution of the Shared World*; Durt, C., Fuchs, T., Tewes, C., Eds.; MIT Press: Cambridge, MA, USA, 2017; pp. 173–192; ISBN 9780262035552.
26. Juliani, A.; Arulkumaran, K.; Sasai, S.; Kanai, R. On the Link between Conscious Function and General Intelligence in Humans and Machines. *Transactions on Machine Learning Research. arXiv* **2022**, arXiv:2204.05133.
27. Koch, C.; Tsuchiya, N. Attention and Consciousness: Two Distinct Brain Processes. *Trends Cogn. Sci.* **2007**, *11*, 16–22. [[CrossRef](#)]
28. Graziano, M.S.A.; Guterstam, A.; Bio, B.J.; Wilterson, A.I. Toward a Standard Model of Consciousness: Reconciling the Attention Schema, Global Workspace, Higher-Order Thought, and Illusionist Theories. *Cogn. Neuropsychol.* **2020**, *37*, 155–172. [[CrossRef](#)] [[PubMed](#)]
29. Mole, C. Attention and Consciousness. *J. Conscious. Stud.* **2008**, *15*, 86–104.
30. Popat, S.; Winslade, W. While You Were Sleepwalking: Science and Neurobiology of Sleep Disorders & the Enigma of Legal Responsibility of Violence during Parasomnia. *Neuroethics* **2015**, *8*, 203–214. [[PubMed](#)]
31. Dijksterhuis, A.; Aarts, H. Goals, Attention, and (un)consciousness. *Annu. Rev. Psychol.* **2010**, *61*, 467–490. [[CrossRef](#)] [[PubMed](#)]
32. Bello, P.; Bridewell, W. Attention and Consciousness in Intentional Action: Steps Toward Rich Artificial Agency. *J. Artif. Intell. Convers.* **2020**, *07*, 15–24. [[CrossRef](#)]
33. Graziano, M.S.A.; Webb, T.W. The Attention Schema Theory: A Mechanistic Account of Subjective Awareness. *Front. Psychol.* **2015**, *6*, 500. [[CrossRef](#)]
34. Velichkovsky, B.B. Consciousness and Working Memory: Current Trends and Research Perspectives. *Conscious. Cogn.* **2017**, *55*, 35–45. [[CrossRef](#)]
35. Friedman, D.A.; Søvik, E. The Ant Colony as a Test for Scientific Theories of Consciousness. *Synthese* **2021**, *198*, 1457–1480. [[CrossRef](#)]
36. Funahashi, S.; Procyk, E. Editorial: Persistent Activity in the Brain—Functions and Origin. *Front. Neural Circuits* **2021**, *15*, 841451. [[CrossRef](#)]
37. Mashour, G.A.; Roelfsema, P.; Changeux, J.-P.; Dehaene, S. Conscious Processing and the Global Neuronal Workspace Hypothesis. *Neuron* **2020**, *105*, 776–798. [[CrossRef](#)]
38. Dehaene, S.; Naccache, L. Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework. *Cognition* **2001**, *79*, 1–37. [[CrossRef](#)]
39. Baars, B.J. *A Cognitive Theory of Consciousness*; Cambridge University Press: Cambridge, UK, 1988.
40. Dehaene, S.; Kerszberg, M.; Changeux, J.-P. A Neuronal Model of a Global Workspace in Effortful Cognitive Tasks. *Proc. Natl. Acad. Sci. USA* **1998**, *95*, 14529–14534. [[CrossRef](#)]
41. Bensemman, J.; Witbrock, M. The Effects of Implementing Phenomenology in a Deep Neural Network. *Heliyon* **2021**, *7*, e07246. [[CrossRef](#)]
42. Long, M.; Cao, Z.; Wang, J.; Yu, P.S. Learning Multiple Tasks with Multilinear Relationship Networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017*; Curran Associates Inc.: Red Hook, NY, USA, 2017; pp. 1593–1602.
43. Tononi, G. An Information Integration Theory of Consciousness. *BMC Neurosci.* **2004**, *5*, 42. [[CrossRef](#)]
44. Friston, K. Consciousness and Hierarchical Inference. *Neuropsychanalysis* **2013**, *15*, 38–42. [[CrossRef](#)]
45. Seth, A.K. From Unconscious Inference to the Beholder’s Share: Predictive Perception and Human Experience. *Eur. Rev.* **2019**, *27*, 378–410. [[CrossRef](#)]
46. Conway, M.A.; Loveday, C. Remembering, Imagining, False Memories & Personal Meanings. *Conscious. Cogn.* **2015**, *33*, 574–581.
47. Seth, A.K. Interoceptive Inference, Emotion, and the Embodied Self. *Trends Cogn. Sci.* **2013**, *17*, 565–573. [[CrossRef](#)]
48. Fotopoulou, A.; Tsakiris, M. Mentalizing Homeostasis: The Social Origins of Interoceptive Inference. *Neuropsychanalysis* **2017**, *19*, 3–28. [[CrossRef](#)]
49. Sajid, N.; Ball, P.J.; Parr, T.; Friston, K.J. Active Inference: Demystified and Compared. *Neural Comput.* **2021**, *33*, 674–712. [[CrossRef](#)]
50. Cooke, J.E. What Is Consciousness? Integrated Information vs. Inference. *Entropy* **2021**, *23*, 1032. [[CrossRef](#)] [[PubMed](#)]
51. Baddeley, A. Working Memory. *Science* **1992**, *255*, 556–559. [[CrossRef](#)] [[PubMed](#)]
52. Baddeley, A. Working Memory: Theories, Models, and Controversies. *Annu. Rev. Psychol.* **2012**, *63*, 1–29. [[CrossRef](#)]
53. Baddeley, A. The Episodic Buffer: A New Component of Working Memory? *Trends Cogn. Sci.* **2000**, *4*, 417–423. [[CrossRef](#)] [[PubMed](#)]
54. Cowan, N. An Embedded-Processes Model of Working Memory. In *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*; Miyake, A., Ed.; Cambridge University Press: New York, NY, USA, 1999; Volume 506, pp. 62–101.
55. Oberauer, K. Access to Information in Working Memory: Exploring the Focus of Attention. *J. Exp. Psychol. Learn. Mem. Cogn.* **2002**, *28*, 411–421. [[CrossRef](#)] [[PubMed](#)]
56. Engle, R.W. Working Memory Capacity as Executive Attention. *Curr. Dir. Psychol. Sci.* **2002**, *11*, 19–23. [[CrossRef](#)]
57. Baars, B.J.; Franklin, S. How Conscious Experience and Working Memory Interact. *Trends Cogn. Sci.* **2003**, *7*, 166–172. [[CrossRef](#)] [[PubMed](#)]
58. Jacobs, C.; Silvanto, J. How Is Working Memory Content Consciously Experienced? The “Conscious Copy” Model of WM Introspection. *Neurosci. Biobehav. Rev.* **2015**, *55*, 510–519. [[CrossRef](#)]
59. Perky, C.W. An Experimental Study of Imagination. *Am. J. Psychol.* **1910**, *21*, 422–452. [[CrossRef](#)]

60. Craver-Lemley, C.; Reeves, A. Visual Imagery Selectively Reduces Vernier Acuity. *Perception* **1987**, *16*, 599–614. [[CrossRef](#)]
61. Pan, Y.; Lin, B.; Zhao, Y.; Soto, D. Working Memory Biasing of Visual Perception without Awareness. *Atten. Percept. Psychophys.* **2014**, *76*, 2051–2062. [[CrossRef](#)]
62. Bona, S.; Cattaneo, Z.; Vecchi, T.; Soto, D.; Silvanto, J. Metacognition of Visual Short-Term Memory: Dissociation between Objective and Subjective Components of VSTM. *Front. Psychol.* **2013**, *4*, 62. [[CrossRef](#)]
63. Bona, S.; Silvanto, J. Accuracy and Confidence of Visual Short-Term Memory Do Not Go Hand-in-Hand: Behavioral and Neural Dissociations. *PLoS ONE* **2014**, *9*, e90808. [[CrossRef](#)] [[PubMed](#)]
64. Baars, B.J. *In the Theater of Consciousness: The Workspace of the Mind*; Oxford University Press: Oxford, UK, 1997; ISBN 9780195102659.
65. Baars, B.J. Global Workspace Theory of Consciousness: Toward a Cognitive Neuroscience of Human Experience. *Prog. Brain Res.* **2005**, *150*, 45–53. [[PubMed](#)]
66. Franklin, S. A Conscious Artifact? *J. Conscious. Stud.* **2003**, *10*, 47–66.
67. Haikonen, P.O. *Consciousness and Robot Sentience*; In Series on Machine Consciousness; World Scientific: Singapore, 2019; Volume 2.
68. Chan, L.-C.; Latham, A.J. Four Meta-Methods for the Study of Qualia. *Erkenntnis* **2019**, *84*, 145–167. [[CrossRef](#)]
69. Ramachandran, V.S.; Hirstein, W. Three Laws of Qualia: What Neurology Tells Us about the Biological Functions of Consciousness. *J. Conscious. Stud.* **1997**, *4*, 429–457.
70. Kriegel, U. Towards a New Feeling Theory of Emotion. *Eur. J. Philos. Sci.* **2014**, *22*, 420–442. [[CrossRef](#)]
71. Overgaard, M.; Mogensen, J. An Integrative View on Consciousness and Introspection. *Rev. Philos. Psychol.* **2017**, *8*, 129–141. [[CrossRef](#)]
72. Descartes, R. *Discours de la Méthode Pour Bien Conduire sa Raison, et Chercher la Vérité Dans les Sciences*; Hachette et cie: Paris, France, 1637.
73. Peters, M.A.K. Towards Characterizing the Canonical Computations Generating Phenomenal Experience. *Neurosci. Biobehav. Rev.* **2022**, *142*, 104903. [[CrossRef](#)]
74. Proust, J. Metacognition and Metarepresentation: Is a Self-Directed Theory of Mind a Precondition for Metacognition? *Synthese* **2007**, *159*, 271–295. [[CrossRef](#)]
75. Brown, R.; Lau, H.; LeDoux, J.E. Understanding the Higher-Order Approach to Consciousness. *Trends Cogn. Sci.* **2019**, *23*, 754–768. [[CrossRef](#)]
76. Cleeremans, A.; Achoui, D.; Beauny, A.; Keuninckx, L.; Martin, J.-R.; Muñoz-Moldes, S.; Vuillaume, L.; de Heering, A. Learning to Be Conscious. *Trends Cogn. Sci.* **2020**, *24*, 112–123. [[CrossRef](#)]
77. Lau, H.; Rosenthal, D. Empirical Support for Higher-Order Theories of Conscious Awareness. *Trends Cogn. Sci.* **2011**, *15*, 365–373. [[CrossRef](#)]
78. Hinwar, R.P.; Lambert, A.J. Anauralia: The Silent Mind and Its Association with Aphantasia. *Front. Psychol.* **2021**, *12*, 744213. [[CrossRef](#)] [[PubMed](#)]
79. Amit, E.; Hoeflin, C.; Hamzah, N.; Fedorenko, E. An Asymmetrical Relationship between Verbal and Visual Thinking: Converging Evidence from Behavior and fMRI. *Neuroimage* **2017**, *152*, 619–627. [[CrossRef](#)] [[PubMed](#)]
80. Lennon, P. Aphantasia and Conscious Thought. *Oxf. Stud. Philos. Mind* **2022**, *3*, 131.
81. Happé, F. Theory of Mind and the Self. *Ann. N. Y. Acad. Sci.* **2003**, *1001*, 134–144. [[CrossRef](#)]
82. Nichols, S.; Stich, S. How to Read Your Own Mind: A Cognitive Theory of Self-Consciousness. *Conscious. New Philos. Essays* **2003**, *1*, 157–200.
83. Michaelian, K. *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*; MIT Press: Cambridge, MA, USA, 2016.
84. Prosser, S.; Recanati, F.; Campbell, J.; Coliva, A.; Folescu, M.; Higginbotham, J.; Ismael, J.; Longuenesse, B.; Morgan, D.; O'Brien, L.; et al. *Immunity to Error through Misidentification: New Essays*; Cambridge University Press: Cambridge, UK, 2012; ISBN 9781139043274.
85. Seth, A. *Being You: A New Science of Consciousness*; Penguin: New York, NY, USA, 2021; ISBN 9781524742874.
86. Kriegel, U. *Subjective Consciousness: A Self-Representational Theory*; OUP Oxford: Oxford, UK, 2009; ISBN 9780191610059.
87. Gennaro, R.J. Consciousness and Self-Consciousness. In *Consciousness and Self-Consciousness*; John Benjamins Publishing Company: Amsterdam, The Netherlands, 1996; pp. 1–230.
88. Carruthers, P. *Consciousness: Essays from a Higher-Order Perspective*; Oxford University Press: Oxford, UK, 2005; ISBN 9780191602597.
89. Korkmaz, B. Theory of Mind and Neurodevelopmental Disorders of Childhood. *Pediatr. Res.* **2011**, *69*, 101–108. [[CrossRef](#)] [[PubMed](#)]
90. Capps, L.; Kehres, J.; Sigman, M. Conversational Abilities among Children with Autism and Children with Developmental Delays. *Autism* **1998**, *2*, 325–344. [[CrossRef](#)]
91. Happé, F.; Frith, U. The Neuropsychology of Autism. *Brain* **1996**, *119 Pt 4*, 1377–1400. [[CrossRef](#)]
92. Senju, A. Spontaneous Theory of Mind and Its Absence in Autism Spectrum Disorders. *Neuroscientist* **2012**, *18*, 108–113. [[CrossRef](#)] [[PubMed](#)]
93. Bora, E.; Bartholomeusz, C.; Pantelis, C. Meta-Analysis of Theory of Mind (ToM) Impairment in Bipolar Disorder. *Psychol. Med.* **2016**, *46*, 253–264. [[CrossRef](#)] [[PubMed](#)]

94. Sprong, M.; Schothorst, P.; Vos, E.; Hox, J.; Van Engeland, H. Theory of Mind in Schizophrenia: Meta-Analysis. *Br. J. Psychiatry* **2007**, *191*, 5–13. [[CrossRef](#)] [[PubMed](#)]
95. Muller, F.; Simion, A.; Reviriego, E.; Galera, C.; Mazaux, J.-M.; Barat, M.; Joseph, P.-A. Exploring Theory of Mind after Severe Traumatic Brain Injury. *Cortex* **2010**, *46*, 1088–1099. [[CrossRef](#)]
96. Yirmiya, N.; Erel, O.; Shaked, M.; Solomonica-Levi, D. Meta-Analyses Comparing Theory of Mind Abilities of Individuals with Autism, Individuals with Mental Retardation, and Normally Developing Individuals. *Psychol. Bull.* **1998**, *124*, 283–307. [[CrossRef](#)]
97. Rathbone, C.J.; Moulin, C.J.A.; Conway, M.A. Autobiographical Memory and Amnesia: Using Conceptual Knowledge to Ground the Self. *Neurocase* **2009**, *15*, 405–418. [[CrossRef](#)]
98. Wilson, B.A.; Wearing, D. Prisoner of Consciousness: A State of Just Awakening Following Herpes Simplex Encephalitis. In *Broken Memories: Case Studies in Memory Impairment*; Campbell, R., Ed.; Blackwell Publishing: Hoboken, NJ, USA, 1995; Volume 444, pp. 14–30.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.