Hybrid Reactive-Deliberative Behaviour in a Symbolic Dynamical Cognitive Architecture

Othalia Larue[1], Pierre Poirier[2], Roger Nkambou[1]
GDAC Research Laboratory
[1] Department of Computer Science
[2] Department of Philosophy
Université du Québec à Montréal
Corresponding author: larue.othalia@courrier.uqam.ca

Abstract - *Sequentiality and reactivity are features that have been deemed important for cognitive architectures[1]and recent emphasis has been put by the community on their development in cognitive architectures. However, the cooperation and competition dynamic between reactivity and sequentiality remains an open issue in the domain [2]. In this paper, we present a three level cognitive architecture for the simulation of human behaviour based on Stanovich's Tripartite Framework [3], which provides an explanation of how reflective and adaptive human behaviour emerges from the interaction of three distinct cognitive levels. We use two classical psychological tasks to study the reactivity/sequentiality dynamic in our architecture. These show that the two features collaborate in interesting and psychologically plausible ways.*

Keywords: cognitive architecure; hybridism; sequentiality; reactivity

1 Introduction

A cognitive architecture is "the overall, essential structure and process of a domain-generic computational cognitive model, used for a broad, multiple-level, multiple domain analysis of cognition of behaviour" [1] (p.4). Psychologists and other cognitive scientists can use these architectures to study the mechanisms responsible for observed behaviour, and engineers can employ them to endow their systems with cognitive (e.g. decision-making) capacities. Theoretical studies of cognitive architectures [1,3, 4] have identified many general features that architectures should have if they are to efficiently play these roles. Two such features, reactivity and sequentiality [1], have proven difficult to integrate in well-unified cognitive architectures, and, emphasis has recently been put on the development of these features (especially in the guise of reactive abilities and of reflective/deliberative). It is difficult to integrate reactivity and sequentiality because the two features appear to be functionally incompatible: designs that favour one almost inevitably hinder the other. As a result, integration of reactive and sequential processing remains a challenge in the domain [2]. In this paper, we address the challenge by introducing a new architecture, one that implements Stanovich's [3] Tripartite Framework: a framework that aims to explain how reflective (characterized by sequentiality) and adaptive (characterized by reactivity) human behaviour emerges from the interaction of three distinct cognitive levels (autonomous/reactive, algorithmic/cognitive control, and reflective). To demonstrate the flexible and coherent behaviour of the resulting architecture, we study its performance on two classical psychological tasks: the Stroop task, which requires perceptual attention and cognitive control, and the Wisconsin card sorting task, which requires cognitive flexibility and efficient collaboration between the reactive and sequential elements of the architecture.

2 Related Work

2.1 Cognitive architectures

Although there are now a variety of cognitive architectures (see for a review [3]), we chose here to focus on three of the most widely used: ACT-R, SOAR, CLARION [1].

ACT-R (Adaptive Components of Thought-Rational) [5] is a cognitive architecture whose development is oriented towards the understanding of human cognition. ACT-R's components are a set of perceptual-motor modules, memory modules, buffers, and a pattern matcher module, which finds productions that match the current state of the buffers. There are two types of memory modules in ACT-R: declarative memory and procedural memory, consisting of chunks or production (for the procedural memory) and associated sub-symbolic values (connectionist hybridism). A long-term memory of production rules coordinates the processing of the modules. Each module has a chunk holding a relational declarative structure. Each chunk has a set of sub-symbolic parameters reflecting its past activity and influencing its future retrieval from long-term memory. Adaptation in ACT-R occurs thanks to a top-down learning approach. ACT-R has been used to simulate a large number of cognitive phenomena but has seldom been used for the simulation of extended metacognitive processes. Further effort has been put into implementing a unified theory of cognition, perception, and action by integrating perceptual and motor modules working in parallel with cognition; however, cooperation between these modules is limited since their content (perceptual, motor, declarative memory) is still processed using distinct buffers.

SOAR (State, Operator And Result) [4] is also a rule-based cognitive architecture aimed at the modelization of general intelligence. Knowledge is in the form of production rules, arranged in terms of operators acting in the problem space (set of states representing the task). Operators provide the system with adaptation since they can externally as well as internally modify the system's state. The primary learning mechanism is chunking, which allows the extraction of rules from problem solving traces. A basic processing cycle repeatedly selects, and applies operators, achieving one decision at a time. In SOAR, different types of learning are applied to different types of knowledge: reinforcement learning to adjust preference values, episodic learning to keep track of the system's evolution, semantic learning for declarative knowledge. SOAR is able to perform high-level reasoning task (planning, problem solving …).

As opposed to ACT-R and SOAR, CLARION requires less a priori knowledge. CLARION (Connectionist Learning Adaptive Rule Induction ON-line) [6] is a hybrid architecture with explicit (symbolic) and implicit (sub-symbolic) processes. CLARION is made of four memory modules, with dual explicit–implicit representation: action-centered subsystem, non-action-centered subsystem, motivational subsystem, and metacognitive subsystem. Action and non-action centered knowledge are stored in implicit form (using neural networks) and in explicit form (using symbolic production rules). Two types of learning support the team work of implicit and explicit processes : bottom-up learning (reinforcement learning methods are used to acquire implicit knowledge, the resulting knowledge is used to modify explicit knowledge at the top level through bottom-up learning mechanism), and top-down learning (extracting knowledge by observing actions guided by these rules). As in dual-process theories of mind, two levels (a meta-cognitive subsystem and a motivational subsystem) cooperate to produce behaviour by combining the action recommendations from the two levels or combining bottom-up and top-down learning. CLARION has often been used for the simulation of higher level cognitive phenomena. . However, CLARION's sensory-motor modules are not as developed as one would wish.

The architectures presented above exhibit many of the desiderata [1] for cognitive architectures, but none support all of them, often because of the theoretical orientation took as foundation (strong symbolism in ACT-R, high modularity in CLARION). Several issues for further research were thus identified; here we highlight specifically two of them [2]: (1) Effectors and perceptual attention: the need for "expanded frameworks that manage an agent's resources to selectively focus its perceptual attention, its effectors, and the tasks it pursues" [2]. In SOAR and ACT-R, perceptual systems are isolated channels providing the information from the environment to the Working memory of the system, but their activity is not regulated by the Working memory, like the activity of other productions. Also, in SOAR, perceptual cues match on separate perceptual working memory elements that are independent of context (matched by the other productions). As we will see below, perceptual information processing in our architecture depends on the current task and the context (environment and long term memory of the system). (2) Combination of deliberative problem solving with reactive control: 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 their situation. CLARION, as well as other architectures (see CogAff, Sloman [7]), has chosen to address this problem by using a dual-process theory of mind [1]. It is this duality that allows CLARION to achieve high level reasoning; however, strong modularity prevents the system from really achieving strong reactive/dynamic processing. In our architecture, thanks to the unified cognitive model we chose, we are able to preserve dynamic processing while maintaining a robust (reflective) behaviour. In this paper, to design an architecture that meets the duality challenge with a clear answer to the interface problem, we have modeled our system after Stanovich's Tripartite Framework [3]. We chose this model, precisely because it provides a good account of dynamic of duality (sequentiality/rule-following and reactivity/dynamicity) in human cognition.

2.2 Cognitive Model

We base our cognitive architecture on Stanovich's tripartite framework [3]. 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. Stanovich's tripartite framework belongs to the "dual-process theories" family of cognitive models, where cognition is characterized by the opposition (duality) between two types of processes or systems [8]. We will follow the latter usage, which often dubs the dual systems "System 1" (fast and automatic reasoning) and "System 2" (abstract and hypothetical reasoning). Stanovich's System 1, which he calls the "Autonomous Mind," includes instinctive behaviours, over-learned process, domain-specific knowledge, emotional regulation and implicit learning. His tripartite framework differs from other dual-process theories in its description of System 2. He divides processes usually ascribed to System 2 in two classes of processes, respectively called the "Algorithmic Mind," responsible for cognitive control, and the "Reflective Mind," responsible for deliberative processes. The Algorithmic Mind acts upon information provided by the Autonomous Mind thanks to two sets of pre-attentive processes (perceptual processes and "processes which access memories and retrieve memories and beliefs" – [9] (p 43), both of which supply content to Working Memory. The Algorithmic Mind is the locus of three processes, each initiated by the Reflective 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). Decoupling, which seems to be supported by the dorsolateral prefrontal cortex (DLPFC) [3], consist in the creation of temporary models of the world, where different alternative scenarios can be tested. The temporary models created through decoupling do not affect the system's current representation of the world but the decoupling process, however, has a cognitive cost: it is difficult for the Algorithmic Mind to perform other processes while decoupling takes place. Decoupling does not occur in every situation where it would be useful, and when it does occur, it is sometimes incomplete. In these cases, subjects simply apply simple models (rules) that *appear* appropriate for the situation. Serial associative cognition supports the implementation of these simple models. The simple model chosen does not in general provide the best solution in a given situation: a better solution could have been found through decoupling, but the cognitive load of decoupling is higher than that of serial associative cognition and thus subjects will often satisfy themselves with less optimal but cognitively easier solution provided by serial association. 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. Individual differences with regard Algorithmic Mind processes are linked to cognitive abilities and fluid intelligence, while Reflective Mind differences are observed in critical thinking skills. We chose the tripartite model Stanovich, precisely because it provides a good account of the diversity in human behaviour.

3 Architecture

Our architecture is implemented in a multi-layer multi-agent simulation platform [10]. As shown in figure 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).

3.1 Reactive level

The Reactive level in our model corresponds to Stanovich's Autonomous Mind. The main roles assigned to agents within this level are "sensor" (C –letters in parenthese in this section appear in figure 1), "effector" (D) and "knowledge" (A).

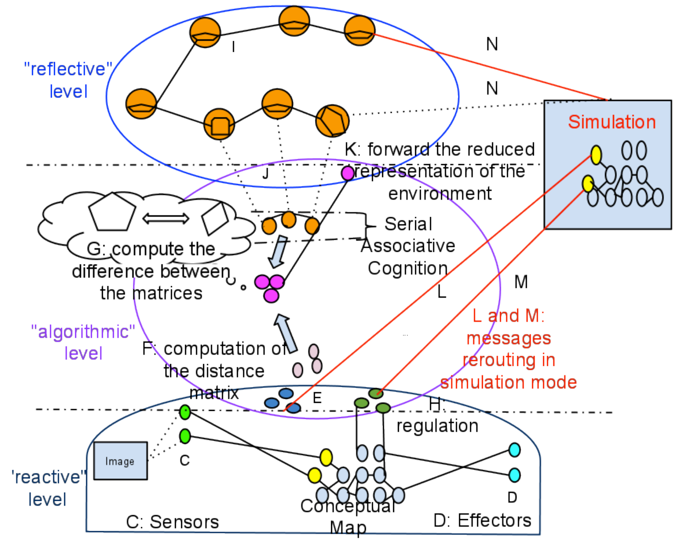

Figure 1: Architecture

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 memories, [11]). 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 (later on, it will also be determined by the activation signals from higher levels – see below). The system's environment is similar to (portions of) human environments. In the Stroop task simulation described below, the system is presented with cards identical to those human subjects see in real Stroop task experiments. 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.

3.2 Algorithmic level

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 [12]. RequestStatus agents (E) belong to both the Reactive and Algorithmic organisation. At regular intervals, they query Knowledge agents about their status (that is, number of messages they 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 a reduced representation of the activity in the Reactive organisation to the Reflective level. 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. Agents in the Algorithmic organisation constitute the system's attention. They activate elements of 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. This implementation of working memory is consistent with the work of Engle [13], in which WM is seen as a set of temporarily activated representations in long-term memory.

3.3 Reflective level

Each agent in this last 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 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 (steps in the plan). The logical and analytical skills of the system will be implemented at this level. A sequence of simple objectives path (J) will be sent to the Algorithmic level, which will take care of its execution. Following Stanovich's Tripartite framework, agents in this last 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 does not lose its dynamicity. Indeed, the reduced representation of the environment are sent on a regular basis by the Status agents so that the Reflective organisation can interrupt serial cognitive association either by: (1) Setting a new starting point in the path, or by, (2) Taking a new branch in the path, based on the current state of the environment

3.4 Simulation

If multiple strategies (thus two or more goal agents) are selected at the algorithmic level, the goal agent that belongs to the algorithmic and reflective level (which usually contains the goal matrix selected at the reflective level) triggers a simulation of the strategies. The simulation capacity as envisioned in Stanovich's Tripartite Framework is implemented at the algorithmic level. When the algorithmic level 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 level 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 level). 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 the Reflective level. Messages from (L) and to

(M) these agents are branched to the SecondaryRepresentations instead of the Reactive level. 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). The WCST realized in this paper illustrates the simulation of an opposite style of thinking. The system first simulates the application of a chosen categorization rule and its negation (negative feedback by the instructor), and is thus able to make a new categorization rule (the alternative) emerge in the simulation.

4 Results and Discussion

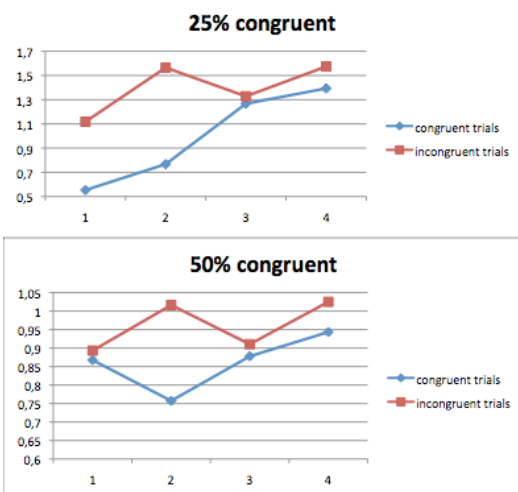

Figure 2: Mean response time per block for congruent and incongruent trials in 25% congruent conditions and 50% congruent condition

4.1 Control and perceptual attention: Stroop task

The Stroop task [14] is used to test attention and inhibitory control. It tests a subject's ability to maintain a goal in mind, suppressing a familiar response in favor of one that is less familiar. The task illustrates the Color–Word Interference effect. The set of trials is a compound of congruent trials (the word "GREEN" written in green) and incongruent trials ("RED" written in green). Two experiments were conducted. In each one, four blocks of 100 cards with a word written in a specific color were shown to the system. 25% of the cards per block were congruent in the first experiment while 50% were in the second.

Control: Set-up this way, we find that the mean response time of the system is longer for incongruent trials, specifically in the 50% congruent condition (Figure 2) (1396 ms) as opposed to the 25% congruent condition (960 ms), a result also found in human subjects. This is because, to achieve the system's goal, Control agents have to send regulation messages to the "recognizecolor" agents to compensate the distribution of agents when the system is initialized for the Stroop task (i.e., with a preponderance of "read" agents to reflect the predominance of the reading ability in normal adult subjects). The system provides an answer once the system has regained its stability, that is, once there is a sufficient difference in activity between the system's two competing responses. During incongruent trials, this stability is harder to achieve (inhibition is a time-consuming operation). Response-time variation is due to the action of the algorithmic level: in incongruent trials, it is particularly called upon to help the system remove perceptual focus on confusing information.

Perceptual attention: the Stroop task especially illustrates the dual nature (orienting and being oriented) of perceptual units in our system. The system gets initial information about its environment through its sensory agents ("recognizecolor" and "read") –the system processes all stimuli in parallel. At first, this leads to an activation of the Knowledge agents that correspond to the external stimulus (e.g., the Knowledge agent that bears the role "Red" will be activated if Sensory agents detect red in the environment). As we mentioned above, when the system is initialized for the Stroop experiment, there is a preponderance of "read" agents. After this initial phase, Control agents, guided by information provided by Status agents, themselves influenced by the system's current goal, can modify the message passing activity between the agents (by increasing the activity of the "recognizecolor" agents). With this experiment, we show that despite an initial setting favoring the "reading " ability and its connected sensor (both belonging to the reactive level), the system is able to change its "natural" tendency (from reading the word to naming the color), thanks to the cognitive control achieved by the algorithmic level. Cognitive control favors the color naming ability by increasing the activation of relevant agents in its goal matrix. Agents connected to them are in turn linked to effectors, therefore achieving perceptual control (perceptual attention) by spreading activation.

4.2 Flexibility: Wisconsin card sorting test

The Wisconsin card sorting test (WCST) [17] is widely used to test executive functioning, especially cognitive flexibility and abstract reasoning. The subject is shown a set of target cards with figures on them, which vary in shape, number and color. The subject has to match stimulus cards to the target cards, one by one. However, he is not told what the sorting rule is and has to discover it. In our experiment, we used material adapted from Dehaene & Changeux [15] representing the context effect of the four reference target cards by adding the following links in the system's conceptual map: Red – triangle – one, Green – star – two, Yellow – square – three, Blue – circle – four. We also linked the shape, color, and number knowledge (already present in ConceptNet) to the sensors. A script provided the system with the series of cards it had to categorize and evaluated the system's answer. The categorization rule was changed after 6 consecutive successes. The experiment consisted of 128 cards with figures varying in shape, number and color. The script attempted to test 6 categorization rules (["shape", "color", "number"] times 2) on the 128 cards. As in human trials, no warning was sent before a rule change: only the answer "no" was given to the system when his answered was wrong.

Figure 3: Wisconsin card sorting task simulation

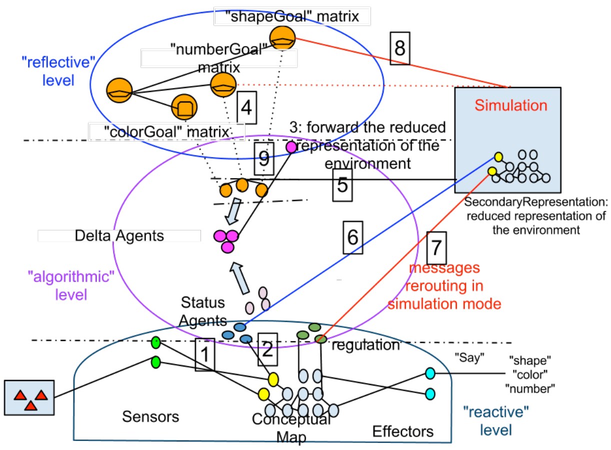

Perceptual attention: (Note that numbers in this section refer to figure 3. Please refer to the figure while the following description of the system's activity) When a new card appears in the environment, Sensors (1) forward the information to Knowledge agents in the reactive level. Activation of the different Knowledge agents is thus influenced by the environment's state. Status Agents then (2) provide the algorithmic level with the status of the agents at the reactive level. This information is used by the Delta Agents in the Algorithmic level to calculate a reduced representation (3) of the environment, which is then forwarded to the Reflective level, leading to the activation of various competing rules (4). When there are more than one winning rule/goal (because the system's working memory is loaded with contradictory contextual information), a process of cognitive decoupling (internal simulation) is launched (5) (by the Decoupling agents). A mini-world (decoupled world) -is thus created: agents in this mini-world are modeled after the reduced representation of the world (6) sent by Status agents. Decoupling agents also recreate an environment, to which the agents of the mini-world will react. For this Experiment, the environment was one in which the "no" response was sent after a categorization rulewas proposed by the mini-world. Regulating and status updating messages are rerouted (7) to act on the mini-world instead of the reactive level. The rules that emerge from the cognitive simulation are sent (8) to the reflective level (with different activations) and then the corresponding matrix is sent (9) to the algorithmic level, which is in charge of regulating the reactive level towards achievement of the goal (encouraging agent activity – even agents involved in perception – according to the system's current goal). The average number (for a hundred simulations) of rules the system (for its 128 trials) was able to discover and apply was 5.33; the maximum number was 6. Normal human subject are able to discover an average of six categories [16]. These results illustrate the system's ability to adapt adequately to the changing situation (new card, error notification) in a bidirectional manner: by orienting its executive control according to the information in the environment but also by orienting its perceptual processing according to the goal.

Equilibrium between deliberative and reactive behaviours: It is the system's design, focused on the interaction between the three cognitive levels, that helps preserve this equilibrium in the architecture to make it efficiently adaptive. Perseverationerror in the system are errors due to the reactive level.

Table 1: A simulation's log.

Serie	Trial	Response	Simulation
Color	1	correct	
	2	correct	
	3	correct	(1)color (2) shape
	4	Incorrect : shape	
	5	correct	
	…		
Shape	1	Incorrect : color	(1) shape (2)number
	2	Incorrect : color	(1) shape (2)number
	3	Incorrect : color	
	4	Incorrect : color, number	(1) number(2) shape
	5	Correct	(1) shape (2)number
	…		

However, the system's ability to achieve five categorizations is the result of a good interaction between reactive, algorithmic and reflective processing. Error notifications are sent from the Reactive level to the Algorithmic level, decreasing the activation of the now wrong classification rule, thereby allowing other rules/goals to take the lead. Activations levels of the goals are a memory of those classification rule that work and those that did not work in the past.. Although, the simulation capacity illustrates the interaction between Algorithmic and Reflective levels, it is also primordial for good cooperation between Deliberative and Reactive behaviour. Decoupling helps generate a prediction of the behaviour at the reactive level and therefore prepare an adaptive plan of action (efficient trial-error adaptation): in the third trial of the "color" series, a cognitive decoupling is run because two competing answers ("color" and "shape") are active. The system creates a simulation of a possible world where the color categorization rule is activated and observed as wrong; in this possible world, the second emerging rule was the shape categorization rule. In the first and second trial of the "shape" series, after a first wrong answer, "color" is selected, a cognitive decoupling is run where "shape" is first activated, since the possible world is an image of the system's environment before the simulation process started and where "color" had been marked as a wrong answer, the selected second answer is "number". In the third trial of the "shape" series, after a first incorrect answer (color), a cognitive decoupling is run where the first rule activated is the “number” categorization rule, and the second is the “shape” rule, leading to a second error. Since the cognitive decoupling had activated the “shape” as second rule, the correct answer is produced.

5 Conclusion and Future Work

Theoretical studies [2,4,1] identified sequentiality and reactivity as two important features cognitive architectures must have if they are to be useful to cognitive scientists and engineers [1]. It has however proven difficult to integrate the two features in coherent architecture due to their functional incompatibility. In this paper, we focused on two such issues: Action attention and perceptual attention and flexibility. To design an architecture that addresses these issues, we sought inspiration from natural minds and modeled our system after Stanovich’s Tripartite Framework [3] (reactive, algorithmic and reflective minds).The initial reaction of our system to stimuli is automatic (reactive level). However, as the task goes on, perceptual information processing is influenced not only the environmental stimuli, but by both activity in the reactive level and at the algorithmic level. The same goes for effectors. Therefore, perceptual information helps orient the system’s behaviour and the system’s gathering of perceptual information is oriented by the its deliberative level (its current plan). The selection of a plan itself is influenced by the information present in the environment. The plan thus depends on perception, but also influences what the system perceives (and of course does): perception and plan selection are dynamically coupled. The hybrid adaptive behaviour of the architecture is achieved through the cooperation of all three levels. Algorithmic cognition allows the system to achieve trial-and-error adaptation and hypothesis testing (e.g. simulation). Although achieved by the algorithmic level, decoupling is launched by the Reflective level from which will emerge the action to be performed by the system. However, achievement of a plan doesn’t cancel the system’s reactivity and dynamicity, since the algorithmic and reflective level stay up to date concerning the internal state of system at the reactive level, and are able to perform operations to adapt its long term behaviour to the constantly evolving environment. Dynamic behaviour emerges from the competition and cooperation between sequential deliberative processes (reflexive and algorithmic level) and reactive processes. In future work, we plan on addressing another issue which we think is also related to problem of flexibly integrating reactive and sequential processes: the emotional modulation of cognitive processes.

6 References

[1] Sun, R., (2004). Desiderata for cognitive architectures. Philosophical Psychology, Vol.17, No.3, pp.341- 373. 2004

[2] Langley, P., Laird, J. E., et Rogers, S. (2009). Cognitive architectures: Research issues and challenges. Cognitive Systems Research, 10(2), 141-160. doi: 10.1016/j.cogsys.2006.07.004

[3] Stanovich, K. E. (2010). Rationality and the reflective mind: Oxford University Press.

[4] Newell, A. (1990). Unified theories of cognition. Cambridge, Mass.: Harvard University Press.

[5] Anderson, J. R., et Lebiere, C. (1998). The atomic components of thought. Mahwah: Erlbaum

[6] Sun, R., Lane, S. M. & Mathews, R. C. (2009). The two systems of learning: An architectural perspective. In Jonathan Evans and Keith Frankish (eds.) *Two Minds: Dual Processes and Beyond (pp.239-264).* New York: OUP.

[7] Sloman, A., et Logan, B. (1999). Building cognitively rich agents using the Sim agent toolkit. Communications of the Association of Computing Machinery, 42, 71--77.

[8] Evans, J.., Newstead, S., Byrne, R. (1993), Human reasoning, Lawrence Erlbaum Associ- ates

[9] Frankish, K., et Evans, J. S. B. T. (2009). The duality of mind: a historical perspective. In K. Frankish & J. S. Evans (Eds.), In Two Minds: Dual Processes and Beyond

[10] Ferber, J., Gutknecht, O., et Michel, F. (2003). From Agents to Organizations : an Organizational View of MultiAgent Systems. Paper presented at the Agent Oriented Software Engineering (AOSE) IV, Melbourne.

[11] Anderson, J. R. (1983). The architecture of cognition. Cambridge, Mass.: Harvard University Press.

[12] Cardon, A. (2005). La complexité organisée : systèmes adaptatifs et champ organisationnel. Paris: Hermès Science .

[13] Engle, R. W. (2010). Role of Working Memory capacity in Cognitive Control. Current Anthropology, 51.

[14] Stroop, J.R. (1935). Studies of interference in serial verbal reactions, JEP 18 task, Can. J. Psych. 36 684–700.

[15] Dehaene, S., & Changeux, J. (1991). WCST: Theoretical analysis and modeling. Cerebral Cortex, 1, 62–79.

[17] Grant, D. A., et Berg, E. A. (1948). A behavioural analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. Journal of Experimental Psychology, 38, 404–411.

[16] Nagahama, Y., Okina, T., Suzuki, N., Nabatame, H., et Matsuda, M. (2005). The cerebral correlates of different types of perseveration in the WCST.J Neuro Psych, 76(2).