Chapter 8
Continuity and the Flow of Time: A Cognitive Science Perspective

Tamas Madl, Stan Franklin, Javier Snaider, and Usef Faghihi

Abstract Modern tools and methods of cognitive science, such as brain imaging or computational modeling, can provide new insights for age-old philosophical questions regarding the nature of temporal experience. This chapter aims to provide an overview of functional consciousness and time perception in brains and minds (Sect. 8.2), and to describe a computational cognitive architecture partially implementing these phenomena (Sects. 8.3, 8.4 and 8.5), and its comparison with data from human behavioral experiments (Sect. 8.6).

8.1 Introduction

The life (existence?) of each of us as human being consists introspectively of a continual flow of conscious or consciously-mediated experience over time. This assertion seems to raise all sorts of questions. What, if anything, "out there" is being experienced? Is this continual flow "really" continual, or do we create the illusion of continuity from a rapid sequence of frames? For the latter, what can we say about the structure of one of these frames? And, what is meant by "over time"? In this chapter we propose possible answers to these questions derived from cognitive neuroscience with the help of an integrated, systems-level cognitive model of how minds work. We hypothesize an answer to the first question above by assuming the existence of a real, physical world that can only be known to us in part through our various senses. We assume that when a tree falls in the forest there are vibrations of the air, but sound would exist only in the mind of some organism (or artificial

---

T. Madl (✉)
School of Computer Science, University of Manchester, Manchester, UK

Austrian Research Institute for Artificial Intelligence, Vienna, Austria
e-mail: thomas.madl@gmail.com

S. Franklin • J. Snaider
Institute for Intelligent Systems, University of Memphis, Memphis, TN, USA

U. Faghihi
Department of Computing and Technology, Cameron University, Lawton, OK, USA

B. Mölder et al. (eds.), *Philosophy and Psychology of Time*, Studies in Brain and Mind 9, DOI 10.1007/978-3-319-22195-3_8

agent?) equipped with an appropriate auditory sense and concomitant cognitive abilities with which to represent and perhaps understand the sound. This process can be thought of as the organism (agent) cognitively modeling its world, at least in part. We say "in part" since the frequency range of the auditory sensory apparatus is typically limited. This view leads us to hypothesize perception as a creative cognitive process at least partially dependent on our senses.

In contrast to other modes of perception, such as taste, color or sound, there is no specific physical sense for time. However, we perceive information from the other senses *over* time; we perceive time in response to change in our sensations. Thus time is viewed here as being fundamental to our cognitive processes. Instead of asking "How can time be perceived?" we will consider "How can a sense of time be produced by a cognitive system?" We hypothesize that our perception of time is constructed by cognitive processes of an organism or other agent. In this chapter we propose to explore possible such processes for producing a sense of time.

Philosophers have proposed that our phenomenal flow of consciousness over time as composed of individual frames (episodes of experiencing), and have given three different accounts of their structure. One of them refers to these three as the cinematic, retentional, and extensional models (Dainton 2010). The *cinematic model* views our introspective flow of time as consisting of a continuous succession of very brief, motion-free frames lacking any (or significant) extension. The *retentional model* takes an entirely similar view, except that the content of each frame is allowed to refer to frames representing intervals. Thus, these contents can represent, though not constitute, temporally extended time intervals. As the name would suggest, the *extensional model* considers each frame to have a brief temporal extension, to comprise an interval of time.

According to our first hypothesis above, what we know of the presumed outside physical world is constructed by us from our conscious perception. Our introspection tells us that this ongoing stream of conscious perception is continuous, extended over time, without gaps, other than those produced by deep sleep. Each of the three models discussed above assume the continuity of our perception of time, But, might it be that our introspection has deceived us as it does when we perceive a sufficiently rapid sequence of still frames in a movie theater as continuous in time? Here we will argue that this is precisely the case, that we in fact construct our apparently continuous flow of conscious perception from a rapid (5–10 Hz) sequence of discrete frames of conscious content (Madl et al. 2011). This view is consistent with recent neuroscience results, which suggest conscious access to arise from periodic phases of information integration (Baars et al. 2013; Dehaene et al. 2014; VanRullen et al. 2014). We will also argue that events in the same frame are consciously perceived as simultaneous (Snaider et al. 2012), and that each frame allows some small amount of motion[1] within its duration (VanRullen and Koch 2003).

---

[1] Confirmed by Christof Koch in personal communication with one of the authors.

8.2 The Cognitive Neuroscience of Consciousness and Time Perception: A Brief Introduction

The cognitive neuroscience of conscious perception is concerned with trying to find minimal neuronal mechanisms which distinguish "conscious" mental states from unconscious ones, as reported by experimental subjects (Crick and Koch 1990; Koch 2004). Consciousness is a difficult phenomenon to study, due to its intrinsically introspective nature; and its experimental investigation is further complicated by some ambiguity as to what exactly is meant by the term.

In this chapter, we will talk only about the functionally relevant aspects of consciousness—"functional consciousness" or "access consciousness" in neuroscience (Block 1995; Baars 2005; Dehaene and Changeux 2003). We will neglect phenomenal consciousness or "qualia" (e.g., what experiences might feel like—such as taste qualia in the case of wine) (Dennett 1988), since it is notoriously difficult to study in a formal, systematic setting. In contrast, functional or access consciousness are described in terms of the availability of mental states to higher-level cognitive processes. If a state or percept enters an agent's (biological or artificial) functional consciousness, it can influence decision making (for example, when a subject correctly presses a button in response to a stimulus, or verbally reports his perception of that stimulus).

The brain mechanisms underlying functional consciousness can be studied in paradigms contrasting conscious and non-conscious brain states. Example conditions in which visual stimuli can be presented such that they cannot be consciously perceived include visual illusions (Kim and Blake 2005), masking[2] (Kouider and Dehaene 2007), or binocular rivalry[3] (Doesburg et al. 2009; Pitts and Britz 2011). Such paradigms help investigate the "neural correlates of consciousness" by identifying which parts of brain activity patterns might correspond to conscious percepts, as opposed to unconscious percepts. Apart from sensory areas such as the visual cortex, brain imaging experiments have indicated that prefrontal and posterior parietal networks exhibit activation strongly correlated with visual awareness (Rees et al. 2002). Unfortunately, there does not seem to be a set of brain areas exclusively involved with conscious processing (Dehaene et al. 2014), casting in doubt the idea of a specific cognitive processor being responsible for consciousness. For example, even areas associated with high-level cognition such as task switching in the prefrontal cortex can be triggered non-consciously (Lau and Passingham 2007; Reuss et al. 2011).

It has been suggested that the difference between conscious and non-conscious processing might be due to differences in temporal coherence or synchronization of

---

[2] Masking involves the elimination of the visibility of one briefly presented stimulus by the presentation of a second brief stimulus (the "mask").

[3] Binocular rivalry involves presentation of different visual stimuli to the left and right eyes of subjects. In this paradigm, conscious perception alternates between the two stimuli—see also Fig. 8.1.

neural activity in the same anatomical substrate (Melloni et al. 2007; VanRullen et al. 2014; Singer 2011) (Fig. 8.1 top). Unlike unconscious perception, which involves local coordination and propagation of sensory information to progressively higher-level representations, conscious perception might require global coordination of widely distributed neurons. This global coordination might be facilitated by long-distance synchronization (Dehaene et al. 2006, 2014), which can temporarily integrate neurons into coherent assemblies and facilitate long-range communication between distant brain areas. There is a large amount of empirical support for this idea—for example, cortical and thalamic neurons discharge synchronously during wakefulness (Steriade 2006) and synchrony is enhanced for consciously perceived

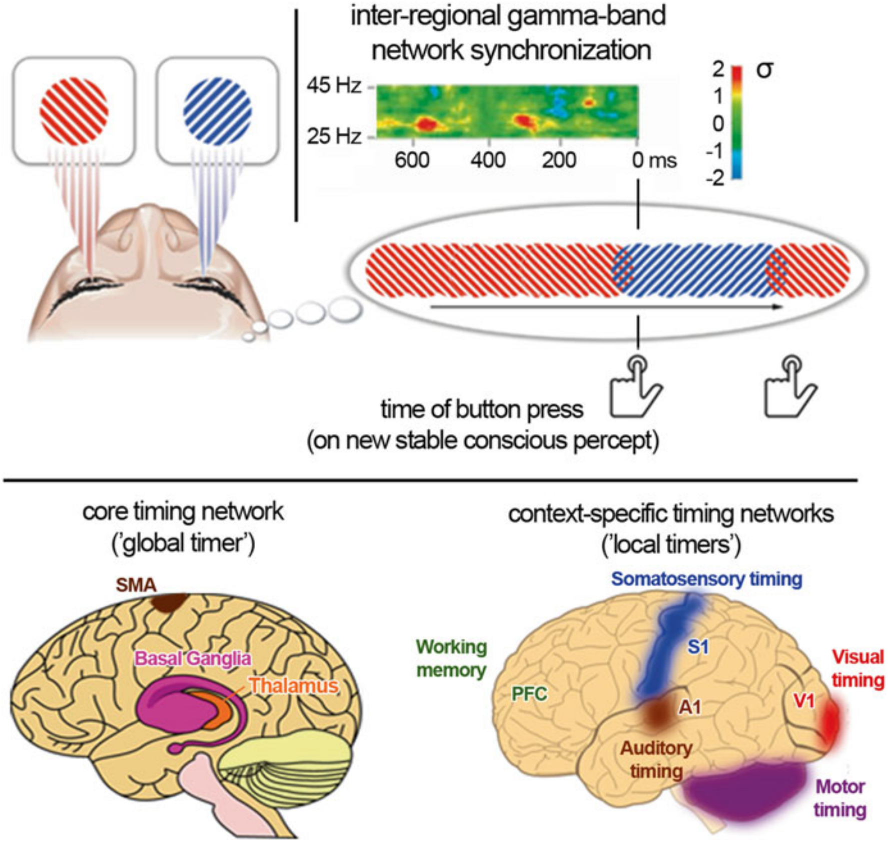

Fig. 8.1 Oscillatory synchrony, and major neural correlates of time perception. *Top*: schematic of a binocular rivalry experiment, and periods of synchrony dissolving and re-forming at each conscious episode (Synchrony data from Doesburg et al. (2009), head image from Dieter and Tadin (2011)). *Bottom*: Neural bases of the core timing network (thalamus, cortex, *BG* basal ganglia and *SMA* supplementary motor area), and example context-specific timing networks including the visual cortex (*V1*), and auditory and somatosensory cortex (*A1* and *S1*), and the cerebellum (Brain images modified from Wikimedia 2009, 2010, based on Merchant et al. 2013)

stimuli (Palva et al. 2005). In masking paradigms, increased gamma frequency band synchrony is induced only by words reported as perceived by subjects (Melloni et al. 2007). Furthermore, in the binocular rivalry paradigm, gamma-synchronous activity locked to an ongoing theta rhythm precedes perceptual switching (as indicated by subjects pressing a button when the stimulus which they are conscious of changes (Doesburg et al. 2009)). Finally, neural activity is globally disintegrated and fragmented in time in unconscious subjects, e.g., those undergoing anesthesia (Lewis et al. 2012), and awake vs. unconscious states can be reliably separated using a measure of the amount of information shared by distant cortical sites (Casali et al. 2013). See Singer (2011) for further evidence.

More recent theories of consciousness are consistent with such empirical results, suggesting consciousness to be a process involving large-scale brain activity, instead of attempting to confine it to one or few brain areas. Prevalent examples include the Global Workspace Theory (which proposes that consciousness is facilitated by a fleeting memory capacity enabling access between spatially separate brain functions (Baars 2005; Dehaene and Changeux 2003; Baars et al. 2013)), and Neural Darwinism (which proposes that conscious experience arises from reentrant neural activity in the thalamocortical system (Edelman and Tononi 2000)).

If periodic large-scale integration via oscillatory synchrony is indeed necessary for conscious processing, then this would have important implications for the structure of experience. Most importantly for the present topic, it would imply that consciousness is a discrete mechanism, since large-scale synchrony in brains is not continuously present, but has been observed to arise and dissolve periodically several times per second (VanRullen et al. 2014; Doesburg et al. 2009; Singer 2011; Madl et al. 2011). Although there is no definite answer to whether consciousness is discrete or continuous, there is substantial neuroscientific (see above) as well as psychophysical evidence supporting the discrete hypothesis, such as the wagon wheel illusion, in which a turning wheel is perceived to rotate in the wrong direction, presumably due to discrete sampling (see VanRullen and Koch (2003) for a review of psychophysical evidence of discrete perception).

Another important consequence of a periodic mechanism facilitating consciousness is that such a mechanism can be used to estimate the durations of events by counting the occurrences of cycles, similarly to pulse accumulator models of time perception in psychology (Grondin 2010). Large-scale oscillatory activity in a cortico-thalamic-basal ganglia circuit has been described as the "core timer" of the brain (Merchant et al. 2013) (the cortico-thalamic system has also been suggested to be involved with conscious perception (Edelman et al. 2011; Steriade 2006)). There is also substantial psychophysical evidence for the existence of a global timing mechanism, e.g., the observation that the variability of interval timing is proportional to the duration of the interval across a large number of tasks, sensory modalities, and species (Gibbon et al. 1997; Buhusi and Meck 2005).

Apart from such a central timing mechanism, there is evidence for "local timers," brain areas with neurons able to measure temporal intervals: see Fig. 8.1 bottom (it has even been argued that timing is a ubiquitous ability of cortical networks, and that a central clock might not be needed (Karmarkar and Buonomano 2007)). Cells associated with temporal processing in the medial premotor cortex are one example confirmed by recording studies (different neurons in this area react most strongly to different time intervals preceding an action such as a button press (Zarco et al. 2009)). Local timers also include several sensorimotor areas with their own local oscillatory cycles, such as the visual, auditory and somatosensory cortices for timing stimuli perceived in these modalities, or the cerebellum for motor timing (see Merchant et al. 2013 for a more comprehensive discussion).

8.3 Models of Time Perception

Here, we will focus on three main aspects of time from the point of view of cognition, namely succession, duration, and temporal perspective (Block 2014). Succession refers to a sequence of events which can be used to perceive temporal order and successiveness. Duration denotes a length of time during which an event might persist, or between events. Temporal perspective in turn addresses the separation of events into past, present, and future. Below, we will discuss a model of time perception focusing on succession and duration, which accounts for these concepts, as well as others including continuity, the duration of the immediate present, perceived length of time.

Many perspectives model time perception. At the end of nineteenth century, William James (1890) developed one of the first, which is relevant to this work. However, most cognitive models that try to explain time are only focus on one or two aspects of it. For example, Michon (1990) studied duration of events, and Block (2014) the sequence of events. Well-known psychological models focusing on duration include the scalar expectancy theory and the pulse accumulator model (Gibbon et al. 1984; Buhusi and Meck 2005). These models use a pacemaker, generating pulses at regular intervals, and a pulse accumulator to estimate event durations. The accumulator facilitates the estimation of event durations by storing the pulses generated by the pacemaker, and comparing them to pulses in a reference memory. Other authors, including Boltz (1995), Grondin (2010), Zakay and Block (1996, 1997), Zakay (1992), and Zakay et al. (1994) describe how the structure of an event influences our perception of its duration. In particular, they consider how the event structure and its complexity affects the accuracy of duration judgments. Most prior work studies event duration perception on the order of magnitude of dozens of seconds or more, whereas this chapter focuses on shorter durations.

In neuroscience literature, time perception is most commonly used to refer to the perception of event duration (Ivry and Schlerf 2008), although some authors including Eagleman (2008) adopt a more general perspective, accounting for duration as well as perception time scale and sequence. Studies on the perception of time abound in both the neuroscience and the behavioral literature. Some are related to memory processes, the order of events as we experience them. They distinguish recalling when an experienced event happened from estimating its duration. Others are related to consciousness, the awareness of subjective time. Still others are concerned with time in relation to sensory processing, for example the processing of speech, music and successive visual images. Grondin (2010) offers pointers to the literatures of each of these, as well as many others. Ivry and Schlerf (2008) contribute a review of dedicated and intrinsic models of time perception.

8.4 Global Workspace Theory and the LIDA Cognitive Architecture

In contrast to most previous models of time, which are limited to one or few cognitive phenomena, our model is based on a general model of cognition: LIDA (Learning Intelligent Distribution Agent), a conceptual and computational cognitive architecture partially implementing and fleshing out the Global Workspace Theory (GWT) of consciousness and a number of other prevalent cognitive science and neuroscience theories, including Anderson (2003), Glenberg and Robertson (2000), Varela et al. (1991), perceptual symbol systems (Barsalou 1999), working memory (Baddeley and Hitch 1974), memory by affordances (Glenberg 1997), long-term working memory (Ericsson and Kintsch 1995), transient episodic memory (Conway 2002), and Sloman's H-CogAff cognitive architecture (Sloman 1999).

8.4.1 Global Workspace Theory

Among different theories of cognition, we choose to work from Baars' (1997) GWT, a prevalent psychological and neurobiological theory of consciousness. According to the GWT, the nervous system is a distributed parallel system incorporating many specialized processes. Various coalitions of these specialized processes facilitate making sense of sensory data currently coming in from the environment. Other coalitions sort through the results of this initial processing and pick out items requiring further attention. In the competition for attention a winner emerges, and occupies the global workspace, the winning contents of which are presumed to be at least functionally conscious. The presence of a predator, enemy, or imminent danger should be expected, for example, to win the competition for attention. However, an unexpected loud noise might well usurp consciousness momentarily even in one of these situations. The global workspace contents are broadcast to processes throughout the nervous system in order to recruit an action or response to this salient aspect of the current situation. The contents of this global broadcast also enable many modes of learning, which explains why it needs to be global. This broadcast provides large-scale integration via access consciousness as discussed in Sects. 8.1 and 8.2 above. We hypothesize that it is accomplished through oscillatory synchrony (Baars et al. 2013). We will argue that Learning Intelligent Distribution Agent (LIDA), which implements Baars' GWT, may be suitable as an underlying cognitive architecture with which to explicate and investigate ideas and hypothesis regarding time.

8.4.2 The LIDA Cognitive Architecture

Autonomous agents (including humans, animals and artificial agents) have to frequently sample (sense) their environments and choose appropriate responses (actions). Agent's "lives" can be thought of as consisting of a sequences of such cycles, which we call cognitive cycles. Each such cycle consists of units of sensing, attending and acting. Cognitive cycles can be thought of as moments of cognition, cognitive "atoms," and are similar to action-perception cycles in neuroscience (Fuster 2002; Freeman 2002). Based on evidence from empirical neuroscience, and consistent with psychophysical paradigms measuring reaction time, we have estimated the duration of cognitive cycles to be approximately 200–500 ms (Madl et al. 2011). However, these cycles can partially overlap (Fig. 8.5b), leading to a rate of 5–10 cycles per second (Baars et al. 2013; Franklin et al. 2013). The LIDA cognitive cycle is not built into the model, but rather, emerges from it. Almost all of the modules as seen in Fig. 8.2, run continuously and asynchronously in parallel.

There are three phases in each cycle: the understanding phase, the attending phase, and the action selection and learning phase. In the understanding phase, the agent tries to make sense of its situation by updating its representations of external entities (perceived through the senses), as well as internally generated features. In the attending phase, the agent selects the most salient, important or urgent part of the constructed representation—the part that needs to be attended to. This part is sent to the rest of the system as the conscious broadcast (and thus becomes the current content of consciousness). In the third phase, internal resources are recruited based on this content—potential actions for the action selection mechanism to choose from. Furthermore, the conscious contents facilitate and modulate learning into multiple different memories. Figure 8.2 shows this process, starting in the upper left and proceeding roughly clockwise. Although the descriptions will be in terms of modules and processes, LIDA makes no commitment regarding whether the neural structure in humans is modular or localized. However, it is possible to tentatively assign neuronal correlates to LIDA's modules based on functional correspondence (Franklin et al. 2013), which we will briefly mention below.

The understanding phase starts with incoming sensory stimuli from the external and internal environments activating low-level feature detectors in Sensory Memory where they are partially interpreted by short term (tens of milliseconds) processes (sensory memory corresponds to sensory brain areas, such as the visual and auditory cortices). Results thereof proceed to LIDA's Perceptual Associative Memory (PAM) (long term associative recognition memory) to be processed by higher-level feature detectors, which can activate more abstract representations, e.g. objects, categories, actions, events, etc., as well as to the preconscious Workspace (a preconscious working memory with duration in tens of seconds). LIDA uses graphical representation,[4] nodes and links, in PAM and in the Workspace to represent

---

[4] More specifically, LIDA often uses directed graphs composed of nodes and links to represent items (nodes) and relationships between them (links).

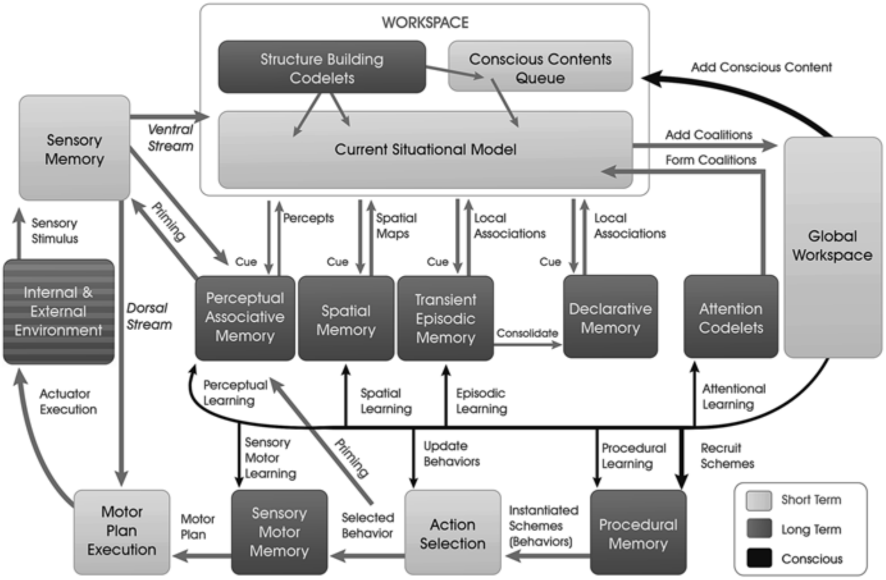

Fig. 8.2 The LIDA cognitive cycle

features, objects, categories, actions, feelings, events, etc. Localizing brain areas functionally corresponding to representations in PAM is not straightforward, as they are distributed and multimodal (Fuster 2004, 2006; Barsalou 2008; Fuster and Bressler 2012). Some areas involved in such representations include the perirhinal and orbito-frontal cortices and the amygdala.

Contents of the Workspace continually cue PAM, Spatial Memory (long term), Transient Episodic Memory (lasting a few hours or a day), and Declarative Memory (long term). Local associations recalled from the cueing of these various memories return to (or perhaps only point to) the Workspace. Neural correlates corresponding to these long-term memories include the hippocampus. On the other hand, the Workspace might correspond to temporo-parietal and frontal lobes and the entorhinal cortex (Franklin et al. 2013).

Workspace contents are operated upon by structure building codelets,[5] with the results being used to update the agent's preconscious Current Situational Model (CSM) within its Workspace. The agent's understanding of events occurring right now (i.e. within the last few cognitive cycles) is represented in this Current Situational Model (Snaider et al. 2012). The understanding phase is concerned with updating this CSM within the Workspace. Representations within the workspace

---

[5] A codelet is a small, single purpose, independently running piece of code, corresponding to a process in Baars' Global Workspace Theory. Structure building codelets build structures of nodes and links.

may persist in subsequent cognitive cycles, until they decay away. Another submodule of the Workspace, the Conscious Contents Queue will be discussed in Sect. 8.5.

For many complex agents with multiple senses, “living” in complex, dynamic environments, the Current Situational Model will contain far too much information to be responded to within a single cognitive cycle (~200–500 ms in humans (Madl et al. 2011)). Some filtering is needed to select the most salient information that must be attended to. In LIDA, attention codelets begin this filtering, or attention, phase of the cognitive cycle by creating coalitions of parts of the Current Situational Model. Each attention codelet looks for items corresponding to its particular concerns. On finding such, it creates a coalition containing their structures, and moves them to the Global Workspace. Subsequently, a competition in the Global Workspace chooses the most salient (the most relevant, important, urgent, novel, unexpected, loud, bright, moving, etc.) coalition, which then becomes the content of consciousness, and is broadcast globally to facilitate action selection and multiple modes of learning (implementing the large-scale integration and broadcasting mechanism suggested to underlie human functional consciousness in Sect. 8.2). The winning contents of the Global Workspace roughly correspond to neurons in different brain areas which are temporarily bound and integrated via oscillatory synchrony (Baars et al. 2013). This broadcast completes the attention phase of LIDA’s cognitive cycle.

The third and final phase of the cognitive cycle is concerned with learning in several modes, and with action selection and its execution. Since these modules and processes play little significant roles in the perception of time, and have been described in detail elsewhere, we will describe them only briefly here. Based on Drescher’s (1991) schema mechanism, data structures in Procedural Memory are called schemes. Each scheme consists of a context, an action, a result, and a base-level activation which measures the likelihood of the result happening should the action be taken in the scheme’s context. Each of the first three components are structures of nodes and links. Schemes whose context and/or results intersect the current conscious broadcast are instantiated as behaviors and passed to the Action Selection mechanism, where one is chosen and sent along for execution. Procedural Memory might correspond to the striatum and anterior cingulate, whereas the action selection mechanism might be grounded in the basal ganglia in brains (Franklin et al. 2013). Learning in the different modes takes place concurrently, completing the final phase of the LIDA cognitive cycle.

LIDA advocates a discrete view of consciousness, in accordance with GWT and the converging neuroscience evidence outlined in Sect. 8.2. As we shall see in the descriptions of the computational LIDA agents reproducing psychological experiments in Sect. 8.6, this view is also consistent with multiple behavioral paradigms investigating consciousness and attention, among others Allport’s (1968) experiments on perceptual simultaneity (which have traditionally been interpreted to require consciousness to be continuous). This view is also consistent with the philosophical conception of streamlikeness provided by Rashbrook-Cooper in this book, which allows for subjective continuity despite of gaps in consciousness. Just like his conception, LIDA’s view can be seen as an extensionalist account of temporal experience.

8.5 A LIDA-Based Model of Time Perception and Production

8.5.1 The Immediate Present Train Model

The LIDA model for time perception and representation is based on ideas from William James (1890). He discussed the "specious present," a term originally coined by E.R. Clay (1882). It has been called "specious" (plausible but wrong) since the present experienced by the human mind, instead of being a duration-less instant, is taken to comprise an interval.

Here we summarize the ideas introduced in Snaider et al. (2012). To describe our model we need first to briefly discuss some basic attributes of time: duration of short events, time duration scale, and succession. These attributes are fundamental for time flow perception, time concepts representation, and for defining what we call the Immediate Present Train (IPT), a more concrete instantiation for James's specious present.

Duration is probably the most well studied property of time. Saint Augustine (Warner 1963) discussed this issue, and argued that because the present is just an instant without duration, memory is required to measure an event's duration. We propose that without *any* memory it is not possible to have any notion of the concepts of past or of event duration. Notice that it is critical what we meant by "any memory" in the previous sentence. Using the LIDA concepts, this refers to the absence of transitive episodic and declarative memories. The workspace would only retain the present percept elements, but no past content is cued. Even in this reduced context, it is still possible to have some functionality, such as reacting to the present perception. However, memory is required to interpret the idea of something past. To evaluate the duration of an interval, some memory for the event (or events) is necessary, or at least some memory of their temporal properties (e.g., its starting time, or an accumulator that counts pulses). If we relax this idea, and we allow some memories of the past few seconds (probably in the preconscious workspace), it would be possible to model the concepts of past and duration.

The relative arrangement of events over time is an acknowledged property required for time awareness. Consider the events perceived by a subject. The arrangement that these events have is, in many cases, a piece of information as important as the events themselves. Processes such as detecting cause and effect situations, planning, and learning a path are possible when the perception and modeling of the sequence of the participant events are available. A nice metaphor for this is a family photo album. The photographs' order is telling us a story. A different arrangement may tell us a completely different one. If instead of an album we have a pile of pictures in no particular order, even the concept of story disappears.

When the intervals and durations of situations are relatively large (i.e. durations of some minutes, hours, days or even longer ones), we assume that an episodic memory module (as described in LIDA) participates in the process of maintaining the chronology. However, when events have durations between a fraction of second to a small number seconds, another entirely different mechanism is required. We consider the sequence of conscious broadcasts as the genesis of the stream of consciousness. Since humans are capable of *perceiving* this succession of broadcasts (Franklin et al. 2013), we introduce a structure that maintains this sequence, called the Conscious Contents Queue (CCQ) (Snaider et al. 2012), in LIDA's Workspace (see Fig. 8.2). James (1890, 606–607) clearly expressed this idea:

> If the present thought is of A B C D E F G, the next one will be of B C D E F G H, and the one after that of C D E F G H I—the lingerings of the past dropping successively away, and the incomings of the future making up the loss. These lingerings of old objects, these incomings of new, are the germs of memory and expectation, the retrospective and the prospective sense of time. They give that continuity to consciousness without which it could not be called a stream.

Although the CCQ name implies a queue's functionality (and in part this is true) it also resembles the behavior of a buffer. Its structure enables random access to its elements, while preserving their order. This allows several time related perceptual operations, such as the measure of an event's duration, or the detection of repeated event sequences.

For James, the specious present was "the prototype of all conceived times… the short duration of which we are immediately and incessantly sensible." James seems to imply a temporal interval, a "short duration," within which perceptions can be viewed to be in the present. For James, this duration could extend up to about 12 s. Latter users of the term "specious present" take it to mean "the (maximal) window through which we are directly aware of change and persistence…"

With this usage, there is considerable controversy as to the timespan of the specious present, but all current contenders are substantially less that James' 12 s (Dainton 2010). Wittmann (2011) reviews evidence suggesting the extent of an experienced moment to span a handful of seconds (mostly suggesting ~3 s; see also the chapter by Wittmann in this book). Block also estimates this duration at about 5 s (Block 2014).

Note that distinguishing events that last less than the estimated timespan of the specious present, and in some degree, modeling their chronology, are still possible. However, as in the case of timespans larger than the specious present that cannot be distinguished directly, some events are too excessively short to be identified as individual events. Images in TV screens are the prototypical example. Although we perceive them as moving images, they are actually static images presented in rapid succession. It is impossible to humans to perceive them as separated events. These ideas suggest that there is a range of event durations that humans can perceive directly, which we name perceptual time-range. Events with durations below this range are represented as a combination (e.g., the frames in the TV screen), or they may not be perceived at all. Events with durations above this time-range are still discernible using other cognitive processes such as episodic memory functionality or reasoning, but direct perception is not possible.

We hypothesize that these limits are not strict or fixed, and vary according to the nature and salience of the events. For example, when riding a rollercoaster, we would perceive a fast succession of stimuli, which may lead to a more than normal fine grain distinction of events.

In LIDA—although we often borrow from and build upon the ideas of James—we do not conceive of the "immediate present" as a fixed, absolute duration. Rather, we define it in terms of the events an agent (biological or artificial) might be currently conscious of. Every event that is broadcast consciously, and can be acted upon (or reported, or introspected upon), becomes a component of the "immediate present" when it is broadcast; and is included into the subjective past only once replaced by different events of a subsequent conscious broadcast.

Keep in mind that motion may be perceived in a single conscious broadcast, making us "directly aware of change and persistence." Thus with the previous estimation of specious present duration, we might claim it is comparable with that of a cognitive cycle, roughly 200–500 ms (Madl et al. 2011). Though motion can be directly perceived, events being perceived within the same cognitive cycle and becoming elements of the same conscious broadcast will be experienced as being simultaneous (e.g., flashes of light separated by a small distance and a few milliseconds), that is, as being a single event.

Snaider and colleagues (2012) combined these attributes of time into their Immediate Present Train where the (specious) present is modeled by a train, in which its extent corresponds to the timespan of the specious present. The cars in the train denote an ordered sequence of time steps, which contains the last few conscious events. In LIDA terms, the cars keep the elements of the recent broadcasts. Notice that the size of a car represents the extent of the shortest interval that can be distinguished directly. In other words, the train models the scope of events' durations discussed previously. We hypothesize that the train receives new conscious content every a few hundreds of milliseconds (for humans), and a car is appended to the front of the train with this content. Correspondingly, cars at the end are removed from the train. The train representation comprises several instants, which allows representing non-simultaneous events as components of the immediate present. In effect, events may be in different cars but still belong to the same train.

Note that an event shorter than the timespan equivalent to a single car may be represented directly as a change event. For example, the movement of a ball can be modeled as a moving-ball event, instead of a sequence of ball-position events.

Although this model may suggest that the duration represented by each car and the number of cars in the train are fixed magnitudes, the model actually allows variations in them. The interval comprising two consecutive conscious events may vary, thus affecting the duration represented by each car. Also, the elements in a few of the cars may decay away, removing these cars from the train (which changes the total duration represented by the train).

As we mentioned previously, the Conscious Contents Queue (CCQ), a submodule in the LIDA's workspace, is a more concrete instantiation of the Immediate Present Train and the specious present (Snaider et al. 2012). It is a combination of a queue and a buffer. It comprises a variable number of cells, similar to the cars described above. CCQ resembles a queue, since it has a head and an end, and the content of one cell is pushed back to the following cell when fresh element arrives

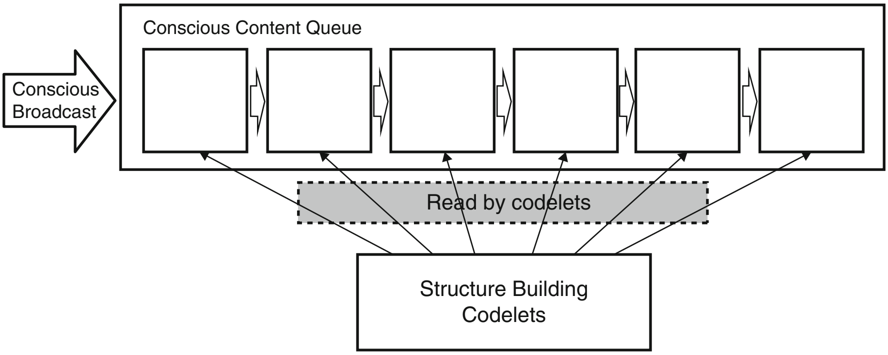

Fig. 8.3 The Conscious Contents Queue (From Snaider et al. (2012) with permission from Elsevier)

(Fig. 8.3). However, unlike queues, the cells can be accessed directly, allowing other process (particularly structure-building codelets) to read several cells simultaneously. With each broadcast, the CCQ receives new content, which is inserted in the head of the CCQ, while the old content is shifted towards the end. As other representations in LIDA, elements in the cells have activation, which decay over time. Some elements loss all their activation and are removed from the queue. Eventually, all elements in one cell may decay away, and then the cell itself is removed from the queue (even if it is not at the end). The frequency of the conscious broadcasts may vary over time as a factor of the many triggers of the Global Workspace module. As a consequence, the size of the CCQ (and the time that it ultimately represents) is not fixed. In other words, the time required to complete LIDA Cognitive Cycle phase defines the frequency of the broadcast and the duration determined by a cell in the CCQ. For humans each phase takes approximately 100 ms (for simple tasks, the understanding phase is estimated to take 80–100 ms, the attending phase an additional 120–180 ms, and the action selection phase 60–110 ms (Madl et al. 2011)). This duration determines the lower limit of the perceptual time-range, and the count of cells in the CCQ defines its maximum.

Structure-building codelets can approximately calculate the duration of short events by simply counting the number of cells that that event spans (see Fig. 8.4). Several factors affect the precision of these calculations. One of these factors is the activation decay of the elements in the CCQ described above. Another factor is the frequency of the conscious broadcast, which can vary. In general, when more stimuli are present, the frequency of conscious broadcast is higher (as in our example of a rollercoaster ride). This has the effect of filling the cells faster, and the 100 ms estimation for the duration represented by each cell becomes inexact. Actually, structure building codelets that inspect the CCQ elements will erroneously consider that each cell still represents 100 ms, and the event duration perception will be distorted producing the effect that the duration is longer than it really is, or in other words, producing the sensation that time flows more slowly, an effect that was reported in several experiments, e.g., Eagleman (2008).

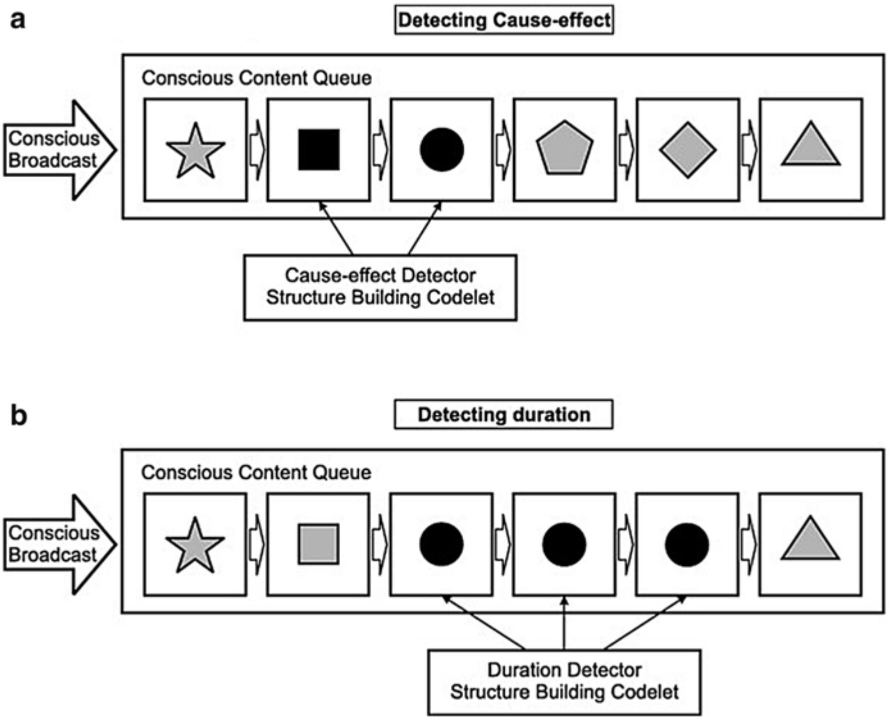

Fig. 8.4 Detecting causes and effects (**a**) and determining durations (**b**) using the Conscious Contents Queue. (**a**) A Cause-effect Detector Structure Building Codelet detects that the "*circle*" content precedes the "*square*" content in the CCQ, and would create a "*circle* before *square*" representation in the Workspace. (**b**) A Duration Detector Structure Building Codelet can select repetitions from the CCQ, and count the number of occurrences. The Codelet can then create a representation of the duration of the selected content in the Workspace, based on this number and on the duration of the cells in the CCQ (From Snaider et al. (2012) with permission from Elsevier)

Codelets may also perform other time related operations, such as determining cause-effect situations (for brief events) using the CCQ, thanks to its sequential order. If one event is present in a cell closer to the head than other event, a codelet could use this a signal for creating a cause-effect relationship between these two events (see Fig. 8.4). Another operation may be the detection of simultaneous or quasi-simultaneous events, depending on the number of cells considered for the tasks.

The current main representation in the LIDA architecture comprises nodes that represent concepts, and links, which denote relationships between these concepts. In the general case, nodes are *grounded* in sensor and motor memories. For example, the node representing the color red is ultimately rooted in the light sensors sensible to that color. However, in LIDA, time-based nodes, such as duration nodes, are grounded by the CCQ. Short duration nodes are instantiated, by codelets when they detect these intervals as we described previously. Other nodes for concepts such as fast and slow, can be derived from these duration nodes. The abstract notions of "duration" and "flow of time" can be created as categorizations of simpler nodes. To explain the creation (and perception) of nodes for larger spans, such as nodes representing minutes, hours, or even longer intervals, we hypothesize that an episodic memory module is required. However, these concepts for longer periods are correctly interpreted and handled thanks to their connection with the simpler ones grounded into the CCQ. In our view, the CCQ mechanism provides the seminal concepts for interpreting and working with time related concepts in LIDA.

8.6 Computational Reproductions of Experiments Involving Time

8.6.1 Consciousness and Continuity: The LIDA Allport Agent

The idea of consciousness possibly being discrete has been strongly criticized and in some cases even outright rejected based on those empirical results in the phenomenal simultaneity paradigm which seem to contradict discrete (e.g., cinematic) models. A number of frequently cited experiments were conducted by Allport (1968), who aimed to compare two prevalent competing theories of consciousness at that time, Stroud's (1967) Discrete Moment Hypothesis (DMH) and the Continuous Moment Hypothesis. The former states that consciousness comprises distinct and not overlapping conscious "moments," within which time-order information is lost, whereas the latter views conscious "moments" as corresponding to continuously moving segments of incoming sensory information. Allport's empirical results contradict the DMH, leading him to reject Stroud's discrete model.

However, although the LIDA model—like Stroud's—also proposes consciousness to be discrete, it can still account for this empirical evidence. To show this consistency, as well as to strengthen the claim that LIDA's GWT-based consciousness mechanism can model human functional consciousness, we have replicated Allport's experiment computationally with a LIDA-based cognitive software agent (Madl et al. 2011).

In Allport's (1968) experiment, participants faced a screen displaying a horizontal line in 1 of 12 possible positions on this screen (see Fig. 8.5a), and rapidly changing position moving upward. Each time the line reached the top position, the screen was first left blank for the amount of time it took for the line to traverse the screen, and then the line reappeared in the bottom position, moving upward. These cycles of the screen alternating between showing the moving line and being blank were repeated. Participants could control the cycle time (τ).

When cycle times were set to be large, participants were able to see the line jumping from one position to the next. When they reduced τ, participants saw multiple lines, moving together. However, at and below a small cycle time S, they reported perceiving an unmoving array of 12 lines which flickered in synchrony, instead of individual lines.

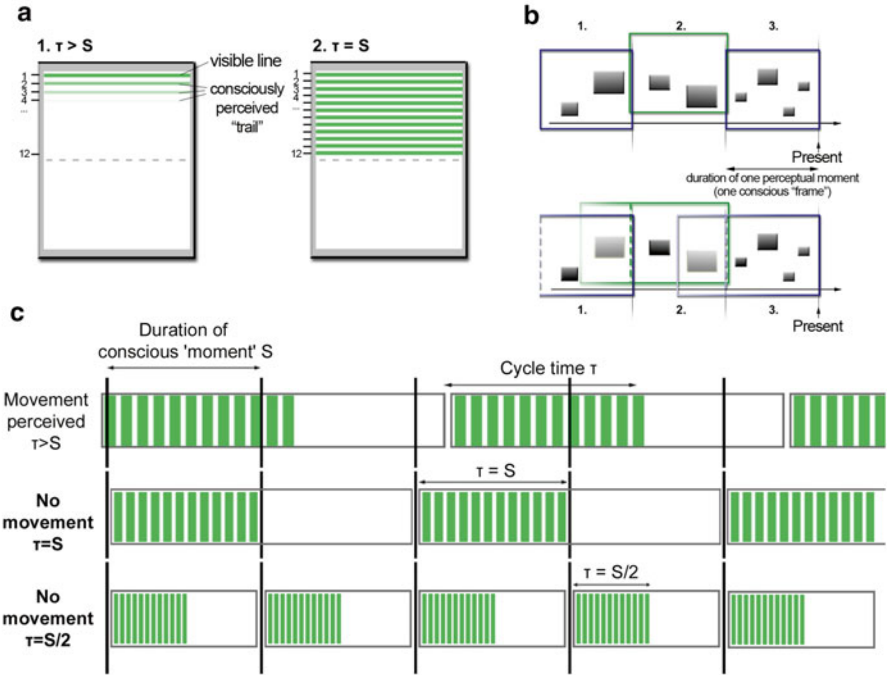

Fig. 8.5 Allport's experiment (**a**), and a comparison of the refuted DMH (**b**, *top*, and **c**) and LIDA's discrete consciousness mechanism (**b**, *bottom*). (**a**) The screen in Allport's (1968) experiment. A visible line was shown in 1 of 12 possible positions, moving upwards. Whenever it reached the top, the *line* vanished for the amount of time it took to reach the top. The cycle time is denoted by τ. When τ >S, participants could see movement (*left panel*). At τ=S, participants perceived all lines at the same time, and saw no movement (*right panel*). (**b**) Schematic comparison of the DMH (*top*) and LIDA's discrete consciousness hypothesis (*bottom*). The frames represent the temporal constraints of a perceptual moment or conscious "frame," and the solid rectangles symbolize incoming percepts. In LIDA, important percepts from previous conscious "frames" can remain conscious (*rectangles left* of the *dashed lines* in the frames in the *bottom* picture). (**c**) Predictions of the DMH. If conscious moments were discrete and distinct, there would be two cycle times at which subjects would perceive no movement (τ=S and τ=S/2). Instead, Allport (1968) reports only one cycle time (From Madl et al. 2011 with permission)

The task of the participants was to keep changing τ until they arrived at cycle time S and stopped perceiving moving lines. They were asked to do this in two types of trials, in which their cycle times were recorded. In the first type, they had to decrease the cycle time from a high value towards S (accelerating the cycles until they reached a cycle time τ_1 at which they saw stationary lines). In the second type, they increased cycle time from a low value towards S (slowing the cycles until they started seeing movement at cycle time τ_2)—see Fig. 8.5a.

The abovementioned hypotheses regarding consciousness make different predictions regarding the cycle times participants should arrive at in these two trial types. The Discrete Moment Hypothesis would predict that they should be different—

there should be two cycle times, τ_1 and τ_2, at which the 12 lines can be seen on the screen without movement. At τ_1=S, subjects would not perceive movement because everything happening on the screen should fall within a single conscious "moment" (all line positions as well as the blank screen). On the other hand, at a time τ_2 which equals S/2 there should also be no movement, since at this cycle time conscious "moments" might alternate between containing all line positions (taking S/2) and between containing a blank screen (also taking S/2)—thus, no moving lines should be seen, only flickering. Therefore, the DMH would predict that in the above experimental setting, subjects will arrive at two distinct cycle times in the two task types, τ_1=S when cycle times are decreased, and τ_2=S/2 when cycle times are increased.

The Continuous Moment Hypothesis, in contrast, would predict only a single cycle time τ_1=τ_2=S at which no movement can be perceived. According to this hypothesis, events are judged to be simultaneous if they fall within one conscious "moment." In this experiment, the lines are perceived to be stationary when all line positions as well as the blank screen fall within a conscious "moment," when the cycle time is S. However, at a cycle time of S/2, there would still be movement—the conscious "moment's" contents would change from containing 12 lines, over containing fewer and fewer lines, to finally only containing the blank screen. Thus, participants should arrive at the same cycle time S in both trial types in this experiment.

Allport (1968) reports that the cycle times in the two trial types were not significantly different. Based on this result, he argued for the implausibility of the Discrete Moment Hypothesis. However, despite LIDA's consciousness mechanism being discrete, we have reproduced Allport's result with a LIDA-based computational cognitive agent.

To simulate Allport's experiment, the LIDA Allport agent used the cognitive cycles outlined in Sect. 8.4. The Allport agent had a pre-defined PAM to model the experimental stimuli, containing a PAM node for each of the line positions on the screen, and feature detectors corresponding to each line passing activation to the respective node corresponding to the currently visible line. The agent also had a pre-defined Procedural Memory (PM) containing two behavior schemes, for the "movement perceived" and "no movement perceived" buttons. The former was activated when the agent perceived no line movement (i.e. when all 12 line positions were present in the conscious broadcast), whereas the latter was pressed by the agent whenever it perceived movement. Cycle times (τ) were adjusted gradually in the environment, and the agent only had to react to whether or not it could perceive movement (this was computationally easier to implement than letting the agent decide the cycle time, but did not make any difference in the implications and predictions of the discrete consciousness mechanism).

The environment first successively decreased the cycle times from a high value, and then successively increased it from a slow value, similarly to the two trial types of the Allport experiment. The button responses of the agent were recorded, and the cycle times at which the agent pressed the "no movement perceived" button compared between the two trial types. The agent pressed this button at the same cycle time in both conditions—at 96 ms (Madl et al. 2011), which matches the results of the human participants described above (in contrast to the predictions of the DMH), and suggests that the durations of conscious “moments” in LIDA approximately match those of humans.

In LIDA, conscious episodes are discrete, but contrary to the DMG as argued by Stroud (1967), not always distinct. Subsequent conscious “moments” might contain percepts from prior moments (symbolized by the rectangles left of the dashed lines in Fig. 8.5b). The duration of older percepts persisting in consciousness is influenced by multiple factors, including when (how long in the past) it was perceived, and on attentional modulation. Here we have an example of a systems level, computational cognitive model providing deeper understanding of an experimental result.

8.6.2 Attention

Two cognitive software agents were developed to reproduce experiments related to attention: the LIDA Attentional Blink (Madl and Franklin 2012) and the LIDA Attention agent (Faghihi et al. 2012).

The first LIDA agent accounted for the attentional blink (AB) (Madl and Franklin 2012), i.e. the observed phenomenon that subjects are frequently unable to report the second of two targets shown within 200–500 ms after the first, within a sequence of target and distractor stimuli (see Fig. 8.6). The AB has a number of observed properties. The second target (T2) can be consciously perceived and reported if it is presented after the first target (T1) but with no distractor in between (“lag-1 sparing”), but not if there are distractors between the targets. Furthermore, the AB effect can be reduced—the likelihood of the second target correctly being reported increased—by increasing its salience (Martens and Wyble 2010) or emotional arousal (Anderson 2005).

Many AB models have been proposed; however, most current models cannot account for all findings and properties in AB experiments (see Dux and Marois 2009 or Martens and Wyble 2010 for reviews). Furthermore, many of these models are specific to the AB, instead of being implemented within a general cognitive architecture.

We have developed a LIDA-based model of the AB (Madl and Franklin 2012) to computationally model the visual attentional blink experiment (Potter et al. 2010), reproducing human behavior data, and to conceptually account for a large number of phenomena. In LIDA, the attentional blink is mainly caused by a temporarily depleted attentional resource (which fully regenerates after ~500 ms), making attending to the second target difficult if it is presented very shortly after the first. Lag-1 sparing arises from both targets entering the same coalition and within the same cognitive cycle, and thus both coming to consciousness.

The second attention agent was based on a modified version of the experiment by Van Bockstaele et al. (2010). Its environment was composed of a screen with two white squares on both sides of a fixation cross (see Fig. 8.6). After a brief delay (fixation period in the original experiment), a colored cue randomly appeared in

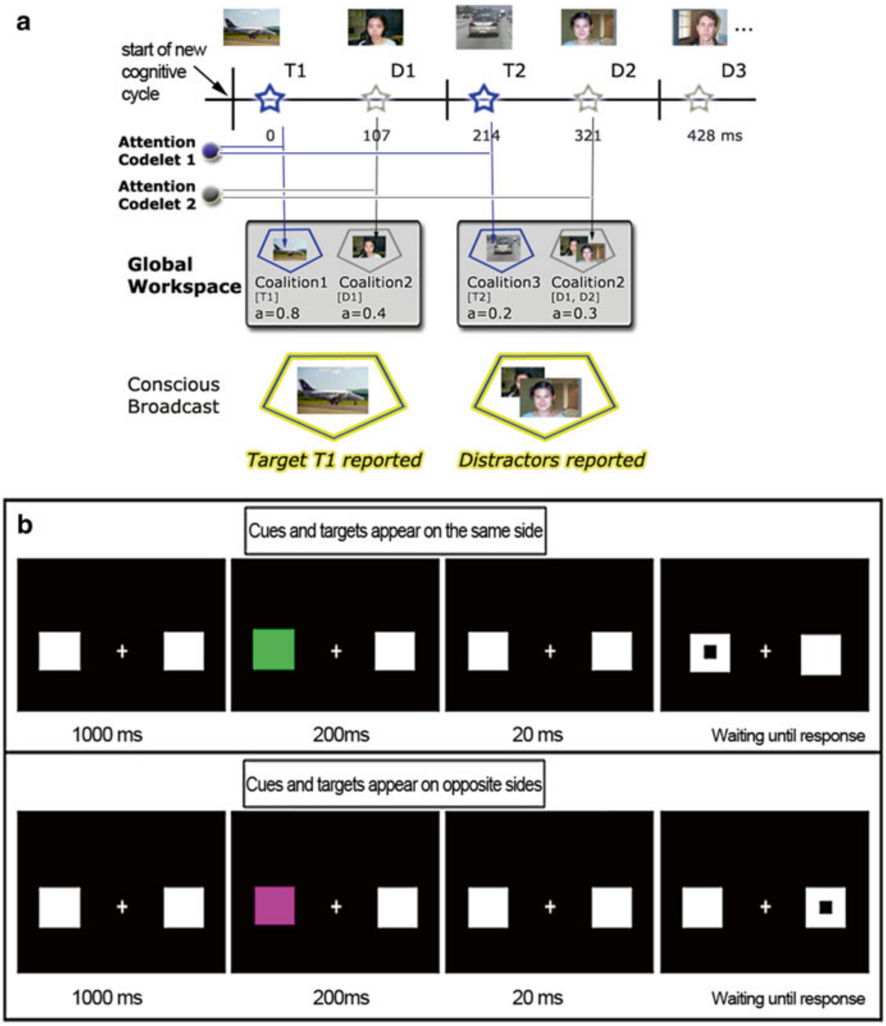

Fig. 8.6 Experimental paradigms reproduced by the LIDA Attentional Blink and LIDA Attention agents. (**a**) Sequentially presented images in the Attentional Blink paradigm. Two targets (vehicles; *T1* and *T2*) are presented with one distractor between them (*D1*) and several distractors after them (faces). In the figure, *T2* cannot be consciously perceived and reported, because in the second cognitive cycle the distractors win the competition for consciousness (starting times of cognitive cycles are marked by *bold vertical lines* on the timeline). If *T1* and *T2* were presented subsequently, they would both be bound within the same coalition, and perceived consciously. (**b**) Timeline of displayed cues in Van Bockstaele's experiment replicated by the LIDA Attention agent, with both target and cue being displayed on the same side in congruent trials (*top*), and on opposite sides in incongruent trials (*bottom*) (Fig. 8.6b from Faghihi et al. (2012) with permission from Elsevier)

either the left or the right square, for 200 ms, followed by the two empty white squares for 20 ms. Subsequently, a small black rectangle (the target) was presented in either the left or the right square, again at random. The agent (just like the participant in the original experiment) had to respond to the target, i.e. press the right one of two buttons, as fast as possible; response times were measured. The experience showed that both participants (Van Bockstaele et al. 2010) and the LIDA Attention agent (Faghihi et al. 2012) were faster in reaction by 20 ms on trials in which the cue and the target were shown on the same side (congruent trials), compared to trials where they appeared on opposite sides (incongruent trials)—average response times were 360 and 380 ms. The reason for this difference in the LIDA agent was the instantiation of the correct behavior scheme. That is, by the time the target arrives to consciousness the cue almost primes a behavior by sending more activation to it. In contrast, in trials with cue and target on opposite sides, different schemes from Procedural Memory needed to be instantiated and then a behavior will be selected and executed. The extra scheme instantiation cost to the Attention agent an additional 20 ms (Faghihi et al. 2012).

8.7 Conclusion

As we argued in the Introduction, the study of mind in all of its aspects, including the perception of time, is best approached from different perspectives. As we have seen throughout this work, it has proved useful to study such difficult questions as the seemingly continual flow of time using the various tools of each of the relevant disciplines, the introspection of the philosopher of mind (e.g., Block 1995; Dainton 2010; James 1890), the behavioral observation of the experimental psychologist (e.g., Buhusi and Meck 2005; James 1890; Michon 1990; Zakay et al. 1994), the brain imaging of the cognitive neuroscientist (e.g., Eagleman 2008; Ivry and Schlerf 2008), and the computational simulation of the cognitive modeler, a computer scientist (e.g., Buhusi and Meck 2005; Madl et al. 2011; Michon 1990; Snaider et al. 2012).

We have reviewed recent neuroscience evidence concerning large-scale integration by oscillatory synchrony as a possible mechanisms underlying functional consciousness, suggesting it to be discrete. We have also briefly reviewed recently suggested neural correlates of global and local time perception mechanisms in brains. After outlining Global Workspace Theory, a prominent theory accounting for functional consciousness, we have described a conceptual and partially computational model of cognition based on GWT—the LIDA cognitive architecture—and argued that it can account for time perception in a cognitively plausible fashion (substantiated by reproduced psychological experiments), and generate concepts such as continuity, immediate present duration, and perceived length of time.

Some of the modelers of time are only concerned with modeling time itself, or even one aspect of it, for example duration (Zakay et al. 1994). Here we have argued for the need to study time in the context of the study of mind, using a broad, systems-level cognitive model such as our LIDA (Franklin et al. 2013). We are not alone. Such arguments have been made earlier by a number of other researchers from disparate fields. Here we support our arguments by quotes from four such. From social psychology, Kurt Lewin says it quite concisely. “There is nothing so practical as a good theory” (1951, 169). A broad, systems level cognitive model is a theory of mind. From computer science, AI pioneer Allen Newell argues against the reliance on modeling individual laboratory tasks saying “You can’t play 20 questions with nature and win” (Newell 1973). Making the same point for his field, psychological memory researcher Douglas Hintzman (2011) writes, “Theories that parsimoniously explain data from single tasks will never generalize to memory as a whole…” Hintzman’s arguments rest precisely on the need for the type of cognitive models that we advocate, and apply broadly beyond memory research. Langley et al. (2008) wrote a review article entitled “Cognitive architectures: Research issues and challenges.” In it they argue for the use of systems-level cognitive architectures such as our LIDA model, asserting that “Instead of carrying out micro-studies that address only one issue at a time, we should attempt to unify many findings into a single theoretical framework, then proceed to test and refine that theory.” Several of the “open problems” described in their review have since been partially or fully solved by our LIDA. The reinterpretation of the Allport experiment provided by the LIDA Allport agent is one example of the value of such an approach. In a table allowing ready comparison of properties of some 26 “biologically inspired cognitive architectures” (Samsonovich 2010), LIDA compares rather well in terms of modeling a complete cognitive system, and also in terms of being truly biologically inspired.

We contend that such systems-level, conceptual and computational modeling can, if it is biologically plausible, integrate findings from the several disciplines, and produce hypotheses that will serve to guide further research.

References

Allport, David A. 1968. Phenomenal simultaneity and the perceptual moment hypothesis. *British Journal of Psychology* 59(4): 395–406.

Anderson, Michael L. 2003. Embodied cognition: A field guide. *Artificial Intelligence* 149(1): 91–130.

Anderson, Adam K. 2005. Affective influences on the attentional dynamics supporting awareness. *Journal of Experimental Psychology: General* 134(2): 258–281.

Baars, Bernard J. 1997. *In the theater of consciousness: The workspace of the mind*. Oxford: Oxford University Press.

Baars, Bernard J. 2005. Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in Brain Research* 150: 45–53.

Baars, Bernard J., Stan Franklin, and Thomas Zoega Ramsoy. 2013. Global workspace dynamics: Cortical “binding and propagation” enables conscious contents. *Frontiers in Psychology* 4: 200.

Baddeley, Alan D., and Graham J. Hitch. 1974. Working memory. *Psychology of Learning and Motivation* 8: 47–89.

Barsalou, Lawrence W. 1999. Perceptions of perceptual symbols. *Behavioral and Brain Sciences* 22(4): 637–660.

Barsalou, Lawrence W. 2008. Grounded cognition. *Annual Review of Psychology* 59: 617–645.

Block, Ned. 1995. How many concepts of consciousness? *Behavioral and Brain Sciences* 18(2): 272–287.

Block, Richard A. (ed.). 2014. *Cognitive models of psychological time*. New York: Psychology Press.

Boltz, Marilyn G. 1995. Effects of event structure on retrospective duration judgments. *Perception & Psychophysics* 57(7): 1080–1096.

Buhusi, Catalin V., and Warren H. Meck. 2005. What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience* 6(10): 755–765.

Casali, Adenauer G., Olivia Gosseries, Mario Rosanova, Mélanie Boly, Simone Sarasso, Karina R. Casali, Silvia Casarotto, Marie-Aurélie Bruno, Steven Laureys, and Giulio Tononi. 2013. A theoretically based index of consciousness independent of sensory processing and behavior. *Science Translational Medicine* 5(198): 198ra105.

Clay, E.R. 1882. *The alternative: A study in psychology*. London: Macmillan.

Conway, Martin A. 2002. Sensory-perceptual episodic memory and its context: Autobiographical memory. In *Episodic memory: New directions of research*, ed. Alan Baddeley, Martin Conway, and John Aggleton. Oxford: Oxford University Press.

Crick, Francis, and Christof Koch. 1990. Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2: 263–275.

Dainton, Barry. 2010. Temporal consciousness. In The Stanford encyclopedia of philosophy. Stanford: The Metaphysics Research Lab, Center for the Study of Language and Information, Stanford University.

Dehaene, Stanislas, and Jean-Pierre Changeux. 2003. Neural mechanisms for access to consciousness. In *The cognitive neurosciences III*, ed. Michael Gazzaniga, 1145–1157. Cambridge, MA: MIT Press.

Dehaene, Stanislas, Jean-Pierre Changeux, Lionel Naccache, Jérôme Sackur, and Claire Sergent. 2006. Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences* 10(5): 204–211.

Dehaene, Stanislas, Lucie Charles, Jean-Rémi King, and Sébastien Marti. 2014. Toward a computational theory of conscious processing. *Current Opinion in Neurobiology* 25: 76–84.

Dennett, Daniel C. 1988. Quining qualia. In *Consciousness in contemporary science*, ed. A. Marcel and E. Bisiach, 42–77. Oxford: Oxford University Press.

Dieter, Kevin C., and Duje Tadin. 2011. Understanding attentional modulation of binocular rivalry: A framework based on biased competition. *Frontiers in Human Neuroscience* 5: 155. doi:10.3389/fnhum.2011.00155.

Doesburg, Sam M., Jessica J. Green, John J. McDonald, and Lawrence M. Ward. 2009. Rhythms of consciousness: Binocular rivalry reveals large-scale oscillatory network dynamics mediating visual perception. *PloS One* 4(7): e6142.

Drescher, Gary L. 1991. *Made-up minds: A constructivist approach to artificial intelligence*. Cambridge, MA: MIT Press.

Dux, Paul E., and René Marois. 2009. The attentional blink: A review of data and theory. *Attention, Perception, & Psychophysics* 71(8): 1683–1700.

Eagleman, David M. 2008. Human time perception and its illusions. *Current Opinion in Neurobiology* 18(2): 131–136.

Edelman, Gerald M., and Giulio Tononi. 2000. *A universe of consciousness: How matter becomes imagination*. New York: Basic Books.

Edelman, Gerald M., Joseph A. Gally, and Bernard J. Baars. 2011. Biology of consciousness. *Frontiers in Psychology* 2: 4.

Ericsson, K. Anders, and Walter Kintsch. 1995. Long-term working memory. *Psychological Review* 102(2): 211–245.

Faghihi, Usef, Ryan McCall, and Stan Franklin. 2012. A computational model for attentional learning in cognitive agents. *Biologically Inspired Cognitive Architectures* 2: 25–36.

Franklin, Stan, Tamas Madl, Sidney D'Mello, and Javier Snaider. 2013. LIDA: A systems-level architecture for cognition, emotion, and learning. *IEEE Transactions on Autonomous Mental Development*. doi:10.1109/TAMD.2013.2277589.

Freeman, Walter J. 2002. The limbic action-perception cycle controlling goal-directed animal behavior. *Neural Networks* 3: 2249–2254.

Fuster, Joaquín M. 2002. Physiology of executive functions: The perception-action cycle. In *Principles of frontal lobe function*, ed. Donald T. Stuss and Robert T. Knight, 96–108. New York: Oxford University Press.

Fuster, Joaquín M. 2004. Upper processing stages of the perception-action cycle. *Trends in Cognitive Sciences* 8: 143–145.

Fuster, Joaquín M. 2006. The cognit: A network model of cortical representation. *International Journal of Psychophysiology* 60(2): 125–132.

Fuster, Joaquín M., and Steven L. Bressler. 2012. Cognit activation: A mechanism enabling temporal integration in working memory. *Trends in Cognitive Sciences* 16(4): 207–218.

Gibbon, John, Russell M. Church, and Warren H. Meck. 1984. Scalar timing in memory. *Annals of the New York Academy of Sciences* 423(1): 52–77.

Gibbon, John, Chara Malapani, Corby L. Dale, and C.R. Gallistel. 1997. Toward a neurobiology of temporal cognition: Advances and challenges. *Current Opinion in Neurobiology* 7(2): 170–184.

Glenberg, Arthur M. 1997. What memory is for: Creating meaning in the service of action. *Behavioral and Brain Sciences* 20(1): 41–50.

Glenberg, Arthur M., and David A. Robertson. 2000. Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language* 43(3): 379–401.

Grondin, Simon. 2010. Timing and time perception: A review of recent behavioral and neuroscience findings and theoretical directions. *Attention, Perception, & Psychophysics* 72(3): 561–582.

Hintzman, Douglas L. 2011. Research strategy in the study of memory: Fads, fallacies, and the search for the coordinates of truth. *Perspectives on Psychological Science* 6(3): 253–271.

Ivry, Richard B., and John E. Schlerf. 2008. Dedicated and intrinsic models of time perception. *Trends in Cognitive Sciences* 12(7): 273–280.

James, William. 1890. *The principles of psychology*. Cambridge, MA: Harvard University Press.

Karmarkar, Uma R., and Dean V. Buonomano. 2007. Timing in the absence of clocks: Encoding time in neural network states. *Neuron* 53(3): 427–438.

Kim, Chai-Youn, and Randolph Blake. 2005. Watercolor illusion induced by synesthetic colors. *Perception* 34(12): 1501.

Koch, Christof. 2004. *The quest for consciousness*. Englewood: Roberts & Company.

Kouider, Sid, and Stanislas Dehaene. 2007. Levels of processing during non-conscious perception: A critical review of visual masking. *Philosophical Transactions of the Royal Society, B: Biological Sciences* 362(1481): 857–875.

Langley, Pat, John E. Laird, and Seth Rogers. 2008. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research* 10: 141–160.

Lau, Hakwan C., and Richard E. Passingham. 2007. Unconscious activation of the cognitive control system in the human prefrontal cortex. *The Journal of Neuroscience* 27(21): 5805–5811.

Lewin, Kurt. 1951. In *Field theory in social science: Selected theoretical papers*, ed. D. Cartwright. New York: Harper & Row.

Lewis, Laura D., Veronica S. Weiner, Eran A. Mukamel, Jacob A. Donoghue, Emad N. Eskandar, Joseph R. Madsen, William S. Anderson, Leigh R. Hochberg, Sydney S. Cash, and Emery N. Brown. 2012. Rapid fragmentation of neuronal networks at the onset of propofol-induced unconsciousness. *Proceedings of the National Academy of Sciences* 109(49): E3377–E3386.

Madl, Tamas, and Stan Franklin. 2012. A LIDA-based model of the attentional blink. In *ICCM 2012 proceedings*, 283–288.

Madl, Tamas, Bernard J. Baars, and Stan Franklin. 2011. The timing of the cognitive cycle. *PloS One* 6(4): e14803.

Martens, Sander, and Brad Wyble. 2010. The attentional blink: Past, present, and future of a blind spot in perceptual awareness. *Neuroscience & Biobehavioral Reviews* 34(6): 947–957.

Melloni, Lucia, Carlos Molina, Marcela Pena, David Torres, Wolf Singer, and Eugenio Rodriguez. 2007. Synchronization of neural activity across cortical areas correlates with conscious perception. *The Journal of Neuroscience* 27(11): 2858–2865.

Merchant, Hugo, Deborah L. Harrington, and Warren H. Meck. 2013. Neural basis of the perception and estimation of time. *Annual Review of Neuroscience* 36: 313–336.

Michon, John A. 1990. Implicit and explicit representations of time. In *Cognitive models of psychological time*, ed. Richard A. Block, 37–58. Hillsdale: Lawrence Erlbaum Associates.

Newell, Allen. 1973. You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In *Visual information processing*, ed. W.G. Chase, 283–308. New York: Academic.

Palva, Satu, Klaus Linkenkaer-Hansen, Risto Näätänen, and J. Matias Palva. 2005. Early neural correlates of conscious somatosensory perception. *The Journal of Neuroscience* 25(21): 5248–5258.

Pitts, Michael A., and Juliane Britz. 2011. Insights from intermittent binocular rivalry and EEG. *Frontiers in Human Neuroscience* 5: 107.

Potter, Mary C., Brad Wyble, Rijuta Pandav, and Jennifer Olejarczyk. 2010. Picture detection in rapid serial visual presentation: Features or identity? *Journal of Experimental Psychology: Human Perception and Performance* 36(6): 1486–1494.

Rees, Geraint, Gabriel Kreiman, and Christof Koch. 2002. Neural correlates of consciousness in humans. *Nature Reviews Neuroscience* 3(4): 261–270.

Reuss, Heiko, Andrea Kiesel, Wilfried Kunde, and Bernhard Hommel. 2011. Unconscious activation of task sets. *Consciousness and Cognition* 20(3): 556–567.

Samsonovich, A.V. 2010. Toward a unified catalog of implemented cognitive architectures. In Proceedings of the 2010 conference on biologically inspired cognitive architectures, 195–244. Amsterdam: IOS Press. ISBN: 978-1-60750-660-7

Singer, Wolf. 2011. Consciousness and neuronal synchronization. In *The neurology of consciousness: Cognitive neuroscience and neuropathology*, ed. Steven Laureys and Giulio Tononi, 43–52. London: Academic.

Sloman, Aaron. 1999. What sort of architecture is required for a human-like agent? In *Foundations of rational agency*, ed. Michael Wooldridge and Anand Rao, 35–52. Dordrecht: Kluwer Academic Publishers.

Snaider, Javier, Ryan McCall, and Stan Franklin. 2012. Time production and representation in a conceptual and computational cognitive model. *Cognitive Systems Research* 13(1): 59–71.

Steriade, M. 2006. Grouping of brain rhythms in corticothalamic systems. *Neuroscience* 137(4): 1087–1106.

Stroud, John M. 1967. The fine structure of psychological time. *Annals of the New York Academy of Sciences* 138(2): 623–631.

Van Bockstaele, Bram, Bruno Verschuere, Jan De Houwer, and Geert Crombez. 2010. On the costs and benefits of directing attention towards or away from threat-related stimuli: A classical conditioning experiment. *Behaviour Research and Therapy* 48: 692–697.

VanRullen, Rufin, and Christof Koch. 2003. Is perception discrete or continuous? *Trends in Cognitive Sciences* 7(5): 207–213.

VanRullen, Rufin, Benedikt Zoefel, and Barkin Ilhan. 2014. On the cyclic nature of perception in vision versus audition. *Philosophical Transactions of the Royal Society, B: Biological Sciences* 369(1641): 20130214.

Varela, Francisco J., Evan Thompson, and Eleanor Rosch. 1991. *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.

Warner, Rex. 1963. trans. *The Confessions of St. Augustine*. New York: A Mentor Book.

Wikimedia. 2009. *Structures of brain: Wikimedia commons*. http://commons.wikimedia.org/wiki/File: Human–brain.SVG. Accessed 17 Apr 2014.

Wikimedia. 2010. *Diagram of human brain: Wikimedia commons*. http://en.wikipedia.org/wiki/File:BrainCaudatePutamen.svg. Accessed 17 Apr 2014.

Wittmann, Marc. 2011. Moments in time. *Frontiers in Integrative Neuroscience* 5: 66.

Zakay, Dan. 1992. The role of attention in children's time perception. *Journal of Experimental Child Psychology* 54(3): 355–371.

Zakay, Dan, and Richard A. Block. 1996. The role of attention in time estimation processes. *Advances in Psychology* 115: 143–164.

Zakay, Dan, and Richard A. Block. 1997. Temporal cognition. *Current Directions in Psychological Science* 6: 12–16.

Zakay, Dan, Yehoshua Tsal, Masha Moses, and Itzhak Shahar. 1994. The role of segmentation in prospective and retrospective time estimation processes. *Memory & Cognition* 22(3): 344–351.

Zarco, Wilbert, Hugo Merchant, Luis Prado, and Juan Carlos Mendez. 2009. Subsecond timing in primates: Comparison of interval production between human subjects and rhesus monkeys. *Journal of Neurophysiology* 102(6): 3191–3202.