CASE FOR SUPPORT:

Computational Modeling of Salience Sensitive Control in Humans and in Artificial Systems

H. BOWMAN, C.G. JOHNSON
(UNIVERSITY OF KENT AT CANTERBURY)
AND
P.J. BARNARD
(MRC – COGNITION AND BRAIN SCIENCES UNIT, CAMBRIDGE)

1 Previous Research Track Record

1.1 The Computing Laboratory, University of Kent at Canterbury (UKC)

UKC contains a diversity of high quality computer science research with groups dedicated to, Networks and Distributed Systems, Software and Systems Engineering, Concurrency, Applied and Interdisciplinary Informatics (AII) and Theoretical Computer Science. Of particular relevance to this proposal are the latter two groups. For example, the AII group contains a large body of research that sits at the junction of computer science and biology. Ongoing research is both focused on importing biological metaphors into computer science, in order to improve computational systems (e.g. applying principles from immunology to data mining) and using computational metaphors in order to explain biological phenomena (e.g. computational simulations of the early stages of cancer formation). The research of this group is supported by funding both from research councils and from industry. For example, the project *"Data mining with antibodies: Artificial Immune Systems for Document Classification and Intelligent Navigation"* is jointly funded by the EPSRC and Sun Microsystems (case award 00802406). The group boasts considerable expertise in neural network modeling, machine learning, adaptive systems in general and self-organising systems in particular. There is also a line of research that is investigating the potential of affective computing in the construction of virtual environments and other computer interfaces [1]. This was reflected in a workshop on the theme recently organized by a member of the group [2].

In addition, the theoretical computer science group is a centre of excellence for formal methods research. Activity in this area has been focused on the application and further enhancement of formal methods technology. With application of these techniques having been made in diverse application areas, such as, distributed systems, safety analysis, testing and Human Computer Interaction. Much of this research has been supported by EPSRC, e.g. the projects *Cross Viewpoint Consistency in ODP* (GR/K13035), *Open Views* (GR/L95878) and *Verification of QOS* (GR/L28890). Furthermore, notations and tools from across the spectrum of formal methods technology are being actively worked on, e.g. Z, Object-Z, Timed Automata, Temporal Logic and, of most relevance to this proposal, process algebra. UKC has considerable expertise in process algebra based specification and verification using the LOTOS notation. This has included work on modeling cognition in HCI using LOTOS [3]; specifying distributed object-oriented systems [4]; subtyping and refinement relations in LOTOS [5] and verification of real-time systems [6]. A large amount of this research has been performed using the CADP verification environment, which is one of the most powerful tool suites available, boasting a spectrum of specification and verification capabilities, e.g. simulators, equivalence and refinement checkers and a model checker. This tool will play a central role in the proposed project.

There is also a considerable body of research at UKC that is directly focused on computational modeling of cognition. A major thread of this work has been towards Human Computer Interaction, focusing on the construction of simulated cognitive architectures in order to perform usability analysis of proposed interfaces, e.g. [3] [7] [8]. This research has been undertaken in collaboration with CNR-Istituto CNUCE in Pisa and was funded by research grants from the European Commission under the ERCIM (contract no. 97-03) and Marie-Curie (grant number Nr ERB4001GT965661) schemes. Another major thread of UKC's cognition related research has been focused on computational modeling of the time course of human attentional processes. This has been performed in collaboration with the MRC's Cognition and Brain Sciences Unit (CBU) in Cambridge, which has supported this research by providing Bowman with a visiting scientist position over an extended period. Output from this research activity has been a computational explanation of one of the key psychological phenomenon in the area – the attentional blink [9]. The proposed project will build directly from this model. A further thread of research has undertaken neural network modeling of perception to action coupling in collaboration with experimental psychology at Cambridge, psychology at Birkbeck College and Psychology at Warwick [10].

1.1.1 Howard Bowman

Howard Bowman is a senior lecturer in the Computing Laboratory at UKC. He received his Ph.D. from Lancaster University in 1991 and then completed a SERC postdoctoral research associateship at Lancaster working on temporal logic based formal methods. Since joining UKC, in addition to continuing his work on formal methods and logic, he has also broadened his interests to computational modeling of cognition. In particular, using his formal methods background, he has developed "specification-level" models of cognitive architectures, e.g. [3] [7] [8] and of cognitive phenomena, e.g. the attentional blink [9] [11]. He has also become interested in connectionism and has developed a neural network model of subliminal priming [10]. Much of his work in this area has been performed in collaboration with the Medical Research Council's Cognition and Brain Sciences Unit (CBU) in Cambridge, where he undertook a six month sabbatical in 2000 and where he remains a visiting scientist.

In the last five years he has run a number of formal methods research projects, three of which were funded by the EPSRC and two were funded by British Telecom. The most recent of these projects to finish (GR/L28890) was very highly rated when it went through the EPSRC post project assessment, receiving feedback such as,

"Exactly what one would expect from the given team: professional and thorough, with excellent dissemination."

Howard Bowman has also been on the programme committees of six different conference and workshop series in the last five years, including the IFIP sponsored FORTE/PSTV, and he is on the editorial board of the journal New Generation Computing. He was the conference chair for FMOODS'97 and FM-Elsewhere'00.

He has published over 70 papers in his career. In particular, in the last five years, he has published 13 journal papers, two book chapters and 29 conference papers. Included in this are invited papers at conferences (FMOODS'99) and in journals (New Generation Computing). In the same time he has also edited two books and one journal special issue. He is co-author of a book on specifying distributed multimedia systems using formal methods and is a member of IFIP TC6 Working Group 6.1.

1.1.2 Colin Johnson

Colin Johnson has been a lecturer at UKC since 1999, having previously worked at Exeter University and Napier University in Edinburgh. His main research interests are in computational and mathematical models of biological systems, and in using ideas from biology to inspire novel algorithms. Of particular relevance to this proposal has been his interest in the newly-emerging field of affective computing, which considers the role that emotions play in human-computer interaction and ways in which we can understand and simulate emotion on the computer. Within this area he has published on the role of affective states in virtual environments [1] [12] and was the organizer of a major affective computing workshop [2] at the 2001 AISB convention in York. Other recent projects have involved working with biomedical scientists at UKC and Cambridge on the development of a model of the molecular processes that underlie cancer, and work on program code induction using genetic programming, in particular tying this in with formal methods of program verification.

1.2 The Medical Research Council's Cognition & Brain Science Unit, Cambridge (CBU)

The MRC Cognition and Brian Sciences Unit is a major centre for UK research in Human Cognitive Processes, their realisation in neural mechanisms and brain systems, as well as computational modeling of these mechanisms and processes. The unit offers expertise in cognitive theory, computational modeling and neuro-imaging and is organised into four large interdisciplinary teams: Memory; Language; Attention; and Cognition and Emotion. Each group contains around 20 staff and graduate students. Their research programmes specialise in different domains but with substantial overlap and interchange among them. Since it was founded (originally as the MRC Applied Psychology Unit), the Unit has maintained an international reputation for outstanding theoretical, empirical and applied research. One of the current initiatives involves the development of advanced computational modeling techniques for cognitive and affective neuroscience.

1.2.1 Phil Barnard

Philip Barnard is a senior member of the research staff at the MRC Cognition and Brain Sciences Unit. He graduated from University College London in 1969, and also completed his PhD there. He joined the MRC Applied Psychology Unit in 1972 where he worked on language, memory, human-computer interaction, and affective disorders until the Unit changed its name and remit in 1998. His most recent research has concentrated on human executive processes and the effects of emotions upon them. He proposed the Interacting Cognitive Subsystems mental architecture and has explored its properties and predictions in core theoretical areas as well as its application to the human use of technologies (e.g. [11] [13]) and to clinical psychology (e.g. [14]). He has collaborated extensively with both clinicians and computer scientists. He coordinated two major ESPRIT Basic Research Actions on human-computer interaction: AMODEUS-3066 and AMODEUS-2 7044. In the last two years he has collaborated with Howard Bowman on the application of formal methods to modeling empirical phenomena linked to cognition and emotion (e.g. [11] [9]).

1. Johnson, C.G. & G. Jones. *Effecting affective communication in virtual environments*. in *Proc of the Workshop on Intelligent Virtual Agents, VR-SIG* 1999.
2. Johnson, C.G., ed. *Proc. AISB'01 Symposium on Emotion, Cognition and Affective Computing, University of York*. 2001.
3. Bowman, H. & G. Faconti, *Analysing Cognitive Behaviour using LOTOS and Mexitl.* Formal Aspects of Computing, 1999. **11**: p. 132-159.
4. Bowman, H., *A LOTOS based tutorial on formal methods for object oriented distributed systems.* New Generation Computing, 1998. **16**(4): p. 343-372.
5. Bowman, H. & J. Derrick. *A junction between state based and behavioural specification (invited paper)*. in *Formal Methods for Open Object-based Distributed Systems*. 1999. Kluwer, p. 213-239.
6. Bowman, H. & J. Derrick. *Extending LOTOS with time: a true concurrency perspective*. in *ARTS'97, AMAST Workshop on Real-Time Systems, Concurrent and Distributed Software*. 1997. Springer-Verlag, LNCS 1231, p. 382-399..
7. Bowman, H. *Modelling concurrent cognitive architectures using process calculi*. in *ECCS'99, European Conf on Cog Sc*.
8. Bowman, H., G. Faconti, & M. Massink. *Towards integrated cognitive and interface analysis*. in *Formal Methods Elsewhere 2000. Electronic Notes in Theoretical Computer Science, H. Bowman (ed).* 2000. Elsevier Science.
9. Bowman, H. & P.J. Barnard, *Computational Modeling of Distributed Executive Control*. 2001, Comp Lab, 12-01, UKC.
10. Bowman, H., et al. *Neural Network Modeling of Inhibition in Visuo-Motor Control*. in *7th Neural Computation and Psychology Workshop on Connectionist Models of Cognition and Perception.* 2001. World Scientific.
11. Barnard, P.J. & H. Bowman, *Rendering Information Processing Models of Cognition and Affect Computationally Explicit: Distributed Executive Control and the Deployment of Attention.* Under Review, 2002.
12. Johnson, C.G. *Affectively tunable environments for the virtual stage*. in *Digital Scenography Conference*. 2000.
13. Duke, D.J., Barnard, P.J. et al., *Syndetic Modelling.* Human-Computer Interaction, 1998. **13**(4): p. 337-393.
14. Teasdale, J. & P.J. Barnard, *Affect, Cognition and Change: re-modelling depressive thought.* 1993, Lawrence Erlbaum.

2 Description of Proposed Research and its Context

Humans are very good at prioritising competing processing demands. In particular, perception of a salient environmental event can interrupt ongoing processing, causing attention, and accompanying processing resources, to be redirected to the new event. A classic example of this is the well-known Cocktail Party Effect. Not only are we easily able to follow just one conversation when several people are speaking, but the occurrence of a salient phrase in a peripheral conversation stream, such as somebody mentioning our name, causes auditory attention to be redirected. It is also clear that emotions, motivation and physiological state in general, play a key role in such prioritisation, e.g. Oatley and Johnson-Laird [15] suggest,

> "… that the function of emotion modes are both to enable one priority to be exchanged for another … and to maintain this priority until it is satisfied or abandoned."

However, in an agent with multiple goals (such as a human) and which is subject to continual environmental stimulus, a compromise needs to be struck between responding optimally according to the order of priorities and efficiency of processing. In the extreme, a system in which interruption is the norm could fail to ever complete any valuable processing. The heart of the conflict is that between the need to respond in a timely fashion and the need to respond optimally according to the salience level of environmental stimuli. The problem is also complicated by the fact that salience is highly context dependent. Hearing a lion roar may be extremely salient if you are walking down the street, but it would be much less salient if you were walking around a zoo. Human's capacity to correctly attribute salience to stimuli in a context dependent manner and interrupt or adjust ongoing processing accordingly is a major reason for our evolutionary success.

In contrast, artificial systems do less well. This manifests itself in two ways. Firstly, they are often bad at adjusting their processing to salient events, especially when assessing salience is context dependent. Thus, they may fail to respond appropriately to a salient event or at the other extreme, they may interrupt processing unnecessarily frequently in response to contextually low salience events. Secondly, when interacting with humans, artificial systems fail to fully utilise salience. In pursuit of a particular goal, interactive systems typically unreel sequences of (in effect) ballistic steps, only being receptive at specific breakpoints to a restricted set of anticipated cues. In contrast, a salience sensitive interface would adapt its behaviour according to the affective state of the user and would provide cues that are sensitive to user's salience schema. See, for example, the affective computing literature [16] for convincing justifications for this position.

A big hindrance to constructing systems that are sensitive in this respect was that it was not fully understood how humans adapted their behaviour according to salience. It was believed that they did it well, but the mechanisms, which achieved that successful outcome, were not well understood. However, through the combination of behavioural experimentation and the recent application of brain imaging, modern cognitive and neural sciences are starting to clarify the underlying mechanism. In particular, a number of experimental paradigms, which fall broadly within the study of human attention, have started to reveal how real-time constraints and sensitivity to salient events are reconciled in humans. Two such experimental paradigms are the *attentional blink* and *psychological refactory period*. Both explore the limits of salience driven control, by determining the constraints (in particular, the time frames) under which salient stimuli fail to be processed as a result of attention being directed at preceding stimuli. Recent work at the CBU has also begun to reveal how context and level of salience, in both a semantic and emotional sense, regulate these effects.

In order to realise the potential of this new understanding it is important that computational models are constructed. These will provide concrete realisations of the mechanisms being revealed and will also act as a bridge to the construction of artificial systems that are sensitive to salience. This is the aim of the proposed project, and it is one that we intend to fulfil through the following objectives.

1. We will construct computational models of the cognitive processes involved in human salience sensitive control. These processes will be modeled within the framework of the Interacting Cognitive Subsystems architecture [11] and will be constructed using formal methods. The latter of these choices has been made since we firstly wish to use our models as general-purpose specifications for artificial system construction and secondly because we feel formal specifications match the abstraction level required in cognitive modeling.
2. We will construct models that reflect the neural substrates that are likely to underlie salience sensitive control in humans. This will involve developing Neural Network (NN) models that are constrained by the biology (e.g. anatomical structure of the brain), brain imaging data and are consistent with the cognitive-level models constructed under objective 1. Relating cognitive and neural models will be one of the major technical issues tackled in this project.
3. Our modeling work will be guided by (and will also guide) the converging evidence now being made available by behavioural studies and brain mapping (both fMRI and EEG). This will be performed through our proven collaborative links with the Cognition and Emotion programme at the CBU.
4. We will inform the construction of artificial systems by using the suite of computational models as general-purpose specifications of such systems. We will do this in three ways, by (a) identifying principles from our models and using them to inform system consturction; (b) using our models to simulate human users in assessing the usability of HCI systems; and (c) investigating the direct construction of systems from our specifications.
5. Finally, we will feed the results of our research into fields of computing for which salience sensitive control is critical, such as affective computing, HCI in general and robotics.

2.1 Background

2.1.1 Attention and the Time Course of Salience Sensitive Control

We pay attention to information that matters, as a result of the cognitive task we are required to perform and as a function of that information's personal salience and of our motivational and emotional state. For example, anxious people preferentially pay attention to external threat [17] and the effectiveness with which humans interact with computers is modulated by the emotional qualities of the interface [18]. In all these domains the key questions concern the dynamic redeployment of attention over time. Empirical paradigms like the psychological refractory period, the Attentional Blink (AB) or task shifting, all illustrate restrictions on our ability to re-deploy attention.

Although the others will be used to inform our computational models, here we focus discussion on the Attentional Blink [19]. In the typical AB, letters are presented using Rapid Serial Visual Presentation (RSVP) at around ten items a second. One letter (T1) is presented in a distinct colour. It is the target whose identity must be reported. A second target (T2) follows after some number of intervening items. For example, the person may have to report whether the letter "X" was among list items that followed T1. Detection of T2 is impaired dependent upon the position in which T2 follows T1. Specifically, if T2 occurs immediately after T1, then its presence is accurately reported (so called, lag 1 sparing). Detection then declines and recovers to baseline when there is a lag of around a half of a second between T1 and T2. This is the basic attentional blink. The empirical literature and alternative theoretical accounts of the AB have all assumed that allocating attention to T1 leaves less attentional resources for processing T2.

As research on the blink has progressed, it is becoming clear that both the semantic and personal salience of items affect the allocation of attention. Indeed, similar blink effects are readily obtained when words are used as list items (and subjects are required to report items from a particular category, e.g. job words). There is also evidence of specific effects of affective variables. [20] reports differences in target detection in the AB paradigm for high and low anxious people. More dramatically, [21] has shown that the blink effect is markedly attenuated when the second target is an aversive word (such as, "rape" or "torture"). This suggests that perception of emotionally salient (T2) stimuli can interrupt the processing of T1 items, causing attentional resources to be redirected to the higher priority T2 item. There is also evidence that patients with damage to specific emotional centres in the brain (viz, unilateral damage to the left Amygdala) show no attenuated blink effect to aversive words [22]. The implication of which is that such brain regions play a central role in the pathway by which affect driven salience is assessed and interruption is initiated. In summary, the AB paradigm is revealing how humans redeploy attentional resources in order to process semantically, personally and emotionally salient stimuli and, moreover, it is clarifying the time course at which such mechanisms operate.

2.1.2 Concurrency, Distribution and Hierarchical Decomposition

It is widely accepted that any realistic architecture or computational model of cognition must, at some level, be *concurrent*. However there is now also a body of opinion that not only is concurrency required, but that distributed control is also essential. That is, cognition should be viewed as the behaviour that emerges from the evolution of and interaction amongst a set of *independently evolving* modules, none of which has access to a complete view of the state of the system. One reason for supporting a "distributed systems" approach is the now substantial evidence that brain damage selectively impairs particular processing capabilities, while leaving others intact. Even under circumstances where "executive" functions are impaired, complex behaviours can still be coordinated [23]. It is very difficult to see how such selective outcomes could arise from fully centralised executive processes.

As a reflection, there are now a number of cognitive architectures based on the distributed control hypothesis, many of which focus on interactions between motivation, affect and cognition. For example, the component process theory of emotion [24] explores interactions among five subsystems: cognitive, autonomic, motor, motivational and a monitor subsystem. The Interacting Cognitive Subsystems (ICS) architecture [14], which will provide a framework for our modeling work, views cognition in terms of the interaction between nine subsystems.

A further element that is essential for the sort of "distributed systems" modeling that we are advocating is that of hierarchy. In particular, a "flat" model, with one level of concurrently evolving modules, is not sufficient. This is because, the same arguments we have used concerning decentralised control at the top-level can be applied at any module level. This combination of distributed control and hierarchical decomposition is reflected in current theories. For example, the subsystems in ICS are themselves systematically decomposed into distinct components.

2.1.3 Computational Models

There are two dominant approaches to computational modeling of cognition: (i) AI approaches based on production systems (e.g. EPIC [25]) and (ii) connectionism [26]. The benefits and drawbacks of both these approaches have been widely aired and both have made profound contributions to our understanding of how mental representations and processes might be realised. We consider the suitability of each with regard to distributed control and hierarchy.

Concurrency can be generated in traditional AI approaches, by allowing multiple productions to fire on each cycle – effectively partitioning the working memory into functional sub-components. However, control remains centralised, being focussed on the working memory. Thus, if one accepts that the brain and the cognitive functions emerging from it are best viewed in terms of distributed control, structure re-aligning maps need to be postulated both when relating "upwards" from production systems architectures to high-level psychological theories and when relating "downwards" to low-level brain models. Although such an approach is certainly illuminating, we feel it is not optimal. An illustration of this is that it is not straightforward to relate data from brain damaged patients to production systems architectures.

Our position with regard to connectionism is less clear cut. For the biologically inspired modeling that we will perform, Neural Networks (NN) are without doubt the most appropriate technique and we will indeed use them. However, for the "coarse-grain" distributed control required of our cognitive-level modeling, we prefer to employ formal methods. There are two reasons for this, (i) the need for hierarchy and (ii) the level of abstraction of the resulting model. Firstly, a level of compositionality is obtained in connectionism through the interconnection of layers, each of which can be viewed as a module. However, there is only one level of composition, i.e., the primitive elements of NNs are neuron-like nodes, not neural networks. Hence it is not possible to nest interacting components within interacting components.

It is also revealing to note that almost all the uses of connectionism in cognitive modeling have been *specialised* in nature. That is, neural networks have been enormously effective in elucidating computational explanations of specific cognitive phenomena. However, in extrapolating from these specific phenomena to the big "architectural" picture, they have done less well. This is in no small part due to the fact that it is hard to construct large architectural models (of the kind required to address salience sensitive control) without hierarchical structuring. Our second reason for not using connectionism concerns abstraction. Modeling based on NNs is, in certain respects, very low-level in nature. In particular, one has to work hard to obtain what are "primitive" constructs and data structures in higher-level notations. However, for a large amount of modeling one would simply like to have high-level constructs at one's disposal.

2.1.4 Formal Methods and Process Algebra

The decision to construct our models using formal methods technology has arisen for two reasons. Firstly, as we have previously stated, we wish to use our models as general-purpose specifications for the construction of artificial systems. This objective suggests abstraction from implementation detail, which would be difficult to obtain with lower-level realisations. Secondly, as we have argued elsewhere [3], we believe non-prescriptive specification fits with the abstraction-level at which many cognitive theories are expressed. The majority of cognitive theories leave much unexplained since complete mechanistic interpretations of cognitive behaviour are unavailable. Thus, in order to construct working simulations large numbers of assumptions have to be made, leaving it unclear what aspect of the behaviour of the simulation corresponds to known cognitive behaviour and what arises from expediency. In the cognitive domain, this problem is called the *irrelevant specification problem* [27]. We believe that limiting such overprescription has much in common with limiting overspecification, which can be addressed using formal methods abstraction devices.

The combination of a desire, on the one hand, not to overspecify and, on the other, to directly reflect distribution of control, has led us to advocate process algebras [28]. These are a class of formal method that was developed to specify and analyse distributed computer systems. A process algebra specification contains a set of top-level modules (called processes in the computing literature) that interact by exchanging messages along channels. Furthermore, process algebra components can be arbitrarily nested within one another, i.e. they allow hierarchical description in the manner desired.

Previous work of ours supports our belief that such notations are relevant to modeling cognitive theories. For example, we have constructed a number of models of the ICS architecture using the process algebra LOTOS [3] [7] [8] and we have undertaken HCI usability analysis using these specifications.

2.1.5 Computational Modeling of Salience Sensitive Control

[11] reports a LOTOS-based computational model of the AB. To our knowledge, this is the only running simulation of the phenomenon. The simulation reproduces a large spectrum of the available data. The model has been constructed within the ICS framework. Accordingly, the distributed control hypothesis is employed, the simulation being constructed using independently evolving subsystems. Most significant of these are subsystems that implement different levels of salience assessment, the first – the *implicational subsystem* – performs a generic (and emotionally charged) assessment of salience, while the second – the *propositional subsystem* – builds a referentially specific representation, which is sufficient for report. An item in an RSVP sequence is only reported if it is found to be both implicationally and propositionally salient. The blink effect is obtained because attentional resources can only be located at one of these subsystems at a time. Thus, if assessing the full salience of a T1 item causes attention to be located at the propositional system, a T2 item can be missed at the implicational system. Such a two-stage model fits with theoretical accounts of the AB, e.g. [29].

Despite the fact that it reproduces a spectrum of AB data, the scope of the model is limited in a number of respects. Firstly, it is currently focussed on the task of reproducing AB data. However, as previously stated, there is other empirical data that it is important to reproduce, e.g. the psychological refactory period. In addition, the current model only implements a fraction of the ICS architecture. However, it is important to situate modeling work within larger architectural frameworks. This contrasts with a lot of the existing simulation work, which reproduces a particular cognitive phenomenon, but has limited applicability beyond it. Architectural frameworks respond to this by postulating a macro-theoretic configuration of cognition, which is applied consistently across the modeling of a spectrum of empirical phenomena. By expanding the scale of our ICS modeling we hope to obtain such a consistent macro-theoretic position.

Not only is it beneficial in a cognitive sense to work within a consistent architectural framework it is also beneficial for our further objective of informing the construction of artificial systems. In the absence of such an architectural influence the resulting artificial systems are unlikely to be suitably general-purpose and of sufficiently broad scope to be applicable across a range of salience sensitive control problems.

The final direction in which we wish to extend our modeling is in respect of neural substrates that could underlie our cognitive-level model. Such an investigation is important because it allows constraints, which arise from brain mapping data to be imposed. We also believe that identifying neural substrates will suggest (NN-based) implementation mechanisms that could underlie the construction of artificial systems.

2.1.6 Salience Sensitive Control and HCI

There are a number of respects in which our modeling work can inform the construction of artificial systems. We concentrate on one such area here – the construction of salience sensitive human-computer interfaces, e.g. interfaces that adapt their behaviour according to the emotional state of the user. Such *affective computing* [16] has a great potential to create "human-centred" systems; having been made possible by a number of emotion recognition technologies, e.g. recognition of facial expressions, voice intonation and galvanic skin response. Many applications of these technologies have been proposed, e.g. systems which learn user preferences or help humans (such as autistic children) to recognize emotions. A typical example of affective computing would be an intelligent tutoring system that modulates its tuition according to the student's emotional state, e.g. curious, fascinated, puzzled, frustrated or anxious. For example, the system might regulate demands on the student according to their level of anxiety; it might make subtle (affect related) changes to the interface in response to user frustration; or it might present emotionally sympathetic responses (e.g. via an avatar).

However, in order to reap the full benefits of affective computing, not only is it important to understand emotions, it is also critical to understand the executive processes (in particular attention) within which emotions function. For example, knowing how our (affective) tutoring system should respond to the anxiety level of the student is dependent upon how anxiety modulates human attention. In addition, knowing when to present emotionally sympathetic responses requires an understanding of how emotional stimuli (e.g. emotional expressions) modulate attention and the time frames over which this modulation functions. These are exactly the kinds of questions that are being answered by empirical studies such as the AB.

As previously indicated, the issue of context is central to our investigation, as can be seen in HCI domains where situational awareness is critical. For example, pilots need to rapidly create an accurate representation of their current situation, by detecting, integrating and interpreting data gathered from a noisy environment. However, for reasons such as task overload, fatigue etc, human operators often miss salient events and thus build an incorrect picture of the situational context in which they are operating. A major goal of HCI research is to build devices that help humans to attend to salient events, to correctly prioritise environmental stimuli and thus to operate with situational awareness. Once again, in order to construct such systems we need to understand how attentional resources are deployed, the time-course of this deployment and how priority levels are assigned to competing stimuli. This is the central objective of our research programme.

Thus, we have argued that many of the goals of HCI research require us to understand human executive processes. Such an understanding will greatly inform interface construction. However it also prompts the question of mechanisms that should be used to construct such systems. Techniques such as user modeling and task analysis have made important contributions to interface construction [30]. However, it is unlikely that they will be sufficient in a domain in which dynamic sensitivity to salience and timing are critical. We believe our computational modeling activities can also help here by providing an abstract specification of user behaviour, which can be placed at the centre of interface usability analysis.

2.2 Programme and Methodology

The aim of the proposed research is, firstly, to further our understanding of how salience drives the time-course of attention deployment, and secondly, to use this understanding to inform the construction of artificial systems. The measurable objectives that we will pursue were highlighted as five points on page 3. We will attain these objectives through the construction of cognitive-level and brain-level models of attentional paradigms. In addition, we believe these paradigms are archetypal of salience sensitive control and thus being able to model them will give strong evidence that the models developed will be broadly applicable. In order to place our work into a consistent macro-theoretic framework, we will work within the context of the ICS architecture. Furthermore, not only will our modeling be responsive to empirical studies but, through our proven collaborative links with the cognition and emotion programme at the CBU, it will guide ongoing empirical studies. Such synergy between modeling and empirical studies is a major strength of our proposal.

We believe these models can serve as general-purpose specifications of salience sensitive control that are applicable across a spectrum of contexts. Thus, we will also explore how the specifications derived from humans can inform the construction of artificial systems such as intelligent agents, avatars in HCI and robotics. The proposed project is both timely and novel. It is timely because of recent breakthroughs in understanding human attention (particularly via brain imaging) and since such an understanding is central to a number of current computer science research initiatives, e.g. affective computing. The project is clearly novel – there has been no previous effort to unite human attention research and construction of intelligent systems in the manner proposed. The following programme of work will be undertaken, where work items 1, 2 and 3 address objectives 1, 2 and 4 (respectively) from page 3. Objective 3 will be addressed through our collaborator, the CBU, at no expense to the EPSRC. Objective 5 will be addressed by a 4th work item, see page 9.

2.2.1 Work Item 1: Cognitive-level Modeling

The simulations presented in [11] [3] [7] will act as prototypes, which we will build upon as follows,

1. The prototype simulations are limited in their modeling of affect. Stimuli are adjudged to be salient on purely semantic grounds, i.e. to what extent they match a target semantic category. Thus, we need to develop mechanisms by which the

affective qualities of stimuli can be assessed and can regulate the deployment of attention. In ICS, attribution of emotional salience arises from interaction of the implicational and body-state subsystems. Furthermore, such interaction can regulate the speed with which the buffer (ICS' metaphor for attention) moves between subsystems. We will enhance our prototype models in respect of these mechanisms in order to reproduce the affective AB data.

2. As previously motivated we will expand the architectural influences on our model. This will be done by developing specifications of a number of further ICS subsystems, i.e. the body-state, object and morphonolexical subsystems. If resources permit, we will also provide implementations of the visual and auditory subsystems.
3. Time course is vital to the computational modeling we are undertaking. Central to which are data structures called delay-lines. These allow integration over time, i.e. data from multiple time instants to be used in building new representations. This use of delay-lines has been very successful, however, there are many issues that still need to be considered, e.g. we need to improve the semantic representation of items in delay-lines.
4. There is data available, which shows that if T1 is semantically related to T2, i.e. it is close to the T2 target category, the blink is deeper. The prototype model produces this effect in a rather basic manner. It is our intention to develop a more refined solution that regulates salience assessment according to semantic distances between concepts.

As previously justified, we will employ process algebra-based formal methods in constructing our cognitive-level computational models. Our specific choice will be the process algebra LOTOS via the CADP tool suite [31]. This choice is justified by our previous work, which has successfully used this combination, e.g. [11].

2.2.2 Work Item 2: Brain-level Modeling

We will also construct computational models of salience sensitive control at the neural-level. These models will be constrained by the biology (e.g. anatomical configuration of the brain), brain mapping data (both fMRI and EEG) and our cognitive-level models. This is a more speculative work item than the previous one and producing implementations that are consistent with all the available levels of constraint is likely to be difficult. However, we believe the potential benefits of this activity are great and warrant the risk. In particular, there is already tantalising evidence, which suggests this activity can succeed. For example, brain structures called synfire chains [32] seem to have much in common with our delay-lines. They are chains of synchonously firing layers of neurons, in which distributed representations are bound via synchrony and stages in the chain correspond to time points. Using principles such as these, we will construct NN models of salience sensitive control. In this activity we will use the SNNS simulator [33], with which we have considerable experience.

Clearly one of the big technical difficulties will be ensuring that the cognitive and brain-level models are consistent. Informally relating the two levels is certain to be useful. However, the benefits of a more formal approach could be even greater. One could, for example, use the process algebra to enforce a hierarchical module structure over the NN. The feasibility of such a combination will be investigated here. Related previous work, e.g. [34], will inform this activity.

2.2.3 Work Item 3: Implications for Construction of Artificial Systems

This work item will extrapolate from our computational models to the construction of artificial systems. It is beyond the scope of the project to implement working prototypes, however, we believe we can inform artificial system construction in terms of specification and abstract architecture, in three ways, by,

1. identifying principles from our models and using them to inform system construction;
2. using our model to simulate human users in assessing the usability of HCI systems; and
3. directly constructing systems according to our specifications.

We believe the first two of these are very feasible undertakings, while the third is more speculative. In particular, we have a good deal of experience of the second level [3] [7]. In fact, there has been a considerable amount of work on using formally specified descriptions of human users in assessing interfaces, e.g. syndetic modeling [13]. A further example is [3], which verifies whether a particular class of multi-modal interaction is possible using a LOTOS specification of ICS. We propose to undertake similar usability verifications using the computational models constructed here.

The final level of activity that we will undertake is the most speculative, but perhaps also of the greatest potential. Specifically our models will realise a particular architecture of how humans undertake salience sensitive control. There will be many aspects to this architecture: division into subsystems and descriptions of how these subsystems behave, use of delay-lines in order to integrate over time etc. All of these could potentially be valuable principles in the engineering of artificial systems. We will investigate how a number of classes of artificial systems can be specified using these principles. Example systems will be taken from the affective computing literature, e.g. an affective tutoring system.

2.3 Relevance to Beneficiaries

This project will benefit both the cognitive and computer science communities. It will benefit cognitive and brain science through an increased understanding of how human's perform salience sensitive control and it will benefit computer science by informing the construction of intelligent artificial systems. The project is fundamentally cross disciplinary, new techniques from computer science will be fed into cognitive science and new theories from cognitive science will be fed into computer science. Specifically, formal methods technology is an important recent development in computer science, which has not yet made a contribution to cognitive science. Through the application of these techniques the cognitive science community will benefit from the provision of running *abstract* models of attentional phenomena.

In addition, modern cognitive science theories and empirical findings on the time course of attention and the influence of affective stimuli, has not yet made a contribution in computing. In particular, it is important that human cognitive processes of the kind to be considered here are understood in order to further the development of HCI systems. The proposed project offers great potential for development of advanced applications that adapt their behaviour according to anticipated user demands, e.g. which anticipate what will become salient in an interface and emphasize it (interfaces for disabled users being a notable application area). Such anticipatory mechanisms can also bring significant performance benefits. In particular, many classes of application that are currently constrained by the availability of processing resources could benefit from predictions of future user demands. A canonical example would be to use the results of our model to direct virtual environment rendering resources towards objects that are about to become salient. In addition, the project offers the potential to predict user's Quality of Service needs, which could be used to inform network resource reservation.

The increased understanding engendered by this project of salience sensitive control will benefit both academia and industry. For example, although working prototypes will not be constructed, the output of work item 3 will benefit the fields of HCI and robotics in respect of specification and architectural-level models of salience sensitive control, which could make an important contribution to UK industrial development in this area.

2.4 Dissemination and Exploitation

The results of this work will be disseminated by publication and also by direct interaction with the CBU, which is one of the major institutes in the world investigating human cognition and is thus the perfect vehicle for dissemination. In addition to the specific deliverables, which are listed in part 3, dissemination will also be made at a number of key forums spanning the project's lifetime. These will include the events, Interact, CogSci and Cognitive and Neural Systems. Dissemination will also be made through publication of articles in learned journals. The computational models constructed by the project and their associated documentation will be made freely available on the WWW. A number of important areas of collaboration will be furthered by the proposed research. These include the central interaction between UKC and the CBU on empirically and computationally inspired theories of salience sensitive control; collaboration between UKC and Pisa on HCI usability analysis; between UKC, Pisa and the CBU on ICS centred modeling; and finally, that between UKC and Inria Rhone-Alpes on use of the CADP tools. Although it is unlikely that there will be widespread short-term commercial exploitation of the research, the availability of our computational models could have a pump priming effect on this area of UK industry. We believe that in the long-term, industrial exploitation of the results of the proposed project can be substantial.

2.5 Justification of Resources

Collaborators. Professor Barnard at the CBU has agreed to contribute 2 hours a week of his time to the project for its full duration. Phil Barnard is a distinguished expert in the field and thus offers considerable added benefit to the research output in addition to that directly supported by EPSRC. The total indirect contribution to the project from the CBU is £14600.

Project RA. Our proposal asks for support for a postdoctoral RA. It is reasonable to expect to find the appropriate mix of skills and experience of cognitive science and computational modeling in someone who has already completed a doctorate.

Support Services. We are also seeking a small amount of technical support from computer technicians to help procure, maintain, update and administer the project computer equipment. Support is also requested for 10% of an administrator for the provision of overall project management and disseminating project information via the internet and other media.

Equipment and Consumables. Simulating cognition is computationally demanding. Thus, we request a high-end PC and a laptop for working offsite, for demonstrations and formal presentations, for which we have budgeted £4503. We request contributions to computing equipment and consumables (maintenance and licenses) at a total cost of £6000.

Travel. We anticipate travel related to the grant in two ways. First, there are sites, which we expect to visit on a regular basis, particularly the CBU in Cambridge. Second, in disseminating the results of the research project we intend to submit papers to and to attend relevant international conferences and workshops in the field.

15. Oatley, K. & P. Johnson-Laird, *Towards a Cognitive Theory of Emotions.* Cognition & Emotion, 1987. **1**(1): p. 29-50.
16. Picard, R.W., *Affective Computing*. 1998, Cambridge, USA: MIT Press.
17. MacLeod, C, A. Mathews, & P. Tata, *Attentional bias in emotional disorders.* J. of Abnormal Psych, 1986. **95**: p. 15-20.
18. Walker, J.H., L. Sproull, & R. Subramasi. *Using a Human Face in an Interface*. in *CHI'94*, ACM.
19. Raymond, J. et al. *Temporary Suppression of Visual Processing in an RSVP Task: An AB.* J. of Exp. Psych: HPP, 1992. **18**(3): p. 849-860.
20. Holmes, A. & A. Richard. *Attentional bias to threat related material in anxiety: a resource allocation or a practice effect..* in *Cog Sec of BPS.* 1999.
21. Anderson, A.K., *Affective influences on attentional dynamics during perceptual encoding: investigations with the AB.* (Under Review), 2001.
22. Anderson, A.K. & E.A. Phelps, *Lesions of the human amygdala impair enhanced perception of emotionally salient events.* Nature, 01. **411**: p. 306-309.
23. Shallice, T., *From neuropsychology to mental structure*. 1988, Cambridge University Press.
24. Scherer, K., *Emotions as Episodes of Subsyst Synch Driven by Nonlinear Appraisal Processes*, in *Emotion, Develop & Self-Org*, 00, CUP. p. 70-99.
25. Kieras, D.E. and D.E. Meyer, *An overview of the EPIC architecture for cognition and performance with application to HCI.* HCI, 1997. **12**: p. 391-438.
26. Rumelhart, D.E., et al, *Parallel Distributed Processing, Volumes 1 and 2*. 1986, A Bradford Book, The MIT Press.
27. Newell, A., *Unified Theories of Cognition*. 1990, Cambridge, Massachusetts: Harvard University Press.
28. Milner, R., *Communication and Concurrency*. 1989: Prentice-Hall.
29. Chun, M.M. & M.C. Potter, *A Two-Stage Model for Multiple Target Detection in RSVP.* J. of Exp. Psych. HPP, 1995. **21**(1): p. 109-127.
30. Preece, J., et al., *Human-Computer Interaction*. 1994, Addison-Wesely.
31. Garavel, H., F. Lang, and R. Mateescu, *An overview of CADP 2001*. 2001, INRIA Technical Report TR-254.

32. Diesmann, M., et al, *Stable Propagation of Synchronous Spiking in Cortical Neural Networks.* Nature, 1999. **402**: p. 529-533.
33. SNNS_Team, *Stuttgart Neutral Network Simulator, http://www-ra.informatik.uni-tuebingen.de/SNNS/*. 2001.
34. Gibson., P., *A LOTOS-Based Approach to Neural network Specification*. 1994, Univ. of Stirling, Comp. Sc., TR-112.

3 Diagrammatic Project Plan

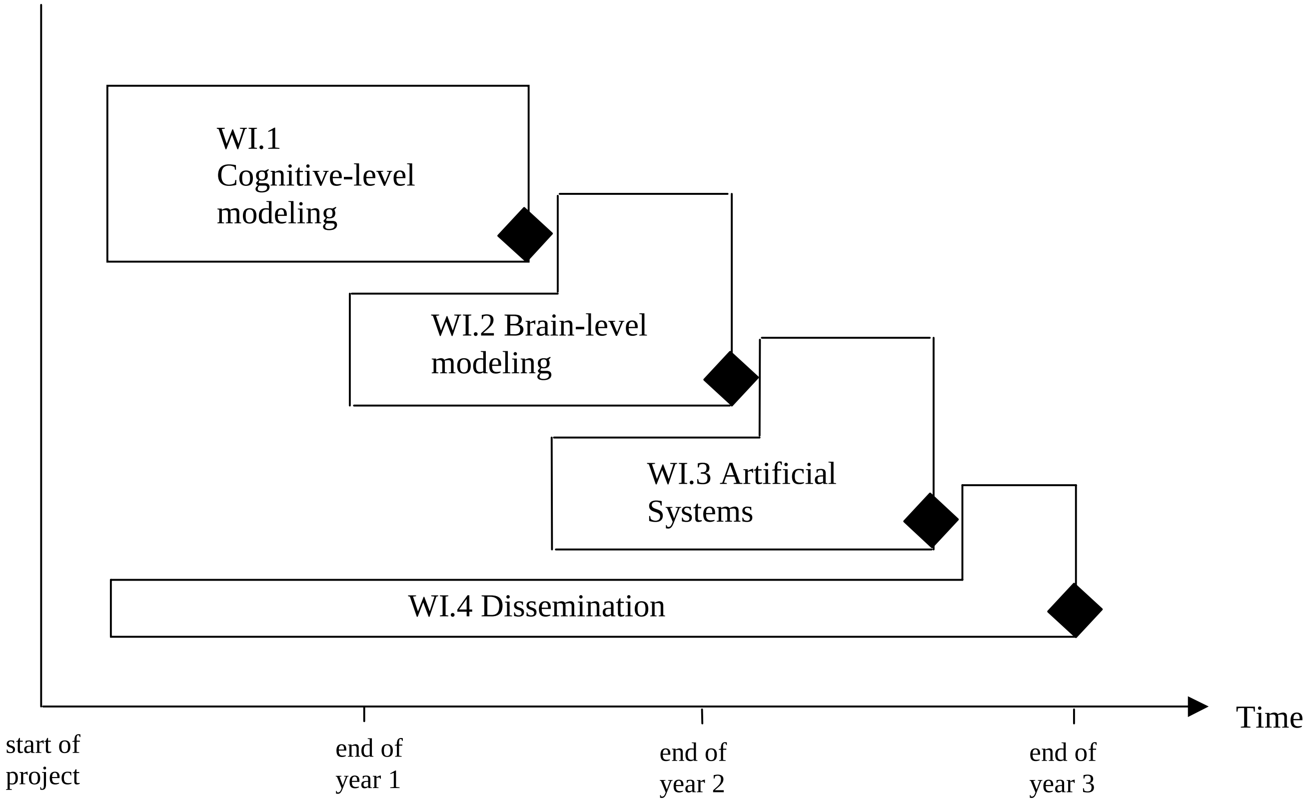

There is a milestone at the end of each work item (denoted by a black diamond) and the project will have the following deliverables.

Work item 1 – Cognitive-level Modeling

- deliverable 1.1 – suite of (CADP implemented) LOTOS specifications of subsystem behaviour and a set of associated global data type definitions;
- deliverable 1.2 – supporting documentation for the specifications; and
- deliverable 1.3 – report describing the cognitive-level model and relating the simulation results to human data.

Work item 2 – Brain-level Modeling

- deliverable 2.1 – neural network models of salience sensitive control, implemented in the Stuttgart Neural Network Simulator;
- deliverable 2.2 – supporting documentation describing the networks, their scope and how to interpret them; and
- deliverable 2.3 – report describing neural network simulations and relating the simulation results to human data.

Work item 3 – Implications for Construction of Artificial Systems

- deliverable 3.1 – report documenting principles arising from our modeling and how they inform construction of artificial systems;
- deliverable 3.2 – running simulations of human interface users and a report on HCI usability simulations; and
- deliverable 3.3 – Exploration of direct system construction. This is the most speculative element of our proposal and we do not wish to commit ourselves to any particular set of outcomes that may or may not be attainable. At the least we will provide a report documenting our exploration into this problem and at the most a prototype implementation of an exemplar system will be delivered.

Work item 4 – Dissemination and Exploitation

- deliverable 4.1 – final report to be made available to the CBU and other interested parties.