



# Cognitive Robots and the Conscious Mind: A Review of the Global Workspace Theory

Alessandro Signa<sup>1,2</sup>  · Antonio Chella<sup>1,3</sup> · Manuel Gentile<sup>2,4</sup>

Accepted: 3 March 2021 / Published online: 6 April 2021  
© The Author(s) 2021

## Abstract

**Purpose of Review** The theory of consciousness is a subject that has kept scholars and researchers challenged for centuries. Even today it is not possible to define what consciousness is. This has led to the theorization of different models of consciousness. Starting from Baars' Global Workspace Theory, this paper examines the models of cognitive architectures that are inspired by it and that can represent a reference point in the field of robot consciousness.

**Recent Findings** Global Workspace Theory has recently been ranked as the most promising theory in its field. However, this is not reflected in the mathematical models of cognitive architectures inspired by it: they are few, and most of them are a decade old, which is too long compared to the speed at which artificial intelligence techniques are improving. Indeed, recent publications propose simple mathematical models that are well designed for computer implementation.

**Summary** In this paper, we introduce an overview of consciousness and robot consciousness, with some interesting insights from the literature. Then we focus on Baars' Global Workspace Theory, presenting it briefly. Finally, we report on the most interesting and promising models of cognitive architectures that implement it, describing their peculiarities.

**Keywords** Global workspace theory · Robot Consciousness · Cognitive Architecture

## Introduction

Despite centuries of interest from philosophers and scholars, an unambiguous definition of the term consciousness remains lacking. Consequently, AI researchers are not unanimous as to the defining characteristics of artificial consciousness and/or what a robot should do to

exhibit conscious behavior [1]. The impossibility to replicate all the aspects of a thinking and conscious subject has led to the conviction that it is essential to know exactly how the brain works in order to implement such a reproduction. It is perhaps due to this conviction that 1994 marked a turning point for research in this field: in April of that year in Tucson, Arizona, the first international conference on the theme of consciousness (“Toward a Science of Consciousness”) took place whereby not only philosophers, but also neuroscientists, physicists, and cognitive scientists were present. Since then, new and multiple criteria have been advanced for the evaluation of consciousness and for the development of theories for describing consciousness. Among them, in this paper, we will focus on one specific theory, namely Global Workspace Theory postulated by Baars [2]. We will summarize its content and analyze some of the architectures inspired by it. Some of them are already widespread and used for several years, such as the LIDA architecture [3–6] and Global Neuronal Workspace theory [7–9], while others are more modern and niche, such as CERA-CRANIUM [10] and Conscious Touring Machine [11•].

---

This article is part of the Topical Collection on *Cognitive Robotics*

---

✉ Alessandro Signa  
alessandro.signa@unipa.it; alessandro.signa@itd.cnr.it

Antonio Chella  
antonio.chella@unipa.it

Manuel Gentile  
manuel.gentile@itd.cnr.it; manuel.gentile@unito.it

- <sup>1</sup> Department of Engineering, University of Palermo, Palermo, Italy
- <sup>2</sup> Institute for Educational Technology, National Research Council of Italy, Palermo, Italy
- <sup>3</sup> Institute for high performance computing and networking, National Research Council of Italy, Palermo, Italy
- <sup>4</sup> Dipartimento di Informatica, Università di Torino, Turin, Italy

## Robot Consciousness

Because it is impossible to define consciousness unambiguously, we are also precluded from defining robot consciousness precisely.

The various attempts at defining consciousness given by scholars, however complete, are often unsuitable as a first step in modeling artificial consciousness. For example, consider, on the one hand, the definition of Consciousness Tetrad given by Singh and Singh [12]. They understand consciousness in four different forms: the “Default,” which distinguishes the living from the non-living, the “Aware,” which distinguishes the state of awake, as opposed to the state of sleep or coma, the “Operational,” which concerns motor, sensory, cognitive, and emotional awareness, and finally, the “Exalted,” which refers to the state of connection with the divine or the soul.

On the other hand, a more recent analysis proposed by Dehaene et al. [13•] offers a definition of consciousness through a model that can be adapted to the language of computer science. Starting from the functioning of the human brain, Dehaene suggests that the term “consciousness” actually corresponds to two different levels of information processing in the human brain: the first, identified as C1, is the level of consciousness responsible for the selection of information of global relevance and worthy of further processing and reasoning; the second, defined as C2, is the level of consciousness related to the autonomous monitoring of the former reasoning, which then leads to a subjective sense of correctness or error. Dehaene argues that although C1 and C2 are the characteristic levels of consciousness, current architectures mainly implement the unconscious processes that occur in the brain, defined as C0. These C0 unconscious processes consist of complex calculations and inferences that occur continuously and in parallel within various areas of the brain. For the survival of the organism, however, it is necessary that these multiple unconscious information flows converge into a single information belonging to the C1 level in order to lead the actions of the organism in a coherent way. To such purpose, Dehaene embraces the theory proposed by Baars [2, 14] relative to the Global Workspace through which it is possible to elevate an unconscious elaboration (C0) up to make it become conscious (C1), sharing it with the other modules of the architecture in such a way that the immediately succeeding actions are executed according to it. Dehaene therefore defines as “conscious” any information that is promoted from C0 to C1 and triggers the decision-making process.

## Global Workspace Theory

An online survey conducted in 2018 on academics and consciousness experts showed that the Global Workspace Theory was perceived as the most promising consciousness theory [15]. The Global Workspace Theory (GWT) is a cognitive

architecture first proposed by Baars in 1988 [2]. It has been over the years and remains a reference point and source of inspiration for further artificial consciousness theories and implementations of modern cognitive architectures [3–11, 16–20].

GWT considers the nervous system as a set of multiple specialized processors working in parallel. The most peripheral processors carry information from sensors and transmit it to other processors responsible for interpreting and processing that sensory information. The processor, or coalition of processors, that carries the most significant information gains access to the Global Workspace that consists of a fleeting memory area in which only one piece of information content can reside at a time. This content is then propagated to all other processors in order to recruit those necessary to act or react on the basis of the current situation. In this theory, the state of consciousness is associated with the memory of the Global Workspace and the information that resides in it. The unconscious processes of the mind instead are associated with all the other specialized processors that elaborate the sensory information, that compete to win the Global Workspace and that subsequently are recruited to perform the actions in response to the current state of the Global Workspace. The criteria for defining the winner of the competition between the unconscious processors are defined by the current context, intended as the achievement of a goal or the fulfillment of an emotion. According to these criteria, the information conveyed by an unconscious processor could be more or less suitable for access to the Global Workspace. The phases of collaboration and competition of the GWT strategy make it dynamic enough to be particularly successful in facing unknown problems that would otherwise be faced in an ineffective or partial way by isolated unconscious processors [16].

In the following section we will present works related to the most relevant theories, models, and cognitive architectures that are inspired by GWT, that is, starting from one or more assumptions in common with Baars’ theory.

## GWT-Inspired Architectures

The first of the models based on GWT we discuss is the Learning Intelligent Distribution Agent (LIDA) model proposed by Franklin [3–6], which implements and flashes out many important aspects theorized by Baars in GWT. This model together with the relative computational architecture constitute the LIDA cognitive cycle. A cognitive cycle is defined as the process by which an agent, living or artificial, perceives through the senses the surrounding environment and creates or updates its representation of it, then processes that information, and finally acts accordingly. Each cycle is therefore divided into three phases. The first phase is the comprehension phase in which the external stimuli activate the functionalities of the sensorial memory of the agent that

generates a percept, such percept is elaborated by an appropriate Workspace that generates local associations correlated to the percept. In this way, the model of the agent relative to the situation and the actual events is produced. The second phase is the consciousness phase in which coalitions of percepts are sent to the Global Workspace where they compete to select the most relevant and worthy to become the content of consciousness to be sent globally to recruit processors and resources. Finally, the third phase is the action selection phase in which action patterns are recruited that are compatible with what the Global Workspace has sent globally, and these possible actions compete in order to bring out the most appropriate one for the current cognitive cycle. This phase is also a learning phase because associations between action patterns, external contexts, and expected and achieved outcomes are learned or reinforced. At this point the selected action pattern is sent to sensory-motor memory in order to activate the behavior, to close the current cycle, and to initiate the next one. These cognitive cycles follow one another continuously throughout the agent's lifetime. The Cognitive Computing Research Group of the University of Memphis led by Dr. Franklin makes available the open source LIDA framework implementation through its website. The LIDA software framework consists of a generic and customizable Java implementation of the LIDA model, easily adjustable to specific needs and domain via XML files or modifying the source code directly. The framework is also provided of several tools, such as a GUI (to let the user monitor running tasks, parameters, etc.) and a logging functionality to record the occurrences of noteworthy actions [6].

In 2005, Shanahan proposed a cognitive architecture that combines the GWT with an internal simulation model [20]. This architecture, which we will refer to later in this paper as Internal Simulation with Global Workspace (IS-GW), is deeply inspired by the human brain and consists of two interacting sensorimotor loops. The first-order loop involves the sensory and motor centers and determines an immediate response to stimuli coming from the external world. This response, however, is mediated by the action of the BG (a component of the architecture that models the functioning of the basal ganglia), which evaluates the relevance of the response by comparing it with other possible responses. These alternative responses are generated (“imagined”) by the other loop, the higher-order loop, which on the basis of current sensory stimuli allows to plan and anticipate actions. Exploiting the previously learned associations, it simulates internally and in parallel the possible scenarios in terms of actions performed and their consequences. Within the higher-order loop, it is the Am (a component of the architecture that models the amygdala) that associates a consequence with each action in the form of reward or punishment. Among the possible actions, the BG chooses and executes the most relevant one for the current context, which is the one exceeding a given threshold of relevance. This is

where the typical GWT competition takes place. The winning action is then broadcasted to all the other elements of the architecture. Despite the unfavorable circumstances in terms of available hardware at that time, Shanahan was also able to provide an implementation of such a system [20]: although only by means of a prototype, the validity of the proposed architecture was demonstrated, which in fact proved the ability of generating a better motor response to the current circumstance (in terms of number of actions carried out) because it was planned due to the intervention of BG and Am. The simulation scenario consisted of a configuration of colored cylinders and a simple robot, which ran the prototypical implementation of IS-GW. The robot, learning from its experience in that scenario, was able to elaborate the outcome of the next hypothetical sensory states; thus, it managed to minimize the required actions to reach the desired cylinder.

Arrabales in 2009 presented CERA-CRANIUM [10], a cognitive architecture whose implementation [21] in 2010 won the 2 K BotPrize Contest, which is a modified version of the Turing Test, to evaluate the ability of computer game playing agents to imitate human gamer behaviors and fool judges into thinking they are human. The game was the multiplayer FPS Unreal Tournament 2004 [22]. There, the Arrabales bot was considered by judges the most human behaving bot. They noticed that, on the one hand, there were several characteristics indicating that a player was a bot (i.e., missing/stupid behavior and low aggression), on the other hand, there were other characteristics indicating that a player was a human (i.e., high aggression, adapting to opponent and good tactics). During that contest, Arrabales' bot fooled the judges 32% of the time and was incorrectly identified as human, while one human player was identified as a machine 35% of time. CERA-CRANIUM (CC) is an architecture for software control and is mainly inspired by GWT. The main difference between CC and other control architectures is that CC does not focus on the details of what the next action to be performed is, but rather it processes the high-level conscious content of the next motor content. The motor context is then responsible for processing this high-level information and deciding how to act. CC consists of CERA, which is a software control architecture divided into four layers, the middle two layers make up CRANIUM, which is the execution environment for creating and managing large amounts of parallel processes in a shared memory. The bottom layer of the architecture is the sensory and motor layer that obtains data from sensors and drives motors to perform actions. The next layer is the physical layer that manipulates the data from sensors and actuators and prepares them for the next layer. Higher up is the mission layer that processes the received data, interpreting them on the basis of the current goals of the agent. In this layer, the information content becomes a percept and is directly related to the mission of the agent, which can be decomposed into different goals. The core layer is the

uppermost layer and is responsible for managing the highest-level cognitive functions. This subdivision into layers gives the architecture a remarkable modularity: the physical layer is specific for the agent that implements it but independent from the current mission, the mission layer is independent from the agent and specific for the mission to be carried out, the core layer is finally independent both from the agent and from the mission, and this allows a straightforward reuse of it. CRANIUM, which is realized in the physical layer and mission layer, has two communicating workspaces in which the processes of competition and collaboration take place, with multiple specialized processors as key players. All processing that occurs in these two layers are modulated by the core layer through special out-of-band commands. These commands indicate which should be the target sensory and motor context, thus influencing the activation of the specialized processors and the results they obtain.

A recently published model by Blum & Blum [11•] aims to study consciousness from the point of view of theoretical computer science. For this purpose, they defined formally a conscious turing machine (CTM). Consciousness in this model is contained in the short term memory (STM), which holds exactly one chunk of information at a time to be sent globally. The long term memory on the other hand is the set of all processors in the model; their processing is unconscious. There are no other processors in the CTM besides these just defined. At each clock tick, each processor generates an information chunk that competes with the others for access to the STM. The process of competition takes the name of up-tree competition and consists of a series of nodes of calculation that according to the criteria of competition let the chunk advance. Then the winning chunk is broadcasted from the short term memory to the long term memory processors via the down tree, which consists of a direct connection between the STM and each processor. In CTM there is also the creation of links between the various processors. These are formed over time if there are sufficiently frequent conscious communications between two processors, that is, requests and responses that pass through the STM involving them. By doing so, unconscious communications are generated by which processors are able to influence each other without passing through the STM. In the CTM, the elaboration and the coding of the information coming from the sensors and directed to the actuators happens in the input map and output map, respectively. The conscious awareness of the CTM consists therefore in the reception by all the processors of the chunk that has won the competition. Consequently, it is possible to define the stream of consciousness as the sequence of chunks broadcasted by the STM in a given time interval.

Dehaene et al. [7–9] start from the same theoretical basis as Baars [2] to postulate a new theory: the Global

Neuronal Workspace (GNW) Theory. The common concept between the two theories is that conscious access is given by the global availability of the informational load and that what is perceived as conscious is the selection, amplification, and global broadcasting of relevant information. The GNW theory proposes two main computational areas in the brain: the first is constituted by the localized and specialized processors that elaborate and evaluate the sensorial and motorial information; the second is precisely the global neuronal workspace, that is, a set of distributed neurons characterized by their ability to communicate with each other even if in distant areas. Given the distribution of these neurons, this theory does not consider the presence of a single central unit in which resides the conscious information to be distributed to the rest of the processors, but the presence of several agglomerates of neurons of the GNW scattered in different areas in order to create a distributed workspace. The consciousness of information is reached when many processors afferent to the GNW converge to a coherent state. The neurons of the GNW are able to receive information from the various processors and transfer to them the results of their elaborations; in this way, they accumulate information in a competitive way in order to eventually make this information reach the state of consciousness. Initially the stimulus goes up the hierarchy of processors (bottom-up) in an unconscious way, then if the stimulus is evaluated as adequate, it is amplified (top-down) and maintained by a subset of the neurons of the GNW. The whole network of the GNW is connected in such a way that only one conscious information at a time can be effective. It is possible that a non-linear activation of some neurons of the GNW occurs; this phenomenon is called “ignition”: the neurons that contain the current conscious content are suddenly and coherently activated, while all the others are deactivated or inhibited. Ignition can be triggered by an external stimulus during the cognitive process or occur spontaneously from within during a state of rest. This state of activation is ruled by appropriate vigilance signals that are strong enough to control transitions between waking (GNW active) and resting (GNW inactive) states. Several concrete computer simulations of the GNW architecture were also proposed by Dehaene et al. [7–9]. Considering for example the simulation regarding the effortful cognitive tasks can help to understand the vigilance signals’ role [7]. In this computer simulation the network is fed with routines made of several versions of the word–color Stroop tasks [23] and the system is provided with specialized units to accomplish the Stroop tasks. During the routines consisting of effortless Stroop tasks, the vigilance was quickly extinguished putting the GNW in the inactive state, due to the absence of errors. Instead, during the routines consisting of effortful Stroop

tasks, the errors and negative rewards led the vigilance to activate the GNW: its aim is to search the rules achieving the stable pattern to the correct performance. After a few effortful tasks trials, the network successfully learned the Stroop test on the basis of realistic neuronal processes.

## Discussion

The theories and models that have been presented are those inspired by the GWT and which represent, according to the authors, excellent starting points for further implementations and/or theories. Table 1 presents a summary of their key characteristics. The column “Year” shows the year in which the respective theory was first published. The column “Maturity” shows the degree of completion of each research and study related to the theory. In that column, “Ready for Implementation” means that the model is only presented theoretically with sufficient details, while “Implementation Available” means that there are working computer implementations related to the architecture. The column “Peculiarities” highlights the key characteristics of each architecture.

However, in some ways it may be partially incorrect to consider these four architectures at the same level. In fact, it should be pointed out that this selection contains the models inspired by Baars’ GWT, which are the most relevant in terms of the results obtained, their popularity, and future potential. Among these, however, the GNW also represents a theoretical model in its own right: its architecture has inspired subsequent theories such as those of Shanahan [24, 25] and those of Whyte related to the predictive global neuronal workspace

(PGNW) [26, 27]. In light of this, it is possible to argue that, although Baars is the founder of GWT which is today the predominant theory in this field, Dehaene’s architecture of GNW, being more specific [28] and less high-level, would be a proper variation of GWT.

## Conclusions

The Global Workspace Theory has been for years, and still is, a reference for those looking for a model of consciousness, whether we are talking about neuroscience or computer science. The appreciation that this theory has received over the years has led the scientific community to pay increasing attention to it, to the point of creating architectures based on it or inspired by it. Despite the fact that more than 20 years have passed since its theorization, GWT continues to inspire new and interesting models that are still waiting to be implemented [11•], while for those which already exist, the field of application is very wide, for example, from the control of NPCs in video games [21] to the medical field through clinical simulations [9]. While the theories remain robust and valid, the remarkable progress made in recent years in the fields of artificial intelligence, neural networks, and deep learning opens the way for new and modern implementations related to them. Therefore, well known and less known architectures have been reported in this paper: the intention is to bring attention to them in order to stimulate more working groups to use the new software and hardware tools to give concrete (and modern) life to these valid models.

**Table 1** Summary of presented GWT-inspired architectures

Name	Year	Maturity	Peculiarities
LIDA	2009	Implementation available	Accurately implements many GWT aspects. It is based on cognitive cycles of comprehension, consciousness, and action.
IS-GW	2005	Implementation available	Two sensorimotor loops by means of which the action to be performed is selected also from those simulated (“imagined”) internally by the model
CERA-CRANIUM	2009	Implementation available	Four hierarchical layers, two workspaces. Consciousness is at the top layer and knows only a high level representation about the next action
CTM - Conscious Touring Machine	2020	Ready for implementation	All processors in the model are unconscious. Local links can be generated between processors to speed up the processing of certain requests
GNW - Global Neuronal Workspace	1998	Implementation available	Several agglomerates of neurons scattered in different areas in order to create a distributed workspace rather than a single central unit



**Funding** Open access funding provided by Università degli Studi di Palermo within the CRUI-CARE Agreement.

## Declarations

**Human and Animal Rights and Informed Consent** This article does not contain any studies with human or animal subjects performed by any of the authors.

**Conflict of Interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

Papers of particular interest, published recently, have been highlighted as:

- Of importance

- Chella A, Cangelosi A, Metta G, Bringsjord S. Editorial: Consciousness in humanoid robots. *Front Robotics AI*. 2019;6. <https://doi.org/10.3389/frobt.2019.00017>.
- Baars BA. *Cognitive theory of consciousness*. Cambridge Univ Press; 1988.
- Baars B, Franklin S, Consciousness I. Computational: the LIDA model of global workspace theory. *Int J Machine Consciousness*. 2009;01:23–32. <https://doi.org/10.1142/S1793843009000050>.
- Franklin S, Strain S, Snider J, McCall R, Faghihi U. Global workspace theory, its LIDA model and the underlying neuroscience. *Biol Inspired Cognitive Architect*. 2012;1:32–43. <https://doi.org/10.1016/j.bica.2012.04.001>.
- Franklin S, Madl T, D'Mello S, Snider JLIDA. A systems-level Architecture for cognition, emotion, and learning. *IEEE Trans Autonomous Mental Develop*. 2014;6:19–41. <https://doi.org/10.1109/TAMD.2013.2277589>.
- Snider J, McCall R, Franklin S. The LIDA framework as a general tool for AGI. *Artificial Gen Intell*. 2011;133–42. [https://doi.org/10.1007/978-3-642-22887-2\\_14](https://doi.org/10.1007/978-3-642-22887-2_14).
- Dehaene S, Kerszberg M, Changeux J. A neuronal model of a global workspace in effortful cognitive tasks. *Proc Natl Acad Sci*. 1998;95:14529–34. <https://doi.org/10.1073/pnas.95.24.14529>.
- Dehaene S, Sergent C, Changeux J. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proc Natl Acad Sci*. 2003;100:8520–5. <https://doi.org/10.1073/pnas.1332574100>.
- Dehaene S, Changeux J, Naccache L. The global neuronal workspace model of conscious access: from neuronal architectures to clinical applications. *Res Perspect Neurosci*. 2011;55–84. [https://doi.org/10.1007/978-3-642-18015-6\\_4](https://doi.org/10.1007/978-3-642-18015-6_4).
- Arrabales R, Ledezma Espino A, Sanchis de Miguel A. CERA-CRANIUM: a test bed for machine consciousness research. 2009.
- Blum M, Blum L. A theoretical computer science perspective on consciousness. *arXiv:2011.09850 [cs.AI]* <https://arxiv.org/abs/2011.09850> 2020 **This study theorise a simple, clean and well designed architecture of consciousness from the theoretical computer science point of view.**
- Singh A, Singh S. Brain-mind dyad, human experience, the consciousness tetrad and lattice of mental operations: and further. The need to integrate knowledge from diverse disciplines. *Mens Sana Monographs*. 2011;9:6. <https://doi.org/10.4103/0973-1229.77412>.
- Dehaene S, Lau H, Kouider S. What is consciousness, and could machines have it? *Science*. 2017;358:486–92. <https://doi.org/10.1126/science.aan8871> **This study correlates the neural science aspect of consciousness with the current machine implementations of it outlining their incongruities.**
- Baars B. Theater of consciousness. 1997. <https://doi.org/10.1093/acprof:oso/9780195102659.001.1>.
- Michel M, Fleming S, Lau H, Lee A, Martinez-Conde S, Passingham R, et al. An informal internet survey on the current state of Consciousness science. *Front Psychol*. 2018;9. <https://doi.org/10.3389/fpsyg.2018.02134>.
- Baars B. The conscious access hypothesis: origins and recent evidence. *Trends Cognitive Sci*. 2002;6:47–52. [https://doi.org/10.1016/S1364-6613\(00\)01819-2](https://doi.org/10.1016/S1364-6613(00)01819-2).
- Newman J, Baars B, Cho SA. Neural global workspace model for conscious attention. *Neural Netw*. 1997;10:1195–206. [https://doi.org/10.1016/s0893-6080\(97\)00060-9](https://doi.org/10.1016/s0893-6080(97)00060-9).
- Bao C, Fountas Z, Olugbade T, Bianchi-Berthouze N. Multimodal Data Fusion based on the Global Workspace Theory. *Proc 2020 Int Conf Multimod Interact* 2020;414–422.
- Ng K, Du Z, Ng G. DSO Cognitive Architecture: Implementation and Validation of the Global Workspace Enhancement. *Artificial Gen Intell* 2018;151–161. [https://doi.org/10.1007/978-3-319-97676-1\\_15](https://doi.org/10.1007/978-3-319-97676-1_15).
- Shanahan M. A cognitive architecture that combines internal simulation with a global workspace. *Conscious Cognit*. 2006;15:433–49. <https://doi.org/10.1016/j.concog.2005.11.005>.
- Arrabales R. Inner speech generation in a video game non-player character: from explanation to self? *Intl J Mach Consciousness*. 2012;04:367–81. <https://doi.org/10.1142/S1793843012400215>.
- Hingston PA. Turing test for computer game bots. *IEEE Trans Computation Intell AI Games*. 2009;1:169–86. <https://doi.org/10.1109/TCIAIG.2009.2032534>.
- MacLeod C. Half a century of research on the Stroop effect: an integrative review. *Psychol Bull*. 1991;109:163–203. <https://doi.org/10.1037/0033-2909.109.2.163>.
- Shanahan M. A spiking neuron model of cortical broadcast and competition. *Conscious Cognit*. 2008;17:288–303. <https://doi.org/10.1016/j.concog.2006.12.005>.
- Connor D, Shanahan M. A simulated global neuronal workspace with stochastic wiring. In *AAAI Fall Symposium: AI and Consciousness*. 2007;43–48. <https://www.aaai.org/Library/Symposia/Fall/2007/fs07-01-008.php>

26. Whyte C. Integrating the global neuronal workspace into the framework of predictive processing: towards a working hypothesis. *Conscious Cogn.* 2019;73:102763. <https://doi.org/10.1016/j.concog.2019.102763>.
27. Whyte C, Smith R. The predictive global neuronal workspace: a formal active inference model of visual consciousness. *Progress in Neurobiology.* 2020;101918. <https://doi.org/10.1016/j.pneurobio.2020.101918>.
28. Prakash R, Prakash O, Prakash S, Abhishek P, Gandotra S. Global workspace model of consciousness and its electromagnetic correlates. *Ann Indian Acad Neurol.* 2008;11:146. <https://doi.org/10.4103/0972-2327.42933>.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.