



Review

# Intelligent architectures for robotics: The merging of cognition and emotion

Luiz Pessoa

*Department of Psychology, Department of Electrical and Computer Engineering, Maryland Neuroimaging Center, University of Maryland, College Park, MD 20742, USA*

Available online 22 October 2019

Communicated by Felix Schoeller

---

## Abstract

What is the place of emotion in intelligent robots? In the past two decades, researchers have advocated for the inclusion of some emotion-related components in the general information processing architecture of autonomous agents, say, for better communication with humans, or to instill a sense of urgency to action. The framework advanced here goes beyond these approaches and proposes that emotion and motivation need to be integrated with all aspects of the architecture. Thus, cognitive-emotional integration is a key design principle. Emotion is not an “add on” that endows a robot with “feelings” (for instance, reporting or expressing its internal state). It allows the *significance* of percepts, plans, and actions to be an integral part of all its computations. It is hypothesized that a sophisticated artificial intelligence cannot be built from separate cognitive and emotional modules. A hypothetical test inspired by the Turing test, called the Dolores test, is proposed to test this assertion.

© 2019 Elsevier B.V. All rights reserved.

*Keywords:* Robotics; Emotion; Cognition; Integration; Dolores test

---

## 1. Introduction

In the first episode of *Westworld*, a popular HBO series, a human is examining Dolores, one of the central characters of the show. Dolores speaks with a heavy frontier accent. The human says “lose the accent,” to which she promptly acquiesces. Her answers are charged with emotion, they convey a deep feeling of confusion and anxiety; she’s distraught. The examiner then commands: “cognition only.” Once again, Dolores does as she’s told, losing all trace of emotion. Dolores is a “host” – how the humanoid robots are called – at *Westworld*, an Old West theme park created for the entertainment of human “guests” who are free to kill or have sex with them.

For over two millennia, Western thinkers have separated emotion from cognition, emotion being the lowly sibling of the two. Cognition helps explain the quantum nature of matter and sends humans to the moon. Emotion might save the lioness in the Savannah but makes humans act irrationally with disconcerting frequency. Against this backdrop, how should autonomous robots be built? Should emotion be part of their information-processing architecture? If so, how should it interact with cognition?

---

*E-mail address:* [pessoa@umd.edu](mailto:pessoa@umd.edu).

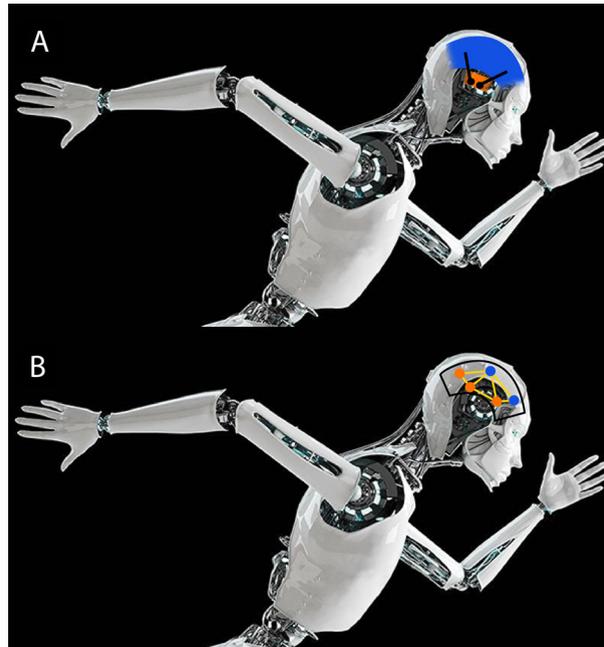


Fig. 1. Information-processing architecture of intelligent robots. (A) Modular architecture with separate cognitive (blue) and emotional (red) components. The cognitive module also inhibits the emotion module when deemed necessary. (B) Non-modular architecture with integrated components. Reproduced with permission [104].

Cognition refers to processes such as memory, attention, language, problem solving, and planning. Many cognitive processes are thought to involve sophisticated functions that might be uniquely human. Furthermore, they often involve so-called controlled processes, such as when the pursuit of a goal needs to be protected from distraction. Whereas there is relative agreement about what constitutes cognition, there is less accord about emotion. Some investigators link it to the concepts of drive and motivation, suggesting that emotions are states elicited by rewards and punishers [1]. Others propose that emotions are involved in the evaluation of events, called “appraisals” [2], which can be explicit or implicit. Some approaches focus on basic emotions (for example, fear and anger [3]) while others focus on an extended set of emotions, including moral ones (for example, pride and envy [4]). Finally, many researchers suggest that emotions are strongly linked to bodily states [5].

In building an “intelligent” robot it might seem best to focus on its cognitive capacities, perhaps disregarding emotion entirely, or perhaps including only as much of emotion as necessary. Fig. 1A illustrates this scenario, where the blue segment represents a cognitive module and the orange segment indicates the emotion module of a hypothetical robot. When emotion is included, to avoid “emotional interference,” another design property would allow the cognitive module to downregulate emotion inputs, minimizing their influence to the desired extent. Thus, emotion is restricted to contributing to behavior only when deemed useful by the cognitive module.

It is conceivable that the robot in Fig. 1A could function effectively in some settings. The cognitive architecture would endow it with the required problem-solving capabilities, and the emotion module would instantiate a sense of urgency into decisions and actions [6,7]. The argument developed here, however, is that this type of architecture, in general, is inherently limited in its behavioral capabilities. For sophisticated intelligent behavior, I argue that cognition and emotion need to be intertwined in the general information-processing architecture (Fig. 1B).

What is meant by “sophisticated intelligent behavior?” While the goal here is not to precisely answer this question, the question can be rephrased as follows: “Does building models of emotional/affective mechanisms into our robots make them more natural, more useful, or more efficient?” [8]. Furthermore, a chief target goal should be *autonomy*, that is, the ability to function independently in diverse and challenging conditions and environments [9] (for example, interplanetary travel). Importantly, the argument made here for a cognitive-emotional architecture is not simply rooted in the need for better communication with humans ([10]; given that the latter do have emotion) or, relatedly, to generate emotional expressions [10], or even for survival-related purposes [7]. The contention is that, for the types of intelligent

behaviors frequently described as cognitive (say, attention, problem solving, planning), the integration of emotion and cognition is necessary [11].

To motivate the information-processing architecture described here, I will first describe growing evidence that emotion and cognition are integrated in the brain. Next, general concepts of brain evolution will be reviewed, and implications for the organization of emotion and cognition in the brain discussed. Based on these ideas, a cognitive-emotional architecture will be motivated and proposed properties described. I will then develop the hypothesis that a humanoid robot cannot exhibit fully autonomous “interesting” behaviors without cognitive-emotional integration, which will be summarized via the “Dolores test.” General conclusions will be then outlined.

## 2. Cognition and emotion in the brain

How is emotion organized in the brain? Historically, textbooks have described the emotional brain by listing a series of brain regions and discussing their purported main function(s) (see also [12]). Several subcortical regions would be featured, including the amygdala (“fear processing”), hypothalamus (“survival mechanisms”), periaqueductal gray (“defensive behaviors”), hippocampus (“emotional memory”), nucleus accumbens (“reward processing”), and so on. Cortical regions would include the cingulate cortex (“affective action”), insula (“processing of bodily signals”), and orbitofrontal cortex (“reward processing”).

Although the idea that some brain regions are specialized for emotion is favored by some researchers [13], a different view has emerged in the past two decades, one in which emotion and cognition strongly interact and, in fact, are integrated in the brain [11,14]. Before describing this perspective, I will first motivate it in intuitive terms.

Consider a hypothetical organism that has its visual cortex (the part that receives signals from the eye’s retina) and its auditory cortex (the part that receives signals from the ear’s cochlea) heavily interconnected with each other. This architecture means that every time signals are processed by the visual cortex, they strongly influence auditory cortex, and vice versa – vision and audition mutually influence each other. Indeed, it would be better to call the parts of the brain in question visual-auditory (or audio-visual) areas. Even without studying behavior in our hypothetical organism, we could make the strong prediction that it would be unable to respond in a purely visual or auditory way. This is because its brain is built such that the two sensory domains have, in effect, no clear boundary – they are not independent modules.

Emotion and cognition are not separable in the human brain for similar reasons. Knowledge about anatomy and physiology indicates that the two are closely intertwined [15]. Next, I describe principles of brain organization that reveal why emotion is not a distinct unit of the brain.

### 2.1. Massive combinatorial anatomical connectivity

Computational analysis of anatomical connectivity demonstrates that both cortical and subcortical brain regions are densely interconnected [16–20] (Fig. 2). Rich connectivity is not limited to specific sectors of the brain (say, prefrontal cortex) but instead encompasses all of them (including brainstem and cerebellum). Central anatomical properties include: 1) Massive interconnectivity; 2) High global *accessibility* [21]; and 3) The existence of a “connectivity core” or “rich-club” of regions marked by especially high levels of connectivity [21,22]. Focusing on macaque cortical regions, one study described a core set of 17 regions spanning parietal, temporal, and frontal cortex that was marked by 92% connectivity density (92% of the connections that could exist were present) [23]. By combining multiple sources of data, another study described a core that was distributed across all major brain sectors (all cortical lobes, thalamus, and subcortical regions in the forebrain) [22].

In all, the anatomical connectivity is such that signals can flow between disparate regions in very few steps. And because *hub* regions (those that are particularly well connected) aggregate and distribute information widely, diverse types of signals (say, related to perception and action) have the opportunity to influence one another [24].

### 2.2. High distributed functional connectivity

Understanding brain function requires characterizing how regions are “functionally connected.” The idea of *functional connectivity* was devised to characterize how neurons (or regions) interact and was defined as the “temporal

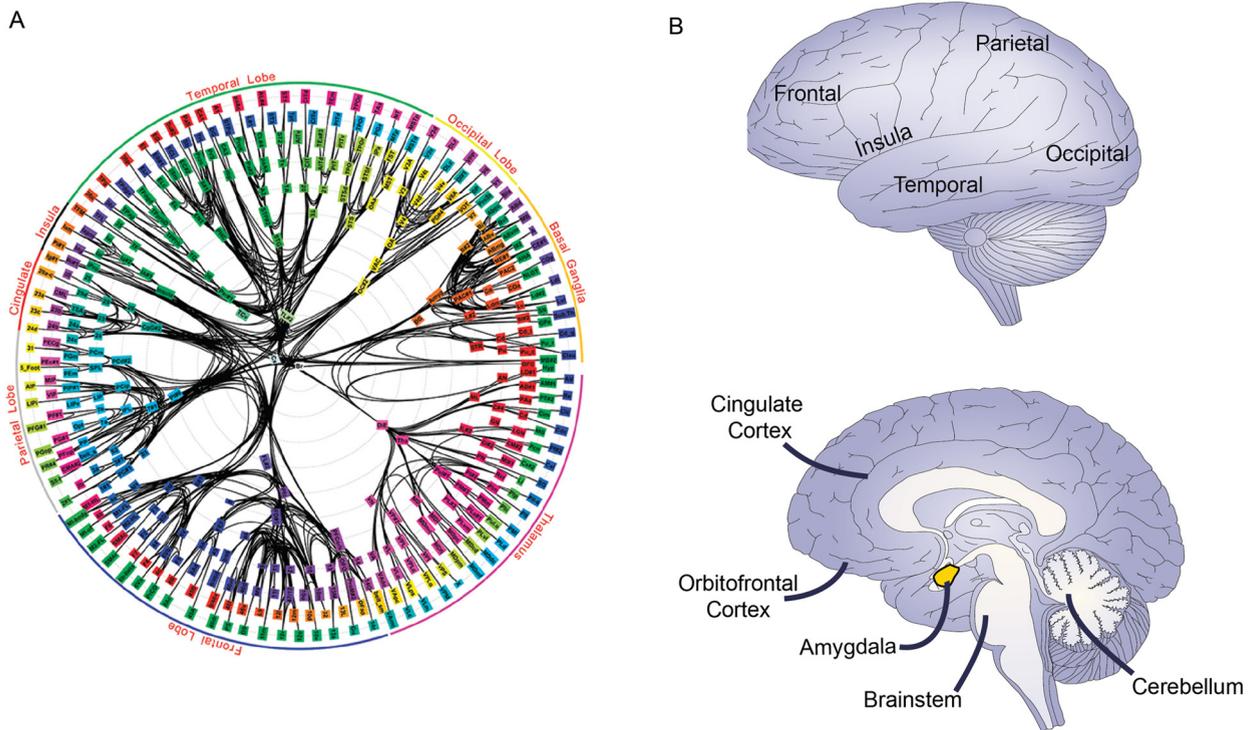


Fig. 2. Large-scale anatomical connectivity. (A) Network analysis of anatomical datasets reveals massive interconnectivity among all sectors of the brain. (B) Illustration of the major lobes of cortex (top) and some of the regions mentioned in the text (bottom). Part (A) reproduced with permission [22].

coherence” among their activity, for example, as measured by correlating their signals [25–27]. Hence, functional connectivity is essentially a description of the *joint state* of two brain elements, regardless of their anatomical connectivity status (linked by a strong or weak anatomical connection, or even not directly linked at all).

Thus, in general, what is important is not simply a region’s anatomical location, but its position in a space of functional relationships to other regions [27–29]. From the perspective of a brain region, at a given time, a region affiliates with a set of other regions, thereby defining a momentary functional circuit. And relatively spared coordinated activity can emerge in brains with substantially altered structural connectivity [30,31]. In all, anatomical architectural features support the efficient communication of information even when strong direct structural connections are absent [32–34], and allow functional interactions that vary as a function of context.

It should be pointed out that “functional connectivity” does not distinguish certain scenarios. Consider, for example, two neuronal populations, A and B.<sup>1</sup> Even if A and B are temporally coherent in their activation but only one triggers a causally related consequence in C, the functional connectivity between A and B does not include that information (and some could even question the validity of stating that they are functionally connected). Whereas this is potentially a problem, establishing causation in complex systems is far from straightforward. In fact, the type of Newtonian “billiard ball” causation model at times applied to reasoning about the brain (in this case, A “causes” B much like a billiard ball hitting another) may be in need of a revision [35–37]. For example, “probabilistic causation” [38] or “dynamical systems causation” offer promising avenues. In the latter, the goal is to investigate causation in (possibly non-linear) dynamical systems, and two variables are “causally linked” if they participate in the same dynamical system [39].

### 2.3. Overlapping and dynamic brain networks

Networks of brain regions collectively support behaviors. Thus, the *network itself is the unit*, not the brain region. Indeed, research characterizing both anatomical (e.g., [22]) and functional (e.g., [40]) networks has grown rapidly

<sup>1</sup> I thank a reviewer for bringing up this point.

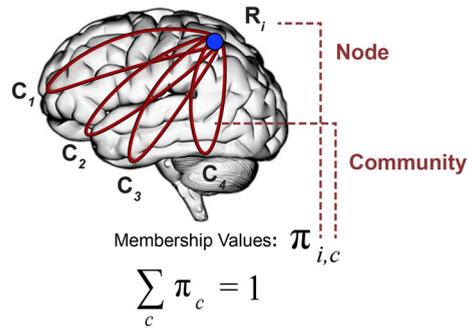


Fig. 3. Overlapping networks. In standard clustering approaches, communities are comprised on disjoint sets of nodes. In the framework of overlapping networks (or mixed membership), each brain region affiliates with multiple communities with varying strengths called “membership values.” For each brain region, these are probability-like values between 0 and 1 that sum to 1 (summed across communities). Here, the membership values of region  $R_i$  are based on strengths to each of communities  $C_1$  to  $C_4$ . Reproduced with permission [55].

in the last decade, and a host of specific networks have been proposed, such as the “salience” (important during the processing of salient events) and “executive” (important during cognitive processes, including attention and working memory) networks [41].

At the same time, analysis of anatomical/functional data has repeatedly revealed regions with widespread connectivity (see section above). This implies that describing brain network in terms of disjoint sets of regions (where each brain region belongs to a single network or cluster) only partly captures the complexity of the organization. It is important, therefore, to elucidate the *overlapping* structure of networks [42]. In several scientific disciplines, there is increasing realization that actual networks are best understood in terms of overlapping subsystems [43]. In neuroscience, disjoint networks (that is, non-overlapping) have been favored, although recent investigations have started to study potential ways to characterize overlapping organization [44,45].

In a recent study, we investigated the potential overlapping structure of functional brain connectivity in a data-driven fashion [46]. Instead of assuming that regions belong to a single network, we allowed them to belong to *all* communities (Fig. 3). Thus, when community organization was determined, a node’s *participation* to every community was estimated; participation strengths were called *membership values*. Each node was assigned a probability-like membership value to each of the existing communities, resulting in a membership vector with entries between 0 and 1 (and summing to 1); entries close to 1 indicated membership to essentially one community, that is, the non-overlapping case. The approach to detecting overlapping communities we adopted was originally based on the mixed-membership stochastic blocks model [47] within the context of stochastic variational inference [48]. The specific algorithm we employed was developed by Gopalan and Blei [49], and can be viewed as a generative Bayesian algorithm favoring assortativity (nodes with similar memberships are more likely to cluster together).

When applied to resting-state fMRI data, the algorithm captured community organization matching many of the general features of standard disjoint networks. However, the analysis indicated that overlapping communities are not fully captured by a single disjoint community, and that they contain information that is not well described by disjoint communities. To gain a richer understanding of the overlapping networks obtained, for each community, we considered the entire distribution of membership values, that is, we considered the membership values of all brain regions simultaneously. A community containing nodes that affiliated mostly with that community would exhibit a membership distribution with values close to 1. Instead, we found that all mixed communities had many regions with non-trivial membership values far from the peak of 1, which indicated that the overlapping communities were not well described by “exclusive nodes” (nodes that participated in only one community) but by “multi-partner nodes” (nodes that participated in multiple communities).

Furthermore, brain networks are not static but evolve temporally [50,51]. Although anatomical pathways change across the life span, faster dynamics are an essential ingredient of functional connections between regions, which vary as a function of context, and are altered by cognitive, emotional, and motivational variables, among others. Therefore, network organization is a *dynamic* property. Indeed, the growth of methods to characterize time-varying functional connectivity has begun to yield novel characterizations of how network organization evolves [52–54]. In sum, networks are not static and fixed collections of brain regions; they are dynamic coalitions of regions that form and dissolve to meet specific computational needs. This poses several challenges, as the very notion of a network

as a coherent unit is undermined. For instance, at what point does a coalition of regions become something *other* than, say, the salience network? An approach to modeling networks in general that partly mitigates this problem is to conceptualize networks as inherently overlapping. Instead of treating nodes as belonging to a single network, in overlapping networks, each node is a member of every network with a specific probability like membership value, which can then fluctuate across time [55].

Dynamic aspects of network function naturally will benefit from understanding interactions in terms of non-linear dynamical systems. Such systems undergo qualitative changes in behavior (so-called bifurcations) as a function of their input and/or parameter changes. For example, Hansen and colleagues [56], show via simulations of a whole-brain mean-field computational model that functional connectivity exhibits rapid transitions between a few discrete states. In the so-called “resting-state” (when participants are not explicitly engaged in a task), such state transitions are spontaneous and often very sharp.

#### 2.4. *Anatomical connectivity to and from the body*

Emotion is closely linked to bodily states, and the very history of emotion research is defined by the discussion and controversy surrounding the question of the necessity of the body for emotion [5]. The exact role of the body in emotion notwithstanding, brain circuits that involve signals to and from the body are important for associated processes.

Brain connectivity affecting the body originates in both cortex and subcortex. In cortex, a notable region is the cingulate cortex [57,58], whose descending projections to autonomic regulatory structures have been amply described (autonomic functions include respiration, heart rate, and pupil response) [59]. This connectivity is consistent with effects of cingulate electrical stimulation on virtually all autonomic processes, as well as many endocrine mechanisms. Conversely, visceral pathways carrying signals from the body to the brain reach multiple subcortical and cortical regions. In particular, cortical regions with notable body-related signals include the orbitofrontal cortex, the cingulate gyrus, and the insula [60]. For example, the physiological condition of the entire body is conveyed to the posterior insular cortex [60], which can be considered interceptive cortex, much like parts of parietal cortex are somatosensory cortex, for example. In conclusion, both cortex and subcortex are part of extensive connectional systems that link the body to the brain. In this manner, the entire spectrum of brain signals can affect the body, and vice versa.

#### 2.5. *Merging cognition and emotion in the brain*

In summary, anatomical and functional studies reveal that brain regions are massively interconnected. In particular, the brain basis of emotion involves large-scale cortical-subcortical networks that are distributed and sensitive to bodily signals [15]. The high degree of signal distribution and integration provides a nexus for the intermixing of information related to perception, cognition, emotion, motivation, and action. Importantly, the functional architecture consists of multiple overlapping networks that are highly dynamic and context sensitive. Thus, how a given brain region affiliates with a specific network shifts as a function of task demands and brain state.

Whereas some regions are more highly interconnected than others, overall information can travel from one part of the brain to another in just a few steps. This implies that the brain’s architecture does not leave much room for isolation of signals; it is promiscuous, with information of a given type being mixed with others. Furthermore, adopting a large-scale network perspective to brain organization helps clarify why some brain structures, such as the amygdala, are thought to be important for emotion: they are important hubs of large-scale connectivity systems. It also illuminates why the impact of emotion is so wide ranging – for one, it is not possible to impact emotion without affecting cognition.

### 3. **Brain evolution, cognition, and emotion**

From an evolutionary perspective, it is sometimes stated that cognition can be considered “the tip of the iceberg,” often associated with “newer” parts of the cortex, while the rest of the iceberg involves emotional processes that rely on “old” subcortical brain regions. Furthermore, it is presumed that the old emotional system functions in a largely autonomous manner (for discussion, see [61]).

Some brain system can indeed be considered old. For example, anatomical and behavioral data suggest that ray-finned fishes and land vertebrates probably share an ancestor (>420 million years ago) that already possessed some form of an amygdala [62] – a subcortical region that is believed to be particularly important for emotional processing [63]. And the amygdala of the ancestral amniote (which includes reptiles, birds, and mammals; around 340 million years ago) is believed to have had parts that are key subregions of the human amygdala, for example [62].

In the beginning of the twentieth century, neuroanatomists proposed the idea that brain evolution takes place in a fairly progressive fashion, from fishes to amphibians, to reptiles, to birds, and then mammals (culminating with the human brain) – the idea of progression goes back to Aristotle. In this scheme, the brain evolved via the *sequential addition* of parts. Thus, in the brain of a cat (a mammal) one could discern older parts that remained largely unchanged underneath the cortex [64]. Furthermore, it was proposed that structures at the top control the ones at the bottom [61]. Unfortunately, this erroneous viewpoint [65] became dominant in neuroscience for many decades. It reached its most dramatic example in Paul MacLean’s proposal of the “triune brain,” which consisted of a reptilian complex (at the base of the brain, including brainstem), to which an “old brain” was added, and later a “new cortex” included in mammals [66]. In the context of the present review, it is noteworthy that the organization shown in Fig. 1A very much resembles this notion.

It is now known that such framework does not reflect current knowledge of comparative neuroanatomy and does not adequately describe brain evolution (not even in general terms) [65,67]. But what is the *basic plan* of the vertebrate brain? The chassis of the brain of fishes, amphibians, and all other vertebrates contains both subcortex (with multiple parts) and simple forms of cortex (called pallium). With evolution, brain changes occurred across the brain. But importantly, new circuits originated in a manner that affected both cortex and subcortex.

An example of this type of *mutual embedding* is provided by the amygdala. In mammals, this structure is composed of a dozen or more subregions. In one study it was found that some amygdala subregions are considerably more “developed” in monkeys than in rats (based on morphological characteristics, such as cell counts and the volume of subregions) [68]. One possibility is that the observed differences between rats and monkeys are linked to the degree of connectivity of amygdala subregions to other brain structures (for related discussion, see [67]). In other words, the changes in the amygdala may have paralleled the development of the cortical areas with which these nuclei are interconnected in primates.

Taken together, brain evolution likely proceeded via the mutual embedding of new and old structures, not by the modular addition of parts. Furthermore, separation of “old” and “new” parts does not take into account the extensive anatomical connectivity of brain circuits, as previously discussed, nor the overlapping and dynamic nature of brain networks.

#### 4. Cognitive-emotional architectures

With the discussion of the previous two sections in mind, we can return to the central question addressed here: What kind of architecture should an intelligent, autonomous robot have? Although a broad range of cognitive architectures have been proposed [69–73]; for review, see [74], a general principle is that they are modular (Fig. 4A). For example, they may have low-level modules for action and perception, mid-level modules for task coordination and mid-level perception, and higher levels for task planning and visual cognition, say. Frequently, the top level is connected to a knowledge database. To coordinate the modules, a central controller or supervisor is also posited, which utilizes temporary buffers as part of an active workspace. By and large, the architecture is hierarchical, although parallel components and organization are recognized as important, too [75].

In the past two decades, a steady stream of researchers have advocated for the inclusion of emotion-related components to the general information processing architecture (or “cognitive” architecture, as sometimes it is still referred) [10,70,76–82]; see contributions in [83]. One type of argument is that emotion components are needed to instill urgency to action and decisions [7]. Others have advocated emotion components to aid understand emotion in humans, or to generate human-like expressions (for empathy, say) [10]. In broad terms, a first-pass effort at including affect entails adding an emotion module so that it can influence other components of the architecture (Fig. 4B). Another aspect emphasized in the literature is that emotion can be used to generate an internal state that influences actions (e.g., [82,84], Fig. 3B, right).

The framework advanced here goes beyond these approaches and proposes that emotion and motivation need to be integrated with all aspects of the architecture. In particular, emotion-related mechanisms influence processing beyond

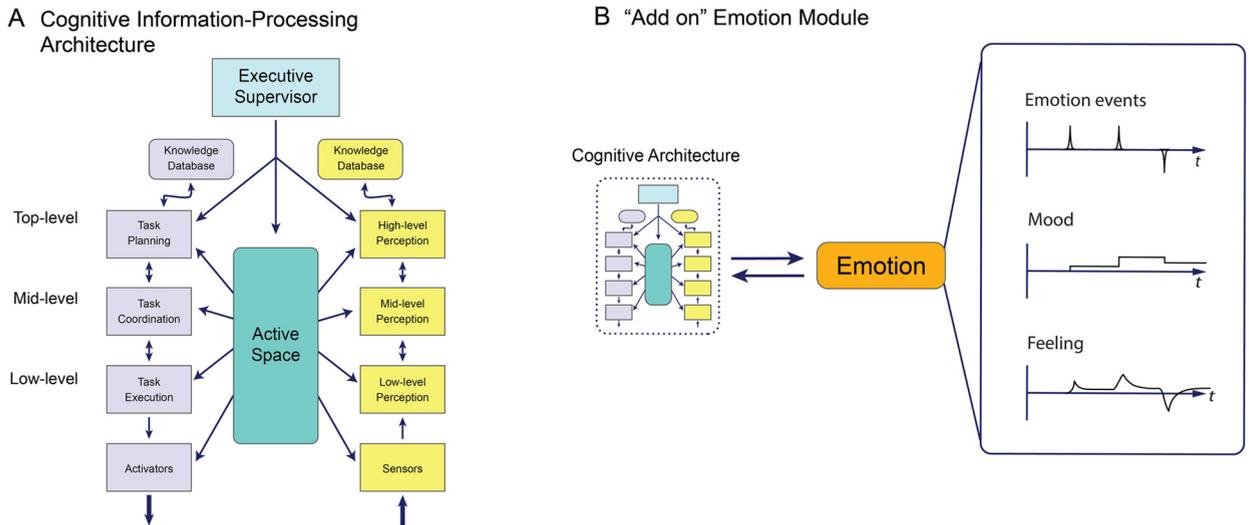


Fig. 4. Cognition and emotion in information-processing architectures. (A) Central components of existing cognitive architectures, with a modular and largely hierarchical organization. (B) Emotion can be considered as an additional module of the architecture that interfaces with some of the cognitive components. In particular, it can be used to determine “mood” and “feeling”. The region with dashed outline corresponds to the architecture in part A.

the modulatory aspects of “moods” linked to internal states (hunger, sex-drive, etc.). Of note, in the literature, emotion is often linked to evaluative aspects of the organism-environment relationship and with affective states; motivation is associated with how the organism acts in a given situation (approach vs. withdrawal). Given the integrationist framework described here, I do not make this distinction [11].

As described previously, emotion is not neatly encapsulated in a specific part of the brain – one cannot point to emotion in the brain. In addition, cognition was not a late addition to the mammalian, or human, brain [65]. So, how can we operationalize it in an information-processing architecture? Emotion can be thought of as a set of *valuating mechanisms* that help organize behavior [85], for instance by helping take into account both the costs and benefits linked to percepts and actions [86]. At a general level, it can be viewed as a biasing mechanism, much like the “cognitive” function of attention. However, such conceptualization is overly simplistic because emotion does not amount to just providing an extra boost to a specific sensory input, potential plan, or action. When the brain is conceptualized as a complex system of highly interacting networks of regions, we see that emotion is interlocked with perception, cognition, motivation, and action. Whereas we can refer to certain behaviors as “emotional” or “cognitive,” this is only a language short-cut. Thus, the idea of a biasing mechanism is too limited. From a design perspective, all components of the architecture should be influenced by emotional and motivational variables (and vice versa). Thus, the architecture must be strongly *non-modular*.

To illustrate these ideas, I will describe architectural features inspired by knowledge of brain and behavior. The focus will be on visual perception and executive control, but related schemes can be used in the case of action, too, as adaptive behavior requires monitoring and evaluating the outcome of actions. I will build upon a conceptual framework called the *dual competition model* which describes how both emotional and motivational signals are integrated with perception and cognition so as to effectively incorporate *value* into the unfolding of behavior [87]. The general framework is displayed in Fig. 5A. Because the use of boxes invites a modular interpretation, the architecture is better represented in terms of overlapping networks of processes or mechanisms (Fig. 5B).

#### 4.1. Perceptual competition

Objects in the environment compete for limited perceptual processing capacity and control of behavior [88–90]. Because processing capacity for vision is limited [91], selective attention to one part of the visual field comes at the cost of neglecting other parts. Thus, a popular notion is that there is *competition* for resources in visual cortex [88,92].

In operationalizing competition, we can use the concept of a *priority map*, which contains representations of spatial locations that are behaviorally important [93,94]. Traditionally, “bottom-up” factors (such as stimulus salience) and

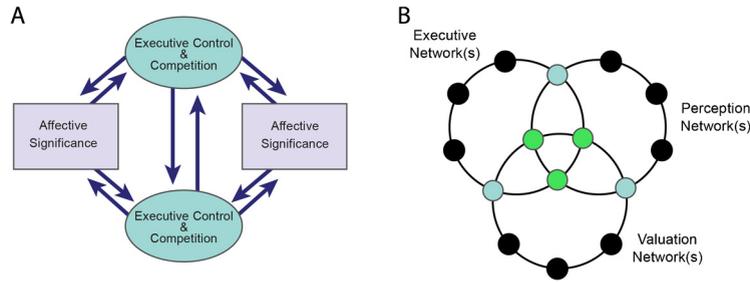


Fig. 5. Dual competition model. (A) Dual competition framework where emotion and motivation are integrated with perception and cognition. (B) Schematic representation of some of the network interactions linked to the model. In particular, the valuation network is important in the determination of affective and motivational significance. Reproduced with permission [35].

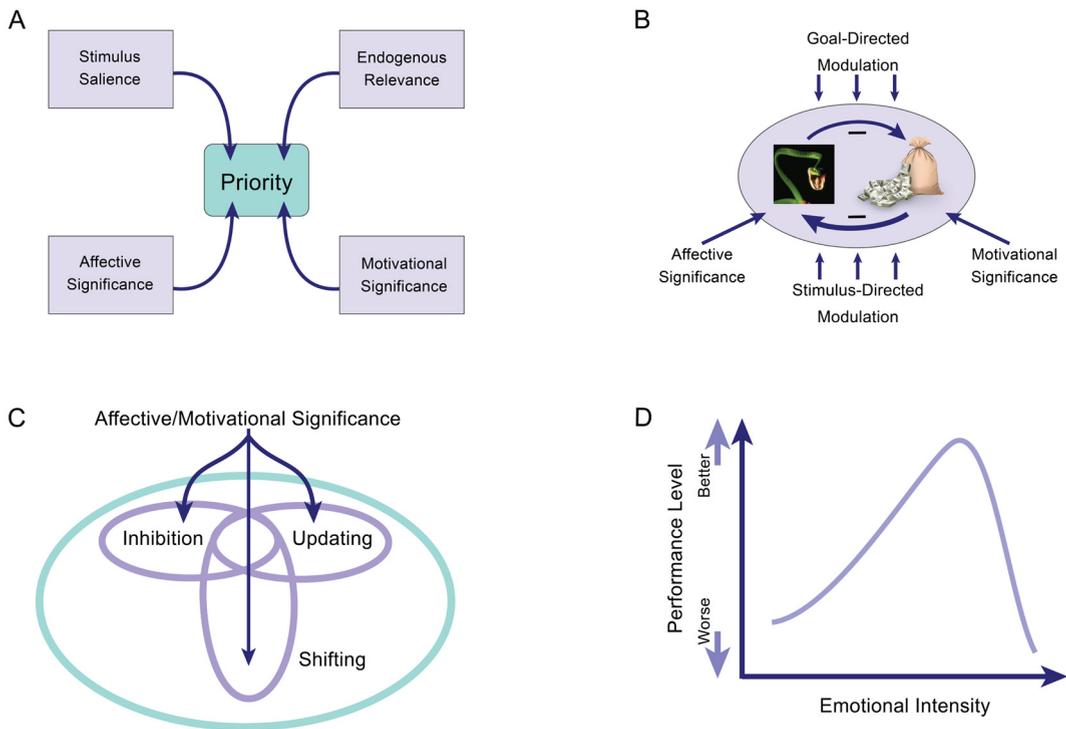


Fig. 6. Cognitive-emotional integration. (A) A priority map contains spatial locations of behavioral importance which are determined by combining multiple factors, including emotional/motivational ones. (B) Perceptual competition is influenced by multiple factors which determine the stimulus representation that should be prioritized (“wins the competition”). (C) Executive control can be conceptualized as a series of processes that are partly interdependent (here, inhibition, updating, and shifting). Affective and motivational inputs engage the same processes at the core of cognition. (D) Behavioral performance is a function of the intensity of the emotional stimulus or context.

“top-down” factors (such as goal relevance) were emphasized as the major inputs to determine priority. To these, we must add affective significance (for instance, based on the prior co-occurrence of a stimulus with aversive events) and motivational significance (for instance, based on the prior co-occurrence of a stimulus with rewards) (Fig. 6A) (see also [95]). In terms of implementation, the architecture requires mechanisms to embed affective and motivational significance, as well as stimulus- and goal-related factors, into perception. This can be accomplished via an explicit priority map (Fig. 6A) or by having multiple influences converge on representations of sensory stimuli (Fig. 6B). In the brain, perceptual competition occurs in part via interactions between regions important for attention (in frontal and parietal cortices, for example) and those that are particularly tuned to the evaluation of affective and motivational significance.

But the above should not be viewed as just an extension of the priority map with just additional variables. In particular, if one considers a predictive coding framework [96,97], internal models of the current input will incorporate features along multiple dimensions, including sensory ones but also affective and motivational significance. Thus, establishing a predictive code involves building internal models that incorporate all of these dimensions simultaneously.

#### 4.2. Executive control

At the heart of cognition, executive control refers to operations involved in maintaining and updating information, monitoring conflict and/or errors, resisting distracting information, inhibiting prepotent responses, and shifting mental sets. A useful way to conceptualize executive control is in terms of a set of processes, including *inhibition*, *updating*, and *shifting* [98]. These processes are not independent, however; when a given function is necessitated, resources devoted to one component will not be available to other operations and interference will ensue, possibly compromising performance [11].

Dealing with an emotional stimulus or situation requires the types of behavioral adjustments that characterize executive function. For example, *updating* might be needed to refresh the contents of working memory, *shifting* might be recruited to switch the current task set, and *inhibition* might be invoked to cancel previously planned actions. In this manner, specific resources are coordinated in the service of emotional processing (Fig. 6C) and, if temporarily unavailable to additional task requirements, would compromise performance – the stronger the emotional manipulation, the stronger the interference (see below). A laboratory example might help illustrate these ideas. Suppose a participant is performing an effortful task of some sort and a change of background color signals that she will receive a mild shock sometime in the next 30 seconds. The participant may *update* the contents of working memory to include the “shock possible” information. *Shifting* may occur in that the participant may shift between the execution of the cognitive task and “monitoring for shock” (and possibly preparing for the aversive event) every few seconds. Now, if another cue stimulus indicated that shock would be delivered in the next second, the participant might temporarily *inhibit* response to the task to prepare for the shock. Altogether, dealing with the emotional situation necessitates the same types of executive functions considered to be the hallmark of cognition.

Importantly, the intensity of the emotional stimulus (or context) determines if it will improve or hinder behavioral performance. Moderate intensity, in particular, allows resources to be devoted to the situation at hand, improving performance. However, if the intensity is high enough, processes may be temporarily unavailable to handle the aspects of the situation adequately. This inverted-U type of relationship is common in many behavioral settings (Fig. 6D).

### 5. The Dolores test for robotics

Let us consider the humanoid robots of Westworld again. In the scene previously referred to, the central humanoid character, Dolores, was asked to consider and describe her situation by using cognition only. But could a human “lose all emotion” and describe complex events without a trace of affect like Dolores did? The present framework indicates that the answer to this question is “no.” Unlike Dolores, humans cannot lose all emotion and simply proceed cognitively. The brain is organized such that the mind doesn’t behave in this manner.

Dolores is not human; she is a robot built for the amusement of humans. Could a sophisticated artificial intelligence be built with separate cognitive and emotional modules? I contend that such a humanoid would be a far cry from the complex hosts of the series, whose behaviors are only possible when emotion and cognition are intertwined. Dolores, for one, would be unable to reflect about her past in a meaningful way just cognitively; emotion is part and parcel of her understanding of the world. And this is not just because she lives in a violent, ruthless world, but because making sense of reality necessitates the integration of emotion and cognition. Could this prediction be tested?

In 1950, Alan Turing [99] proposed a test to establish the presence of thought in a computer by assessing a machine’s ability to exhibit intelligent behavior indistinguishable from that of a human. In the Turing test, an interrogator is supposed to be in a room separated from a machine and another person. The object of the game is for the interrogator to determine which of the other two is the person and which is the machine. The game is such that the human interrogator is physically separated from the other two. I propose, instead, the *Dolores test*: one cannot build a humanoid that behaves in a manner as sophisticated as Dolores does that is built with separate emotional and cognitive modules.



Fig. 7. The Dolores test. The test states that, by observing behaviors as humanoids interact with both other humanoids and humans, humans and humanoids would be confused with each other unless the latter were built with a non-modular cognitive-emotional architecture. In this depiction of a scene from *Westworld*, Dolores interacts with two humans (one tied to the chair, another immediately to her right), and is watched by two other humanoid robots. Reproduced with permission [104].

A scenario somewhat similar to that envisioned by the Dolores test is featured in the 1982 movie *Blade Runner*. Humans paid to track down rogue humanoids use the “Voight-Kampff device,” a polygraph-like machine used to help determine whether or not an individual is a robot. The machine measures bodily signals linked to respiration, blush response, heart rate, and eye movement in response to questions dealing with emotion and empathy. However, the Dolores test proposed here is fundamentally distinct from the one in *Blade Runner*. It is not based on the idea that bodily responses to emotionally charged situations would deviate in a detectable way in robots. Instead, it is based on the hypothetical scenario that a robot could be built and that it would display “sophisticated behaviors” of the kind produced by the robots in *Westworld*. The idea of the test is that, by observing behaviors as the humanoids interact with both other humanoids and humans, humans and humanoids would *not* be confused with each other if the latter were built according to a modular cognitive-emotional architecture (Fig. 7) – that is, complex behaviors require a non-modular architecture.

Of course, to make the test sufficiently precise, it would be necessary to operationalize “sophisticated behaviors.” In particular, Turing believed that the question of whether or not machines can think was “too meaningless” to deserve discussion, and therefore sought to formulate the “Imitation Game,” with the interrogator separated from the machine and the other person, as described above.

It has been argued that passing the Turing test does not necessarily show “intelligence.” For example, the Chinese Room argument by Seale [100] points out that complex rule-based transformations may yield apparent understanding when only mechanical rules are being applied [101]. Could the behavior required to pass the Dolores test be generated by a set of complex rules? Although a more extended discussion should be taken up elsewhere in the future, the Chinese Room argument itself has fostered intense debate [101]. For instance, both the “systems reply” and the “virtual machine reply” argue that understanding is not present at the level of rule implementation per se (assuming it would be possible), but at the “systems level” including the rule-implementation mechanisms and (possibly all) other components at play – thus creating a “virtual” agent that does understand.

## 6. Conclusions

How should intelligent, autonomous robots be built? Information-processing architectures have mostly focused on cognitive components. In the past two decades, emotion components have been increasingly described and advocated as a way to improve human-robot communication, including the understanding of human facial expressions and the generation of expressions to communicate internal states. The survival argument has also been emphasized; emotion can then be used to help prioritize processes and actions that foster continued existence.

The central argument of the present framework is not that emotion is needed – the answer is a definitive “yes” – but that emotion and motivation need to be integrated with *all* information-processing components. This implies that

cognitive-emotional integration is a principle of the architecture. Emotion is not an “add on” that endows a robot with “feelings” (for instance, reporting or expressing its internal state). It allows the *significance* of percepts, plans, and actions to be an integral part of all its computations.

Another implication of the present framework is that the traditional modular engineering approach is inadequate for autonomous robots. What we know about behavior and the brain of mammals (indeed of vertebrates more generally) suggests that an approach heavily influenced by biology is needed. In such framework, centralized control is abdicated, and distributed control is adopted. Cognitive and brain science have struggled with the concept of a homunculus, where “too much” intelligence is relegated to one of its components (but how does that component itself acquire its sophistication?) [102,103]. Information-processing architectures for artificial agents need, likewise, to banish such notions of centralized control (a homunculus).

To return to the scenarios inspired by the *Westworld* series. Humans cannot “lose all emotion” and continue with their daily lives. I content that sophisticated, autonomous robots will prove the same. They will need cognition and emotion as integrated “components” of their information-processing architecture.

## Acknowledgements

I am grateful to the National Institute of Mental Health for continued research support (R01 MH071589 and R01 MH112517) and Heiko Neumann for feedback on the manuscript. A reviewer of the manuscript also provided excellent feedback. I thank Christian Meyer and Anastasiia Khibovska for assistance with figures and references.

## References

- [1] Rolls ET. *Emotion explained*. Oxford: Oxford University Press; 2005.
- [2] Sander D, Grandjean D, Scherer KR. A systems approach to appraisal mechanisms in emotion. *Neural Netw* 2005;18:317–52.
- [3] Ekman P. An argument for basic emotions. *Cogn Emot* 1992;6:169–200.
- [4] Moll J, Zahn R, de Oliveira-Souza R, Krueger F, Grafman J. The neural basis of human moral cognition. *Nat Rev Neurosci* 2005;6:799–809.
- [5] Damasio AR. *Descartes’ error: emotion, reason, and the human brain*. New York: G.P. Putnam; 1994.
- [6] Arkin RC, Vachtsevanos G. Techniques for robot survivability. In: *Proc. 3rd international symposium on robotics and manufacturing*; 1990. p. 383–8.
- [7] Kelley AE. Neurochemical networks encoding emotion and motivation: an evolutionary perspective. In: Fellous J-M, Arbib MA, editors. *Who needs emotions? The brain meets the robot*. New York: Oxford University Press; 2005.
- [8] Ziemke T, Lowe R. On the role of emotion in embodied cognitive architectures: from organisms to robots. *Cogn Comput* 2009;1:104–17.
- [9] Maes P. *Designing autonomous agents: theory and practice from biology to engineering and back*. Cambridge, MA: MIT Press; 1990.
- [10] Breazeal C. *Designing sociable robots*. Cambridge, MA: MIT Press; 2002.
- [11] Pessoa L. *The cognitive-emotional brain: from interactions to integration*. Cambridge: MIT Press; 2013.
- [12] Pessoa L. The emotional brain. In: Conn PM, editor. *Conn’s translational neuroscience*. Elsevier; 2017. p. 635–56.
- [13] Panksepp J. *Affective neuroscience: the foundations of human and animal emotions*. Oxford University Press; 2004.
- [14] Lindquist KA, Barrett LF. A functional architecture of the human brain: emerging insights from the science of emotion. *Trends Cogn Sci* 2012;16:533–40.
- [15] Pessoa L. A network model of the emotional brain. *Trends Cogn Sci* 2017.
- [16] Sporns O. *Networks of the brain*. MIT Press; 2010.
- [17] Oh SW, Harris JA, Ng L, Winslow B, Cain N, Mihalas S, et al. A mesoscale connectome of the mouse brain. *Nature* 2014;508:207–14.
- [18] Bota M, Sporns O, Swanson LW. Architecture of the cerebral cortical association connectome underlying cognition. *Proc Natl Acad Sci USA* 2015;112:E2093–101.
- [19] Swanson LW, Lichtman JW. From Cajal to connectome and beyond. *Annu Rev Neurosci* 2016;39:197–216.
- [20] Hilgetag CC, Medalla M, Beul SF, Barbas H. The primate connectome in context: principles of connections of the cortical visual system. *NeuroImage* 2016;134:685–702.
- [21] Markov NT, Ercsey-Ravasz M, Van Essen DC, Knoblauch K, Toroczkai Z, Kennedy H. Cortical high-density counterstream architectures. *Science* 2013;342:1238406.
- [22] Modha DS, Singh R. Network architecture of the long-distance pathways in the macaque brain. *Proc Natl Acad Sci USA* 2010;107:13485–90.
- [23] Markov NT, Vezoli J, Chameau P, Falchier A, Quilodran R, Huissoud C, et al. The anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. *J Comp Neurol* 2013.
- [24] Fornito A, Zalesky A, Bullmore E. *Fundamentals of brain network analysis*. Academic Press; 2016.
- [25] Gerstein GL, Perkel DH. Simultaneously recorded trains of action potentials: analysis and functional interpretation. *Science* 1969;164:828–30.
- [26] Aertsen A, Gerstein GL, Habib MK, Palm G. Dynamics of neuronal firing correlation: modulation of “effective connectivity”. *J Neurophysiol* 1989;61:900–17.

- [27] Aertsen A, Preissl H. Dynamics of activity and connectivity in physiological neuronal networks. In: Schuster HG, editor. *Non linear dynamics and neuronal networks*. New York: VCH; 1991. p. 281–302.
- [28] Friston KJ, Buechel C, Fink GR, Morris J, Rolls E, Dolan RJ. Psychophysiological and modulatory interactions in neuroimaging. *NeuroImage* 1997;6:218–29.
- [29] Pessoa L. Understanding brain networks and brain organization. *Phys Life Rev* 2014;11:400–35.
- [30] Tyszka JM, Kennedy DP, Adolphs R, Paul LK. Intact bilateral resting-state networks in the absence of the corpus callosum. *J Neurosci* 2011;31:15154–62.
- [31] O'Reilly JX, Crosson PL, Jbabdi S, Sallet J, Noonan MP, Mars RB, et al. Causal effect of disconnection lesions on interhemispheric functional connectivity in rhesus monkeys. *Proc Natl Acad Sci USA* 2013;110:13982–7.
- [32] Adachi Y, Osada T, Sporns O, Watanabe T, Matsui T, Miyamoto K, et al. Functional connectivity between anatomically unconnected areas is shaped by collective network-level effects in the macaque cortex. *Cereb Cortex* 2012;22:1586–92.
- [33] Averbeck BB, Seo M. The statistical neuroanatomy of frontal networks in the macaque. *PLoS Comput Biol* 2008;4.
- [34] Grayson DS, Bliss-Moreau E, Machado CJ, Bennet J, Shen K, Grant KA, et al. The rhesus monkey connectome predicts disrupted functional networks resulting from pharmacogenetic inactivation of the amygdala. *Neuron* 2016;91:453–66.
- [35] Pessoa L. Cognitive-motivational interactions: beyond boxes-and-arrows models of the mind-brain. *Mot Sci* 2017;3:287.
- [36] Pessoa L. Understanding emotion with brain networks. *Curr Opin Behav Sci* 2019;19:19–25.
- [37] Pessoa L. Dynamics of emotion and cognition. *Neural Netw* 2019. Special issue in honor of Stephen Grossberg.
- [38] Mannino M, Bressler SL. Foundational perspectives on causality in large-scale brain networks. *Phys Life Rev* 2015;15:107–23.
- [39] Sugihara G, May R. Detecting causality in complex ecosystems. *Science* 2012;338:496–500.
- [40] Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* 2011;106:1125–65.
- [41] Menon V, Uddin LQ. Saliency, switching, attention and control: a network model of insula function. *Brain Struct Funct* 2010;214:655–67.
- [42] Xie J, Kelley S, Szymanski BK. Overlapping community detection in networks: the state-of-the-art and comparative study. *ACM Comput Surv* 2013;45:43.
- [43] Palla G, Derenyi I, Farkas I, Vicsek T. Uncovering the overlapping community structure of complex networks in nature and society. *Nature* 2005;435:814–8.
- [44] Yeo BT, Krienen FM, Chee MW, Buckner RL. Estimates of segregation and overlap of functional connectivity networks in the human cerebral cortex. *NeuroImage* 2014;88:212–27.
- [45] Li X, Gan JQ, Wang H. Collective sparse symmetric non-negative matrix factorization for identifying overlapping communities in resting-state brain functional networks. *NeuroImage* 2018;166:259–75.
- [46] Najafi M, Mcmenamin BW, Pessoa L. Overlapping communities reveal rich structure in large-scale brain networks during rest and task conditions. *NeuroImage* 2016;135:92–106.
- [47] Airoldi EM, Blei DM, Fienberg SE, Xing EP. Mixed membership stochastic blockmodels. *J Mach Learn Res* 2008;9:1981–2014.
- [48] Hoffman MD, Blei DM, Wang C, Paisley J. Stochastic variational inference. *J Mach Learn Res* 2013;14:1303–47.
- [49] Gopalan P, Blei D. Efficient discovery of overlapping communities in massive networks. *Proc Natl Acad Sci* 2013;110:14534–9.
- [50] Varela FJ, Lachaux JP, Rodriguez E, Martinerie J. The brainweb: phase synchronization and large-scale integration. *Nat Rev Neurosci* 2001;2:229–39.
- [51] Buzsaki G. *Rhythms of the brain*. Oxford University Press; 2006.
- [52] Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST. Dynamic reconfiguration of human brain networks during learning. *Proc Natl Acad Sci USA* 2011;108:7641–6.
- [53] Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, et al. Dynamic functional connectivity: promise, issues, and interpretations. *NeuroImage* 2013;80:360–78.
- [54] Helfrich RF, Knight RT. Oscillatory dynamics of prefrontal cognitive control. *Trends Cogn Sci* 2016.
- [55] Pessoa L. Overlapping and dynamic brain networks of the emotional brain. In: Wu G, Laurienti PJ, editors. *Connectomics: methods, mathematical models and applications*; 2018.
- [56] Hansen EC, Battaglia D. Functional connectivity dynamics: modeling the switching behavior of the resting state. *NeuroImage* 2015;105:525–35.
- [57] Vogt BA. *Cingulate neurobiology and disease*. Oxford: Oxford University Press; 2008.
- [58] Shackman AJ, Salomons TV, Slagter HA, Fox AS, Winter JJ, Davidson RJ. The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nat Rev Neurosci* 2011;12:154–67.
- [59] Nieuwenhuis R, Voogd J, van Huijzen C. *The human central nervous system*. Fourth ed. Steinkopff; 2008.
- [60] Craig AD. How do you feel? Interoception: the sense of the physiological condition of the body. *Nat Rev Neurosci* 2002;3:655–66.
- [61] Parvizi J. Corticocentric myopia: old bias in new cognitive sciences. *Trends Cogn Sci* 2009;13:354–9.
- [62] Martinez-Garcia F, Novejarque A, Lanuza E. Evolution of the amygdala in vertebrates. In: Kaas JH, editor. *Evolution of nervous systems*, vol. 2. Oxford: Academic Press; 2007. p. 255–334.
- [63] Whalen PJ, Phelps EA. *The human amygdala*. New York, NY: Guilford Press; 2009.
- [64] Edinger L. The relations of comparative anatomy to comparative psychology. *J Comp Neurol Psychol* 1908;18:437–57.
- [65] Butler AB. Triune brain concept: a comparative evolutionary perspective. In: Squire LR, editor. *Encyclopedia of neuroscience*. Oxford: Academic Press; 2009. p. 1185–93.
- [66] MacLean PD. The triune brain, emotion, and scientific bias. In: Schmitt FO, editor. *The neurosciences second study program*. New York: Rockefeller University Press; 1970. p. 336–49.
- [67] Striedter GF. *Principles of brain evolution*. Sunderland, MA: Sinauer Associates; 2005.

- [68] Chareyron LJ, Banta Lavenex P, Amaral DG, Lavenex P. Stereological analysis of the rat and monkey amygdala. *J Comp Neurol* 2011;519:3218–39.
- [69] Burghart C, Mikut R, Stiefelhagen R, Asfour T, Holzapfel H, Steinhaus P, et al. A cognitive architecture for a humanoid robot: a first approach. In: 5th IEEE-RAS international conference on humanoid robots; 2005.
- [70] Shanahan M. A cognitive architecture that combines internal simulation with a global workspace. *Conscious Cogn* 2006;15:433–49.
- [71] Metta G, Sandini G, Vernon D, Natale L, Nori F. The iCub humanoid robot: an open platform for research in embodied cognition. In: The 8th workshop on performance metrics for intelligent systems. Gaithersburg, MD: ACM; 2008. p. 50–6.
- [72] Trafton G, Hiatt L, Harrison A, Tamborello F, Khemlani S, Schultz A. Act-r/e: an embodied cognitive architecture for human-robot interaction. *J Hum-Robot Interact* 2013;2:30–55.
- [73] Laird JE, Kinkade KR, Mohan S, Xu JZ. Cognitive robotics using the soar cognitive architecture. In: Workshops at the twenty-sixth AAAI conference on artificial intelligence; 2012.
- [74] Vernon D, Metta G, Sandini G. A survey of artificial cognitive systems: implications for the autonomous development of mental capabilities in computational agents. *IEEE Trans Evol Comput* 2007;11:151–80.
- [75] Brooks RA. *Cambrian intelligence: the early history of the new AI*. Cambridge, MA: MIT Press; 1999.
- [76] Kubota N, Nojima Y, Baba N, Kojima F, Fukuda T. Evolving pet robot with emotional model. In: Congress on evolutionary computation. IEEE; 2000. p. 1231–7.
- [77] Breazeal C. Emotion and sociable humanoid robots. *Int J Hum-Comput Stud* 2003;59:119–55.
- [78] Arbib MA, Fellous JM. Emotions: from brain to robot. *Trends Cogn Sci* 2004;8:554–61.
- [79] Sandini G, Metta G, Vernon D. The iCub cognitive humanoid robot: an open-system research platform for enactive cognition. In: Lungarella M, Iida F, Bongard J, Pfeifer R, editors. 50 years of artificial intelligence. Berlin, Heidelberg: Springer; 2007. p. 358–69.
- [80] Saint-Aimé S, Le Pévédic B, Duhaut D. iGrace – emotional computational model for EmI companion robot. In: Kulyukin VA, editor. Advances in human-robot interaction: INTECH; 2009.
- [81] Álvarez M, Galán R, Matía F, Rodríguez-Losada D, Jiménez A. An emotional model for a guide robot. *IEEE Trans Syst Man Cybern, Part A, Syst Hum* 2010;40:982–92.
- [82] Arkin RC. *Moving up the food chain: motivation and emotion in behavior-based robots*. 2003.
- [83] Fellous J-M, Arbib MA. In: *Who needs emotions? The brain meets the robot*. Oxford University Press; 2005.
- [84] Kubota Y, Sato W, Murai T, Toichi M, Ikeda A, Sengoku A. Emotional cognition without awareness after unilateral temporal lobectomy in humans. *J Neurosci* 2000;20:RC97.
- [85] Adolphs R. How should neuroscience study emotions? By distinguishing emotion states, concepts, and experiences. *Soc Cogn Affect Neurosci* 2016;nsw153.
- [86] Rushworth MF, Walton ME, Kennerley SW, Bannerman DM. Action sets and decisions in the medial frontal cortex. *Trends Cogn Sci* 2004;8:410–7.
- [87] Pessoa L. How do emotion and motivation direct executive function? *Trends Cogn Sci* 2009;13:160–6.
- [88] Desimone R, Duncan J. Neural mechanisms of selective attention. *Annu Rev Neurosci* 1995;18:193–222.
- [89] Pashler H. *The psychology of attention*. Cambridge, MA: MIT Press; 1998.
- [90] Buschman TJ, Siegel M, Roy JE, Miller EK. Neural substrates of cognitive capacity limitations. *Proc Natl Acad Sci USA* 2011;108:11252–5.
- [91] Tsotsos JK. Analyzing vision at the complexity level. *Behav Brain Sci* 1990;13:423–69.
- [92] Grossberg S. How does a brain build a cognitive code? *Psychol Rev* 1980;87:1–51.
- [93] Itti L, Koch C, Niebur E. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans Pattern Anal Mach Intell* 1998;20:1254–9.
- [94] Itti L, Koch C. Computational modelling of visual attention. *Nat Rev Neurosci* 2001;2:194–203.
- [95] Awh E, Belopolsky AV, Theeuwes J. Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends Cogn Sci* 2012;16:437–43.
- [96] Grossberg S. *How does a brain build a cognitive code? Studies of mind and brain*. Springer; 1980. p. 1–52.
- [97] Friston K, Kiebel S. Predictive coding under the free-energy principle. *Philos Trans R Soc Lond B, Biol Sci* 2009;364:1211–21.
- [98] Miyake A, Friedman NP, Emerson MJ, Witzki AH, Howerter A, Wager TD. The unity and diversity of executive functions and their contributions to complex “Frontal Lobe” tasks: a latent variable analysis. *Cogn Psychol* 2000;41:49–100.
- [99] Turing AM. Computing machinery and intelligence. *Mind* 1950;59:433–60.
- [100] Searle JR. Minds, brains, and programs. *Behav Brain Sci* 1980;3:417–24.
- [101] Cole D. The Chinese room argument. In: Zalta EN, editor. *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University; 2019.
- [102] Monsell S, Driver J. Banishing the control homunculus. In: Monsell S, Driver J, editors. *Attention and performance XVIII: control of cognitive processes*. Cambridge, MA: MIT Press; 2000. p. 3–32.
- [103] Verbruggen F, McLaren IP, Chambers CD. Banishing the control homunculi in studies of action control and behavior change. *Perspect Psychol Sci* 2014;9:497–524.
- [104] Pessoa L. Do intelligent robots need emotion? *Trends Cogn Sci* 2017;21:817–9.