

Body Schema in Autonomous Agents

Zachariah A. Neemeh

*Department of Philosophy and Institute for Intelligent Systems, The University of Memphis,
3720 Alumni Avenue, Memphis, Tennessee 38152, United States of America
zaneemeh@memphis.edu*

Christian Kronsted

*Department of Philosophy and Institute for Intelligent Systems, The University of Memphis,
3720 Alumni Avenue, Memphis, Tennessee 38152, United States of America
christian.k@memphis.edu*

Sean Kugele

*Computer Science Department and Institute for Intelligent Systems, The University of Memphis,
3720 Alumni Avenue, Memphis, Tennessee 38152, United States of America
skugele@memphis.edu*

Stan Franklin

*Computer Science Department and Institute for Intelligent Systems, The University of Memphis,
3720 Alumni Avenue, Memphis, Tennessee 38152, United States of America
franklin@memphis.edu*

Received (Day Month Year)

Revised (Day Month Year)

A body schema is an agent's model of its own body that enables it to act on affordances in the environment. This article presents a body schema system for the LIDA cognitive architecture. LIDA is a conceptual and computational implementation of Global Workspace Theory, also integrating other theories from neuroscience and psychology. This article contends that the 'body schema' should be split into three separate functions based on the functional role of consciousness in Global Workspace Theory. There is 1) an online model of the agent's effectors and effector variables (*Current Body Schema*), 2) a long-term, recognitional storage of embodied capacities for action and affordances (*Habitual Body Schema*), and 3) "dorsal" stream information feeding directly from early perception to sensorimotor processes (*Online Body Schema*). This article then discusses how the LIDA model of the body schema explains several experiments in psychology and ethology.

Keywords: Body Schema; Embodied Cognition; Global Workspace Theory; LIDA.

1. Introduction

An autonomous agent is a "system situated within, and a part of, an environment, that senses that environment and acts on it, over time, in pursuit of its own agenda, and so as to

affect what it senses in the future” [Franklin & Graesser, 1997, p. 25]. An agent that is situated or embedded in an environment must perceive affordances that are relevant to its own embodied capacities. For example, a cat does not perceive a graspability^a affordance when it sees a mug with a handle, while a human can. An agent must also pursue an agenda that is appropriate to its embodiment. A human cannot pursue a task of flying around the sky and catching prey to eat, while a bat can. Although humans have had the option of using aircraft or gliders to fly since the 20th century, this is a highly costly, restrictive, and inefficient means of procuring food. Humans do not have the right kind of embodiment or effectors to make flying a reasonable method of finding nourishment, yet this comes naturally to bats who do have the right bodies for the task.

In order to perceive affordances, and pursue an agenda appropriate to the agent’s embodiment, the agent uses a *body schema*, which is a model of its own embodiment [Haggard & Wolpert, 2005]. The model of the body includes its embodied capacities, properties of its body, and the current positions of its body, limbs, and digits (or whatever the agent’s effectors may be). It is not the entire body that is modeled. The agent does not need to model its digestive system, or its own nervous system functions. It needs a model of everything relevant to its performance of actions in the world, such as its posture, effectors and size. “The primary function of an autonomous agent is to continually and iteratively answer the question, ‘What do I do next?’” [Franklin, 1995, p. 412]. The body schema is a central tool by means of which an agent answers that question. The affordances that the agent perceives, the actions it takes, and the agenda it pursues are all relative to this body schema. This is consistent with the enactive approach, where the agent actively brings forth or enacts its world, based on its autonomy but is structurally coupled to that world, rather than passively representing it [Franklin, 1995; Varela *et al.*, 1991/2016].

The body schema must not be confused with the body image, especially given the deep historical conflation of the two terms [Gallagher, 2005]. Body image can refer to two related ideas. First, body image includes semantic representations of one’s attractiveness, strength (e.g., being muscular), weight (e.g., overweight, underweight), race, and similar socially-relevant judgments about one’s own body. It is in this sense that the term is typically used in social psychology. Disorders of this body image include anorexia, bulimia, and body dysmorphic disorder, and muscle dysmorphia. This sense of body image has been primarily studied in humans. Second, body image can refer to the conscious perception of the body. Body schema, on the other hand, is largely kinesthetic or proprioceptive. Its disorders include deafferentation, phantom limb syndrome, and alien hand syndrome.

While the literature on body schema tends to speak about it as a unified construct, we believe the body schema serves three quite different functions (see Fig. 1). The *Current Body Schema* is a model of the agent’s effectors and effector variables, such as current effector positions. For example, for a human agent, the Current Body Schema would model

^a Affordances have traditionally been marked off linguistically by the postposition *-ability*: e.g., grasp-ability, turn-ability, or bite-ability [Gibson, 1979/2015].

bodily posture, orientation, limb, and digit positions, as well as those of effectors used in language production such as the tongue and lips. This model is continually updated and, although it is typically not conscious, if the agent attends to it, parts of it can be brought to consciousness. The *Habitual Body Schema*, on the other hand, is a unified collection of the agent’s innate and learned bodily capacities for action and affordances. For example, a human agent may know it has the ability to ski or turn a door handle, and it may know several properties of its body such as how heavy it is (e.g., “Can I jump across these two rocks?”) or how tall it is (e.g., “Can I fit through this doorway?”). This collection of embodied capacities for action is a form of stored long-term memory that can be cued and regularly updated. Finally, the *Online Body Schema* is realized as “dorsal” stream information in early perceptual processes, which is never-conscious [Franklin & Baars, 2010], and directly modulates sensorimotor processes.^b *Body image* refers to other semantic and visual aspects of the body, and is not part of the body schema.

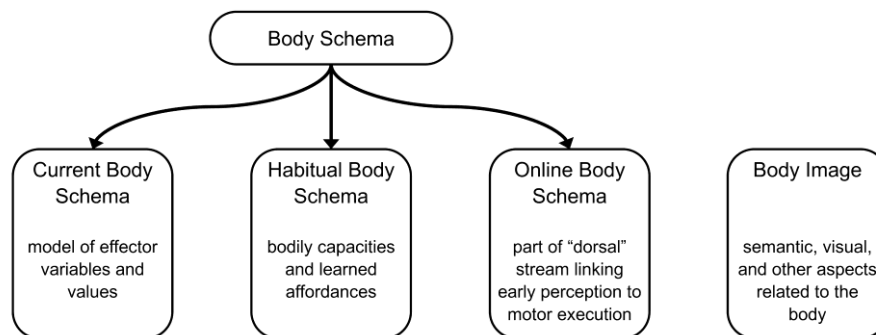


Fig. 1. The three functions of the body schema, and the body image. The literature on the body schema does not typically distinguish between these different functions of the body schema. This fine-grained analysis allows us to account for a broader range of cognitive phenomena.

Our primary contribution in this paper is to outline a conceptual implementation of the body schema in LIDA, an embodied and situated cognitive architecture. Additionally, we distinguish between three separate functions for what the literature identifies as ‘*the* body schema.’ This distinction is based on the functional role of consciousness in Global Workspace Theory, and we believe our distinction ought generally to hold beyond the specific modeling concerns of LIDA. We also discuss how the LIDA model of the body schema helps to explain several experiments in psychology and ethology.

In contrast to cognitive models that seek to explain specific aspects of cognition, such as working memory [Baddeley, 2007], LIDA is a comprehensive, systems-level, cognitive

^b In humans, this processing pathway is not literally dorsal, but it is analogous to the dorsal visual stream (see Sec. 1.1). LIDA’s dorsal stream is a sensory pathway from Sensory Memory to the Sensory Motor System (see Fig. 3).

model and agent architecture that strives to be a “unified theory of cognition” [Newell, 1990]. While a disadvantage of doing so is that it makes the model highly complex, it has the advantage of integrating a broad range of cognitive phenomena, while also specifying how they function interdependently. LIDA implements and fleshes out significant portions of the Global Workspace Theory of consciousness [Baars, 1988, 2019], as well as several other psychological and neuroscientific theories including the perception-action cycle [Cutsuridis *et al.*, 2011; Dijkstra *et al.*, 1994; Freeman, 2002; Neisser, 1976], perceptual symbol systems [Barsalou, 1999], appraisal theory [Scherer, 2001], long-term working memory [Ericsson & Kintsch, 1995], among others. It is also broadly consistent with embodied cognition [Gallagher, 2005; Varela *et al.*, 1991/2016], nonlinear, self-organizing neural dynamics [Freeman, 2000; Kelso, 1995; Kozma & Freeman, 2016], and Dynamic Field Theory [Erlhagen & Schöner, 2002; Schöner, 2008; see Franklin *et al.*, 2013, 2014, 2016].

We continue the introduction with a brief overview of the concept of *body schema*, including the neuropsychological evidence (Sec. 1.1). Next, we briefly describe the LIDA cognitive architecture and its *cognitive cycle* (Sec. 2), followed by our conceptual implementation of the body schema in LIDA (Sec. 3). We discuss the relation of the body image to the body schema (Sec. 3.3). We examine the plasticity of the Current Body Schema and its relation to the phenomenon of phantom limbs (Sec. 4). Finally, we describe several experiments related to the body schema, and show how LIDA helps to explain their results (Sec. 5).

1.1. *The History and Neuropsychological Underpinnings of the Body Schema*

The concept of body schema is often traced to Head and Holmes’ [1911] early work on brain lesions, although it was prefigured in the work of Bonnier [1900]. Bonnier argued that a stable sense of posture (*sens d’attitudes*) persisted despite the rapid changes in somatosensory information. Munk [1890] earlier posited the existence of “images” of posture and movement recorded in the brain, allowing for the smooth, syncopated coordination and movement of the body. Head and Holmes reject the notion that visual images of the body (which is today one of the meanings of the term *body image*) serve to guide sensorimotor action. Referring to a patient who appears to be deafferented, they remark that when this person closes their eyes, they lose all sense of where their arm may be if you change its position. The sense that is lost in this patient is proprioception or kinesthesia. But proprioception is simply a sense modality. Just as the visual modality helps to build up a model of the world, they thought the proprioceptive information helped to build a model of the body.

“By means of perpetual alterations in position we are always building up a postural model of ourselves which constantly changes. Every new posture or movement is recorded on this plastic schema, and the activity of the cortex brings every fresh group

of sensations evoked by altered posture into relation with it” [Head & Holmes, 1911, p. 187].

While today we may associate the term “posture” with the way one carries their spine, Head and Holmes were referring to the position and movement of the body and its limbs. The subsequent literature uses the term *body schema*. Head and Holmes emphasize that it is only in virtue of this model that the agent is able move, act, and manipulate objects. Mere proprioceptive feedback is not in itself enough, in the same way that visual information from the retinas is not by itself enough to produce a visual scene. The agent needs a knowledge of its embodied capacities and a model of its body and limb positions. This kind of knowledge is what philosophers call *know-how* [Ryle, 1949/2009]. It is not propositional, but is instead practical, embodied, and intuitive. For example, knowing that I can fit through a doorway is a type of know-how, while knowing that I am 172 cm tall is not. In Head and Holmes’ view, the body schema is nonconscious, in contrast to the body image which is conscious.

Merleau-Ponty [1945/2012] first emphasized the aspect of habit or know-how in the body schema. “[T]he subject does not weld together individual movements and individual stimuli but acquires the power to respond with a certain type of solution to situations of a certain general form” [p. 164]. These habits, as well as a sense of bodily size, undergirds for example the way we are able to move through a doorway “without checking the width of the doorway against that of my body” [p. 165]. Merleau-Ponty also emphasizes that the body schema is marked by egocentrically-coded peripersonal space. Following a suggestion from Head and Holmes, Merleau-Ponty proposes that phantom limb syndrome is underlain by the undue persistence of the body schema. A phantom limb arises when a limb is amputated but the body schema erroneously continues to model that limb despite the objective absence of proprioceptive feedback.

Gallagher argues that the body schema is a “system of sensory-motor capacities that function without awareness or the necessity of perceptual monitoring” [2005, p. 24]. Although it is largely nonconscious, it can be brought to consciousness. For example, we may not be aware of the movement of our legs while walking. When not in motion, we may have no awareness at all of our body and limb positions [Jeannerod, 2010]. We can, however, shift our attention to them and gain awareness of their movement and positions.

The neuropsychological underpinnings of the body schema include separate action and perception streams in somatosensory perception [de Vignemont, 2010; Hoffmann *et al.*, 2010; Paillard, 1999]. The action stream is martialed in service of movement, which is ultimately the purpose of the body schema. The perception stream is used for identifying areas of the body. Vision is well-known to have two visual pathways: 1) the dorsal stream, or “how” pathway, a nonconscious stream processing action-guided perception, and 2) the ventral stream, or “what” pathway, a largely conscious stream for object categorization and scene perception [Goodale & Milner, 1992; Ungerleider & Mishkin, 1982]. Other sensory modalities have an analogous organization, including hearing [Kaas & Hackett, 1999] and somatosensory perception, including touch and proprioception. The equivalent of the

ventral visual stream, a pathway from the anterior parietal cortex (APC) to the secondary somatosensory cortex (SII) to the posterior insula, underlies object recognition through touch, and processing of the tactile aspects of body image. Body image is also underlain by the ventral visual stream itself. The equivalent of the dorsal visual stream, a pathway from the thalamus to the posterior parietal cortex, underlies the body schema [Dijkerman & De Haan, 2007]. Although the sensory mechanisms underlying the body schema are known, it is not fully understood how the brain converts the contents of the action stream into a model of the body [Longo *et al.*, 2010].

Bimodal neurons in the premotor cortex, intraparietal cortex, and the putamen code for both a body part's somatosensory receptive field *and* the visual receptive field immediately around that part [Duhamel *et al.*, 1998; Fogassi *et al.*, 1996; Graziano *et al.*, 1994]. These neurons play an important role in the plasticity of the body schema. Head and Holmes say, "a woman's power of localization may extend to the feather in her hat" [1911, p. 188]. Although a rather dated example, what they mean is that such a woman would have an immediate sense of her *extended* bodily height, having no difficulty negotiating doorways. Merleau-Ponty [1945/2012] gives the example of a blind man with a cane. The blind man, Merleau-Ponty claims, does not experience the cane as an intermediary between himself and the world. Rather, he has the sense of a direct perception of the world that is being tapped and touched in front of him, as if the cane were an extension of his body. Maravita and Iriki [2004] find that neural representations of the body extend to incorporate tools in macaques. They recorded bimodal neurons in the intraparietal cortex, where visual and somatosensory modalities are combined. After two weeks of training with a tool for grabbing food, the visual and somatosensory receptive fields of the macaque bimodal neurons expanded to integrate habitual tool use. In Fig. 2 below, distal neurons code for the hand (and later the hand-tool extension), while proximal neurons code for arm reach (and later the arm-tool extended reach).




Fig. 2. Evidence of the plasticity of the body schema. Reprinted from Maravita and Iriki [2004, p. 80] with permission.

In humans, the body schema integrates proprioception and the vestibular sense of balance, also integrating visual, tactile, and auditory information, producing a model of the agent's own body [Hoffmann *et al.*, 2010]. Despite a sharp dissociation between body image (as conscious visual percepts of the body) and body schema [Gallagher, 2005], the latter integrates multimodal information to help it update.

LIDA distinguishes between three different functions that subservise what in the broader literature is unified under the label of 'body schema': 1) Current Body Schema, 2) Habitual Body Schema, and 3) Online Body Schema (see Fig. 1). Henceforth, we will use the term 'body schema' to mean the set containing these three components. Our distinction is not merely a modeling decision relevant to the LIDA cognitive architecture. We base this distinction on the functional role of consciousness in Global Workspace Theory.

According to Global Workspace Theory, consciousness in humans and some other animals [Boly *et al.*, 2013] broadcasts contents to distributed networks in the brain. In this way, information that is not otherwise available to regional brain networks is made available to them [Baars, 2019; Dehaene, 2014; Mashour *et al.*, 2020]. Global Workspace Theory proposes the Conscious Learning Hypothesis, viz. that "significant learning takes place via the interaction of consciousness with the various memory systems" [Franklin *et al.*, 2013, p. 10]. Unconscious "learning," such as priming, is not labeled as 'learning' in this context because it is not considered a route to significant learning, but it is still accounted for [Franklin *et al.*, 2013, 2016; Kugele & Franklin, forthcoming]. The broadcast

is mediated by theta-gamma coupling of neural assemblies across the brain [Baars *et al.*, 2013; Strain *et al.*, in preparation; Varela *et al.*, 2001]. Essentially, local neural assemblies distributed across regional networks are entrained by the mediation of long-fiber pyramidal neurons, causing a nonlinear bifurcation of brain activity into a state of global synchronization, and simultaneously inhibiting other neural assemblies [Kozma & Freeman, 2016]. The Current Body Schema—i.e., the proprioceptive and vestibular model of the body—is *continually* updated, however, meaning that it cannot rely on periodic conscious broadcasts to update. Furthermore, most of its updates are unconscious. If it were updated through conscious broadcasts, that would imply that we would be continually phenomenally conscious of the entirety of our bodily movements, positions, and microadjustments, which is plainly false. The Habitual Body Schema, on the other hand, does not require continuous updating, and can learn in the typical way, through conscious broadcasts (see Sec. 2). For example, when we learn to drive, we learn to recognize new affordances. To learn these new affordances, however, we must first pass through a learning stage involving intense concentration and conscious focus. Finally, “dorsal” stream information is never-conscious, but it allows agents to make fine, online adjustments as an action is being performed [Gallagher, 2005]. In sum, we believe that our distinction of the ‘body schema’ into three separate functions ought to hold beyond the specific case of LIDA, as it is based upon the general functional role of consciousness in Global Workspace Theory.

2. The LIDA Cognitive Architecture

LIDA, or Learning Intelligent Decision Agent, is an embodied, situated, and functionally conscious cognitive architecture. It is based on IDA, a software agent developed for the US Navy [Franklin *et al.*, 1998]. LIDA is a conceptual and partly computational, implementation and fleshing out of Global Workspace Theory [Baars, 1988, 2019]. LIDA models minds, both biological and artificial, which are defined as “control structures for autonomous agents” [Franklin, 1995, p. 412]. LIDA agents are autonomous agents with control structures based on the LIDA architecture and its conceptual commitments. LIDA agents have been created for both scientific (e.g., attempting to replicate a psychological experiment) [Madl & Franklin, 2012; McCall, 2014] and engineering (e.g., controlling a robot) [Dong & Franklin, 2015] pursuits. Depending on the task domain and needs, a LIDA agent may not implement all elements of the LIDA modules. LIDA can be used to model both humans and non-humans and, in this paper, we will often refer to implementations of former as “human agents.” In this section, we provide a brief overview of the LIDA model, focusing on the emergence of cognitive cycles, the role of consciousness, and the modules^c most relevant to the body schema. For a fuller overview, see Franklin *et al.* [2016].

In LIDA, cognition arises from *cognitive cycles*, which are akin to the action-perception cycles referred to by many psychologists and neuroscientists [Cutsuridis *et al.*,

^c The modular structure of LIDA does not necessarily imply a modular view of the brain.

2011; Dijkstra *et al.*, 1994; Freeman, 2002; Neisser, 1976]. In this context, ‘cognition’ refers broadly not only to higher-order cognition, but also to basic perceptual and motor processes. Cognitive cycles emerge as LIDA’s cognitive modules receive, process, and send information, where each cycle conceptually begins with perception and ends with the execution of an action. Most of LIDA’s modules operate asynchronously, that is, they do not explicitly synchronize their operations with other modules; therefore, cognitive cycles typically overlap. They are divided into three phases: 1) the perception and understanding phase, 2) the attention phase, and 3) the action and learning phase. A single cognitive cycle in a human typically takes place in on the order of 260–390 ms, with conscious broadcast happening around 200–280 ms after the onset of the stimulus [Madl *et al.*, 2011]. The timing involved is consistent with neuroscientific evidence of theta-gamma coupling of neuronal populations from sensory to corticothalamic regions [Doesburg *et al.*, 2009; Varela *et al.*, 2001].

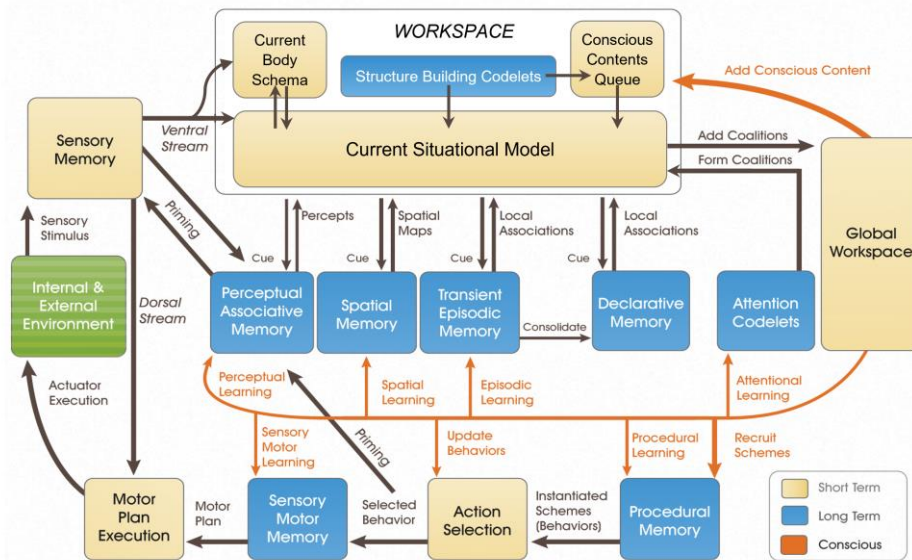


Fig. 3. Schematic of the LIDA cognitive architecture.

2.1. The Perception and Understanding Phase

A single cognitive cycle begins with the perception and understanding phase, which starts along with stimulus onset. For a LIDA agent modeling humans, this typically takes 80–100 ms to initiate. A sensory stimulus is first processed in Sensory Memory, which is a rapidly decaying memory, but with high storage capacity. It has only the most basic of feature detectors. Both neural networks [He *et al.*, 2015; Kugele & Franklin, 2020; Madl *et al.*, 2014] and cortical learning algorithms [Agrawal & Franklin, 2014; Hawkins *et al.*,

2010] have been explored as implementations of Sensory Memory. Ventral stream visual information from Sensory Memory is sent to the preconscious Workspace, as well as to Perceptual Associative Memory. The preconscious Workspace is a form of short-term memory; it is not conscious *per se* but some of its contents can be brought to consciousness [Franklin & Baars, 2010]. Structure building codelets monitor the preconscious Workspace for items that are relevant to their concerns, and they assemble structures. In the Current Situational Model (CSM), a submodule of the preconscious Workspace, structure building codelets assemble structures that model the environment and certain aspects of the agent's current situation, such as options and emotions. It is a form of short-term memory, and the model of the agent's situation often changes as its environment changes.

Structure-building codelets build structures using input originating from Sensory Memory, memories cued from long-term memory modules (Perceptual Associative Memory [PAM], Spatial Memory, Transient Episodic Memory [TEM], and Declarative Memory), and representations constructed by other structure-building codelets. PAM, based on the CopyCat architecture [Hofstadter & Mitchell, 1994], receives ventral stream^d information from Sensory Memory as percepts are activated, instantiated, and sent to the CSM. PAM's primary task is recognition. One of its primary data structures is the event, typically comprised of agent, action, and object thematic links, and sometimes others. Spatial Memory stores allocentric spatial maps (egocentric spatiality is modeled within the CSM). TEM stores autobiographical memories that decay away within the span of a few hours or a day [Conway, 2001; Franklin *et al.*, 2005]. Declarative Memory stores autobiographical memories that have been consolidated from TEM (for example, during REM sleep), as well as semantic memories that are autobiographical memories that have lost their place and time contexts.

2.2. *The Attention Phase*

The next phase of the cognitive cycle is the *attention phase*, wherein attention codelets form coalitions, and compete for a conscious broadcast, which in humans typically begins around 200–280 ms after stimulus onset. Like structure building codelets, attention codelets monitor the CSM for their specific concerns, such as loud, bright, dangerous, relevant, urgent, and important objects [Faghihi *et al.*, 2012]. Attention codelets contribute to consciousness' function as a saliency filter, and allow only the most salient structures to compete for the global broadcast. The structures built by structure building codelets can be further included in coalitions by attention codelets. Coalitions contain one or more structures that are advocated for by one or more attention codelets, which are assigned an overall activation based on the activation of the attention codelet, the activation of its component structures, the degree to which the structure matches the codelet's concern, as

^d In LIDA, the ventral stream of information is defined as multimodal information used for representation. The dorsal stream is defined as multimodal information used for online adjustment of the Sensory Motor System. They loosely correspond to the ventral and dorsal visual streams of the human brain, but they are not limited to one modality and they are not necessarily neural pathways (e.g., they can be properties of artificial agents).

well as the modulatory effects of a refractory period [Madl & Franklin, 2012]. Coalitions are sent to the Global Workspace, where they compete for the conscious broadcast [Baars, 2019; Baars *et al.*, 2013]. While most modules are asynchronous, the competition for the global broadcast within the Global Workspace is synchronous, and it is initiated by one of several triggers [Kaur, 2011]. For example, when a coalition over a certain activation threshold arrives, a global broadcast is triggered.

Content in the global broadcast is considered “functionally consciousness” content; that is, it’s made globally available to all LIDA’s modules and processes, though each module and process is free to ignore that content based on its own concerns. This corresponds to the human evidence of a phase shift in neural dynamics during conscious cognition [Kozma and Freeman, 2016], wherein distributed neural assemblies across the brain temporarily synchronize through theta-gamma coupling [Baars *et al.*, 2013; Strain *et al.*, in preparation]. Overall, the agent’s model of the world and of its body, as well as the plans, memories, and imaginings surrounding those models, remain overwhelmingly unconscious (i.e., in the preconscious Workspace). The agent is only conscious of a very small portion of that material in the form of the contents of the winning coalition at any given point in time.

2.3. *The Action and Learning Phase*

The final phase of the cognitive cycle is the *action and learning phase*, wherein the conscious broadcast leads to a process of action selection, as well as profligate learning in a majority of the modules. In humans, this begins 60–110 ms from the initiation of conscious broadcast [Madl *et al.*, 2011]. The choice of an action begins as conscious content is used to recruit schemes in Procedural Memory. Schemes are Procedural Memory’s primary data structure, and they consist of a context, an action, and a result. Upon receipt of a conscious broadcast, schemes can be instantiated using the conscious content to fill out the details of the context, action, and result. Instantiated schemes, called behaviors, are then sent to Action Selection. Action Selection then chooses a behavior. Under some circumstances, the selection of a behavior may require *deliberation*, such as when two possible behaviors have equal weighting; this pushes behavior selection to further cognitive cycles (multi-cyclic deliberation is achieved through ideomotor theory [Franklin, 2000; James, 1890/1983]). Once a behavior is selected, it is sent to Sensory Motor Memory, which generates a motor plan to realize the action of the behavior. The motor plan is then sent to Motor Plan Execution. Profligate learning also occurs as a result of the conscious broadcast: attentional, procedural, episodic, spatial, perceptual, behavioral, sensory motor, and other types of learning [Franklin *et al.*, 2013, 2016; Kugele & Franklin, forthcoming]. This accords with the conscious learning hypothesis of Global Workspace Theory, which states that significant learning (excluding priming) occurs only during conscious broadcasts [Baars, 1988, 2019]. Information from the Current Body Schema, when broadcast, can be used by Procedural Memory, Action Selection, Sensory

Motor Memory, and Motor Plan Execution—potentially, by any of LIDA’s modules and processes. In this way, the body schema supports action and movement.

3. The Body Schema in LIDA

LIDA makes a distinction between three functions subserving what in the broader literature are labeled “body schema” (see Fig. 1 and Fig. 4).

- (1) The *Current Body Schema*. It is a continually-updated model of postural and effector positions, effected in terms of effector variables and their values. In humans, it is primarily built out of proprioceptive and vestibular feedback, but visual feedback can influence its states. It is typically preconscious but parts of it can be brought to consciousness. The module is a system composed of a long-term memory component holding effector variables, and a short-term memory component holding rapidly-updated values for those variables.
- (2) The *Habitual Body Schema*. It is a unified collection of bodily and effector capacities for action, learned affordances, accustomed weight, and accustomed size. It is a form of long-term memory. It is preconscious and stored largely in PAM, from whence it can be cued into the CSM and from there possibly brought to consciousness.
- (3) The *Online Body Schema*. It consists of “dorsal” stream information sent directly from Sensory Memory to the Sensory Motor System. Since it bypasses conscious broadcast, taking a shortcut straight from early sensory processes to late motor performance, it allows for the rapid sensorimotor adjustments that are critical to smooth motor performance in the world.

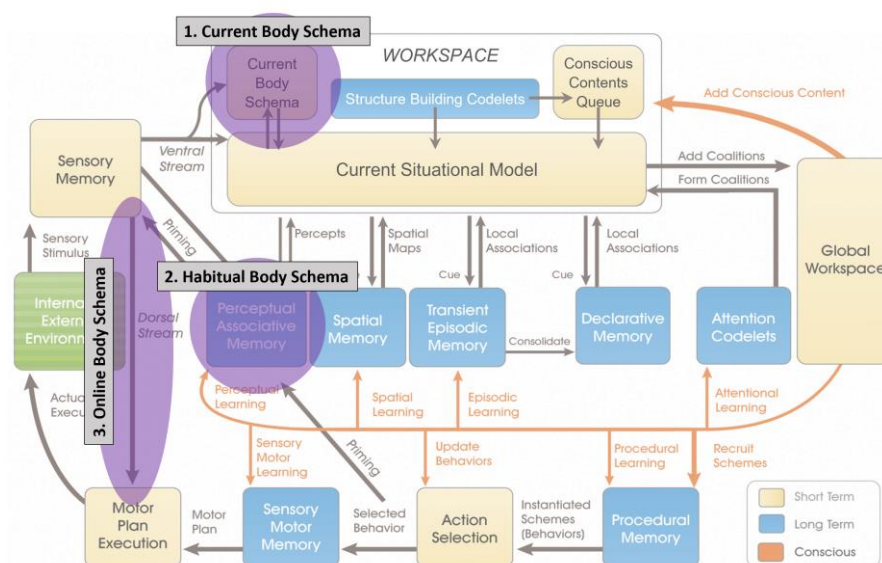


Fig. 4. The place of the three functions of the body schema in LIDA's cognitive cycle.

Body image, on the other hand, is an umbrella term capturing the variegated aspects of the body that are represented in linguistic, visual, auditory, and other modalities. Nonhuman LIDA agents may have only a limited body image composed of visual and other sensory information. This is because many nonhuman animals and artificial agents have little capacity for semantic knowledge of their attractiveness and other bodily attributes that are social in nature. LIDA agents *without* a body, such as IDA [Franklin *et al.*, 1998], may need neither a body image nor a body schema. In humans, the body image is composed of ventral stream visual information, the somatosensory “ventral” stream from the anterior parietal cortex (APC) to the secondary somatosensory cortex (SII) to the posterior insula (see Sec. 1.1), semantic knowledge of the body, and socially-relevant representations of the body. For humans, the semantic elements of the body image are part of the self-concept in Declarative Memory [Ramamurthy *et al.*, 2012]. The body image can influence and modulate the body schema (see Sec. 3.3 below), but they remain distinct systems, and human studies have shown that they are underlain by different neurological streams [Dijkerman & De Haan, 2007; Gallagher, 2005; Paillard, 1999].

The philosophical, psychological, and neuroscientific literature on the body schema does not make a distinction between the short-term and long-term memory aspects of the body schema, nor does it make a place for the functional role of consciousness. LIDA, however, implements functional consciousness (based on Global Workspace Theory), and a model that integrates consciousness will need to distinguish several separate aspects of the body schema. The body schema in humans contains not only an online model of posture

and effector positions, but it also stores the habitual capacities of the body and its effectors [Gallagher, 2005]. This is how, for example, a phantom limb can persist even in the absence of proprioceptive feedback. Any long-term memory must be learned through a conscious broadcast. The Current Body Schema's online model of effector positions, however, does not require consciousness to model movement, although it can be brought to consciousness. It does, however, require consciousness in order to learn.

The body schema guides an embodied agent's interactions with the world. It models aspects of the body that are relevant to acting in the world [Graziano & Webb, 2015]. Without a body schema, simple movements such as walking, for example, would be rendered exceedingly difficult, because the agent would not know the positions of its limbs. LIDA agents with effectors or simulated effectors (i.e., in a simulated environment) have a body and hence will sometimes benefit from a body schema. An effector is an actuator that "assert[s] physical forces on the environment" [Russell & Norvig, 2016, p. 196]. Software agents may have actuators that do not realize physical force, such as generating text. Basically, a LIDA agent that acts on a physical environment (and hence is embodied), or on a simulated environment [Dong & Franklin, 2015; Madl *et al.*, 2016], may need a body schema. But software agents that do not act on the physical environment or a simulated equivalent, such as Medical Agent X (MAX) [Strain & Franklin, 2011; Strain *et al.*, 2014], do not need a body schema.

3.1. *The Current Body Schema in the Preconscious Workspace*

The Current Body Schema is part of the preconscious Workspace, interacting with the Current Situational Model (CSM) (see Fig. 3). The CSM can represent aspects of a LIDA agent's external environment as well as its internal environment, including memories, desires, plans, and other items. External items are labeled 'real,' while many internal items are labeled 'virtual.' Most of the objects represented in the CSM are about environmental objects, such as a visual scene in humans. For an agent that moves or locomotes, its environment continually changes from its perspective. For example, the contents of consciousness continually shift as a human agent progressively moves by a chair, a sofa, and a table in a room. If it closes its eyes, most of the visual contents of consciousness will very rapidly disappear. The corresponding representations in the CSM, however, remain until they decay away. However, when the agent moves out of the room, those items will likely decay out of the CSM in the absence of reinforcement. The content may decay within a few seconds for humans, but that duration ultimately depends on the kind of agent in question and the nature of the environment. The Current Body Schema shares these features of the CSM, but is marked by two important differences:

- (1) Most of its contents are persistently represented. The agent's body and effectors largely remain the same throughout its existence, while the contents of the CSM typically change with changing environments and situations.

- (2) Its contents are automatically labeled ‘self.’ As a result, it rarely confuses its own body with the rest of the world.

The Current Body Schema is continually updated with incoming proprioceptive and vestibular information. It is a system composed of both a long-term and short-term memory component. The long-term memory component stores effector variables such as the angles of joints. Effector variables may be innate (or programmed), or they may be learned through continual reinforcement by proprioceptive information from conscious broadcasts. The second component of the Current Body Schema system is a form of short-term memory, storing the values of effector variables.^e These values continually change as effector positions shift in space. Variables may model the angles between joints (e.g., between human synovial joints) or directionality for swivel joints (e.g., human cartilaginous joints), but some effectors do not have joints at all. For example, the tongue and lips, used in speech production, are modeled with effector variables without joint angles. Octopus tentacles are another example of effectors modeled without joint angles. Variables’ values change rapidly as an agent interacts with the world, and they also rapidly decay. In Fig. 3, we label the Current Body Schema as short-term memory to emphasize the continual update of effector variables that occurs as it models the agent’s body, but the reader should keep in mind that it is a system composed of both a long-term memory for effector variables and a short-term memory for the values of these variables (see Fig. 5).

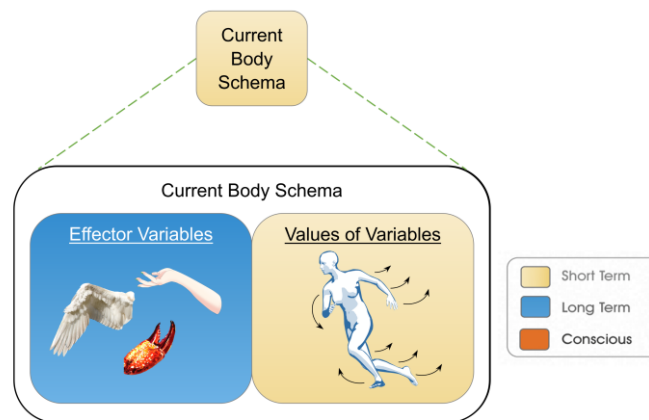


Fig. 5. The Current Body Schema is a system composed of long-term and short-term components. The Current Body Schema is labeled simply as short-term memory in Figs. 2 and 6 for simplicity, pointing out that its most important aspect is modeling the agent’s body as it moves.

^e Every autonomous agent must have actuators, but immobile agents (whether in a virtual or real world) have no effectors. Because of this, they may not need a body schema.

In the sustained absence of proprioceptive feedback, the Current Body Schema may no longer be able to track the agent's body, or some subset of joints or variables. For example, it is common for humans to experience a limb "going to sleep." This occurs when blood circulation is temporarily cut off from a limb, such as by sitting on it. The variables representing the angles between the limb's joints decay, and the Current Situational Model temporarily ceases to model the affected part of the limb. Visual feedback (body image) may be needed in order to coordinate walking until sensation has been restored.

A more radical decay of *both* variable *and* joint representations in the Current Body Schema can be seen in the case of Ian Waterman. Waterman was affected by a neuropathy that caused him to lose proprioceptive and tactile feedback in most of his body [McNeill *et al.*, 2010]. He gradually regained the ability to walk, use his arms and hands, and engage in many everyday activities. However, this was a functional rehabilitation that operated using different cognitive structures from those typically undergirding bodily motion. Waterman used the body image, in the sense of visual information about his body, to guide his actions [Gallagher & Cole, 1995]. In LIDA terms, the joint variables decayed out of the Current Body Schema, since they were no longer reinforced by incoming proprioceptive feedback. In this sense, Waterman had a "missing" or "lost" body schema [Gallagher, 2005]. This case highlights the utility of the Current Body Schema. Although Waterman regained the ability to walk, he is only able to do so concentrating full attention on his bodily movements through visual feedback. Typically, walking is an automatized process requiring infrequent or little attention at all.

The Current Body Schema can be built from the same kinds of content as structures in the CSM, such as percepts and category links. Despite the content being the same in kind, the agent does not confuse its own body with its outside environment. Everything in the Current Body Schema is recognized as 'self,' for example, by associations with a "self node" as in PAM. LIDA integrates three levels of selfhood: the protoself, the minimal self, and the extended self (see Sec. 3.3) [Ramamurthy *et al.*, 2012; Ryan *et al.*, 2020]. The Current Body Schema is part of the protoself. The protoself in human agents involves neural representations of the body [Damasio, 1999].

The content of the Current Body Schema includes the rapidly-updated positions, size, and pressure information of the agent's effectors. This information is cued into the CSM when a relevant item appears. For example, a human agent may receive a visual stimulus from a doorknob. The sensory information generated by this stimulus is sent to the preconscious Workspace. If the object is recognized as a doorknob in PAM, an object node for the doorknob can be sent into the CSM as part of a percept. An action node may also be activated in PAM, e.g., a 'turning' action. This turning action is part of the Habitual Body Schema (see Sec. 3.2). The action is not a full specification of the scheme or motor plan of the turning motion. Rather, it is simply the knowledge that the agent is able to perform that action. Additionally, the agent's hand is activated in PAM as part of the body image by association with the doorknob node. As a result, affordances related to the doorknob are recognized in PAM, suggesting possible bodily actions. Next, the Current Body Schema can be cued for effector variables and their values, as they are rapidly

changing online. We have presented this as if these are sequential steps, but this is only for the convenience of explanation. The activation of the object, action, and hand in PAM, and the activation of the current joint positions and variables of the hand and its digits from the Current Body Schema, are all cued from the CSM.

An event in a LIDA agent is a structure usually composed of three or more nodes, often with the following thematic role links: 1) agent, 2) action, 3) object, and (for embodied actions) 4) instrument [McCall *et al.*, 2010]. In our example, the agent is ‘self,’ the action is ‘turn,’ the object is the doorknob, and the instrument is ‘hand.’ Incoming ventral stream information from Sensory Memory activates an event node that was previously learned into PAM, including the agent, action, object, and instrument. The entire event is part of the Habitual Body Schema (see Sec. 3.2).^f If the event’s activation is high enough, it can be copied into in the CSM. The hand that is part of this event then cues the Current Body Schema. As a result of this cueing process, the unbound variables that are part of the hand node in the CSM can be associated with the corresponding values from the part of the Current Body Schema that could be identified in PAM as a hand. This specifies the exact context, position, and articulation of the hand (see Fig. 6). Once in the CSM, these values now associated with the hand node will be used to cue other long-term memory modules. For example, for those who use sign language, the specifics of hand orientation can be associated with the recognition of words in PAM, and it can cue associated information from Declarative Memory.

^f Not all events in PAM are part of the Habitual Body Schema; only those incorporating bodily effectors in the instrument role are.




Fig. 6. An event is recognized in PAM. The event consists of an agent role (self), an action role (turn), an object role (doorknob), and an instrument role (hand). This event is part of the Habitual Body Schema (marked in green). The event recognized in PAM is then sent as part of a percept into the CSM. In the CSM, the variables in the event structures are then bound with values, including those for the hand. Depending on how the programmer wishes to implement a specific LIDA agent, this could be done by structure building codelets, or the module could do the binding. The values for the hand used to bind the hand node in the CSM are cued from the Current Body Schema.

For a human agent, who has the embodied capacity to turn a doorknob, the doorknob affords turnability. This requires that the agent has previously learned the affordance, such

as through exploration or cultural learning. The affordances an agent perceives are always relative to the embodied capacities of that agent. For example, a cat typically does not perceive a turnability affordance in a doorknob. This is because the structure of their paw is not suitable for the grasping, twisting, and pulling that is required to open with it. The cat's Habitual Body Schema has no stored memory of such bodily capacities, and so they cannot be cued to the CSM. Nevertheless, this does not absolutely prevent an agent from recognizing affordances for other kinds of agents.

The agent's current size is not measured in centimeters or inches. For example, if a human agent knows it can fit through a door, it is not because it is cued from Declarative Memory the fact that it is e.g. 170 cm tall, and compares that to the height in centimeters of the doorway, which anyway is usually unknown. Rather, it has an embodied sense of its size in relation to other objects. Typically, current size and accustomed size (which is a long-term memory held in the Habitual Body Schema) will match up. However, discrepancies will arise when e.g. a human agent wears platform shoes or high heels, or when a tool becomes incorporated into its Current Body Schema (see Sec. 4).

Structures in the Current Body Schema are not capable of directly being broadcast because the module is not part of the CSM, and so attention codelets do not see it. Structures that are cued into the CSM, however, may form parts of coalitions that compete for conscious broadcast. If a model of an effector wins the competition for broadcast, the effector variables and their values can be used by Procedural Memory and Sensory Motor Memory. Procedural Memory is the first stage of the action phase of the cognitive cycle, and its data structure is the scheme. An important aspect of schemes is that they are contextually sensitive (the scheme consists of a context, action, and result). Schemes that are bound with variables from conscious broadcasts are copied into Action Selection as behaviors. Procedural Memory may use broadcast effector values to help determine the scheme's context, action, and result. For example, the position of the hands is important in determining whether an agent can turn a doorknob.

In each cognitive cycle, Sensory Motor Memory receives the winning behavior from Action Selection, which can also utilize broadcast effector values to help determine the motor plan template. The motor plan template is a skeletal outline towards a concrete motor plan. The effector values, which ultimately come from the agent's model of its own body in the Current Body Schema, are fully specified in the behavior. However, by the time the cognitive cycle has moved from early processing of the stimulus in Sensory Memory to Sensory Motor Memory, for many LIDA agents many milliseconds (hundreds in the case of humans) will have already passed. The object in the environment that the agent intends to act upon may have slightly changed its position or orientation in the interim. To prevent a disconnect between action and the environment due to this lag, dorsal stream information feeds directly into Motor Plan Execution. The motor plan generated by Sensory Motor Memory can have some of its variables updated within Motor Plan Execution by the dorsal stream information of the Online Body Schema. This is consistent with neuropsychological evidence that patients with damage to object categorization processes are still able to adjust for changes in orientation when performing tasks [Goodale *et al.*, 1991].

The dorsal stream information about effectors differs from the information modeled by the Current Body Schema insofar as the former 1) does not form a full bodily model, but instead consists of localized and unintegrated information, and 2) is not capable of being consciously broadcast. As we saw in Sec. 1.1, the ‘body schema’ in humans is underlain by two separate neural pathways, and the unconscious pathway (analogous to the dorsal visual stream) runs from the thalamus to the posterior parietal cortex in humans. LIDA does not label this “dorsal” stream information as part of the body schema, but it is important to keep in mind that LIDA does account for this aspect of the body schema. In LIDA, the ‘body schema’ similarly uses both dorsal and ventral stream information (which in LIDA are not necessarily neural in nature). In other words, the ‘body schema’ as it is presented in many scientific and philosophical accounts [e.g., Gallagher, 2005] is realized in LIDA partially through aspects that are capable of becoming preconscious (Current Body Schema, Habitual Body Schema) and partially through the dorsal stream.

3.2. *The Habitual Body Schema*

While the Current Body Schema continually tracks effector variables and their values including the angles between joints in human agents, the Habitual Body Schema is a unified collection of the body’s capabilities and capacities for action, learned affordances, and accustomed weight and size. While the Current Body Schema has a long-term memory component, it only stores effector variables. Embodied capacities for action are recognitional. Upon seeing an affordance, an agent may recognize that it has the embodied capacity to act on that affordance. PAM is LIDA’s recognitional memory, and the Habitual Body Schema is a label for that part of PAM that stores embodied capacities and affordances.

The Habitual Body Schema helps an agent to effectively maneuver its body during various tasks by storing its actions and capacities for use in modeling within the CSM. The knowledge that is stored in the Habitual Body Schema does not consist of detailed plans of how to implement an action with an effector. Actions are specified and become progressively more detailed as they develop from schemes, to behaviors, to motor plans (see Sec. 2.3). Just as the Current Body Schema does not model every minute part of the agent’s real body, but instead only aspects of the body related to the control of its effectors, the Habitual Body Schema does not model the detailed implementations of actions, but only the broad capacities for action. The actual moving of the effectors is accomplished in the action and learning phase, and the capacities stored in the Habitual Body Schema are used to help with the simulation of that action prior to it being carried out.

The Habitual Body Schema instead contains events which include a bodily effector in the instrumental thematic role (see Fig. 6). Actions such as grasping, climbing, or turning, are learned into the Habitual Body Schema as parts of events through conscious broadcasts. For example, if an agent learns a new skill such as opening a door, the Habitual Body Schema will be updated to reflect this new embodied capacity for action. Modules related to action, including Procedural Memory and Sensory Motor Memory, may also be updated

to reflect the new skill, resulting in modified or new behaviors and motor plans. A single conscious broadcast may include an agent's trying of a new action for the first time, such as turning a doorknob. The conscious broadcast becomes available to other LIDA modules, including PAM, where the Habitual Body Schema resides. Many skills require multiple trials in order to reinforce a new action node in the Habitual Body Schema. In this way, the Habitual Body Schema builds a repertoire of skills as the agent engages with the world over time.

Weight is an example of something that is realized by the interaction of the Current Body Schema and the Habitual Body Schema in the CSM. Weight is not directly measured in the Current Body Schema, but it can instantiate a copy of the pressure sensations it receives into the CSM, where it can be interpreted as weight. The current size and weight of the body and effectors are not measured in standard units of measurement. For example, we can put on a heavy backpack. We will feel encumbered by the weight, and have a sense that our body is heavier than it usually is. We will rightly attribute this difference to the backpack. This intuitive sense of a change in weight is not measured in kilograms or pounds. We may know our weight, e.g., 73 kg. It is possible that we may also know the combined weight of ourselves plus the backpack, e.g., 93 kg. This semantic knowledge is not primarily what makes us feel encumbered. Rather the current weight we are referring to here is a bodily *feeling*. Although semantic knowledge can influence our sense of how heavy something is, the linguistic token '93 kg' is different in kind from the feeling of heaviness. The feeling of weight is only approximately related to our weight in kilograms.

Pressure information from Sensory Memory feeds directly into the CSM, the Current Body Schema, and PAM. PAM, as a recognitional memory, is where weight is recognized. However, PAM is not directly able to determine weight from pressure information it receives from Sensory Memory. This is because weight is a contextual feeling, and pressure can be interpreted in other ways depending on context. For example, pressure in a fast vehicle or spaceship may be interpreted as acceleration, or pressure on an arm when a human punches a wall may be interpreted as force. As part of its recognition, PAM interprets the pressure as weight (the context is part of the recognitional process). Once PAM recognizes weight, it can then be sent as part of a percept back up to the CSM. The feeling of weight (or acceleration, or however the pressure information is interpreted) can then be consciously broadcast if it is selected by an attention codelet into a coalition that wins the competition for global conscious broadcast. That corresponds to our conscious sense of weight. The difference should be intuitive: our conscious awareness of how heavy we feel when we put on a heavy backpack is not the same thing as our awareness of our saying, or subvocalizing, the sentence, "I weigh 73 kg."

3.3. *Body Image*

As previously mentioned, 'body image' is an umbrella term referring to representations of the agent's body that are not proprioceptive, vestibular, or affective, and some of its diverse functions are biologically realized by two separate neural pathways in humans [Dijkerman

& De Haan, 2007]. It is not a model or unified collection like the body schema. Many semantic representations of the body, for humans, are deeply embedded in social distinctions, hierarchy, competition, and mating, and can be a significant source of pride or distress. These include representations, on a positive to negative scale, of attractiveness, weight, height, athleticism, and so forth. Several psychological disorders, such as anorexia, bulimia, and body dysmorphic disorder, are essentially disorders of the body image. They may incorporate both negative semantic body image and negative visual or imaginative body image.

In LIDA, following Global Workspace Theory, the semantic elements of the body image are part of the self-concept. The self-concept is an aspect of the Extended Self. The self-concept, or selfplex [Blackmore, 2000], includes an agent’s beliefs about themselves [Baars, 2019; Ramamurthy *et al.*, 2012; Ryan *et al.*, 2020]. There are also visual and imaginative aspects of the body image [Gallagher, 2005], and these are likely widely distributed across the kingdom *Animalia*, more so than semantic aspects. These can include sensory and perceptual experiences of their own bodies, such as how their bodies visually appear, smell, taste, and feel, as well as mental simulations of those experiences. Any agent that has the capacity to sense its own body through modalities other than proprioceptive and vestibular means can have at least a minimal body image. By contrast, the Current Body Schema is marked as part of the protoself.

4. The Plasticity of the Current Body Schema and Phantom Limbs

As we noted (Sec. 1.1), the body schema is, under the right conditions, plastic and able to extend out into external objects. This only happens in situations where an object, such as a tool, is repeatedly utilized by an agent over a period of time. Maravita and Iriki [2004] documented the neuroplastic shift in body schema representations for macaques when they were trained to use a reaching tool, over a period of two weeks, to grab and eat fruit. A more classic example is of a blind man, whose cane has become an extended part of his body schema [Merleau-Ponty, 1945/2012]. In LIDA, the Current Body Schema is a component that is plastic. The Habitual Body Schema can also incorporate novel capacities or the affordances that tools enable, and in that sense can also be “extended.” However, this is not a literal extension, since the Habitual Body Schema is not a model (since it does not include any processes), but rather is a collection of capabilities.

When a LIDA agent uses a tool, e.g., a food-grasping tool [as in Maravita & Iriki, 2004], the Current Body Schema begins to incorporate that tool into its model of the body. Upon using a tool once, that tool can be learned into the Current Body Schema as a new variable (in the module’s long-term memory component) through a conscious broadcast, and will have a very low activation. This is an instance of the profligate learning that occurs in LIDA agents [Franklin *et al.*, 2016]. Old variable values become associated with this new variable. For example, when a cane is learned into the blind man’s Current Body Schema, there are no new sensors. Values generated by sensors (such as pressures in the

hand) get associated with the cane. Without reinforcement, however, it rapidly decays out of the Current Body Schema. Once tools or other objects in the CSM become incorporated into the Current Body Schema, they are automatically labeled as ‘self.’ The model of the agent’s body and effectors incorporates the extraneous implement. The tapping sounds of the cane, used for spatial navigation, are not proprioceptive, and are processed in the CSM.

As Merleau-Ponty notes, “The blind man’s stick ...has become an area of sensitivity, extending the scope and active radius of touch, and providing a parallel to sight” [1945/2012, p. 165]. What Merleau-Ponty describes is the shift that the tool makes from a pure object into an extension of the agent’s own hand or arm. In a LIDA agent, not only does the stick become an extension of the agent’s body, but it also retains its status as an object, as well. The agent can shift its disposition towards the cane, and look at it purely as an object, because it is still represented in the long-term memory modules as an object. Although a tool, such as a food grasper or a cane, has no proprioceptive sensors, information about its position is readily available through the sensors of the effector that is manipulating the tool, and possibly other sensors. As previously mentioned (Sec. 3), effectors in the Current Body Schema can be learned into it—specifically, into its long-term memory component (see Fig. 5). The plastic incorporation of tools into the Current Body Schema is an instance of such learning.

An object node for a cane in PAM can be cued into the CSM, where it can be linked with existing effector variables by structure building codelets, and so become a new effector structure. If it is selected by attention codelets and eventually wins the competition for conscious broadcast, the Current Body Schema may access it and write a new variable for it in its long-term memory component. At this point, the structure becomes a new effector. This is a form of instructional learning, which occurs when new structures are learned into a module as a result of a conscious broadcast. That new variable will then either decay out or be reinforced. Reinforcement occurs through selectionist learning, wherein the base-level activation of structures are strengthened through further conscious content. If the variable is reinforced enough, it may become a permanent or at least long-enduring entity within the Current Body Schema. For example, the macaques studied by Maravita and Iriki [2004] only incorporated the food grabber into their body schemas after a full two weeks of repeated trials of using the tool to grab food.

When the tool is integrated into and augments the Current Body Schema, it only extends the latter while it is actively being used. Once the tool is dropped, the Current Body Schema stops modeling the extended effector. The variables associated with the tool cease to have values in the short-term memory component of the Current Body Schema, but they remain in the long-term memory component. If the tool is picked back up again, the Current Body Schema once again models the extended effector. The variables do not typically decay out of the long-term memory component, possibly unless it has not been reinforced in a very long time. The short-term memory component, however, registers no incoming information concerning the variables’ current positions. In the absence of this information, it is unable to model those variables in action. Hence, the extended effector seems to

“disappear” when the tool is no longer in use, but it can rapidly “reappear” once incoming information is again received.

There is typically no proprioceptive information that directly comes from an object that is incorporated into the Current Body Schema. What changes is the way sensations are interpreted. A cane, for example, has no sensors. The effector or effectors that are directly in contact with the cane (viz., the hand), does have sensors, however. It is through these sensors that proprioceptive feedback related to the artificial effector comes into the Current Body Schema. For example, when the blind man’s cane bumps against an external object, the physical force is transmitted through the cane and to the pressure sensors in the hand. The changes in pressure in the hand used to only be associated with the hand effector in the Current Body Schema. However, some of them now also become associated with the extended cane effector once it has been learned into the Current Body Schema. Additionally, the sounds produced by the cane’s tapping also become associated with the new effector.

The Habitual Body Schema is also affected and shaped by the repeated utilization of a certain tool or object over time. New embodied capacities or affordances are learned in PAM, or existing affordances become associated with novel objects. For example, a human agent may become accustomed to using a virtual reality system for gaming. Many such virtual reality systems utilize manual controllers, which the user grasps and manipulates in order to interact with objects and events in the virtual world. Once a gaming controller is used enough, it should become integrated with the Current Body Schema [see Maravita & Iriki, 2004]. Controllers often have one or more buttons, in addition to position sensors. The human agent may want to open a virtual elevator, or pick up a virtual object. To do so, they would place their hand near the virtual elevator panel, or near the virtual object. Next, they would press the button on their controller. These virtual objects have affordances, but the affordances are slightly different because they do not involve biological effectors. Rather, they involve the use of the extended, technological effector (the controller).

In our model, the phenomenon of phantom limbs typically arises by means very similar to extended effectors, although there may be alternate routes as well, given the diversity of the phenomenon. A phantom limb is a proprioceptive model of the variables and values associated with a limb or other body part that has been amputated, or even one that was never there to begin with (as in congenital aplasia). The phenomenon occurs for the overwhelming majority of amputees [Halligan, 2002] (96% of amputees, in one study; [Fraser *et al.*, 2001]). For example, a person may have the feeling of a hand that has been missing for decades. In one study, over half of participants reported that they could move their phantom limb like other limbs, and over a quarter reported that they felt they could manipulate and pick up things with it [Fraser *et al.*, 2001]. Many cases of phantom limbs seem to be associated with pain, making them a medical and neuropsychological issue.

Phantom limbs are diverse, with some fixed, some movable, and many even occasionally involving pain [Giummarra *et al.*, 2007]. Some phantoms are not even limbs at all, such as phantom breasts [Dijkstra *et al.*, 2007]. There may be several routes to phantom limbs depending on the specificities of the case; we consider a few of these below,

but we do not claim to have exhaustively accounted for every possible phantom phenomenon. Whereas in the plastic extension of the Current Body Schema, extended effectors variables are learned into the long-term memory component of the Current Body Schema, the effector variables for phantom limbs are mostly innate [Halligan, 2002]. That the model of limbs in the long-term component of the Current Body Schema is largely innate is shown by the existence of people who report phantom limbs, but who were born with a congenital lack of the limb [Brugger *et al.*, 2000]. Phantom limbs resemble plastically-extended effectors, however, insofar as they both associate already-existing proprioceptive information with a different effector. Phantom limbs get their proprioceptive information from adjacent areas of the human brain's homunculi or somatotopic maps. For example, "[i]n some patients, stimulation of selected areas on the face, surviving upper limb or trunk, reliably produce referred sensations to corresponding parts of the phantom limb" [Halligan, 2002, p. 263; see also Ramachandran *et al.*, 1992]. In addition, some information from the contralateral limb may affect the phantom limb [Giummarra *et al.*, 2007]. In a LIDA human agent, the effector values of such bodily regions would become associated with the already existing and innate limb variables once they stop receiving proprioceptive information of their own, in the same way that extended effectors like canes become associated with the values of preexisting variables.

The phenomenology of the phantom limb would arise when that effector and its associated values are cued into the CSM, for example in order to bind the variables of an instrument node in an event (which itself originates in PAM). If that hand node were to be selected by an attention codelet and become part of a coalition, and that coalition were to win the competition for conscious broadcast, then the agent would experience that effector and its values. For example, they may experience them as a hand opened up, bunched into a fist, and even sensations of pain. The phenomenology of the phantom will remain proprioceptive, and will lack a visual phenomenology.

In some cases, the body image may further modulate the phantom limb phenomenon. This is particularly salient in cases where visual feedback has an effect on the phantom limb. In such cases, visual feedback triggers a "fusion" of the body image (the perceived limb) with the phantom limb. This occurs frequently with prosthetic limbs, where the phantom and the prosthesis become merged [Giummarra *et al.*, 2007]. It has also occurred in surgery. For example, in a study by Melzack and Bromage [1973], participants had the brachial plexus temporarily anaesthetized. Nearly all reported a phantom arm within 20 to 40 minutes. The proprioceptive location of the phantom arm typically lay close to the torso, while their real arm was placed 90 degrees away from the body for surgery. The mismatch between the location of their phantom arm and their real, anaesthetized arm was corrected when the two "fused" as participants looked at their real arm. Since the Current Body Schema is proprioceptive and vestibular, it is not affected by visual information. The "fusion" would occur within the CSM. The Current Body Schema would be cued for the variables and values associated with the phantom arm in the CSM. The body image, particularly in the form of visual information about the arm within the CSM, will in this case trigger an update of the values of the phantom arm.

5. The Body Schema in LIDA in Relation to Several Experiments

Having laid out the foundations of the Current Body Schema and the Habitual Body Schema in LIDA, we will close this paper with a few examples of difficult cases from psychological experiments that can be well explained by LIDA. Some researchers have argued that the body image in humans can in some cases influence the body schema [Gallagher, 2005], especially by virtue of the dominance of the visual modality, but also through semantic elements. For example, in the rubber hand illusion, a subject is visually presented with an artificial hand in front of them [Botvinick & Cohen, 1998]. To begin with, the subject's real arm is hidden from view. The researcher physically stimulates the subject's arm, and makes the same stimulating motions over the artificial arm in the subject's visual field. Many subjects experience a feeling that their body ownership is of the rubber hand, and a common explanation is that, through visual dominance, a transfer of the body schema to the artificial arm was made. As Gallagher [2013] notes, while the feeling of the body seems to transfer to the rubber hand, the motor aspects of the hand are not typically subject to the same illusion. The variables in the Current Body Schema do not undergo an alteration. The subject still receives proprioceptive information from Sensory Memory. It is within the CSM, rather than in the Current Body Schema, that the rubber hand illusion would take effect. The activated hand in the Current Body Schema is interpreted in the CSM as the visually recognized hand activated in PAM. Generally, structure building codelets attempt to build hypotheses to best interpret the situation within the CSM. The interpretation in this case is, of course, mistaken, giving rise to the illusion. The structure building codelets build an interpretation that the hand the agent is seeing is the agent's own hand. This is a hypothesis that has probably always held true prior to the agent's entering into the artificial laboratory conditions. The Current Body Schema only receives pressure and vestibular information from Sensory Memory. Feelings are activated in PAM and can be interpreted, based on contextual factors, as being feelings about the agent's own hand. In other words, the body image does not truly modulate the body schema in this case.

Another case of the purported influence of body image on body schema was first discovered by Nielsen [1963] [see Gallagher, 2005]. In this experiment, subjects were instructed to draw a straight line. Their vision was altered in a device such that they actually saw an experimenter's hand, which they mistakenly identified as their own (similar to the rubber hand illusion). The experimenter then began to veer and curve the line drawn. Subjects experienced a mismatch between their proprioceptive sense of drawing a straight line and their visual body image of their hand (or what they misidentified as their hand). Several subjects

“described that they felt something strange and queer outside themselves, which pressed their hand to the right or resisted their free mobility. They suggested that ‘magnets’, unidentified forces’, ‘invisible traces under the paper’, or the like, could be the cause” [Nielsen, 1963, p. 228].

The Current Body Schema, of course, does not receive any pressure or other information from Sensory Memory related to being forced or moved, since that feeling is an illusion. Rather, an event is recognized in PAM. That event consists of:

- force or natural cause thematic role: magnet, unidentified force, etc.
- action thematic role: push
- object thematic role: self (I am pushed)
- subject thematic role: hand

In natural language, this event is ‘My hand is being pushed by a magnet.’ The Current Body Schema is a model of the variables and values associated with the hand, and is cued from the CSM. PAM sends the event as part of a percept to the CSM, where its remaining unbound variables can be bound. This includes the variables for the hand that takes the subject thematic role, variables which are bound using the copy of the hand-associated variables and values from the Current Body Schema (cf. Fig. 6). Meanwhile, PAM is receiving pressure information from Sensory Memory. The interpretation it makes of this pressure information depends on the event that is activated. If the event is ‘My hand is being pushed by a magnet,’ the pressure information may be interpreted as the feeling of an external force. The event that is activated, however, depends on the cultural background of the agent. If the event is ‘My hand is being controlled by a demon,’ a culturally plausible interpretation in some societies, then the pressure information would still be interpreted as the feeling of an external force, but that feeling would have a negative affective valence (it would be scary). If the feeling, as part of the percept, becomes part of a winning coalition, it can be consciously broadcast. In this way, the agent will consciously feel and be able to report being pushed by an outside force, even though no related information is being received in the Current Body Schema that would indicate such a force.

New Caledonian crows (*Corvus moneduloides*) have an ability to manufacture basic hook-like tools [Hunt, 1996; Taylor *et al.*, 2007]. In one experiment, researchers originally provided differently shaped wires with the aim of seeing whether a crow would figure out how to bend a wire to manufacture a similar tool to that which they make in the wild [Weir *et al.*, 2002]. The crow took the straight wire and bent it to make a hook. Although these corvids are known to make hook-like tools, this case demonstrated remarkable intelligence due to the novelty of the environment and the strangeness of the material (wire). This was later replicated with larger samples [Rutz *et al.*, 2016]. In the LIDA model, the crow would need to recognize its embodied capacities and the affordances available in the wire. It may learn the affordances of the wire through brief experimentation, but its embodied capacities are stored in the Habitual Body Schema. These capacities include the ability to bite and manipulate things with its beak and feet. Both the object’s affordances (once learned) and its embodied capacities (and their associated effectors) are recognized in PAM based on incoming information from Sensory Memory, and cues from the CSM. Once activated, they can be copied into structures in the CSM (*viz.*, percepts). Structure building codelets may assemble the embodied capacities and their associated effectors, the wire (also

recognized in PAM), and the affordance link from the wire to the embodied capacity or solicitation to action (e.g., a graspability affordance) all together into a structure. The Current Body Schema contributes with an updated model of the beak's and foot's current positions. If attention codelets pick this structure out as salient, it may win the competition for a conscious broadcast, and eventually lead to actions being selected for and performed in the action and learning phase of a cognitive cycle.

6. Conclusion

This paper has laid out the conceptual foundations of the body schema, as well as the body image, in LIDA, a functionally conscious cognitive architecture that computationally implements and fleshes out Globally Workspace Theory. Body schema has been an object of inquiry since the early 1900s and recent advances have been made in our knowledge of the underlying neural structures supporting it. What LIDA can contribute is an understanding of how an agent might computationally implement such a structure. Additionally, LIDA, as a computational implementation of Global Workspace Theory, can contribute to the study of the body schema by refining that concept into three separate functions, based on the functional role of consciousness in cognition. The Current Body Schema models the agent's effectors, whose positions are continually updated, while the Habitual Body Schema is a unified collection of the embodied capacities for action that the agent has, either through programming, evolution, or learning. Additionally, the Current Body Schema does not solely rely on a conscious broadcast for updating, whereas the Habitual Body Schema relies on conscious broadcasts for updating. The final function is the Online Body Schema, which incorporates "dorsal" stream, never-conscious aspects of the 'body schema,' such as highly granular effector positions that are used in online control [cf. Gallagher, 2005]. This information is not integrated into a model, and feeds directly from Sensory Memory to the Sensory Motor System.

We believe these distinctions within LIDA ought to hold for biological agents, as well. The embodied capacities of an agent must be functionally different from the online modeling of their effectors because of the different ways they are updated. Since the Habitual Body Schema updates through conscious broadcasts, it can only update at most once per cognitive cycle. The Current Body Schema's long-term memory component, which stores variables such as joints and effectors, can similarly be updated through conscious broadcasts. However, the bulk of the Current Body Schema's work is in the online modeling of current effector positions (viz., the values of variables). This updating is continuous and does not solely rely on conscious broadcasts. In sum, the idea that there is 'a body schema' is untenable, according to a Global Workspace Theory of consciousness. There is a continually-updated proprioceptive and vestibular model of the body, a long-term memory of embodied capacities and effectors, and an online coupling of the sensorimotor system from early perception.

References

- Agrawal, P., and Franklin, S. [2014] "Multi-layer cortical learning algorithms," in *2014 IEEE Symposium on Computational Intelligence, Cognitive Algorithms, Mind, and Brain (CCMB)* (IEEE), pp. 141–147.
- Baars, B. J. [1988] *A Cognitive Theory of Consciousness* (Cambridge University Press).
- Baars, B. J. [2019] *On Consciousness: Science and Subjectivity* (Nautilus, New York).
- Baars, B. J., Franklin, S., and Ramsøy, T. Z. [2013] Global workspace dynamics: Cortical “binding and propagation” enables conscious contents, *Front. Psychol.* **4**, 200, doi:10.3389/fpsyg.2013.00200.
- Baddeley, A. [2007] *Working Memory, Thought, and Action* (Oxford University Press).
- Barsalou, L. W. [1999] Perceptual symbol systems, *Behav. Brain Sci.* **22**(4), 577–660.
- Blackmore, S. [2000] *The Meme Machine* (Oxford University Press).
- Boly, M., Seth, A. K., Wilke, M., Ingmundson, P., Baars, B., Laureys, S., Edelman, D., and Tsuchiya, N. [2013] Consciousness in humans and non-human animals: Recent advances and future directions, *Front. Psychol.* **4**, 625, doi:10.3389/fpsyg.2013.00625.
- Bonnier, P. [1900] *L'orientation* (Carré et Naud, Paris).
- Botvinick, M., and Cohen, J. [1998] Rubber hands ‘feel’ touch that eyes see, *Front. Psychol.* **391**, 756, doi:10.1038/35784.
- Brugger, P., Kollias, S. S., Müri, R. M., Crelier, G., Hepp-Reymond, M.-C., and Regard, M. [2000] Beyond re-membering: Phantom sensations of congenitally absent limbs, *P. Natl. Acad. Sci. USA* **97**(11), 6167–6172, doi:10.1073/pnas.100510697.
- Conway, M. A. [2001] Sensory–perceptual episodic memory and its context: Autobiographical memory, *Phil. Trans. Roy. Soc. B* **356**(1413), 1375–1384, doi:10.1098/rstb.2001.0940.
- Cutsuridis, V., Hussain, A., and Taylor, J. G. [2011] *Perception-Action Cycle: Models, Architectures, and Hardware* (Springer).
- Damasio, A. [1999] *The Feeling of What Happens: Body and Emotion in the Making of Consciousness* (Harcourt Brace, New York).
- de Vignemont, F. [2010] Body schema and body image—Pros and cons, *Neuropsychologia* **48**(3), 669–680, doi:10.1016/j.neuropsychologia.2009.09.022.
- Dehaene, S. [2014] *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts* (Penguin, New York).
- Dijkerman, H. C., and De Haan, E. H. [2007] Somatosensory processing subserving perception and action: Dissociations, interactions, and integration, *Behav. Brain Sci.* **30**(2), 224–230.
- Dijkstra, P. U., Rietman, J. S., and Geertzen, J. H. [2007] Phantom breast sensations and phantom breast pain: A 2-year prospective study and a methodological analysis of literature, *Eur. J. Pain* **11**(1), 99–108, doi:10.1016/j.ejpain.2006.01.002.
- Dijkstra, T. M. H., Schöner, G., and Gielen, C. [1994] Temporal stability of the action-perception cycle for postural control in a moving visual environment, *Exp. Brain Res.* **97**(3), 477–486, doi:10.1007/BF00241542.
- Doesburg, S. M., Green, J. J., McDonald, J. J., and Ward, L. M. [2009] Rhythms of consciousness: Binocular rivalry reveals large-scale oscillatory network dynamics mediating visual perception, *PLOS ONE* **4**(7), doi:10.1371/journal.pone.0006142.

- Dong, D., and Franklin, S. [2015] A new action execution module for the learning intelligent distribution agent (LIDA): The sensory motor system, *Cogn. Comput.* **7**(5), 552–568, doi:10.1007/s12559-015-9322-3.
- Duhamel, J.-R., Colby, C. L., and Goldberg, M. E. [1998] Ventral intraparietal area of the macaque: congruent visual and somatic response properties, *J. Neurophysiol.* **79**(1), 126–136, doi:10.1152/jn.1998.79.1.126.
- Ericsson, K. A., and Kintsch, W. [1995] Long-term working memory, *Psychol. Rev.* **102**(2), 211, doi:10.1037/0033-295X.102.2.211.
- Erlhagen, W., and Schöner, G. [2002] Dynamic field theory of movement preparation, *Psychol. Rev.* **109**(3), 545, doi:10.1037/0033-295X.109.3.545.
- Faghihi, U., McCall, R., and Franklin, S. [2012] A computational model of attentional learning in a cognitive agent, *Biologically Inspired Cognitive Architectures* **2**, 25–36, doi:10.1016/j.bica.2012.07.003.
- Fogassi, L., Gallese, V., Fadiga, L., Luppino, G., Matelli, M., and Rizzolatti, G. [1996] Coding of peripersonal space in inferior premotor cortex (area F4), *J. Neurophysiol.* **76**(1), 141–157, doi:10.1152/jn.1996.76.1.141.
- Franklin, S. [1995] *Artificial Minds* (MIT Press, Cambridge, MA).
- Franklin, S. [2000] Deliberation and voluntary action in ‘conscious’ software agents, *Neural Netw. World* **10**, 505–521.
- Franklin, S., and Baars, B. [2010] Two varieties of unconscious processes, in E. Perry, C. Daniel, L. Fiona, and A. Heather (eds.), *New Horizons in the Neuroscience of Consciousness* (John Benjamins, New York), pp. 91–102.
- Franklin, S., Baars, B. J., Ramamurthy, U., and Ventura, M. [2005] The role of consciousness in memory, *Brains, Minds, & Media* **1**.
- Franklin, S., and Graesser, A. [1997] Is it an agent, or just a program? a taxonomy for autonomous agents, in J. Müller, M. J. Wooldridge, and N. R. Jennings (eds.), *Intelligent Agents III: Agent Theories, Architectures, and Languages* (Springer, Berlin), pp. 21–35.
- Franklin, S., Kelemen, A., and McCauley, L. [1998] "IDA: A cognitive agent architecture," in *SMC'98 Conference Proceedings. 1998 IEEE International Conference on Systems, Man, and Cybernetics* (IEEE), **3**, 2646–2651.
- Franklin, S., Madl, T., D’Mello, S., and Snider, J. [2014] LIDA: A systems-level architecture for cognition, emotion, and learning, *IEEE T. Auton. Ment. De.* **6**(1), 19–41, doi:10.1109/TAMD.2013.2277589.
- Franklin, S., Madl, T., Strain, S., Faghihi, U., Dong, D., Kugele, S., Snider, J., Agrawal, P., and Chen, S. [2016] A LIDA cognitive model tutorial, *Biol. Inspir. Cogn. Arc.* **16**, 105–130, doi:10.1016/j.bica.2016.04.003.
- Franklin, S., Strain, S., McCall, R., and Baars, B. [2013] Conceptual commitments of the LIDA model of cognition, *JAGI* **4**(2), 1–22, doi:10.2478/jagi-2013-0002.
- Fraser, C. M., Halligan, P. W., Robertson, I. H., and Kirker, S. G. B. [2001] Characterising phantom limb phenomena in upper limb amputees, *Prosthet. Orthot. Int.* **25**(3), 235–242, doi:10.1080/03093640108726607.
- Freeman, W. [2000] *Neurodynamics: An Exploration in Mesoscopic Brain Dynamics* (Springer).
- Freeman, W. J. [2002] "The limbic action-perception cycle controlling goal-directed animal behavior," in *Proceedings of the 2002 International Joint Conference on Neural Networks. IJCNN'02* (IEEE), **3**, 2249–2254.

- Gallagher, S. [2005] *How the Body Shapes the Mind* (Oxford University Press).
- Gallagher, S. [2013] The enactive hand, in Z. Radman (ed.), *The Hand, an Organ of the Mind: What the Manual Tells the Mental* (MIT Press, Cambridge, MA), pp. 209–225.
- Gallagher, S., and Cole, J. [1995] Body image and body schema in a deafferented subject, *J. Mind Behav.* **16**, 369–389.
- Gibson, J. J. [2015] *The Ecological Approach to Visual Perception* (Classic Ed.) (Psychology Press, New York), (Original work published 1979).
- Giummarra, M. J., Gibson, S. J., Georgiou-Karistianis, N., and Bradshaw, J. L. [2007] Central mechanisms in phantom limb perception: the past, present and future, *Brain Res. Rev.* **54**(1), 219–232, doi:10.1016/j.brainresrev.2007.01.009.
- Goodale, M. A., and Milner, A. D. [1992] Separate visual pathways for perception and action, *Trends Neurosci.* **15**(1), 20–25, doi:10.1016/0166-2236(92)90344-8.
- Goodale, M. A., Milner, A. D., Jakobson, L. S., and Carey, D. P. [1991] A neurological dissociation between perceiving objects and grasping them, *Nature* **349**(6305), 154–156, doi:10.1038/349154a0.
- Graziano, Michael S., Yap, G. S., and Gross, C. G. [1994] Coding of visual space by premotor neurons, *Science* **266**(5187), 1054–1057, doi:10.1126/science.7973661.
- Graziano, Michael SA, and Webb, T. W. [2015] The attention schema theory: a mechanistic account of subjective awareness, *Front. Psychol.* **6**, 500, doi:10.3389/fpsyg.2015.00500.
- Haggard, P., and Wolpert, D. M. [2005] Disorders of body scheme, in H.-J. Freund, M. Jeannerod, M. Hallett, and R. Leiguarda (eds.), *Higher-Order Motor Disorders* (Oxford University Press), .
- Halligan, P. W. [2002] Phantom limbs: the body in mind, *Cogn. Neuropsychiatry* **7**(3), 251–269, doi:10.1080/13546800244000111.
- Hawkins, J., Ahmad, S., and Dubinsky, D. [2010] *Hierarchical temporal memory including HTM cortical learning algorithms* (Numenta, Inc., Palto Alto, CA).
- He, K., Zhang, X., Ren, S., and Sun, J. [2015] "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1026–1034.
- Head, H., and Holmes, G. [1911] Sensory disturbances from cerebral lesions, *Brain* **34**(2–3), 102–254, doi:10.1093/brain/34.2-3.102.
- Hoffmann, M., Marques, H., Arieta, A., Sumioka, H., Lungarella, M., and Pfeifer, R. [2010] Body schema in robotics: A review, *IEEE T. Auton. Ment. De.* **2**(4), 304–324, doi:10.1109/TAMD.2010.2086454.
- Hofstadter, D. R., and Mitchell, M. [1994] The Copycat project: A model of mental fluidity and analogy-making, in J. A. Barnden and K. J. Holyoak (eds.), *Advances in Connectionist and Neural Computation Theory* (Intellect Ltd., Chicago), Vol. 2, pp. 205–267.
- Hunt, G. R. [1996] Manufacture and use of hook-tools by New Caledonian crows, *Nature* **379**, 249–251, doi:10.1038/379249a0.
- James, W. [1983] *The Principles of Psychology* (Harvard University Press, Cambridge, MA), (Original work published 1890).
- Jeannerod, M. [2010] De l’image du corps à l’image de soi, *Rev. Neuropsychol.* **2**(3), 185–194, doi:10.3917/rne.023.0185.
- Kaas, J. H., and Hackett, T. A. [1999] “What” and “where” processing in auditory cortex, *Nat. Neurosci.* **2**(12), 1045–1047, doi:10.1038/15967.
- Kaur, S. [2011] *When to pay attention? Asynchrony requires a trigger* (The University of Memphis).

- Kelso, J. S. [1995] *Dynamic Patterns: The Self-Organization of Brain and Behavior* (MIT Press, Cambridge, MA).
- Kozma, R., and Freeman, W. J. [2016] *Cognitive Phase Transitions in the Cerebral Cortex-Enhancing the Neuron Doctrine by Modeling Neural Fields* (Springer).
- Kugele, S., and Franklin, S. [forthcoming] Learning in LIDA, *Cogn. Syst. Res.*
- Kugele, S., and Franklin, S. [2020] "'Conscious' multi-modal perceptual learning for grounded simulation-based cognition," in S. Denison, M. Mack, Y. Xu, and B. C. Armstrong (eds.), *Proceedings of the 42nd annual conference of the Cognitive Science Society* (Cognitive Science Society, Austin, TX), pp. 2459–2465.
- Longo, M. R., Azañón, E., and Haggard, P. [2010] More than skin deep: body representation beyond primary somatosensory cortex, *Neuropsychologia* **48**(3), 655–668, doi:10.1016/j.neuropsychologia.2009.08.022.
- Madl, T., Baars, B. J., and Franklin, S. [2011] The timing of the cognitive cycle, *PLOS ONE* **6**(4), doi:10.1371/journal.pone.0014803.
- Madl, T., and Franklin, S. [2012] "A LIDA-based model of the attentional blink," in *ICCM 2012 Proceedings*, pp. 283–288.
- Madl, T., Franklin, S., Chen, K., Montaldi, D., and Trapp, R. [2014] Bayesian integration of information in hippocampal place cells, *PLOS ONE* **9**(3), doi:10.1371/journal.pone.0089762.
- Madl, T., Franklin, S., Chen, K., Montaldi, D., and Trapp, R. [2016] Towards real-world capable spatial memory in the LIDA cognitive architecture, *Biol. Inspir. Cogn. Arc.* **16**, 87–104, doi:10.1016/j.bica.2016.02.001.
- Maravita, A., and Iriki, A. [2004] Tools for the body (schema), *Trends Cogn. Sci.* **8**(2), 79–86, doi:10.1016/j.tics.2003.12.008.
- Mashour, G. A., Roelfsema, P., Changeux, J.-P., and Dehaene, S. [2020] Conscious processing and the global neuronal workspace hypothesis, *Neuron* **105**(5), 776–798, doi:10.1016/j.neuron.2020.01.026.
- McCall, R. J. [2014] *Fundamental motivation and perception for a systems-level cognitive architecture* (The University of Memphis).
- McCall, R. J., Franklin, S., and Friedlander, D. [2010] "Grounded event-based and modal representations for objects, relations, beliefs, etc.," in *Twenty-Third International FLAIRS Conference*.
- McNeill, D., Quaeghebeur, L., and Duncan, S. [2010] IW—"The man who lost his body," in S. Gallagher and D. Schmicking (eds.), *Handbook of Phenomenology and Cognitive Science* (Springer), pp. 519–543.
- Melzack, R., and Bromage, P. R. [1973] Experimental phantom limbs, *Exp. Neurol.* **39**(2), 261–269, doi:10.1016/0014-4886(73)90228-8.
- Merleau-Ponty, M. [2012] *Phenomenology of Perception* (D. A. Landes, Trans.) (Routledge, London), (Original work published 1945).
- Munk, H. [1890] *Über Die Functionen Der Grosshirnrinde* (A. Hirschwald).
- Neisser, U. [1976] *Cognition and Reality: Principles and Implications of Cognitive Psychology*. (WH Freeman, San Francisco).
- Newell, A. [1990] *Unified Theories of Cognition* (Harvard University Press, Cambridge, MA).
- Nielsen, T. I. [1963] Volition: A new experimental approach, *Scand. J. Psychol.* **4**(1), 225–230.

- Paillard, J. [1999] Body schema and body image: A double dissociation, in G. N. Gantchev, S. Mori, and J. Massion (eds.), *Motor Control, Today and Tomorrow* (Academic Publishing House, Sofia), pp. 197–214.
- Ramachandran, V. S., Rogers-Ramachandran, D., Stewart, M., and Pons, T. P. [1992] Perceptual correlates of massive cortical reorganization, *Science* **258**(5085), 1159–1159, doi:10.1126/science.1439826.
- Ramamurthy, U., Franklin, S., and Agrawal, P. [2012] Self-system in a model of cognition, *International Journal of Machine Consciousness* **4**(02), 325–333, doi:10.1142/S1793843012400185.
- Russell, S., and Norvig, P. [2016] *Artificial Intelligence: A Modern Approach* (3rd ed.) (Pearson).
- Rutz, C., Sugawara, S., Van der Wal, J. E., Klump, B. C., and St Clair, J. J. [2016] Tool bending in New Caledonian crows, *Roy. Soc. Open Sci.* **3**(8), 160439, doi:10.1098/rsos.160439.
- Ryan, K., Agrawal, P., and Franklin, S. [2020] The pattern theory of self in artificial general intelligence: A theoretical framework for modeling self in biologically inspired cognitive architectures, *Cogn. Syst. Res.* **62**, 44–56.
- Ryle, G. [2009] *The Concept of Mind* (60th Anniversary Ed.) (Routledge, New York), (Original work published 1949).
- Scherer, K. R. [2001] Appraisal considered as a process of multilevel sequential checking, in K. R. Scherer, A. Schorr, and T. Johnstone (eds.), *Appraisal Processes in Emotion: Theory, Methods, Research* (Oxford University Press), pp. 92–120.
- Schöner, G. [2008] Dynamical systems approaches to cognition, in R. Sun (ed.), *Cambridge Handbook of Computational Psychology* (Cambridge University Press), pp. 101–126.
- Strain, S. F., Franklin, S., Baars, B. J., and Jeuniaux, P. [in preparation] *Brain rhythms, cognitive cycles and mental moments*.
- Strain, Stephen, and Franklin, S. [2011] Modeling medical diagnosis using a comprehensive cognitive architecture, *J. Healthc. Eng.* **2**(2), 241–258, doi:10.1260/2040-2295.2.2.241.
- Strain, Steve, Kugele, S., and Franklin, S. [2014] "The Learning Intelligent Distribution Agent (LIDA) and Medical Agent X (MAX): Computational intelligence for medical diagnosis," in *2014 IEEE Symposium on Computational Intelligence for Human-like Intelligence (CIHLI)* (IEEE), pp. 1–8.
- Taylor, A. H., Hunt, G. R., Holzhaider, J. C., and Gray, R. D. [2007] Spontaneous metatool use by New Caledonian crows, *Curr. Biol.* **17**(17), 1504–1507, doi:10.1016/j.cub.2007.07.057.
- Ungerleider, L. G., and Mishkin, M. [1982] Two cortical visual systems, in D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (eds.), *Analysis of Visual Behavior* (MIT Press, Cambridge, MA), pp. 549–586.
- Varela, F. J., Thompson, E., and Rosch, E. [2016] *The Embodied Mind: Cognitive Science and Human Experience* (Revised Ed.) (MIT Press, Cambridge, MA), (Original work published 1991).
- Varela, F., Lachaux, J.-P., Rodriguez, E., and Martinerie, J. [2001] The brainweb: Phase synchronization and large-scale integration, *Nat. Rev. Neurosci.* **2**(4), 229–239, doi:10.1038/35067550.
- Weir, A. A., Chappell, J., and Kacelnik, A. [2002] Shaping of hooks in New Caledonian crows, *Science* **297**(5583), 981–981, doi:10.1126/science.1073433.