# An integrative pluralistic approach to phenomenal consciousness

Rick Dale, Deborah P. Tollefsen & Christopher T. Kello
Cognitive and Information Sciences, University of California, Merced / Department of Philosophy, The University of Memphis / Cognitive and Information Sciences, University of California, Merced

We review theories and empirical research on underlying mechanisms of selfhood, awareness, and conscious experience. The mechanisms that have been identified for these phenomena are many and multifarious, lying at many levels of space and time, and complexity and abstractness. Proposals have included the global workspace for conscious information, action and its centrality to self awareness, the role for social information and narrative, and more. We argue that phenomenal experience, whatever it "really is," is probably dependent upon all of these levels simultaneously. We end with two challenges for consciousness research. Both are couched in terms of the dynamics of phenomenal experience. The first is to investigate the sustained dynamics of phenomenal experience; the second is to unveil the way that multi-scale processes in the cognitive system interact to produce that richness of experience. We do not aim to solve the hard problem, but argue that any solution will require this plural characteristic.

## 1. Introduction: A plural approach

So-called "qualia" remain a mystery. They are a mystery partly because we are not *really* sure if they exist. Some proclaim that they are the very things about which we should be most certain; others proclaim that they are a fanciful illusion, supported only by questionable "intuition pumps" in the form of thought experiments (Dennett 1988). Even if we grant their existence, they seem to be tucked away into first-person, private experience, inaccessible to third-person science. This first-/third-person chasm discussed by Nagel (1974) is still seen by many as an obstacle to progress (e.g. Dietrich & Hardcastle 2004; McGinn 2000); others have forcefully endorsed *both* first- and third-person perspectives as crucial for making progress in a science of consciousness (e.g. Velmans 1990, 2009). Other disagreements abound.

As a neural approach to consciousness was emerging two decades ago, Crick and Koch (1990) sought to overcome such extensive disagreement through a set of conditions. These conditions were intended to keep their discussion from being "frittered away in fruitless argument." (p. 264). They may have helped kickstart a neurobiology of consciousness, but they did little to stem broader debate. For example, Crick and Koch recommended that researchers should avoid getting stuck in particular definitions of consciousness. Yet some have argued that a clear definition is exactly what we need (e.g. Velmans 2009). Others still argue that a precise definition, or so-called "semantic ascent," could impede progress when we remain still so confused about consciousness (e.g. Antony 2001). Some have fractionated the concept, identifying more than one kind of "consciousness" – now associated with a large literature itself (e.g. Block 1995, 2005, 2007). As we'll see below, significant progress has come in our neurobiological understanding by designing basic tasks, based on rudimentary conceptions for now, and taking an incremental approach. The stipulations of Crick and Koch (1990) are instructive to read, as practically each of them remains a continuing issue, from the definition and function of consciousness, to the problem of "qualia."[1]

Despite the continuing debate, the groundwork laid by Crick and Koch (1990) and other researchers has led to progress in the science of consciousness. Part of our goal in this chapter is to review some of this progress, found in impressive work in neuroscience and cognitive science. This large and still-growing literature has the distinct flavor of scientific progress: nascent consensus, thorough experimental investigation, and well-articulated theoretical and quantitative models. The problem, as we will argue, is that our glance at this literature does not support one particular theory of what remains a significant puzzle: phenomenal experience. This is because phenomenal experience is brought about by *diverse* mechanisms. Our background review, we will argue, recommends a *pluralistic* approach to experience. After all, consciousness is a complex biological function underlain by a variety of mechanisms, at a variety of spatial and temporal scales. We will therefore argue that any solution

1. The following is a paraphrase of some key stipulations of their discussion (on their p. 264): (1) Don't define consciousness precisely, it may mislead us; (2) it is too premature to speculate on the function of consciousness; (3) other animals probably have consciousness (and ergo: language is not necessary); (4) there may be many other forms of consciousness, but few of these (e.g. lucid dreaming) help with scientific headway on the issue; (5) no neural theory will explain everything about consciousness; (6) let's leave qualia aside for now, though they may be assailable at some point.

to the "hard problem" (Chalmers 1997) will require a *pluralistic* approach to experience.

Admittedly, the observations we offer are not terribly new. They are not meant to be, because they will derive from our background review of such topics as the Global Workspace Theory (Baars 1988; Dehaene et al. 1998), social modulation of cognitive processing and action (Castiello et al. 2010), and the relationship between action and awareness (Jeannerod 2006). We aim, instead, to emphasize that further progress could be made by assembling a more integrative account of phenomenal experience. Our review suggests that phenomenal experience is underlain, sustained, and modulated by *diverse controlling variables* in and outside the organism. We highlight the pluralistic nature of this explanatory strategy: The multi-scale property of phenomenal experience allows *numerous theories* to have a role in the explanation, depending on one's meaning of "experience," and what measurement grain size is chosen in analysis. An argument for such "pluralistic" stances can be found in many domains of cognitive science (e.g. Bechtel 1990; Dale 2008; Dale, Dietrich & Chemero 2009; Dove 2009; Edelman 2008; McCauley & Bechtel 2001; Weiskopf 2009). This leads us to a multi-scale account of experience that, we argue, is best articulated through the language of dynamics and complexity science. This is similar in spirit to, and is influenced by, multiple-constraint accounts pioneered by, for example, Metzinger (2004; Wiese and Metzinger, this volume), Merker (2007, this volume), Shallice (1972), Tulving (1985), and of course many more.

## 2. Outline of the chapter

In what follows, we first cover three key themes in the science of awareness and consciousness, in order of relative grain size, from basic mechanisms to larger-scale social organization: (1) Global Workspace Theory and related discoveries, (2) the role of action in awareness, and (3) the central role of social experience. All three of these sections identify aspects of our nervous system, and more abstract cognitive characterizations, that are part of the story of phenomenal experience. Motivated by the role of all these levels, we then discuss the multi-scale nature of phenomenal experience, and identify two outstanding issues that may be pursued through dynamics and multi-scale complexity. We end by discussing how a more integrative approach would lead to new avenues of investigation into phenomenal experience. We do not aim to solve the hard problem, but argue that any solution would require such a pluralistic agenda.

## 3. Brains and workspaces

Recent research has made significant headway into what happens in the brain when a person becomes aware of a stimulus. An oft-cited "emerging consensus" derives from Baars' (1988, 1997) Global Workspace Theory of consciousness, which now has a well-developed neurophysiological description (e.g. Dehaene et al. 1998; Dehaene & Naccache 2001), along with carefully developed behavioral experiments to test it (e.g. Dehaene et al. 2006; Sergent et al. 2005; Sergent & Dehaene 2004), and even large-scale computational models that implement it (e.g. Baars & Franklin 2003, 2009; Franklin & Graesser 1999).

Global Workspace Theory identifies consciousness with a process that resembles working memory. At durations of several seconds in length, information becomes broadly accessible to the cognitive system, propagating widely throughout the system to influence action: "a fleeting memory capacity whose focal contents are widely distributed ('broadcast') to many unconscious specialized networks." (Baars & Franklin 2003, p. 166) This has, for example, become associated with a distributed pattern of activity that occurs across regions of the brain when a stimulus is detected – the information rendered by a stimulus propagates throughout the nervous system, engaging it in a kind of "tidal wave" of recurrent activation between low- and high-level systems (among many related proposals, see: Dehaene & Naccache 2001; Del Cul et al. 2007; Tononi & Edelman 1998; Rodriguez et al. 1999; Singer 2001).

A problem that this approach faces, as often noted, is the "hard problem" of consciousness: Why on earth do these processes result in delightful first-person experiences anyway? (Chalmers 1997) For example, if such patterns of neural synchrony, coherence, workspace-neuron connectivity, etc. are *correlated* with consciousness, then a natural question that arises is *why* does consciousness occur through these processes? The query seems intuitive. However, it could cause researchers to further stray. As Dennett (2001) warns, one must remember that the proposal is that this neuronal coherence *is* consciousness. Consciousness is no more than just that (above-threshold) wave of activity that a stimulus engenders when it participates in a conscious experience. This sophisticated (type-type) identity theory aims to avoid a category mistake through what some have called modern cognitive science's phlogiston (e.g. Churchland 1998): "qualia," or phenomenal experience, as something else *beyond* these dynamic properties of the nervous system.

But one can still be wary of the overall story so far. This worry does not stem from identifying detection or awareness of some stimulus event with the global neuronal "workspace" (because the data are quite compelling). The problem is that these experiments, and even the theories that motivate them, are based on very simple conceptions of phenomenal experience. Yoshimi (2004) provides an elegant discussion of this problem, which he calls the "mereological dogma." Our everyday experience is not constituted by a staccato sequence of stimuli detected flittingly amidst a complex array of background information (see also Spivey 2007). It is not best exemplified by punctuate moments of seeing red patches, or isolated experiences like the taste of beer alone. Everyday experience is suffused in event structure that has a high dimensionality and complex dynamic properties – subtleties that are very difficult to describe in natural language or, at present, any theory. As Yoshimi (2004) argues, "if one begins with *parts* of mental states, one has to outline what might be called a 'structural neuro-phenomenology,' which takes account of how conscious states divide into parts, how brain-states divide into parts, and how these two kinds of parts relate to one another." (See Blackmore 2002 and papers in the same issue for a related discussion, and perhaps disagreement with this point.)

In other words, one could, in a cynical mood, note that most modern cognitive neuroscience of "consciousness" is simply stimulus detection associated with certain kinds of voluntary motor control. The rub is whether you're happy referring to this as phenomenal consciousness (in a form of identity theory), or believe there's something left to explain. Many of these researchers would probably accept the basic thrust of this latter concern (see, e.g. Dehaene & Naccache 2001, p. 29). Still, the idea of the global workspace provides clues to how it might work. We revisit this below.

## 4. Action, awareness, and consciousness

Whatever these internal patterns of activity might be, the nervous system is devoted to engaging its external environment. It does not always sit passively and enjoy conscious experiences of red patches. It acts, and the way it does has important links to phenomenology. Our sense of volitional action relates to our perception of selves, and the distinction between self and other (see Jeannerod 2006 for review). Willful action also brings about changes in experience, such as our sense of relative time between successive events (Haggard et al. 2002). The sense of having caused something in the environment is itself based on a variety of variables that, when satisfied, can experimentally induce illusions of will (Wegner 2003). O'Regan and Noë (2002) and Noë (2005) have laid out convincing arguments and review showing that action is fundamentally related to perceptual processes and experience.

An interesting recent account that further seeks to explain why action is central to consciousness is found in Morsella (2005). In this "supramodular theory," phenomenal states emerge in cross-modal and integrative contexts that converge to control body plans. When multimodal constraints converge to bring about executed actions, the accompanying internal states have phenomenal properties that derive from binding these diverse information sources. This is hypothesized to produce a quality to experience that we typically call "consciousness," "awareness," and so on.

For example, recent evidence from the first author's laboratory has shown that explicit predictive actions tend to accompany awareness of learning (Dale et al. in press). Predictive actions – reaching into the world in anticipation of where it is going – may reflect a convergence of information from prior perception and action experiences that, in Morsella's (2005) terms, interfere with ongoing body plans. Phenomenal states reflect the cognitive system's integration or binding of these experiences for overall skeletomotor control. In this research on learning, Dale et al. (in press) found that overt anticipatory behavior significantly correlated with a sense of awareness of that learning. In this sense, the strategy of explicit, stimulus-specific prediction must draw the motor system away from other possible actions and guide it towards a single coordinated activity for a time. Action binds the mind into phenomenal states.

Prediction need not always be conscious, however. There is indeed evidence for unconscious prediction and learning (e.g. Turk-Browne et al. 2010). And overt decision-making and valuation may proceed just fine using a surprisingly flexible "unconscious" part of our cognitive system (Bargh & Morsella 2008). Even in perception, there have been proposals for rapid assemblage of "predictive associations" underlying vision (e.g. Bar 2009). Nevertheless, in a variety of studies, aspects of phenomenal experience have been systematically related to motoric activity, as Morsella (2005) would hypothesize. There is a long literature on distinguishing between "explicit" and "implicit" processes in cognitive psychology that we do not have space to consider here, but it may shed further light on some properties of phenomenal experience, action, and learning (e.g. Cleeremans et al. 1998; Cleeremans & Jiménez 2002; Haggard et al. 2002; Hurley 2002; Reber 1992; Sarrazin, Cleeremans & Haggard 2008; Jordan 1998; Morsella et al. 2008). It is important to note that action, as a specific variable underlying aspects of our experience, may functionally relate to the workspace account: Voluntary action may be related to the process of "broadcasting," as it could serve the function of coordinating the cognitive system to interact with its world (e.g. Franklin 2000).

## 5. Sensitive dependence on social conditions

But we do not always act alone. A frequent part of our day-to-day experience has a social dimension. Our thought processes seem to be highly sensitive to these social variables, such as who is around us when we speak loudly about something, or who is with us during a movie (e.g. Crosby et al. 2008). In the past several years, experimental research in cognitive science has shown a pervasive influence of the social on our thought processes. For example, in the domain of language, when two people converse, they become loosely coupled across many behaviors, from eye movements to gestures to speech patterns and all the way to aspects of discourse, like sarcasm (see Shockley et al. 2009 for a review). This coupling occurs across a variety of spatial and temporal scales, from basic perception all the way to higher cognitive functions. It may be that these interpersonal processes weave together during day-to-day activities, and fundamentally shape human experience.

A popular example of a low-level bridge between persons is the hypothesized mirror neuron system (Gallese et al. 2004; Rizzolatti & Craighero 2004), and research on the network involved in empathy has found similar kinds of patterns (Decety & Jackson 2004) – namely, that humans have emotional circuitry that mediates both the experience of emotion in oneself, and comprehension of emotions in others. There is of course debate about whether processes of behavioral and emotional comprehension are "innate" (Heyes et al. 2005), but one thing is for certain: Human cognition rapidly integrates information to judge states of others.

As an example of this rapidity of processing, inferences based on the beliefs and knowledge of others can sometimes be faster than other types of judgment, such as inferences from signs and symbols (Cohen & German 2010). Though there is debate about whether this skill innate and automatic (Apperly et al. 2006), there seem to be many robust contexts in which humans are inclined (through learning or otherwise) to carry out fairly complicated judgments of others' cognitive states. Even basic orienting responses can be influenced by emotion-relevant facial stimuli, especially when they involve potential threat (Friesen & Kingstone 1998; Frischen et al. 2007; Tipples 2006).

So from the perception of action and emotion, to the judgment of "mental states" in other agents, humans are capable of rapidly assessing the social relevance of stimuli at various levels of complexity (Meltzoff & Decety 2003). Cognitive processes involve rapid integration of social information to make, for example, social judgments (Freeman & Ambady 2010), and this process of social judgment and interpretation may work very quickly and robustly with even a minimum of information (Ambady & Rosenthal 1993).

From this necessarily selective review of a vast literature, it may be said that humans are sensitively dependent on social conditions. There is now a growing movement in the cognitive, social, and neurosciences investigating how deeply social variables penetrate mental processes (e.g. see for review: Amodio & Frith 2006; Adams et al. 2010; Balcetis & Lassiter 2010; Castiello et al. 2010; Frith & Frith 1999; Galantucci & Sebanz 2009; Sebanz et al. 2006; Tomasello 2009).

Given this pervasive influence of "the others," it may therefore be unsurprising that many have proposed that social information and reasoning, language, and discourse figure into our sense of selves and consciousness. If our cognitive system is actively modeling ourselves in the context of others, then the domains of narrative and discourse may indeed be central to the way we experience everyday life, and conceptualize ourselves in it. Work by Gallagher and others argues that our narrative process helps create the very experiences we have when we navigate the world (Gallagher 2000; see also Harré 2000; Humphrey 1982; and discussion in Dennett 1992).

Narratives help frame our expectations and reactions to events, which partly provide, quite possibly, the richness of human experience itself. They may help us to overcome the problem referred to as the "mereological dogma" described above (Yoshimi 2004) by hinting at answers to "What is the nature of this sense of a continuous self? Is it carried by a succession of momentary minimal selves that are tied together by real connections?" (Gallagher 2000, p. 18). Indeed, Dennett refers to our sense of selves as a kind of "center of narrative gravity" (Dennett 1992), and highlights the potential convenient *fictions* that that we construct for ourselves in these narrative structures (because, after all, centers of gravity themselves are fictions, of an important sort).

In contrast to Dennett's narrative fictions, Gallagher (2000) reviews work that points to the importance of the narrative structure as "decentered, distributed, and multiplex" (p. 20), granting us insights into mysterious moments of the self like "conflict, moral indecision and self-deception." (p. 20). Drawing from phenomenology (see also Yoshimi, this volume), Gallagher argues that this decentered conception of self – a self which takes part in a narrative that is heterogeneous, with diverse clusters of narrative structure and activity, containing a variety of important players – is more consistent with empirical evidence and everyday experience than Dennett's fictional account.

In a recent example that is likely to have an impact on this discussion, Graziano and Kastner (2011) develop a framework within which human consciousness is fundamentally related to social information. The core hypothesis put forth by the authors is that awareness "is a product of social perception" (p. 99), and is supported by a wide range of behavioral and neuroscientific studies. A central part of their explanation is the role of *simulation* in understanding other minds (e.g. more recent reviews: Gallese et al. 2004; Vogeley et al. 2001; and earlier discussion: Goldman 1992). Perhaps most interesting in this paper is the development of explicit, testable predictions about this hypothesis. Graziano and Kastner (2011) offer three such predictions. First, damage to areas associated with social processing (i.e. constructing perceptual models of others) should cause problems with conscious awareness. Second, these same areas should correlate with activity in brain structures associated with attentional control. And third, the same areas involved in building models of others should also be involved in sustaining models of the self. The authors review current evidence in favor of these predictions, and offer insight on further investigation that would support them.

The idea that consciousness is socially constructed has a simple interpretation that can be easily discarded as false (see discussion in Block 1996; Kurthen et al. 1999), since sociality can be seen as something simply built "on top" of more fundamental processes that engender our phenomenology. But basic processes that distinguish self from other seem foundational even to simple forms of consciousness and selfhood (e.g. Damasio 2010, for recent discussion). And if it is true that any kind of consciousness is somehow emergent from the collection of growing experiences and memories that a person constructs in his or her lifetime, then the social should be central to this, too. Indeed, there are well-developed conceptions of autobiographical memory as driven, centrally, by (the often social) events of our day-to-day lives (Nelson & Fivush 2004). If this is true, then human conscious experience has *as a foundation* social experience and its associated narrative structures.

## 6. Plural processes underlie phenomenal consciousness

We reviewed a series of proposals for consciousness that range from relatively simple experiences propagated in the global workspace, to processes that lie on longer timescales such as social information and experience. From this review alone, it appears that an account of *human* phenomenal experience is not going to be a simple affair of identifying some key characteristic – some strict sufficiency criterion. The Global Workspace Theory provides some understanding of how cascading interactivity among systems produces experiences of individuated stimulus events; but the extended conscious experience of self in the world may be underlain by narrative structures that depend upon a social timescale. Put simply, all of these proposals help us understand how we experience objects in our world, maintain our sense of agency, and frame it in some broader social and cultural context that provides further hues to our experience.

As Crick and Koch (1990) argued, it should be fairly clear that no theory alone can account for all aspects of phenomenal consciousness. In a discussion of explanatory proposals for consciousness, Dietrich and Hardcastle (2004) provide an impressive list of proposed foundations of conscious experience, and state that all "the items on the list have two properties: they either are necessary (at best), but not sufficient, for consciousness, or are as puzzling as consciousness itself." (p. 8) The possibilities reviewed here, while now based on growing empirical evidence and sophisticated theoretical apparatus, may still face these problems too.

Consider, for example, the Global Workspace Theory. In Dehaene et al.'s (2006) elegant threshold account, a considerable amount of local processing of a stimulus can take place "sub-threshold," and only those stimuli that coordinate a fairly distributed cluster of brain areas are "brought to awareness." Are these conditions *sufficient* for consciousness? If we granted sufficiency, it would have to be limited in the range of experiences it accounts for. Many acknowledge these limitations, sometimes remarking that any theory must, at present, be provisional.

So the problem is not that any theorist urges their account to be all-encompassing, or to be the necessary and sufficient conditions for consciousness (perhaps some do: Rosenthal 1990; Lau & Rosenthal 2011). We might argue, from the vantage point of admitted dilettantes of this literature, that proposing strict *necessary and sufficient* conditions for complex biological functions is not an enterprise that is likely to be successful. Complex biological functions are underlain by an assemblage of processes, lying at multiple scales, and subject to diverse range of scientific vantage points (see, e.g. Mitchell 2003). An explanation for the complex biological function of consciousness can buck the trend of seeking pure "demarcation" conditions for *phenomenal* consciousness, and instead simply specify how these various processes work together. This presents two challenges that we face to achieve a fuller understanding of phenomenal experience. Both can be framed as an investigation into the *time-extended* dynamics of that experience. The first is articulating the longer-timescale *sustained dynamics* of phenomenal experience. The second is to articulate its many interacting parts.

## 7. Challenge 1: The sustained dynamics of phenomenal experience

Most of the experimental evidence we have reviewed above is based on brief conscious events in the laboratory. These are, of course, the most easily controlled and induced experimentally. They are punctate – hearing a tone, noticing a word, seeing one interpretation of an ambiguous figure, and so on. One needs such delimiting moments to identify when conscious experiences are occurring so that their effects and neural correlates can be found.

But phenomenal experience does not seem like the staccato sequence of punctate experiences of the laboratory. It has a kind of structure that extends over broader scales of space and time (Yoshimi 2004). How does our cognitive system sustain this continuity, and produce more complex phenomenal experience?[2] The science of consciousness will eventually, one would suppose, go beyond these individuated moments, and into the sustained dynamics of day-to-day phenomenal experience. There are two ways this might occur. One could be through studying the "intrinsic dynamics" of the neural system – the continuity of experience arises from some dynamical process that can be identified in the brain. A second is that the dynamics of the brain must be investigated as coupled to continuous flow of energy present in the environment – sustained consciousness is an organism-environment affair.

Taking the second position, Noë (2011) argues that conscious experience must be explained by reference to the way the brain is coupled to the external world. He specifically attacks the notion that our conscious experience can be, fundamentally, identified with processes of the brain. Soft versions of his proposal would be straightforwardly endorsed by many researchers: Ultimately, how a system is coupled to its environment is important to understanding how conscious experience and related phenomena work. However, a much stronger version of this argument is presented in his book, and readers might come away with the impression that recent cognitive neuroscience is so misguided that little insight can be gleaned from more mainstream analyses (such as those described in previous sections of this chapter).

Perhaps among the more interesting issues considered by Noë is that of dreaming. If it is true that conscious experience can be had in dreams, then it seems that the external world is not *purely* necessary in stretches of conscious experience. In other words, the brain is capable of engaging in "closed" dynamics that give way to conscious experience, fully internally in the system itself. But he gives dreams short shrift: "So the appeal to dreams, like the appeal to neuroscientific interventions, leaves us more or less where it starts: with unspecific Cartesian intuitions about the interiority of our experience." (Noë 2011, p. 180)

The problem with Noë's discussion is that it assumes opposing positions few researchers would espouse. The appeal to dreams by "brain advocates" need not be for the purpose of showing the *pure and universal* interiority of experience, but rather the fundamental importance of the brain's own "intrinsic dynamic" (Kello et al. 2007; Van Orden et al. 2003) that permits real conscious experiences (of

2. This question could be answered by some by noting that this phenomenal experience is an illusion (see Blackmore 2002 and the papers in that special issue). But, as sometimes acknowledged by some of these theorists, it still remains a problem to explain the quality of that illusion and why it is there at all. This section could be phrased in these terms too; though the authors of the current chapter have all agreed with each other that such accounts do not succeed in deflating what seems like a "real illusion."

particular kinds). And indeed, as Noë also observes, when the brain is "recoupled" to the external world, that external world fundamentally changes the landscape of experience through dynamic exchanges with the organism. Yoshimi (2010, 2011, this volume) lays out a mathematical framework for understanding this, and much like Velmans' (2009) more pluralistic, perspectival approach, sees value in both levels of analysis. In a perspective he calls "active internalism," Yoshimi argues that intrinsic brain dynamics can be subjected to its own analysis and mathematical modeling. However, in order to understand how it operates in ecological context, we require a conception of how these intrinsic dynamics "fold into" the dynamics of the world itself. He proposes model systems to get at these conceptual questions, focusing on dynamic neural network models.

So Yoshimi would predict that conscious experience can emerge from the intrinsic dynamics of the central nervous system. Indeed, there has been some astonishing evidence gathered in the past few decades about the phenomenal quality of a particular kind of dream referred to as *lucid dreaming*. As observed in Crick and Koch (1990), discussed above, conscious experiences of this kind have not typically figured into scientific discussion of consciousness (but see, e.g. Revonsuo 1995; Metzinger 2003). There is now very strong evidence that lucid dreams are real, and the development of methods by LaBerge et al. (1981) and others has allowed the detection of lucid dreams, and an exploration of their qualities (see Erlacher & Chapin 2010, for a review). In the original work (e.g. LaBerge et al. 1981), subjects identified as lucid dreamers were instructed (before falling asleep) to move their eyes in particular ways during a dream. By measuring oculomotor activity during identifiably REM sleep, researchers were capable of detecting the onset of the (conscious) dream state.

Recent evidence suggests that brain imaging may reveal activities carried out by lucid dreamers. Dresler et al. (2011) instructed lucid-dreaming participants to carry out particular actions during the dream state. They had 6 lucid dreamers first carry out a pre-instructed left-right-left eye-movement so that the onset of the lucid dream could be identified. The participants were then instructed to clench one of their fists 10 times. After this first clenching, they were instructed to do the left-right-left eye movement again, then switch hands and clench once more. Researchers used the eye-movement signal to demarcate regions of brain-imaging data (with fMRI and NIRS) and found the expected lateralized pattern in sensorimotor cortex: They detected the specific activity that was being executed in the dream state.

Other concepts related to conscious experience can also be explored in lucid dreams. In some earlier studies, researchers have investigated the way that lucid dreamers experience time. By having participants count from 1001 to 1010, LaBerge (reviewed in Erlacher & Chapin 2010) found that dream time approximates real time. Studies since this one have found that there may be modality-specific effects of time perception in lucid dreams. Erlacher and Schredl (2004) found that one particular activity (squats), when performed in dreams, appeared to stretch time relative to the same activity during wakefulness.

These findings show that consciousness during dream states may be similar to that of wakefulness in important ways (e.g. sensorimotor activity and time perception), but that, as Noë would certainly predict, there are important differences. The latter point may seem obvious, given reports by many lucid dreamers that activities such as flying can be used to confirm being in the dream state (see the fundamental role of metacognition in developing and controlling lucid dreaming: Kahan & LaBerge 1994). These studies of cognitive activity in dreams, and the sorts of experiences that occur in them, surely raise many questions, but they serve as evidence, *at least*, that dream consciousness is rich with structure.

Rich conscious experience – a sense of will, and experience of events that can be remembered – all *can* take place in the brain in a relatively self-contained way. One may respond that these phenomenal experiences depend on other already-had interactions with the outside world (e.g. for their content). But this is not the same thing as saying that consciousness *requires* coupling to the external world. The intrinsic dynamics of the human brain are capable of producing distinct qualities of phenomenal experience, with structure that resembles that of world-coupled waking. So, while "distal" explanations of conscious experience must employ active interaction with the external environment, there may be "proximal" accounts of conscious experience that can use the intrinsic dynamics of the human brain as a unit of analysis. As we further discuss in the next section, this relates to the *timescales* that we choose in the explanatory agenda. The slower timescale of brain-world coupling (e.g. social interactions, extended perceptual and event experiences, etc.) serve to constrain the landscape of experiences that occur in the "phenomenological now," occurring at a faster timescale.

Anderson and colleagues (Anderson 2000; Anderson et al. 2006) have argued that an account of sustained conscious functioning must look to dynamical systems and a property of complex dynamic systems called pink or fractal noise (see also Van Leeuwen & Smit, this volume, for some discussion of this). According to Anderson and colleagues, the operation of the brain gives way to these intrinsic dynamics during conscious processing, which can be identified both in waking states and in REM sleep. Kello (under review; Kello et al. 2010) has developed a model that suggests that sustained operation of a complex nervous system requires this "echo" of noise, as it reflects the balance between two dangerous equilibria: zero activity in the brain (akin to death) and saturation of activity (akin to debilitating seizure) (cf. self-organized criticality: Bak et al. 1987; also its relationship to homeostasis for cognitive function: Parvizi & Damasio 2001). Modes of the nervous system give way to distinct dynamic patterns, and measurable signatures may reflect the sustained dynamics of conscious experience, from single-neuron firing to system-level activity (for recent discussion see Bieberich 2002; Sevush 2006; Werner 2010).

## 8. Challenge 2: Multi-scale phenomenal consciousness

So if we suppose that the sustained dynamics of conscious functioning is underlain by a kind of pink-noise-inducing interaction-dominant dynamics, what does it mean? While, as mentioned above, we do not advocate seeking a strict sufficiency demarcation point for conscious vs. non-conscious states, we should still be wary of generic properties that may inadvertently render our explanation relevant to non-biological entities, like traffic jams (Helbing 2001), that violate important explanatory intuitions (e.g. Block 1978). A second challenge we wish to describe is that the dynamics of phenomenal experience involve a richness across levels: They are integrative across space and time, between internal states and the world, and so on.

This is not merely the problem of "binding," the experimental work of which still focuses on relatively simple phenomenal states like binding features into a single visual percept. Instead, we mean that interaction-dominant dynamics giving way to *sustained* phenomenal experience involves a systematic interplay among different emergent levels of organization (Jordan & Ghin 2006; Van Orden & Holden 2003). It is also not merely a problem of specifying the structure of the nervous system and how dynamic interaction takes place within it (e.g. Buzsáki 2007; Damasio et al. 2000). This is because such exploration of structure does not, by itself, reveal the *functional* properties of that integration and how it gives way to forms of phenomenal experience. Instead, one must combine what can be seen in the interactive dynamics of the brain, in conjunction with behavioral and theoretical explorations of how they sustain our day-to-day experience: What are the *contents*, so to speak, of the various systems-level characterizations, and how are they integrated?[3]

3. The rapidly proliferating structure-function proposals has been lamented recently, along with efforts to seek integration of these proposals and associated findings (e.g. Morin 2006). See also Saygin et al. (2011) for a promising example of mapping structure to function using advanced brain imaging methods, which may provide future assistance with this general problem.

A range of brain structures and processes have been implicated in levels of awareness and self-consciousness (see Morin 2006 & Seth et al. 2005 for some concise review). The role of the cortex in the global workspace has already been discussed. But a case may even be made for the brainstem, which may play an important role in sustaining a rudimentary form of awareness (see Merker 2007, and this volume, for review). It may do so by *deciding and guiding* for an organism that faces complex probabilistic processing in the cortex, and by integrating that processing with a model of the self as the center of a world with which it is interacting. As Damasio reviews (2010; see also Merker 2007 for review), hydranencephalic children, with almost no cortex but preserved brainstem, indeed appear to show awareness of their environment, despite lacking the rich interconnected neocortex supposedly required by other theories, such as the global workspace.[4] To theorists such as Merker, *sufficiency* for conscious experience may simply lie in functioning brainstem that can integrate basic awareness of the self with the external environment. Indeed, Sevush (2006) has argued that the dynamics of consciousness are already present in the activity of a single neuron!

But, as Morin responds in discussion of Merker (2007), basic awareness is one component of phenomenology, but does not represent the complete picture. It is likely not sufficient, under other contexts of inquiry, to unveil the nature of phenomenal experience. This could include complex social cognition, which may involve building models of self or other (e.g. Frith & Frith 1999); rapid computation of timing between events of one's actions and those in the world to determine agency (e.g. Wegner 2003); a complex and probabilistic layering of computations to extract perceptual information, and feedback connectivity from conceptual knowledge (e.g. Balcetis & Dale 2007); and so on. These levels of organization in the nervous system, associated with particular functions, become integrated under normal circumstances, and bring about the subtle tinges that accompany any complex, extended event.

One way to pursue this strategy is to use more ecological experiential tasks, and collect behavioral and brain data. For example, in widely cited work by Hasson et al. (2004), intersubjective correlation between brains was explored to find what areas are activated during a viewing of *The Good, The Bad, and The Ugly*. They found that, during extended viewing of this film, subjects exhibited a surprising intersubjective consistency in the patterns of activity of the brain. These kinds of

4. The Global Workspace Theory is still consistent with this case, because the dynamic inter-region communication suggested by the global workspace may still be preserved in the case of hydranencephaly. Such a debate is outside of current discussion, but a challenge may be made to *cortex-centered models* of the workspace from these cases.

methods may be employed to explore the *differences* between states of the brain during extended events. Differences in phenomenal experiences may be determined by verbal report (e.g. from what aspects were focused on, accompanied by emotional states, etc.), and these differences may be associated with differing patterns of integration at the systems-level: a kind of "neurocinematics" (Hasson et al. 2008; see also Schier 2009, for a related discussion).

Another means of exploring the sustained and multi-scale dynamics of conscious experience is through computational modeling. In a cleverly entitled paper "Drinking from the firehouse of experience," Kuipers (2008) describes the development of dynamical systems models for capturing experience. The key computational quality of the proposed model is that the cognitive system must cope with an overwhelming amount of *possible* information from the real-world. The continuous flow of sensory data must be filtered, in some way, to guide behavior. In his model, some of that filtering is done through "trackers" that rely on the external world for such information; other filtering may take place through processes that function like a Kalman filter, which can construct predictive models of internal and external states. The result is, again, a multi-part theory of phenomenal experience, which integrates mind with world, but also proposes a series of internalized computational mechanisms that support this.

We therefore conclude that an agenda to uncover the sustained dynamics of phenomenal experience in all its complexity requires a research agenda that extends the experiencer in time, and can seek an exploration of the interaction among neural events, along with the appropriate behavioral (e.g. eye-movements) and linguistic reports. The methods are diverse. And the discoveries, in our opinion, are likely to reveal pluralistic theoretical relevance: from brainstem, to global workspaces, and social narratives.

## 9. The hierarchy of human activity: Which scales?

We've described the study of consciousness as requiring an integration of multiple scales. Scales in science run the entire gamut of possible measurement, so one may ask which *among* the multiple scales of measurement are the ones relevant to conscious experience? As we consider the multiple scales of consciousness, we should also consider whether a bigger picture is emerging. The parable of the blind men and the elephant fits naturally in a discussion of consciousness (e.g. Sloman & Chrisley 2003), and our concern with scales is no exception. The parts of the elephant are not just different. Importantly, the parts sit at different scales of a hierarchy. As Herb Simon famously wrote, hierarchical structures appear universal to all natural and artificial systems of sufficient complexity (Simon 1973). Humans are complex in this sense because, for example, neural and other physiological activities at smaller/faster scales hierarchically combine to form behavioral actions at larger/slower scales, which further combine to form social, cultural, and other group phenomena at even larger/slower scales. This description of human activity does not entail reductionism, because phenomena at each scale may require their own ontologies and explanations. This is what is intended by the moniker "interaction-dominant dynamics" (Van Orden et al. 2003), and the recognition that smaller/faster scales are nested within larger/slower scales.

So where does the elephant of phenomenal consciousness sit within the bigger picture of the hierarchy of human activity? If one first considers the temporal extent of phenomenological experiences, then there appears to be a privileged *band* of timescales (in the sense of Newell 1990), roughly on the order of seconds to minutes, at which the elephant sits. People do not experience time in nanoseconds or in years. To be clear, one can hold in mind abstractions of events on any timescale – from the cycling of an atomic clock (over 9 billion cycles per second for the caesium-133 atom; International Systems of Units, 2006) to the deep time of planets and beyond. However, the duration with which humans hold thoughts in mind has a characteristic timescale. The boundaries can be pushed somewhat. For example, some baseball players report consciously perceiving the spin of a major league pitch; some meditators concentrate on a single thought for an hour or more. They still do so under the constraints of the characteristic timescales of phenomenal experience.

So what of our blind men? Each has a piece of the elephant in hand, and the multi-scale, hierarchical perspective gives us a frame in which to relate those pieces. For starters, neural dynamics at the scale of spiking patterns unfold on roughly the same timescale as conscious awareness. Thus one could say that the Global Workspace Theory gives us a functional view of consciousness at the scale of phenomenology (seconds to minutes), and Tononi and Edelman's (1998) dynamic core hypothesis gives us a neuroscientific view at the same scale.

From here we can go up and down the hierarchy. At a smaller/faster scale, theories of reservoir computing (Maass, Natschläger & Markram 2002) may tell us something about the way that action potentials (spikes) are stitched together to form the dynamic core of phenomenological experience. For the present purposes, reservoir computing refers to the basic idea that recurrent networks with nonlinear dynamics have a generic "fading memory" with a capacity for computation, in the sense that patterns of activity transform as they fade. Studies have shown that the relatively fast timescales of membrane potentials (i.e. their rates of decay, which are on the order of tens of milliseconds) lead to spiking patterns in neural networks that fade on the order of seconds and minutes, in line with phenomenology (see Mauk & Buonomano 2004). Computational studies of homeostatic stability in spike dynamics further reinforce the *faster* phenomenological timescales, showing how simple neural mechanisms can maintain flexible, ever-changing (metastable) spike patterns, and even relate them to behavior and action (Kello, under review).

Behavior unfolds on multiple timescales, including the timescale of phenomenology. But as we move up the hierarchy, we encounter the slower timescales of learning, long-term memory, culture, and evolution. We leave phenomenology behind, but not consciousness writ large, because slower dynamics shape and constrain faster dynamics. Thus a comprehensive view of consciousness encompasses not only activity on the order of seconds and minutes, but also the larger/slower constraints on that activity. In other words, the larger/slower scales of human activity make spike patterns and behavioral patterns more than just patterns – recall that traffic jams are patterns too, and such patterns often exhibit the same earmarks of complexity (e.g. scaling laws; Kello et al. 2010) as neural and behavioral activity. What traffic jams and many other complex patterns lack, however, is a hierarchical nesting within the dynamics of learning, long-term memory, culture, and evolution.

This challenge posed by a hierarchical, multi-scale perspective is this: What is it about the larger/slower constraints of human activity that endow patterns at the scale of phenomenology with consciousness? In view of complex, self-sustaining and replicating biological organisms, goal-directedness may be an important piece of the puzzle. The basis for this hunch starts with the very definition of pattern formation, which occurs when a physical system runs counter to the second law of thermodynamics, that is, against the universal backdrop of ever-increasing entropy (Deacon 2012). Pattern formation occurs in open thermodynamic systems by virtue of energy moving along a gradient, and gradients are created by placing relatively larger/slower constraints on thermodynamics. For instance, the classic Rayleigh-Bénard preparation uses a stable temperature gradient to induce heat transport in fluid molecules rising from the bottom to the top of a heated pan. Patterns of convection rolls form under the right constraints – that is, a certain gradient coupled with a certain fluid viscosity in a certain container.

The analogy may seem far-fetched, but one conjecture is that the larger/slower constraints of human activity are somehow special, in that they create gradients for neural and behavioral pattern formation that endow them with goal-directedness. And this cascading process of running counter to the second law of thermodynamics predicts a layering of scales, with ever more complex patterns allowing diverse goals, under flexible conditions and constraints, to be carried out (Jordan & Ghin 2006). Human conscious experience may emerge from this; and if it does, it has at its foundation "proliferated" scales of organization, instantiated in the phenomenological band, while being continually constrained by even faster and slower timescales at its fringes. How we accomplish such an integration, and what theories can be integrated, are topics that conclude this chapter.

## 10. Models of an integrative pluralistic solution

Hierarchical, multi-scale dynamics, expressed in one form or another, have become the centerpiece in many discussions of consciousness (e.g. Chemero 2009; Edelman 2008; Jordan & Ghin 2006; Perlovsky & Kozma 2007; Silberstein & Chemero 2011; Thompson 2007). Though this perspective is gaining recognition, it has not shaped and defined the debate. Instead, even amongst theoretical discussion by many philosophers, consciousness is still frequently discussed in terms of punctate moments, with theories aiming to explain "conscious mental *states*" rather than "conscious mental *processes*." What we have aimed to do in this chapter is highlight the extraordinary range of processes supposed to be involved: from some core cortical dynamics to the role of higher-level social narratives and related variables. These diverse processes should help us approach two open puzzles about consciousness: the dynamics of sustained conscious experience, and how it is constituted by multiple interacting scales.

At the beginning of this chapter, and after our background review, we argued that a solution to these challenges should utilize an integration of *multiple* perspectives. This strategy is often referred to as "explanatory pluralism." It is not relativism ("all perspectives are correct"), nor is it a brand of theoretical nihilism ("let's give up this debate"). Pluralism is inspired by the observation that a complex system, measured in diverse ways, will admit of emergent patterns that "exist" in some (at-minimum observer-centered) way, and that theories capturing these patterns can play a powerful role in our overall understanding of that system (see Dale et al. 2009, for a brief summary). The set of all relevant theories will not necessarily be obviously consistent because, after all, emergent patterns by definition have properties not easily inferable from the properties of the parts that bring about the pattern: More is different (Anderson 1972). The challenge is to have these theories mutually inform one another, rather than seeing them *only* as competitors (McCauley & Bechtel 2001).

So what diverse families of theories are relevant here, and how should they be integrated? We discussed three in this chapter: the Global Workspace, action-centered consciousness, and social constitutivity. But we chose to discuss a particular range of complexity, and used these as examples. Other theories are also relevant. For example, Rosenthal's well-known higher-order thought (HOT) account may be related to the workspace, because the workspace suggests there is broad informational redundancy in activation across regions of the brain when one becomes conscious of a particular stimulus. HOT would state that these broad patterns of neural communication serve to generate higher-order referential encoding (Lau & Rosenthal 2011). Gathering evidence for this may reveal that HOT is relevant to the Global Workspace, as it may specify one functional role of cascading activation across the cortex. Comparisons have been drawn between these theories in the past, but it is not obvious to us that they must be seen as mutually exclusive, despite some simplistic comparisons (see Lau and Rosenthal 2011, Table 2).

Another theoretical example is Chalmers' (1997) well-known discussion of protophenomenal panpsychism. In this account, one supposes that matter has both physical *and* phenomenal properties. This perspective is unlikely to convince many people about the phenomenal constitution of the universe's basic physical material. But it may be possible to explore a kind of *neural panpsychism.* Sevush (2006) has recently argued that the neural dynamics that appear to be intrinsic to waking conscious states can be attributed to the firing properties of *single neurons*. In this account, he sees each neuron as having a particular distribution of activity over its dendritic extensions that may constitute a simple form of "consciousness"; as layers of these neurons fire together, their collective dynamics of dendritic activity may be what brings about higher-order, complex phenomenal states. The proposal is interesting, but at the very least it reveals a problem that panpsychism both "suffers from" (Chalmers 2002) yet succeeds in highlighting: There may be levels of complexity cascading across neural activity, each of which may be casually attributable to consciousness, and understanding their *composition* is an open puzzle (what Yoshimi 2004, calls a 'structural neuro-phenomenology').

If consciousness is a complicated biological function underlain by diverse processes at multiple scales of complexity, then we should marshal the tools of dynamics and complexity science to explore it. It therefore seems *very unlikely* that some unitary "theory of consciousness" is going to emerge that renders all other theories irrelevant. Complex biological functions are produced by an assemblage of (only approximately separable) components, each functioning in sophisticated, context-dependent ways. But this may be said for many concepts in cognitive science. Perhaps labeling consciousness the last frontier of cognitive science is getting a bit ahead of ourselves. After all, cognitive scientists are still grappling with fairly vague concepts such as "representation" and "affordances,"[5] central to many explanations of other cognitive processes, such as language and perception. To move consciousness along, the goal could be to situate it amongst the normal kind of confusion and controversy, characteristic of our young science. One way

5. Anti-representational advocates of "affordances" often state that affordances are somehow better defined than "representations." In the experience of at least one of the authors of this chapter, one can with great facility find as much debate about the nature of affordances as representations; there is, therefore, no *specific* theoretical consensus on either of these ideas from either camp advocating for them.

to accomplish this, in cognitive science, is to go full bore on producing *models* of what we propose to underlie consciousness. The literature on consciousness is replete with verbiage, but relatively short on implemented mechanisms (admittedly, the current chapter is no exception).

Models of how the various emergent capacities of our cognitive system, coupled to the environment, hang together to produce our subjective human experience requires theoretical tools that can carry out such a multi-scale integration. We have described some of these already, including Kuipers (2008) dynamical model. Other hybrid systems may be relevant, such as the large-scale cognitive architecture called LIDA (see Franklin et al. 2007). We might also take a foundational approach, assembling the smaller/faster scales to observe their emergent properties, such as in models of reservoir computing with critical branching (Kello, under review; Kello & Mayberry 2010), K-set models (Kozma et al. 2007), or large-scale models of thalamo-cortical dynamics (Izhikevich & Edelman 2008). Other approaches, such as sequential dynamic systems or "simfrastructures" (Barrett et al. 2006) would allow dynamical modeling that cuts between these bottom-up and top-down strategies. These models permit simulation of large network structures, specified at various scales, producing a generative model of complex system behavior that is still amenable to system-level analysis. The development of adaptive models, with similar goals, is just beginning (e.g. Hernandez et al. 2009; Sanz et al. 2009).

## References

Adams, R.B., Ambady, N. & Nakayama, K. (2010). *The science of social vision* (Vol. 7). Oxford University Press.

Ambady, N. & Rosenthal, R. (1993). Half a minute: Predicting teacher evaluations from thin slices of nonverbal behavior and physical attractiveness. *Journal of Personality and Social Psychology, 64*, 431–441. doi:10.1037/0022-3514.64.3.431.

Amodio, D.M. & Frith, C.D. (2006). Meeting of minds: The medial frontal cortex and social cognition. *Nature Reviews Neuroscience, 7*(4), 268–277.

Anderson, C.M. (2000). From molecules to mindfulness: How vertically convergent fractal time fluctuations unify cognition and emotion*. *Consciousness & Emotion, 1*(2), 193–226.

Anderson, C.M., Lowen, S.B. & Renshaw, P.F. (2006). Emotional task-dependent low-frequency fluctuations and methylphenidate: Wavelet scaling analysis of 1/f-type fluctuations in fMRI of the cerebellar vermis. *Journal of Neuroscience Methods, 151*(1), 52–61.

Antony, M.V. (2001). Is "Consciousness" ambiguous? *Journal of Consciousness Studies, 8*(2), 19–44.

Apperly, I.A., Riggs, K.J., Simpson, A., Chiavarino, C. & Samson, D. (2006). Is belief reasoning automatic? *Psychological Science, 17*(10), 841–844. doi:10.1111/j.1467-9280.2006.01791.x.

Baars, B.J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.

Baars, B.J. (1997). *In the theater of consciousness: The workspace of the mind*. Oxford University Press, USA.

Baars, B.J. & Franklin, S. (2009). Consciousness is computational: The LIDA model of global workspace theory. *International Journal of Machine Consciousness, 1*(1), 23–32.

Bak, P., Tang, C. & Wiesenfeld, K. (1987). Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters, 59*(4), 381–384.

Balcetis, E.E. & Dale, R. (2007). Conceptual set as a top-down constraint on visual object identification. *Perception, 36*(4), 581–595.

Balcetis, E.E. & Lassiter, G. (2010). *Social psychology of visual perception*. New York, NY: Psychology Press.

Bar, M. (2009). The proactive brain: Memory for predictions. *Philosophical Transactions B, 364*(1521), 1235.

Bargh, J.A. & Morsella, E. (2008). The unconscious mind. *Perspectives on Psychological Science*, *3*(1), 73–79.

Barrett, C., Eubank, S. & Marathe, M. (2006). Modeling and simulation of large biological, information and socio-technical systems: An interaction based approach. In D. Goldin, S. Smolka & P. Wegner (Eds.), *Interactive Computation: The New Paradigm* (pp. 353–394). Springer Verlag.

Bechtel, W. (1990). Multiple levels of inquiry in cognitive science. *Psychological Research, 52*(2–3), 271–281.

Bieberich, E. (2002). Recurrent fractal neural networks: A strategy for the exchange of local and global information processing in the brain. *Biosystems, 66*(3), 145–164.

Blackmore, S. (2002). There is no stream of consciousness. What is all this? What is all this stuff around me; this stream of experiences that I seem to be having all the time? *Journal of Consciousness Studies, 9*, *5*(6), 17–28.

Block, N. (1978). Troubles with functionalism. In C. Savage (Ed.), *Perception and cognition: Issues in the foundations of psychology*. University of Minnesota Press.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences, 18*(02), 227–247. doi:10.1017/S0140525X00038188.

Block, N. (1996). On a confusion about a function of consciousness.

Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences, 9*(2), 46–52.

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences, 30*(5–6), 481–499. doi:10.1017/S0140525X07002786.

Buzsáki, G. (2007). The structure of consciousness. *Nature*, *446*(7133), 267.

Castiello, U., Becchio, C., Zoia, S., Nelini, C., Sartori, L., Blason, L., D'Ottavio, G. et al. (2010). Wired to be social: The ontogeny of human interaction. *PLoS ONE, 5*(10), e13199. doi:10.1371/journal.pone.0013199.

Chalmers, D.J. (1997). *The conscious mind: In search of a fundamental theory*. Oxford University Press, USA.

Chalmers, D.J. (2002). Consciousness and its place in nature. In S. Stich & T. Warfield (Eds.), *The Blackwell guide to philosophy of mind* (pp. 102–142). Blackwell Publishing Ltd. Retrieved from http://onlinelibrary.wiley.com/doi/10.1002/9780470998762.ch5/summary.

Chemero, A. (2009). *Radical embodied cognitive science*. The MIT Press.

Cleeremans, A., Destrebecqz, A. & Boyer, M. (1998). Implicit learning: news from the front. *Trends in Cognitive Sciences, 2*(10), 406–416.

Cleeremans, A. & Jiménez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. *Implicit learning and consciousness: An empirical, computational and philosophical consensus in the making*, 1–40.

Cohen, A.S. & German, T.C. (2010). A reaction time advantage for calculating beliefs over public representations signals domain specificity for "theory of mind." *Cognition, 115*(3), 417–425.

Crick, F. & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences* (Vol. 2, p. 203).

Crosby, J.R., Monin, B. & Richardson, D. (2008). Where do we look during potentially offensive behavior? *Psychological Science, 19*(3), 226–228.

Dale, R. (2008). The possibility of a pluralist cognitive science. *Journal of Experimental & Theoretical Artificial Intelligence, 20*(3), 155–179.

Dale, R., Dietrich, E. & Chemero, A. (2009). Explanatory pluralism in cognitive science. *Cognitive Science, 33*, 739–742.

Damasio, A. (2010). *Self comes to mind: Constructing the conscious brain*. Pantheon Books.

Deacon, T.W. (2012). *Incomplete nature: How mind emerged from matter*. WW Norton & Co Inc.

Decety, J. & Jackson, P.L. (2004). The functional architecture of human empathy. *Behavioral and Cognitive Neuroscience Reviews, 3*(2), 71–100. doi:10.1177/1534582304267187.

Dehaene, S., Changeux, J.P., Naccache, L., Sackur, J. & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences, 10*(5), 204–211.

Dehaene, S., Kerszberg, M. & Changeux, J.P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences, 95*(24), 14529.

Dehaene, S. & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79*(1), 1–37.

Dennett, D.C. (1988). Quining qualia. *Consciousness in modern science*.

Dennett, D.C. (1992). The self as a center of narrative gravity. *Self and consciousness: Multiple perspectives*.

Dennett, D.C. (2001). Are we explaining consciousness yet? *Cognition, 79*(1–2), 221–237. doi:10.1016/S0010-0277(00)00130-X.

Dietrich, E. & Hardcastle, V.G. (2004). *Sisyphus's Boulder: Consciousness and the limits of the knowable* (Vol. 60). John Benjamins.

Dove, G. (2009). Beyond perceptual symbols: A call for representational pluralism. *Cognition, 110*, 412–431.

Dresler, M., Koch, S.P., Wehrle, R., Spoormaker, V.I., Holsboer, F., Steiger, A., Sämann, P.G. et al. (2011). Dreamed movement elicits activation in the sensorimotor cortex. *Current Biology, 21*(21), 1833–1837. doi:10.1016/j.cub.2011.09.029.

Edelman, S. (2008). On the nature of minds, or: truth and consequences. *Journal of Experimental & Theoretical Artificial Intelligence, 20*(3), 181–196.

Erlacher, D. & Chapin, H. (2010). Lucid dreaming: Neural virtual reality as a mechanism for performance enhancement. *International Journal of Dream Research*, 3(1), 7–10.

Erlacher, D. & Schredl, M. (2004). Time required for motor activity in lucid dreams. *Perceptual and Motor Skills*, 99(3 Pt 2), 1239–1242.

Franklin, S. (2000). Deliberation and voluntary action in "conscious" software agents. *Neural Network World*, 10, 505–521.

Franklin, S. & Patterson Jr, F. (2006). The LIDA architecture: Adding new modes of learning to an intelligent, autonomous, software agent. *pat, 703*, 764–1004.

Franklin, S., & Graesser, A. (1999). A software agent model of consciousness. *Consciousness and Cognition, 8*(3), 285–301. doi:10.1006/ccog.1999.0391.

Freeman, J.B. & Ambady, N. (2010). When two become one: Temporally dynamic integration of the face and voice. *Journal of Experimental Social Psychology, 47*, 259–263.

Friesen, C.K. & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic Bulletin & Review, 5*(3), 490–495.

Frischen, A., Bayliss, A.P. & Tipper, S.P. (2007). Gaze cueing of attention: Visual attention, social cognition, and individual differences. *Psychological Bulletin, 133*, 694–724. doi:10.1037/0033-2909.133.4.694.

Frith, C.D. & Frith, U. (1999). Interacting minds – A biological basis. *Science, 286*(5445), 1692–1695. doi:10.1126/science.286.5445.1692.

Galantucci, B. & Sebanz, N. (2009). Joint action: Current perspectives. *Topics in Cognitive Science, 1*(2), 255–259.

Gallagher. (2000). Philosophical conceptions of the self: Implications for cognitive science. *Trends in Cognitive Sciences*, *4*(1), 14–21.

Gallese, V., Keysers, C. & Rizzolatti, G. (2004). A unifying view of the basis of social cognition. *Trends in Cognitive Sciences, 8*(9), 396–403. doi:10.1016/j.tics.2004.07.002.

Graziano, M.S.A. & Kastner, S. (2011). Human consciousness and its relationship to social neuroscience: A novel hypothesis. *Cognitive Neuroscience, 2*(2), 98–113.

Haggard, P., Clark, S. & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience, 5*(4), 382–385.

Harré, R. (2000). Social construction and consciousness. *Investigating phenomenal consciousness*. Amsterdam, Netherlands: John Benjamin Publishing Company, 233–253.

Hasson, U., Landesman, O., Knappmeyer, B., Vallines, I., Rubin, N. & Heeger, D.J. (2008). Neurocinematics: The neuroscience of film. *Projections, 2*(1), 1–26.

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G. & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science, 303*(5664), 1634–1640. doi:10.1126/science.1089506.

Helbing, D. (2001). Traffic and related self-driven many-particle systems. *Reviews of Modern Physics, 73*(4), 1067–1141. doi:10.1103/RevModPhys.73.1067.

Hernandez, C., Lopez, I. & Sanz, R. (2009). The Operative mind: A functional, computational and modeling approach to machine consciousness. *International Journal of Machine Consciousness, 1*(1), 83–98.

Heyes, C., Bird, G., Johnson, H. & Haggard, P. (2005). Experience modulates automatic imitation. *Cognitive Brain Research, 22*(2), 233–240.

Humphrey, N. (1982). Consciousness: A just-so story. *New Scientist*, 95, 473–477.

Hurley, S.L. (2002). *Consciousness in action*. Harvard University Press.

Izhikevich, E.M. & Edelman, G.M. (2008). Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences, 105*(9), 3593–3598. doi:10.1073/pnas.0712231105.

Jeannerod, M. (2006). *Motor cognition: What actions tell the self*. Oxford University Press.

Jordan, J.S. (1998). Recasting Dewey's critique of the reflex-arc concept via a theory of anticipatory consciousness: Implications for theories of perception. *New Ideas in Psychology, 16*(3), 165–187.

Jordan, J.S. & Ghin, M. (2006). (Proto-) consciousness as a contextually emergent property of self-sustaining systems. *Mind and Matter, 4*(1), 45–68.

Kahan, T.L. & LaBerge, S. (1994). Lucid dreaming as metacognition: Implications for cognitive science. *Consciousness and Cognition, 3*, 246–264.

Kello, C.T. (under review). Critical branching neural networks.
Kello, C.T., Beltz, B.C., Holden, J.G. & Van Orden, G.C. (2007). The emergent coordination of cognitive function. *Journal of Experimental Psychology: General, 136*(4), 551.
Kello, C.T. & Mayberry, M.R. (2010). Critical branching neural computation. Neural Networks (IJCNN), *The 2010 International Joint Conference on* (pp. 1–7).
Kozma, R., Aghazarian, H., Huntsherger, T., Tunstel, E. & Freeman, W.J. (2007). Computational aspects of cognition and consciousness in intelligent devices. *Computational Intelligence Magazine, IEEE, 2*(3), 53–64.
Kuipers, B. (2008). Drinking from the firehose of experience. *Artificial Intelligence in Medicine, 44*(2), 155–170. doi:10.1016/j.artmed.2008.07.010.
Kurthen, M., Grunwald, T. & Elger, C.E. (1999). Consciousness as a social construction. *Behavioral and Brain Sciences, 22*(01), 197–199. doi:null.
LaBerge, S., Nagel, L.E., Dement, W.C. & Zarcone, V.P. (1981). Lucid dreaming verified by volitional communicaton during REM sleep. *Perceptual and Motor Skills, 52*(3), 727–732. doi:10.2466/pms.1981.52.3.727.
Lau, H. & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences, 15*, 365–373.
Maass, W., Natschläger, T. & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation, 14*(11), 2531–2560.
Mauk, M.D. & Buonomano, D.V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience, 27*(1), 307–340. doi:10.1146/annurev.neuro.27.070203.144247.
McCauley, R.N. & Bechtel, W. (2001). Explanatory pluralism and heuristic identity theory. *Theory & Psychology, 11*(6), 736–760.
McGinn, C. (2000). *The mysterious flame: Conscious minds in a material world*. Basic Books.
Meltzoff, A.N. & Decety, J. (2003). What imitation tells us about social cognition: A rapprochement between developmental psychology and cognitive neuroscience. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 358*(1431), 491–500.
Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences, 30*(1), 63–80.
Metzinger, T. (2003). *Being no one: The self-model theory of subjectivity*. Bradford Books.
Mitchell, S.D. (2003). *Biological complexity and integrative pluralism*. Cambridge University Press.
Morin, A. (2006). Levels of consciousness and self-awareness: A comparison and integration of various neurocognitive views. *Consciousness and Cognition, 15*(2), 358–371. doi:10.1016/j.concog.2005.09.006.
Morsella, E. (2005). The function of phenomenal states: Supramodular interaction theory. *Psychological Review, 112*(4), 1000–1021.
Morsella, E., Bargh, J.A. & Gollwitzer, P.M. (2008). *Oxford handbook of human action* (1st ed.). Oxford University Press, USA.
Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review, 83*(4), 435–450.
Nelson, K. & Fivush, R. (2004). The emergence of autobiographical memory: A social cultural developmental theory. *Psychological Review, 111*(2), 486.
Noë, A. (2005). *Action in perception*. MIT Press.
Noë, A. (2011). Out of our heads: Why you are not your brain, and other lessons from the biology of consciousness.
O'Regan, J.K. & Noë, A. (2002). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*(5), 939–973.

Parvizi, J. & Damasio, A. (2001). Consciousness and the brainstem. *Cognition, 79*(1–2), 135–160. doi:10.1016/S0010-0277(00)00127-X.

Perlovsky, L.I. & Kozma, R. (2007). *Neurodynamics of cognition and consciousness*. Springer.

Reber, A.S. (1992). An evolutionary context for the cognitive unconscious. *Philosophical Psychology, 5*(1), 33–51.

Rizzolatti, G. & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience, 27*, 169–192. doi:10.1146/annurev.neuro.27.070203.144230.

Rodriguez, E., George, N., Lachaux, J.P., Martinerie, J., Renault, B. & Varela, F.J. (1999). Perception's shadow: Long-distance synchronization of human brain activity. *Nature, 397*(6718), 430–433.

Rosenthal, D.M. (1990). *A theory of consciousness*. Zentrum für interdisziplinäre Forschung.

Sanz, R., HernáNdez, C., GóMez, J., Bermejo-Alonso, J., RodríGuez, M., Hernando, A. & SáNchez, G. (2009). Systems, models, and self-awareness: Towards architectural models of consciousness. *International Journal of Machine Consciousness, 01*(02), 255–279. doi:10.1142/S1793843009000359.

Sarrazin, J. C., Cleeremans, A. & Haggard, P. (2008). How do we know what we are doing? Time, intention and awareness of action. *Consciousness and Cognition, 17*(3), 602–615.

Saygin, Z.M., Osher, D.E., Koldewyn, K., Reynolds, G., Gabrieli, J.D.E. & Saxe, R.R. (2011). Anatomical connectivity patterns predict face selectivity in the fusiform gyrus. *Nature Neuroscience*.

Schier, E. (2009). Identifying phenomenal consciousness. *Consciousness and Cognition, 18*(1), 216–222. doi:10.1016/j.concog.2008.04.001.

Sebanz, N., Bekkering, H. & Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in Cognitive Sciences, 10*(2), 70–76.

Sergent, C., Baillet, S. & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience, 8*(10), 1391–1400.

Sergent, C. & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science, 15*(11), 720–728. doi:10.1111/j.0956-7976.2004.00748.x.

Seth, A.K., Baars, B.J. & Edelman, D.B. (2005). Criteria for consciousness in humans and other mammals. *Consciousness and Cognition, 14*(1), 119–139.

Sevush, S. (2006). Single-neuron theory of consciousness. *Journal of Theoretical Biology, 238*(3), 704–725.

Shallice, T. (1972). Dual functions of consciousness. *Psychological Review, 79*(5), 383.

Silberstein, M. & Chemero, A. (2011). Complexity and extended phenomenological-cognitive systems. *Topics in Cognitive Science*. doi:10.1111/j.1756-8765.2011.01168.x

Simon, H.A. (1973). The organization of complex systems. In H. Pattee (Ed.), *Hierarchy theory: The challenge of complex systems* (pp. 3–27). George Brazilier.

Singer, W. (2001). Consciousness and the binding problem. *Annals of the New York Academy of Sciences, 929*(1), 123–146.

Sloman, A. & Chrisley, R. (2003). Virtual machines and consciousness. *Journal of Consciousness Studies, 10, 4*(5), 133–172.

Spivey, M.J. (2007). *The continuity of mind*. Oxford, UK: Oxford University Press.

Tipples, J. (2006). Fear and fearfulness potentiate automatic orienting to eye gaze. *Cognition & Emotion, 20*, 309–320. doi:10.1080/02699930500405550.

Tomasello, M. (2009). *Why we cooperate*. The MIT Press.

Tononi, G., & Edelman, G.M. (1998). Consciousness and complexity. *Science, 282*(5395), 1846–1851. doi:10.1126/science.282.5395.1846.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie Canadienne, 26*(1), 1–12. doi:10.1037/h0080017.

Turk-Browne, N.B., Scholl, B.J., Johnson, M.K. & Chun, M.M. (2010). Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience, 30*(33), 11177–11187.

Van Orden, G.C. & Holden, J.G. (2002). Intentional contents and self-control. *Ecological Psychology, 14*(1), 87–109.

Van Orden, G.C., Holden, J.G. & Turvey, M.T. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology. General, 132*(3), 331–350. doi:10.1037/0096-3445.132.3.331.

Velmans, M. (1990). Consciousness, brain and the physical world. *Philosophical Psychology, 3*(1), 77–99.

Velmans, M. (1991). Is human information processing conscious? *Behavioural and Brain Sciences, 14*(4), 651–726.

Velmans, M. (2009). *Understanding consciousness*. Taylor & Francis.

Wegner, D.M. (2003). The mind's best trick: How we experience conscious will. *Trends in Cognitive Sciences, 7*(2), 65–69.

Weiskopf, D.A. (2009). The plurality of concepts. *Synthese, 169*(1), 145–173.

Werner, G. (2010). Fractals in the nervous system: Conceptual implications for theoretical neuroscience. *Frontiers in Physiology, 1*.

Yoshimi, J. (2011). Phenomenology and connectionism. *Frontiers in Psychology*, 2.

Yoshimi, J. (2004). Field theories of mind and brain. In L. Embree (Ed.), *Gurwitsch's relevancy for cognitive science* (Vol. 52, pp. 111–129). Dordrecht: Springer Netherlands. Retrieved from http://www.springerlink.com/content/g568600qt55j1640/

Yoshimi, J. (2011a). Active internalism and open dynamical systems. *Philosophical Psychology*, 1–24. doi:10.1080/09515089.2011.569919.

Yoshimi, J. (2011b). Phenomenology and connectionism. *Frontiers in Psychology*, *2*(288). doi:10.3389/fpsyg.2011.00288.

Yoshimi, J. (2007). Mathematizing phenomenology. *Phenomenology and the Cognitive Sciences, 6*(3), 271–291. doi:10.1007/s11097-007-9052-4.