A Cognitive Architecture Based on Cognitive/Neurological Dual-System Theories

Othalia Larue, Pierre Poirier, and Roger Nkambou

GDAC Research Laboratory, Université du Québec à Montréal
Montréal, QC, Canada
larue.othalia@courrier.uqam.ca, {poirier.pierre,nkambou.roger}@uqam.ca

Abstract. In previous work [11], we introduced a cognitive architecture for the simulation of human behaviour based on Stanovich's tripartite framework [1], which provides an explanation of how reflexive and adaptive human behaviour emerge from the interaction of three distinct cognitive levels (automatic/reactive, algorithmic and reflective) capable of dynamically influencing each other in real time. Dynamic control between the levels in the architecture was then enhanced with a diffuse control system, based on the physiology of neuromodulations. In this paper, we study the physiological, functional and gross (mesoscopic) anatomical parallels between our architecture and the human brain, in order to understand how reactive yet sequential and planned behavior can emerge from physical systems of interacting simple processors such as our agents or low-level neural processors.

1 Introduction

Human behaviour is both reactive, manifesting dynamical sensitivity to multiple parallel factors, and sequential, manifesting planned analytic concatenation of basic or high-level actions. When uttering a sentence, for example, one respects the order in which sounds must be produced while adapting the vocal apparatus to many simultaneous constraints, both internal (need to breathe, to swallow, and even chew if one is talking with gum in his mouth) and external (observed response of interlocutor, unexpected events, etc.) The same is true of higher-level behaviour. In making breakfast, one must toast bread before buttering it but both toasting and buttering can only be achieved properly if they are dynamically responsive to a number of simultaneous parallel factors. This fact, which we will call the dynamical/sequential nature of behaviour, imposes a major constraint both on cognitive theories of the mind and on the design of cognitive architectures.

In this paper, after surveying previous attempts to account for the dynamical/sequential nature of cognition both in cognitive science (theories of human cognition) and artificial intelligence (design of cognitive architectures), we present our own proposal, based on a specific theory of human cognition, Stanovich's Tripartite Framework [1]. We have presented concrete results elsewhere [11][12] and, given the stated aims of the present conference, we shall focus

F.M. Zanzotto et al. (Eds.): BI 2012, LNCS 7670, pp. 288–299, 2012.

here on a detailed account of how processing in our architecture parallels known cognitive and neurological processing and how this allows us to realize hybrid reactive/deliberative behaviour in a system.

2 Accounting for the Dynamical/Sequential Nature of Behaviour: Related Work

2.1 Dual-System Theories of Cognition and Stanovich's Tripartite Framework

Dual-system (or dual-process) theories are a major current trend in cognitive theories to account for the dynamical/sequential nature of behaviour. We shall not review the variety of dual-system theories here (for such a review, see [2]). Suffice to say that such theories posit a host of opposed (dual) traits. System 1 is said to be automatic, fast, implicit, heuristic, associative, domain-specific, unconscious and so on. It should be noted that although one speaks of "System 1," it is believed that System 1 is actually a number of system, that all possess System 1 traits (so a better name would be "Type 1 system"). System 2 by contrast is said to be under voluntary control, slow, explicit, logical, sequential, domain general, conscious, and so on. Dual process theories account for the dynamical/sequential nature of behaviour by claiming that the reactive and dynamical aspects of behaviour result from the action of System 1 while the sequential aspects of behavior reflect the involvement of System 2. There are a number of problems with this way of accounting for the dynamical/sequential nature of behaviour, but most researchers in the field believe that these problems ask for refinements of the theory rather than it outright rejection.

Stanovich's tripartite framework is one such refinement of a dual-process theory. Stanovich's System 1, which he calls the "Autonomous Mind," includes instinctive behaviours, over-learned process, domain-specific knowledge, emotional regulation and implicit learning. However, contrary to other dual process theories in Stanovich's tripartite framework, System 2 is divided in two classes of processes: the "Algorithmic Mind," (cognitive control), and the "Reflective Mind" (deliberative processes).

The Algorithmic Mind acts upon information provided by the Autonomous Mind by means of two sets of pre-attentive processes which supply content to Working Memory : perceptual processes and memory access and retrieval.Three processes initiated by the Reflective Mind are implemented by the Algorithmic Mind: (1) inhibition of Autonomous Mind processes, (2) cognitive simulation (decoupling), and (3) serial associative cognition. Performance of these processes leads to an activation of the anterior cingulate cortex (ACC) [3].

Decoupling , is the creation of temporary models of the world upon which different alternatives can be tested. It is not a mandatory step in the decision making process, and when decoupling occurs, it can be incomplete. Subjects may apply simple models that appear appropriate for the situation, and the simple model chosen is not necessarily the best solution for the given situation. In fact,

through decoupling a better solution could have been found; however, since the cognitive load of decoupling is higher than that of serial associative cognition, subjects will often satisfy themselves with less optimal but cognitively easier solution provided by serial association. Serial associative cognition supports the implementation of the models. This operation is supported by the dorsolateral prefrontal cortex (DLPFC) [3].

Operations supported by the Reflective Mind define the subject's cognitive style. The Reflective Mind performs three processes: (1) initiation of inhibition of Autonomous Mind processes by the Algorithmic Mind (i.e., it tells the Algorithmic Mind: "Inhibit this Autonomous Mind process") and (2) initiation of decoupling in the Algorithmic Mind (i.e., it tells the Algorithmic Mind: "Start Decoupling") and (3) interruption of serial cognition, either by sending a new sequence to the Algorithmic Mind or by initiating a full simulation of the situation through decoupling. According to Stanovich, the division of human cognition into three sets of processes, instead of the traditional two of dual-process theories, provides a better account of individual cognitive differences. We chose the tripartite model Stanovich, precisely because it accounts of this diversity can help us to provide a good account of the hybrid nature of human behaviour.

2.2 Cognitive Architectures That Generate Dynamical/Sequential Behaviour

Many cognitive architectures have been designed over the years to explicitly exhibit behaviour that manifest a dynamical/sequential nature. We review two of the most influential here and one that theoreticallly resembles ours.

CLARION [4] reproduces the implicit/explicit duality. The architecture is divided into different subsystems: 1) oriented action subsystem (decision making level), 2) non action-oriented (general knowledge of the world), 3) motivational subsystem: ('why the action is performed'), and 4) the meta-cognitive subsystem (close to the 3) that monitors and regulates cognitive processes in the system). Each of these subsystems is split into two representational levels: implicit (distributed) and explicit (symbolic).

LIDA [14] is also a hybrid cognitive architecture approach, but unlike CLARION it models the duality that lies between conscious and unconscious processing. In LIDA's cognitive cycle, transition from conscious to unconscious processing is made through five steps: 1) Interpretation of sensory stimuli by perception units - unconscious stage; 2) Transmission of percepts to preconscious buffer (LIDA's working memory and emotions); 3) Local Associations between percepts and episodic, declarative memories (past emotions/situations); 4) Competition for attention (between coalitions of memories and perception; 5) Access of the winning coalition to the global workspace and broadcast of its content conscious; 6) Resource mobilization; 7) Production of a new hierarchy of objectives according to the context; 8) Action selection based on the new goal or the previously activated one; 9) Action execution.

HCogAff [5] was directly constructed in the context of dual-process theories. HCogAff is a tripartite architecture, with a reactive level (sensory and

perceptual processes, constantly changing with environment), a deliberative level (motor processes, effectors action to modify the internal state of the system), and a thought processing level (meta-management units, planification, decision making). Higher levels exert control upon those below, but this control can be disrupted (alarm signals from the reactive level, environmental stimuli, and motivations).

This last architecture is probably the one architecture that is structurally closest to our own; however, as we will see, the design focus in our architecture is on the dynamics of interactions between three distinct levels for the production of a coherent behaviour. Various open issues that remain to be properly addressed have been identified [6]. One of them is a better combination of deliberative problem solving with reactive control, that is, the need for architectures that can combine deliberative problem solving with reactive control by changing their location on the deliberative vs. reactive behaviour spectrum dynamically-based on the situation in which they find themselves. Our proposal, we believe, clearly addresses this issue.

3 Accounting for the Dynamical/Sequential Nature of Behaviour: Our Proposal

We base our cognitive architecture on Stanovich's tripartite framework [1]. This allows us a complete model of the cognitive mind, from automatic and implicit processes to explicit processes involving control (attention and executive functions) to more abstract planning and reasoning. Our architecture is implemented in a multi-layer multi-agent simulation platform [16]. As shown in Fig. 1, each level presented is composed of groups of agents acting in parallel, each agent having one or more role (an abstract representation of their functionality). In the following sections, we present the three level of our architecture (inspired by Stanovich's tripartite framework) from the most deliberative to less deliberative (most reactive). Interaction between those three levels and inside each level is achieved thanks to the message passing activity between agents. We cannot claim fine-grained neurological plausibility for this system. Parallels can however be drawn at the higher level of gross neurological structure and mesoscopic dynamical activity, allowing us to claim that a measure of neurological plausibility for the architecture. The neurological plausibility of dual process theories has been extensively studied [8][9]. Stanovich's tripartite framework, individually, is also supported by neurological data [1]. Since the design of our architecture is based on this model, it naturally inherits its neural plausibility.

3.1 Reflective Level

Structure. The logical and analytical skills of the system will be implemented at this level. Each agent in this group has a shape (a distance matrix) which represents the state that the system must be in to achieve a simple goal. Goal agents (I) are organized in a direct graph. A path in this graph represents a

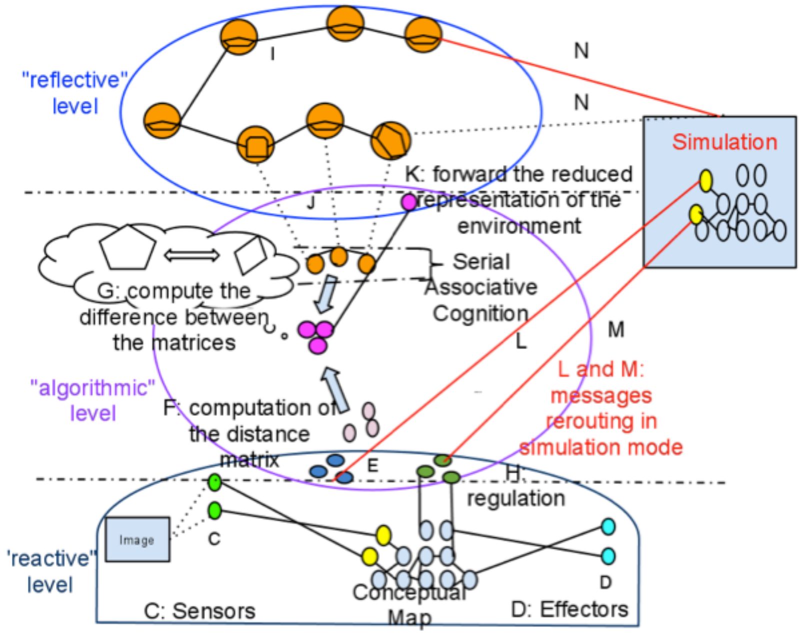

Fig. 1. Our architecture

plan that can be applied to achieve a complex behaviour A set of Goal agents represents a graph of several complex plans or strategies decomposed into a sequence of simple objectives. A succession of simple objectives (J) will be sent to the Algorithmic level, which will take care of its execution. Following Stanovich's Tripartite framework, agents in this group will have access to a reduced representation of the environment. This representation is provided by the Status agents of the Algorithmic Group to other status agents (K) that carry the reduced representation and announce themselves to the goal agents, which in turn compute their similarity to this representation. The activation of the Goal agents will be determined by the computed similarity between these two matrices. Activation propagates from the Goal agent most matching the reduced representation to those that follow in its path. The last agent in the path will send the parsed path to the Algorithmic level. Thus, the shortest path and the most active (with the most messages exchanged) will be sent first to the Algorithmic level. The shortest path (simplest model) or the one the most activated (model used more recently or more often) will prevail over the other paths. The limited serial associative cognition of the Algorithmic level will execute this path step by step. The path executed by serial associative cognition provides the system with the sequentiality necessary to achieve complex goals. However, the system is still a dynamical one. Indeed, a reduced representation of the environment is sent on a regular basis by the Status agents so that the Reflective organisation can interrupt serial cognitive association either by setting a new starting point in

the path, or by taking a new branch in the path, based on the current state of the environment. Decision-making at the Reflective level is therefore dynamically influenced by the current strategy and the state of the environment. Gross structural and mesoscopic dynamical parallels. The DLPFC, which provides support for goal directed behaviour is implemented in our system thanks to the Goal Agents.

3.2 Algorithmic Level

Structure. Corresponding to Stanovich's Algorithmic Mind, the Algorithmic group is responsible for the control of the system. Control is achieved with the help of morphology [7]. RequestStatus agents (E) belong to both the Reactive and Algorithmic organisation. At regular intervals, they query Knowledge agents about their status (number of messages sent during that interval to each of the Agents to which they are connected). Status agents (F) represent the system's activity at a given time in the form of a distance matrix that describes the message passing activity of the system at that time. The distance between two concepts in the conceptual map is measured by the number of messages sent between the Knowledge agents bearing these two concepts as their role. Status agents also send the reduced representation of the activity in the Reactive organisation to the Reflective level as we described in the previous section. Globally, this matrix thus represents a form or shape, and it is this form that will be transformed to reach the shape describing the goal assigned to the system. At the Algorithmic level, we thus find the short-term goals of the system in the form of a graph of Goal agents sent by the Reflective level. Each Goal agent (I) contains a distance matrix that specifies the distance necessary between each Knowledge agents (that is, the number of messages that must be sent between Knowledge agents) if the system is to reach goal. Graphs of short-term goals in our architecture correspond to Stanovich's serial associative cognition. Delta agents (G) compute the difference between the matrix provided by the Status agents and the one provided by the Goal agents. The resulting difference (another matrix) is provided to Control agents (H), which in turn send regulation messages to agents in the Reactive organisation to modify (i.e., increase) their activation so that their global activity more closely matches the shape describing the current short-term goal. Attention and Working memory in the system are implemented thanks to agents of the Algorithmic organisation. Control agents regulate the activation of elements in the system's semantic memory in relation to its current goal. The system's long term memory is made up of the Knowledge agents in the Reactive organisation, and the system's working memory (WM) at a given time is made up of the Knowledge agents that are activated in the Reactive group at that time. Decoupling, an operation initiated by agents of the Reflective Group, is supported by agents of the Algorithmic Group. When multiple strategies (meaning two or more GoalSet Agents) are selected at the algorithmic level, the Goal Agent that belongs to the Algorithmic and Reflective Group triggers a simulation of the strategies. When the Algorithmic Group is in simulation mode, a possible world is created thanks to the reduced representation sent by

the Delta agents. This secondary representation is realized with a limited number of agents (20). These agents are assigned dynamically the same roles and links as those agents from the Reactive Group they are replicating, as indicated by the reduced representation. Since this possible world is carried out thanks to distinct agents (SecondaryRepresentation Agents instead of Knowledge agents) and a distinct group (Algorithmic instead of Reactive), we can be sure that this secondary representation is totally independent from the current representation of the world (i.e., Knowledge Agents from the Reactive Group). To reproduce the cognitive cost of the simulation operation, the cognitive operations (goal inhibition and selection) are carried out by the Control Agents, Delta Agents, and agents from the Reflective Group. Messages from (L) and to these agents (M) are branched to the SecondaryRepresentation Agents of the Algorithmic Group instead of the Knowledge Agents of the Reactive Group. Once the simulation is completed, the activation of Goal Agents is regulated accordingly at the Reflective level (N), therefore potentially replacing the next action carried out by the Algorithmic level (by the first rule simulated).

Gross Structural and Mesoscopic Dynamical Parallels. Regulation of posterior brain regions is implemented by the regulating messages sent by the Control Agents to agents of the Reactive Group Cognitive tasks supported by the Algorithmic Mind lead to an activation of the Anterior Cingulate cortex (ACC) [1]. There is furthermore evidence that decoupling is achieved by the Dorsolateral PreFrontal Cortex (DLPFC) [1]. Furthermore, the different roles ascribed to the agents in the architecture correspond to functional roles that has been mapped to specific anatomical structures see Fig. 2 [15]. It must be noted that the ACC has been identified as the response conflicting monitoring system [15] in the human brain, regulating control's engagement. Conflict monitoring is achieved in our system by the collaboration between the Delta Agents and the Control Agents.

4 Reactive Level

Structure. The Reactive level in our model corresponds to Stanovich's Autonomous Mind. The main roles assigned to agents within this level are "sensor" (C), "effector" (D) and "knowledge" (A). The network of Knowledge agents (agents assigned with the "knowledge" role) is initialized with a knowledge base that makes up the system's declarative knowledge (semantic memory): a conceptual map made up of concepts and the semantic links between them. Knowledge agents therefore have two attributes: "knowledge" and a word from the knowledge base (e.g., "Red"); knowledge agents are also connected together according to the links in the conceptual map. Upon receiving a message from a Sensor agent or from another Knowledge agent, Knowledge agents send a message to those Knowledge agents they are connected to, therefore spreading activation in the network (a process similar to that of semantic memory, [13]). The number of messages exchanged between the agents, and therefore their activation, is

at first determined by the distance between them in the conceptual map. The system's environment is similar to (portions of) human environments. Each Sensor agent is sensitive to some particular type of information in the environment (colors, sounds, texts, etc.). If the type of information to which they are sensitive to is present in the environment, Sensor agents will (at short intervals) extract it and send messages to Knowledge agents with a role associated with the sensor's function (”read” for Knowledge agents connected to Sensor agents reading characters, ”recognizeColor” for Knowledge agents connected to Sensor agents recognizing colors). Activation in the network therefore depends on the number of messages sent by the Sensor agents and the activation of the Knowledge agents in the conceptual map. Taken together, the action of Sensor and Knowledge agents make up the system's sensory motor level. This means that the system's sensory abilities are always a function of the Sensor agents' information extracting capacities and of the system's knowledge about the environment: the system is fully situated. Effectors agents work similarly: a knowledge agent associated to the function of the effector (“sayRed”, “sayBlue”) sends messages to Effector agents with a similar role, which will then act on the environment.

Gross Structural and Mesoscopic Dynamical Parallels. Knowledge Agents are linked to Sensor Agents as the Medial Temporal Lobe (MTL in Fig. 2) is known to mediate sensory memory. The medial temporal lobe is also identified as the functional locus of semantic memory. “Effector” and “Sensor” agents are associated with distinct roles in the Reactive Group since their functional role is achieved by distinct anatomical structures (PreCentral Gyrus, PostCentral Gyrus and Visual Cortex in Fig. 2). Through nested sensorimotor and goal-directing loops, we are therefore able to implement the cognitive dynamics of a goal-sensitive sensory-motor architecture.

5 Neurological Parallels

Signals in dynamical systems are called diffuse if they rule secondary variables in large areas. In our system, neuromodulations are diffuse signals that control the overall state of our system. In the system, the processing occurs thanks to message passing. Messages are exchanged between group of agents inside the levels, and thanks to intermediary agents which belongs to both levels to pass information between levels. Messages are exchanged at regular intervals, those intervals varying from groups to groups (for example: Goal agents which send the future objective to the system at regular intervals, Control agents also send regulation messages at regular intervals to agents at the reactive level, this interval being inferior to the one used by Goal agents, and this interval (the one use by Control agents) is superior to the one used by Knowledge agents at the reactive level) Neuromodulations are messages sent by agents belonging to another group. Agents of the system (of the three levels) upon reception of those messages count the number of activation messages received, and , through a transfer function representing the sensitivity of the agents to this modulation (mirroring

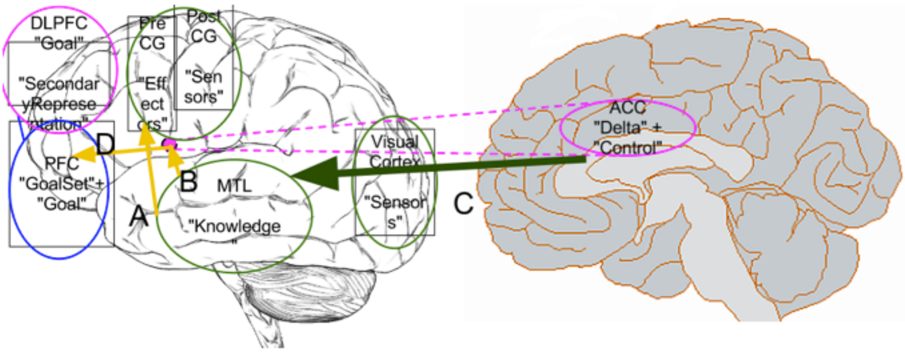

Fig. 2. Neurological parallels

the sensitivity of neurons to neuromodulations), compute the "modulated" interval that the agent will now be using. Neuromodulations help us modify the system's general dynamics (allowing us finer-grained control on the dynamics), giving specific profile to the system (reactive oriented, reflective oriented,). We intend on reproducing the four following neuromodulatory systems [10].

Dopamine is mainly present in two important brain pathways: the mesocortical pathway (frontal lobe and prefrontal cortex), involved in the update of working memory, and the nigrostriatal pathway (basal ganglia), involved in motor skills (A). In our system, the mesocortical pathway is implemented thanks to the messages sent by agents of the reactive group to the StatusUpdate agents of the algorithmic group (B). The nigrostriatal pathway, for its part, is implemented by the messages sent to the effector agents. The reward and motivation effect occur thanks to the reinforcement or inhibition of goal agents that have led to a success in the past (D).

For the *Serotonin* neuromodulation, we will replicate the pathway projecting from the raphea nuclei to the thalamus, subthalamus, limbic system and cerebral cortex, especially its role in long-term memory and executive control. The influence on long-term memory will be implemented through the modulation of the messages exchanged between agents with the "knowledge" role (D). The influence of executive control will be reproduced by modulating messages between control agents and agents belonging to the reactive group (E).

Norepinephrine (Noradrenergic system), acting on the pathway from the locus coreuleus to the the forebrain and brainstem, it can create arousal; in the system, this action will be represented by a global increase in the message passing activity at the Reactive level. The pathway from the locus coreulus to the posterior parietal lobe, responsible for the integration of sensory information will be represented in our system by the modulation of the sensitivity to sensory stimuli (E) (sensitivity of Sensor agents).

Concerning *Acetylcholine*, the cholinergic system also affects the regulation of specific sensory relay (by modulating activity in the basal forebrain). We will reproduce this by modulating the message passing activity from the sensor

agents. The action of acetylcholine on the pathway from frontal cortex to nucleus basalis.impacts divided attention (attention to different tasks at the same time). In our system this would be represented by the message passing activity occurring with Decoupling agents.

6 Current Status and Future Directions

For now, we have validated the architecture for the first two levels (reactive and algorithmic) and part of the third levels, thanks to two psychological tasks (the Stroop task [11] and the Wisconsin card sorting task [12] for which we obtained results comparable to human results. The Stroop task is a task aimed at the evaluation of a subject's cognitive control: subject is confronted to two stimuli resulting in the competition of two responses, with the one required to fulfill the task not being the habitual response. With the Stroop task[11], we were able to evaluate the system perceptual's attention (achieved thanks to the regulation of knowledge agents and linked effector agents by the control agents) and cognitive control (tregulation of knowledge agents by control agents). With a variant of the Stroop task, Semantic Stroop task, we were able to evaluate its ability to process context and retrieve information in its long term memory (conceptual map formed by knowledge agents). In the Wisconsin Card Sorting Task[12], a subject has to discover and use different sorting rules on a set of cards. With this task, we were able to validate the system's cognitive flexibility, and the equilibrium between deliberative and reactive behaviours. Cognitive flexibility is made possible by the interaction between the three cognitive levels, that helps preserve this equilibrium in the architecture to make it efficiently adaptive : the rule discovery is made in conjunction by the reflective level (rules firing) and the algorithmic level (decoupling), the application of the rule is made thanks to a mechanism similar to the one involved in the Stroop task (regulation of the knowledge agents of the reactivel level by the Control agents of the algorithmic level). We have worked for now on the implementation of two neuromodulations : dopamine and serotonine. In the future, we plan on adding two other neuromodulations : acetylcholine (sensors and executive functions) and noradrenaline (arousal). In future work, we plan on addressing the physical embodiment and emotional modulation of cognitive processes. Neuromodulations that influence communications between organisations, and therefore modify the general dynamic of the system would be the basis of physiological dimension of emotions in the system (their conceptual/cognitive dimension is already present in the conceptual map). Neuromodulation agents could then be linked to emotional concepts/knowledge ("happy"), Allowing for an interaction between the physiological activation (neuromodulation) and conceptual interpretation of emotions.

7 Conclusion

We modeled our cognitive architecture on Stanovich's Tripartite Framework, which explains how cognition can be at once reactive and sequential by positing

three levels of interacting cognitive processes. Brains are the physical systems we know that have best managed to dynamically integrate real-time reactivity with sequentiality; and human brains are perhaps the only system that has perfected this integration to allow such cognitive feats as sequential problem solving that is sensitive to the evolving problem space as actions are applied or long-term planning that can self-adapt as environments, knowledge and goals change. Our views on Artificial Intelligence thus recommend that we complement our functional-level model of cognition by neural-level knowledge of the implementation of the cognitive processes posited by such theories. To conclude, we therefore want to look at how processing may be organized in physical systems of interacting simple agents or processors to allow reactive and sequential processes to influence each other in order to fully take advantage of this interaction. The cognitive processes described here result from an interaction between the medial temporal lobe (MTL), the anterior cingulate cortex (ACC) and the dorsolateral prefrontal cortex (DLPFC). The MLT-based autonomous sensorimotor agent is sensitive to the meaning of inputs and outputs and to the semantic relations that make (declarative) knowledge into a spreading activation semantic memory. Because no perception or action is uninfluenced by knowledge, this sensorimotor agent always acts in an informed (or meaningful) way. However, in itself, it is not sensitive to internal goals or plans: it is condemned to live in the present. The ACC-based Goal agents provide internal motives to the reactive-level agents by slightly, but constantly, altering the activity of Knowledge agents (which in turn can affect the activity of Sensor agents). If this model is right, then goals merely influence the dynamics of a sensorimotor system already capable of meaningful action in the world. The emerging picture of MLTACC interaction at this point is one where the ACC forces MLT-processing into a preferred direction. This is not the extant of processing in our agent, however. We saw that DLPFC-based agents receive a reduced representation of the environment (through a reduced representation of the activity of Knowledge agents), which determines which path in the goal structure (the directed graph of Control agents) will be sent to the Algorithmic level for execution. Since these DLPFC agents can halt current execution of a goal-path (plan or strategy) or force a switch to a new goal-path, MLT-processing can influence, albeit indirectly, the pressures ACC imposes on its own processing. The resulting picture is one of two linked dynamical loops: the meaningful sensorimotor loop through the world and the goal-setting MLT-DLPFC-ATT-MLT loop. Our current agent has two additional control structures. The first is an ACC-based ability to simulate counterfactual situations, initiated by DLPFC-based agents, the result of which may determine which goal-path is sent to the ACC for execution. This first additional control structure is therefore part of the agents goal control loop, allowing the selection of goals that would be chosen in different yet related environments. The second is a lower-level modulatory influence (neuromodulations) diffusely controlling the activity of Knowledge agents (serotonin) and of Goal agents (dopamine). The effect of neuromodulation is either to make the agent more sensitive to the environment (serotonin) or more sensitive to its goals (dopamine). Serotonin neuromodulation

is thus part of the sensorimotor loop, whereas dopamine neuromodulation is part of the goal control loop. At this point, however, neuromodulations are point quite crude in our system (we basically switch them on or off individually). We plan in future work to give neuromodulation its own dynamics, which would determine the modulatory state of the agent at any time. The neuromodulation dynamics would be influenced by low-level physiological parameters (drives, excitation, wakefulness) as well as by the system's other two loops that characterize, that is, by events in the environment (through the sensorimotor loop) and by current or coming goals (through the goal control loop).

References

1. Stanovich, K.E.: Rationality and the reflective mind. Oxford University Press (2010)
2. Evans, J.S.B.T.: Dual-Processing Accounts Of Reasoning, Judgment, and Social Cognition. Annual Review of Psychology 59, 255–278 (2008)
3. Frankish, K., Evans, J.S.B.T.: The duality of mind: a historical perspective. In: Frankish, K., Evans, J.S. (eds.) In Two Minds, pp. 1–29 (2009)
4. Sun, R., Lane, S., Matthews, R., et al.: The two systems of learning. In: Frankish, K., Evans, J.S. (eds.) In Two Minds. Oxford University Press (2009)
5. Sloman, A., Chrisley, R.L.: Virtual machines and consciousness. Journal of Consciousness Studies 10(4-5), 133–172 (2003)
6. Langley, P., Laird, J.E., Rogers, S.: Cognitive architectures: Research issues and challenges. Cognitive Systems Research 10(2), 141–160 (2009)
7. Cardon, A.: La complexit organise. Herms Science Publications, Paris (2005)
8. Goel, V.: Cognitive Neuroscience of Thinking. In: Berntson, G., Cacioppo, J. (eds.) Handbook of Neuroscience for the Behavioral Sciences. Wiley (2009)
9. Lieberman, M.: What zombies can't do: a social cognitive neuroscience approach to the irreducibility of reflective consciousness. In: Frankish, K., Evans, J.S. (eds.) In Two Minds: Dual Processes and Beyond, pp. 293–316. Oxford University Press (2009)
10. Bear, M.F., Connors, B.W., Nieoullon, A., Paradiso, M.A.: Neurosciences Additions Pradel (2007)
11. Larue, O., Poirier, P., Nkambou, R.: A Three level Cognitive architecture for the simulation of cognitive phenomena. In: CAI 2012 (2012)
12. Larue, O., Poirier, P., Nkambou, R.: Hybrid Reactive-Deliberative Behavior in a Symbolic Dynamical Cognitive Architecture. In: International Conference on Artificial Intelligence 2012 (2012)
13. Anderson, J.R.: The architecture of cognition. Harvard University Press, Cambridge (1983)
14. Franklin, S., Ramamurthy, U., D'Mello, S., McCauley, L., Negatu, A., Silva, R., Datla, V., et al.: LIDA: A computational model of global workspace theory and developmental learning. AAAI Fall Symposium on AI and Consciousness. Arlington (2007)
15. Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S., Cohen, J.D.: Conflict monitoring and cognitive control. Psychol. Rev. 108(3), 624–652 (2001)
16. Ferber, J., Gutknecht, O., Michel, F.: From Agents to Organizations: An Organizational View of Multi-agent Systems. In: Giorgini, P., Müller, J.P., Odell, J.J. (eds.) AOSE 2003. LNCS, vol. 2935, pp. 214–230. Springer, Heidelberg (2004)