

The Interacting Cognitive Subsystems (ICS) project: background to, and evolution of a macro-theory of mental architecture.

With a list of key ICS papers and abstracts 1985-2023

Philip Barnard

What is ICS?

Interacting cognitive subsystems, or ICS, is a macro-theory of the how the components of our minds are arranged and function. It proposes that the human mind is organised into an “architecture” of nine subsystems. As a theory, it simply specifies all the “processing” resources that are in the architecture and how they interact, or communicate, one with another. This includes how they receive information from sense receptors and control somatic, visceral responses and skeletal musculatures. It also includes resources to support bodily, verbal, spatial and semantic imagery. It works at a high level of abstraction. Rather like, say a diagram of a railway system that involves nine “hubs” with two-way tracks interlinking them, with the hubs themselves decomposed into station platforms, sidings, maintenance sheds and offices, it shows how the stuff of cognitive-affective ideation passes around a complex network and what happens within each subsystem. The architecture itself does not specify all the minutiae of what happens in all circumstances, – although it does provide a platform for such detail to be added in a principled manner using evidence and task analysis. Importantly, each of the nine subsystems share an identical internal architecture. Rather like how cells in the human body differentiate to fulfil different roles, so over the course of cognitive evolution new subsystems emerged and were added to a simpler mammalian mental architecture to support our more advanced cognitive capabilities.

What can it be used for?

ICS was developed be a theory of broad scope that can be used to help think about issues that arise in broad range of practical settings – how people use instructions manuals; how they make and use tools; how they use information technology; how they create art, listen to music or watch movies, what happens in a range of mental health conditions and their treatments, what goes on in mindful contemplation, how we should reason about and study the mental capabilities of other animals and so on. It was developed to be a theory of broad scope and applicability. While primarily directed at practical applications, the architecture itself is grounded in the very substantial range of empirical evidence developed in the basic psychological and cognitive sciences over the last seven decades.

How can the ICS project contribute to studies of spirituality and contemplation?

There are four key topics for which ICS may provide some support for broader debates concerning religious or spiritual experience and any associated contemplative practices. These are its analyses of 1) Subjective experience/consciousness; (2) The support it provides for thinking about attention – to stuff out there in the world and to stuff in the mind itself (thought); (3) the presence in ICS of two levels of meaning – propositional and implicational; and (4) How thinking with language may support the interrogation and differentiation of implicational models of spirituality. The first three of these topics have already been quite

extensively explored and written about. Fraser Watts himself has focussed on the idea that implicational meanings provide support for thinking about spiritual beliefs and experiences. John Teasdale and I have thought, talked, and published on and around how attention to meanings in meditational techniques can support the remodelling of implicational meanings about the self, the world, and others. Using ICS there are now several substantial contributions to our understandings of contemplative practices and suggestions for future research.

Why this document?

For novice users of this kind of macro-theory, it must be acknowledged that it requires a significant degree of effort to learn about ICS and how to use it. First, the human mind is complicated and capable of supporting many patterns or streams of mental activity. Those patterns can shift moment to moment – we can after all walk, talk, chew gum and avoid bumping into an oncoming pedestrian while being excited an event we are walking to. Rather like the title of a 2022 movie “everything is going on everywhere (in the mind) all at once.” Second, the ICS project has focussed on a rather broad range of practical applications of psychology and the range and depth of the analyses it supports have themselves evolved over the last four decades. The outputs of the project have also been addressed to quite different communities of practice. When talking about theory to clinicians, technologists, computational theorists or artists we often need to call upon very different ways of describing specifying or graphically depicting ICS. This document outlines the history of the ICS project from my own personal perspective, and is intended to illustrate how many features of this diverse research trajectory could be of relevance to other projects.

The background

The appendix contains abstracts to some 60 key ICS papers on the theory which are available in my electronic archives. A few pieces of background knowledge, some of it personal, may help to frame an understanding of the rather unusual theoretical trajectory of the research on ICS as well as some of the key features of this rather diverse set of papers. I did my PhD research with Philip Johnson-Laird at University College London on memory for narrative and descriptive prose. I was interested in meaning and how memory for meaningful material was constrained by both by its structuring and by its content hence the title of my thesis was Structure and Content in the Retention of Prose. That long-standing core interest in meaning systems has remained with me ever since and over the last couple of decades that interest has deepened. When I arrived in Cambridge at the then MRC Applied Psychology Unit (APU) in 1972, I was tasked to work on and around issues associated with answering questions and understanding instructions about how to fill in forms or operate equipment – the topic was not research questions but quite literally questions you had to answer on government forms – driving licence applications, benefit forms etc. At the time the APU was a hotbed of theoretical discourse with allied concerns not just about “good” cognitive theory but about how theory and evidence could *usefully be applied* in real world contexts. I had two personal demons at the time.

First, what I had learned as an undergraduate and postgraduate gave me a firm grounding in laboratory theory and practice. Most of this I struggled to apply in the intricate real-world contexts I was investigating. Laboratory output was all very clever stuff, but it was mostly paradigm bound. This seemed to me to be too simple minded when clearly multiple factors were involved in most practical situations where natural language was used to real effect. I knew a lot about language, semantic memory and reasoning with affirmative and negative statements from my PhD mentors. I was well versed in theories of reaction times and errors

in judging the truth of statements like “Birds have wings” or “Canaries can fly” as well as deciding if pictures matched affirmative or negative contingencies like “*If* the envelope has a second class stamp on one side, *then* it must *not* be sealed.” A theory that accounted for 90% of the variance in a very restricted reaction time paradigm was, for example, deemed better than a theory that accounted for 84% of the variance. The most valued theories were often those that had given rise to a counter-intuitive prediction that was subsequently supported by evidence. What was valued in terms of theory in cognitive psychology was by no means of *that much value* outside of a laboratory, journal or conference. In sum, I did not know how to use what I had been taught and knew. I doubted that it was all that relevant to my practical concerns. Indeed, Donald Broadbent (a very celebrated theorist and Director of the APU when I arrived) once remarked to a group of us that, faced with a new applied problem, it was often quicker to get straight on and carry out some experiments without much recourse to reading relevant published literature. I also struggled to make paradigms on real practical issues deliver publishable data.

Second, the more senior people around me *did appear to be able to use what they knew very effectively*. I attended many meetings where senior colleagues like Donald Broadbent, Patricia Wright, Christopher Poulton, John Morton and Alan Baddeley, faced with some practical problem, were all able to develop all sorts important practical implications of their knowledge of basic psychology theory and evidence. I had not got a clue how they were doing it. During my first three-year research contract (1972-1975), I got my first taste of lived experience of anxiety and depression – in this instance about my own career and capabilities as an applied scientist. I was pretty good at masking it and I only discovered when became assistant director of the Unit around fifteen years later that quite a lot of my colleagues were equally affectively challenged and perhaps even better than I at camouflaging it. Some years before that revelation, I had myself recovered enough poise to figure out that my ‘elders and betters’ relied on their scientists’ *craft skills* in mapping theory and data onto problems in the real world. I myself gradually came to acquire a useful subset of those craft skills while at the same time became keen to expose this aspect of practice.

The culture of the then Applied Psychology Unit was very much framed in terms of how the examination of applied problems worked not just to the benefit of society but also to help the science of psychology evolve. The stimulus of new applied problems was a way of directing research towards novel theoretical problems and experimental paradigms. As many people know the stimulus of wartime or of emerging technologies in telecommunications led APU scientists to innovate substantive developments in, for example, attention and memory theory, the development of numerous forms of psychological tests, the first flight simulator for pilot training and much more. Within that culture, there was something of an intellectual divide. Some researchers focussed primarily on applications and rarely touched on or published evidence and theory in the core community of practice we call “basic psychology”. These researchers’ home community of practice was most closely allied to the discipline of Ergonomics/Human Factors. There were effectively three communities of practice. The third community of practice was typified by Broadbent, Baddeley and Morton. They moved relatively seamlessly between basic and applied research. Even this is a somewhat oversimplification since no two people every really saw the link between pure and applied research in exactly the same way. One feature of this research landscape fascinated me. Once applied research gave rise to a powerful new topic and/or theory that captured attention with the basic research community, it most typically transitioned into a standard laboratory paradigm for testing or validating competing theories. Research “guidance” went “paradigm internal” and basic rather than applied. Other than an occasional glance, they rarely took a

good look back in their rear-view mirrors at the applied origins of the topic. Often these paradigms gained a life of their own, with diminishing returns, which only expired when the respective theory owners retired. Why were theoretical developments *not* folded back and assessed for breadth and applicability? Why were theories themselves *not* evolved to better suit answering applied as well as basic questions? Clearly this is not a black and white issue. In the research world of today, for example, there are bodies of cognitive theories in clinical psychology that address the first of those two questions but not the second.

This experience as an early career researcher prompted a clear personal goal for theory development that framed many of the twists and turns in what actually happened in the subsequent development of ICS. That goal was to think about how we might *develop theories of broad scope and applicability*. Allied to that major goal were two further associated subgoals that followed from the remarks so far. First, to be useful the theory would need to be developed and evaluated with reference not just to laboratory phenomena. It would also need to be *constantly folded back to different real-world tasks and applications to assess whether or not it was fit for purpose*. The second associated subgoal was trickier. How can we address and overcome the *craft skill* aspects inherent in the challenges of applying theory in practice? It is perhaps apposite to note that simply producing an AI architecture is not in and of itself an answer because craft skills also underlie implementation and the way in which a particular architecture is seeded. It is well known in psychology that “theories do not make predictions, theorists do” (Allen Newell, William James lectures, 1987). But in other disciplines the elements of deduction are often supported by technique, they have formal proof schema – logic or maths or at least systematised semi-formal ways of exposing their reasoning as is evident in the presentation of legal judgements. In psychology, theorists could often only make predictions by adding a raft of additional (often implicit) assumptions about the tests being used and then if and only if a significant subset of members of that community of practice all agree that the prediction is grounded (e.g. referees). While psychology is a science and relies on scientific methods, the reasoning of its theorists is often pragmatic. The theories of psychology are tried by a mix of evidence *and* the equally pragmatic judgements of juries of peers.

Although these goals are, to say the least, intellectually challenging. Adopting such track was also challenging from the perspective of the sociology of science. In many domains of science there can be tensions between the more macroscopic and the more microscopic analysis – as today with tensions between General Relativity and Quantum Mechanics. Even though theories in psychology were far from as powerful and accurate, there were similar divisions. “Grand” theories of the sort proposed in the pre-war years by Hull were regarded as having failed to deliver and the post-war climate favoured more specific theories. To “rule” the mind it was considered sensible to divide the problems up into tractable chunks and study component mental faculties and interrelationships on a more local scale. This became the core research paradigm for cognitive psychology and was most valued. In the 1970’s, we had distinct theories of language, visual perception, auditory perception, attention, thinking and problem solving, knowledge representation, pre-categorical acoustic storage, iconic memory, short-term memory, long-term memory, autobiographical memory, semantic memory, learning, perceptual-motor skill and much else. Although research in artificial intelligence moved toward more “Unified Theories of Cognition,” they were rooted in problem solving and were really only unified theories of a subset of our actual mental capabilities with, among other topics, sensation, perception, bodies and emotions mostly absent or, at best, in peripheral intellectual vision. In the case of many of my contemporary colleagues who leaned more toward empirical psychology than AI, they felt more

comfortable in empirical paradigms and with theories of local scope. Attempts to move beyond the local to more ‘macro’ theoretical concerns were often described with words like “ambitious” – where the pragmatics of the use of that particular term in this context conveyed negative connotations associated with being over-ambitious. Whilst moving towards a more integrated perspective on cognition was accepted in AI, it was swimming against the sociological tide in most experimental psychology laboratories. While the advocates of local theory would acknowledge the argument that a mind was a mind for all seasons and all tasks, a more unified form of theory was a far-off long-term vision. In the here and now, they clearly favoured a research strategy that focussed on tractability and assumed that like a jigsaw puzzle the pieces would eventually somehow fit together to yield an image of a whole mind – magical thinking?

There were of course plenty of arguments in favour of both the “macro” and “micro” approaches to theory development and I was influenced by many of them. Of the less standard criticisms, I was often amused by the thought that if you assembled the many information processing models of the various cognitive faculties listed earlier that used some form of box and arrows notation you would land up with a wide a ranging mess. You would have hundreds of different boxes and arrows many of which overlapped in function and where what the boxes and arrows signified or meant was not the same between the different models. While the models may have been fit for purpose in the pragmatic processes of predicting or explaining phenomenon of local scope, they were decidedly not fit for purpose as ingredients of a more integrated theory of the whole mental architecture. There was no systematic vision for the cumulative development of a body of more integrated theory of broader scope. The computer metaphor and its allies within signal processing and information theory created elements of a common culture, as did the rules, production systems and knowledge representations of AI. However, these were more like dialects where speakers of one preferentially communicated with those who shared the same dialect, and this in turn inhibited more integrative theoretical thinking. There were also few constraints on the proliferation of theories within a domain. Indeed, the culture in psychology valued most highly independent theoretical contributions and marked down much work as derivative rather than original if that work simple piggy-backed off the work of others. Information processing models of bits and pieces of the mind proliferated in ways that made it difficult to see the wood for the trees. Even grasping the details of individual trees was a problem because the “contents” of many boxes or stages of processing were not always clearly specified and there was more than a little wriggle room for ambiguities of interpretation.

ICS was developed as a reaction to all of these concerns. The immediate stimulus was a discussion with my then intellectual mentor, John Morton, about the nature of Short-Term Memory. I argued that having “a box” labelled STM was maybe a bit of a red herring, since most of the material was verbal, we should be thinking more that STM was most likely a by-product of the collective requirements of language comprehension and production. John challenged me to produce a model that expressed those ideas. At around the same time I had been reading a working paper that John had written while on a period of sabbatical leave at the University of Michigan. The paper was called Consideration of Grammar and Computation in Language Behaviour (1968) and in it he used formal logic to specify key aspects of his influential model of aspects of language processing called the logogen model. This paper was never published and existed only as a progress report to the US Department of Health, Education and Welfare. It remains a mimeograph created ‘cognitive-archaeological artifact’ buried in the archives of the Cognition and Brain Sciences Unit. What interested me was that this version of the logogen model was articulated in formal logic with

rules (principles) and abstract properties of representations. It looked to me very much like what a “proper” theory should be striving to do in the sense that it went beyond loosely formulated theories that simply described in words what theorists thought the component boxes of their information processing models did. By using formal logic Morton was able to encapsulate the essence of the theory uncluttered by the type of implementational detail that was involved with AI simulations of the time.

My full response to Morton’s challenge took many, many months which resulted in a very lengthy working paper of my own. I assembled a list of as many short-term memory phenomena that were in vogue at the time and set an objective to find an architecture that would allow the evidence associated with all of them to be accommodated. Since this was to be an account of STM in terms of language processing I, of course, also had notes about a whole range of phenomena that informed our then understanding of language comprehension and production, cross-modal processing (e.g. reading, McGurk effects) and memory phenomena more widely. This was the era when neuropsychology was offering up a whole raft of evidence from deficits in the use of spoken and written language as well as memory that effectively put constraints on the structuring of the wider architecture for cognition. Issues associated with cross-modal phenomena (like reading) proved the most problematic. Although I only learned about requirements analyses in computer science several years later, I was effectively doing a requirements analysis based both on evidence and on what followed from the concerns for theory noted earlier.

The requirement for broad scope meant that I could not allow a proliferation of boxes while the requirement for explicitness meant that I needed to specify not only the stuff inside of the boxes but also to be explicit about the principles that governed the operation of that stuff. There were other concerns that also arose out of the requirement for broad scope. One was unevenness of coverage. So, for example, Morton’s logogen model was very explicit about the activation of lexical units but far less explicit about other aspects of cognition – these were assigned to boxes sometime labelled “general cognitive system” or “context system”. Likewise, Baddeley’s influential model of working memory was reasonably explicit about the phonological loop and visuo-spatial scratch pad components of his system, but the all of the clever stuff was assigned to a central executive component that was far less well specified. Although its functions were later clarified it remained, as Baddeley himself acknowledged, a bit of a rag bag entity not unlike the general cognitive system of Morton’s model. In both cases the larger fraction of mental functions was allocated to just one undifferentiated box while tiny fractions of mental functionality were given the prominence of their own boxes!




Figure 1: The eight-subsystem architecture and the internal arrangement of processes in 1985 [1].

An important lead for me at the time was the idea that different subsystems of mind worked on different “codes.” At the time, there was powerful evidence that short term memory relied on acoustic coding and that long-term memory relied on a more semantic form of encoding. After a number of specification cycles, the settled form of the first version of ICS was composed of eight subsystems (Figure 1a). Each subsystem was defined as organised around one *domain* of mental coding of information in which it (a) could represent over time in an image record (b) extract regular patterns that were then held in the image record, and (c) transform into one or more other mental codes that were then passed to other subsystems over a data network. There was no central executive, general cognitive system or any homunculus at all. There were just processes and exchanges between two meaning systems fulfilled what were considered “higher level” functions. Importantly, these two systems were specified in exactly the same way as all other subsystems. In the background, this provided economy of theoretical expression as well as being based (at the time) on a rather loose evolutionary argument in favour of a common internal architecture. The overall justification and accounts of evidence are laid out in the 1985 book chapter [1].

There was something of a lag in between model creation and its final emergence in published form in 1985. I was working at the time most extensively on Human-Computer Interaction (HCI), and through that exploring how to address more detailed issues associated with the development of a model of broad scope and applicability, including the tricky issue of the use of craft skills in applying psychology. In the fifteen years following the first publication of ICS, we continued the development of the theory by folding back onto more and more applications. In the case of HCI we started out with work on command language and menu dialogues [2] [5] followed by applications to visual/iconic interfaces [12] [14] [15] and finally multimodal ones [19] [20] [21] and even touched on affective influences on interactions [34]. We addressed the craft skill issues by first creating an expert system to create approximate models of cognitive activity [2] [3] [4] [7] [23]. These did not simulate or automate ICS. Rather the expert system provided an explicit and executable formalisation of how ICS resources were organised together with the rules and assumptions required to apply it. The reasoning from theory to application was laid out and automated in an AI production rule base and shown to have the capability to generalise to new problems. We were, in effect, being explicit both about the model and about the expertise needed to apply it. The approximate models described in a space of attributes of abstract properties of the workings of subsystems and their internal interactions and interactions with stuff in the world. We also developed tutorial material for human factors/user experience professionals to help them conduct a form of task analysis, namely cognitive task analysis, for how the mind would work on intricate interactions with advanced technologies and how to support design decision making based upon theory [13] [15] [10]. In the more usual forms of task analysis carried out by human factors specialists, the objective is to analyse a practical task, often a task using some form of technology, to determine what needs to be done, how it can be done, what you need to do it and so. Cognitive task analysis is essentially the same idea but applied to how the resources of the mind, including knowledge, are being used in a practical setting [2] [11] [23] [25] [32].

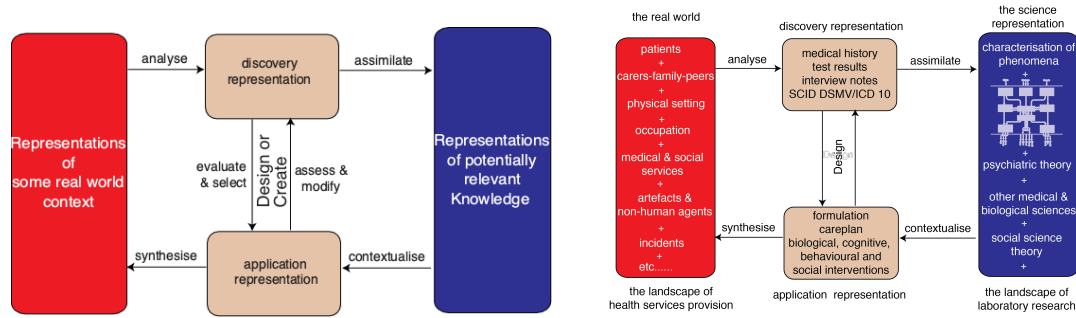


Figure 2 Bridging model – left hand side identifies generic craft skills of analysing real-world settings, assimilating the products of discovery methods, such as experiments, to knowledge; contextualising that knowledge for application to a new setting; and the synthesising all relevant concerns for use in some actual real-world context. The design and development of tools for discovery and application techniques sits in the centre and requires evaluation and selection processes plus identification of candidate modifications to a design. Right hand side shows an illustration for how this kind of model can be used to support dialogue about the use of theory in clinical psychology – from [36].

An additional strand of work looked in abstract terms at the use of theory in tackling applied problems and identified six kinds of thought processes that collectively constituted those craft skills. John Long, a long-standing collaborator and colleague, had developed a way of thinking about the discipline of human factors and how knowledge was deployed in different paradigmatic practices of applied science. I meddled with, and modified, his framework and renamed it *a bridging model* that I subsequently applied to craft skills in technology design [7], clinical psychology [36], and much later to creative domains such as drama, choreography [56] and cognitive archaeology. I have included in Figure 2 a generic form of this model alongside an example characterisation for clinical psychology [36]. The core notion is that discovery representations and applications representations bridge between the real world and the world of knowledge about that world. Creative processes of design work with those representations when creating and drawing upon on relevant knowledge. Inspection of the variant for clinical psychology will help ground how to think what is in ‘a box.’ The collective of six labelled arrows signify where the craft skills of practitioners come into play.

In the context of the projects directed at understanding contemplative or spiritual practices, for example, the “discovery” representation would include whatever means are selected to interrogate contemplative practices so that systematic theoretical accounts can be developed to be assimilated to the science representation and to support design cycles in specifying and creating a spiritual assistant. It is important to note that in any research endeavour there may well be more than one type of discovery representation or application representation. Different application representations would come into play for publications, pedagogy and the development of software for the cognitive modelling of spiritual or mediative practices [61].

A prerequisite for work on the expert system for HCI was the need for a schema or framework to support the systematic development of detail within ICS theory itself and, obviously, to support the specification of executable rules and the attributes they were to work with. In this research we really made a significant move from a core concern with box and arrow notations in information processing models, to a wider concern with how to reason about behaviours of systems. This framework considered the behaviour of any system to be a complex function of:

- (1) *the identity and organisation of the entities that interact within it (the configuration of interactors);*
- (2) *the capabilities of the individual interactors;*
- (3) *the requirements that must be met for the capability to be exhibited; and*
- (4) *a specification of how the system is dynamically controlled and co-ordinated.*

Military systems provide an easy example to follow. Large armies are usually made up of units – infantry, artillery, tanks, air support, logistics supply, field hospitals with doctors and nurses, and so on. If we describe the units as interactors (or subsystems that can themselves be decomposed into smaller units) then they will be deployed in particular places and able to communicate along particular paths (radio, courier, pigeon). The *interactors are configured to interact* in a specific arrangement (e.g. senior commanders give order to unit commanders who in turn issue orders to each soldier). We can talk about and define attributes of capabilities. Artillery may have hardware that enables them to fire projectiles over a specific number of miles, the further they fire the sooner they can engage an approaching threat. In order to fire them certain requirements must be met. They need soldiers to work them, ammunition to be available at the point of use and the co-ordinates of what they are firing at. The whole system is dynamically controlled and co-ordinated over time by a hierarchy of commanders and by rules of engagement. As a battle progresses all of these attributes are subject to change – forces move, run out of supplies, radio communication may be broken commanders may be injured and so on. Many factors and how they change over time affect what we might want to infer about system behaviour. Modelling a battle or an economy has a lot in common with modelling a mind. We once used the Battle of Cannae as an illustration of how a real military engagement could be shown to have similar underlying systemic principles governing interactions as the learning of a word processing system by novice users! [26]. It takes a lot of abstract properties to be specified to support an analysis or predict outcomes. Our AI production rules encapsulated in the expert system for analysing human-computer interactions were organised according this wider four-component schema for constraints on system behaviour and effectively automated the process of reasoning with ICS.

Earlier it was noted that, for the wider discipline, I could not see an engaging and systematic vision for the cumulative development of a body of more integrated theory of broader scope. This four-component framework was the response to that issue. It could be applied to any system level and was potentially able to support adding detail to theory in comparable ways across the entities that interacted within a system, subsystem or supra-system.

Why bring it up at this point in a discussion of the background to ICS research and development? Many in the community of psychological/cognitive scientists have understandable concerns that ICS is too complicated a system to work with and cannot figure out how to use the model to make predictions. The four-component system analysis is of value in helping understand what ICS is as a theory and what it isn't. The ICS architecture *is a comprehensive specification of the entities that interact in the mind and of how they can be configured to interact one with another and the external world.* The architecture also inherently provides a high-level specification of process *capabilities*. For example, one type of processes maps one mental code to another. The $ac \rightarrow mpl$ (morphonolexical) process transforms an acoustic code onto a phonological one and the $mpl \rightarrow prop$ process generates a referentially specific proposition from a phonological surface structure. ICS itself has a bunch of principles that constrain how that capability is exercised – they can only handle a single stream of information and the incoming information must have the right temporal properties

(not too fast or too slow). The image record enables pattern completion from past experience if the input is incomplete and so on. However, ICS is not a theory of any given language. Attributes of the mappings are a product of learning so the repertoire of mappings for understanding or producing Cantonese, in which pitch contour can alter the meaning of words, would be different for those of English where pitch contour does not change meanings in the same way. The formulation of the model allows for such variation in the properties of process capability that is broad reaching – there are many sources of such variation in psychology from different forms of knowledge, education, personality characteristics, age, brain damage, alcohol or drug intake, tiredness and more.

Such diverse variation means that ICS in and of itself cannot make predictions without additional assumptions laid out in a supporting four component specification outlined earlier. Another example is perhaps helpful for understanding that certain requirements need to be met for normal capability to be expressed. I am a protanope. This is a condition where I lack one of three normal retinal systems for encoding colour (the red one), so all the processes in my visual subsystem do not meet *the requirements for normal colour vision* and any down line processes will react differently from people with normal colour vision. My ability to discriminate colours is severely compromised particularly in low light. I used to make terrible fashion decisions, but now rely on a restricted range of colours and avoid multi-coloured patterns. In the case of dynamic control and co-ordination of process actions, *a key feature of ICS is that there is no central executive or homunculus*. Our HCI expert system had a first pass at set of attributes to characterise the dynamic control and co-ordination such as demands on memory, the extent of processing exchanges required to learn new mappings or resolve problems as well as a characterisation of the pattern of “buffering” over time.

The details of the four-component framework for system and cognitive task analysis itself appears in a number of rather technical papers on HCI achieved only through the sterling work of Jon May and Ann Blandford. It re-appears in a later consideration of how mentation transits between depressed, euthymic and manic states in Bipolar conditions [36] and this is reproduced in Figure 3. Notice that in papers on ICS *depressive interlock* is actually *a state of dynamic control and coordination of a sub-configuration of interactor capabilities* in the two meaning subsystems and Morphonolexical subsystem in ICS. We can hypothesise about the *capabilities* of the constituent processes and *requirements* that must be met for interlock to establish (negative self-models in the Implicational code mapping to propositions AND the inverse propositions mapping back to negative self-models; AND a preponderance of *dynamic control and co-ordination* with buffering in the verbal sub-configuration (usually at PROP) [44]. We need a wider attribute space for the whole ICS system when dealing with issues associated with transdiagnostic symptomatology. While rumination in depression is allied to low body state activation, worry in anxiety shares similar properties of exchanges within three central subsystems but differs in that worry is more likely to be associated with inability to control autonomic and physical arousal with consequences for what the Body-state subsystem feeds back to the Implicational subsystem. Likewise, those experiencing Anorexia Nervosa are also prone to negative ruminative thoughts. In such cases it can be argued that schematic models of self actually encapsulate high levels of agency that contrast markedly with low levels of agency more typical of low self-esteem in many other forms of depressed state [48]. Figure 3 shows property lists such as those just outlined with letters to index them to give a sense of four component parameter spaces for this type of explicit modelling process. Rules operate on the parameters to derive a characterisation of behaviour for a given segment in a phase. Do not worry about the detail, the point is that modelling intricate patterns of mentation, how those patterns change over time and prediction requires a

lot of parameters. The physical sciences are often dealing with a small number of parameters (think $E=MC^2$) and psychology is infused with that mentality. Information systems involve large parameter spaces and that is why the expressive power of the languages of computer science are both relevant and useful.




Figure 3. A family of interrelated models for different phases and segments of cognitive affective mentation in Bipolar conditions from [36].

In the more intellectual cognitive tasks, the dynamic control and co-ordination of the mental mechanism tends toward having particular properties. In the case of problem-solving or agonising about something one of which can be that there is a lot of dynamic reconfiguration in the central engine and a lot of switching between different streams of information. In mindful and other varieties of contemplative processing stream switching is likely to have very different properties as well as displaying a very different balance of reconfigurations centred on Implicational processing. The attentional score concept [57] is a notation that supports a description of dynamic control and co-ordination with reference to both reconfigurations and stream switching.

There are two further points to draw out from more than two decades of work on ICS in human-computer interaction. We worked collaboratively with a whole range of computer scientists and in the course of this work I encountered a very different community of practice who worked with what are known as ‘formal methods.’ This community created and studied models of complex systems. They used a whole range of logics and mathematical formalisms, including process algebras, not to simulate these systems, but to reason about the ways in which they would behave. Given a formal model they could assess whether or not a system design could, for example, runaway or deadlocked states that could lead to catastrophic failures. Such modelling is time consuming and expensive and was only done by a limited group of people in the UK mostly connected to the Programming Research Group founded at Oxford University in 1965. My own colleagues were heavily influenced by the work of Tony Hoare who ran the group at the time we were working on these issues. Modelling of this kind was generally directed at big issues like nuclear power stations or software systems in aircraft. My collaborators on a large EU project were doing such modelling and the question arose about whether we could not only produce a formal model of computers but also formal models of mental systems and a “supra-model” of the interactions between mental systems and computer systems. We did exactly that but only because very

clever computer scientists, mostly David Duke, did the work. Using deontic extensions of modal action logic, a set of axioms was produced for the workings of a particular multimodal computer system, a set for a particular cognitive model – ICS, and a set relating to their interaction. We called this approach Syndetic modelling from the Greek term *syndesis* meaning to bind together [16] [20] [21] [26] [28].

For myself, the development of Syndetic models was a major milestone because it did for ICS what Morton had done for the logogen model in 1968 but with many enhancements. These were models that enabled a theorist to test conjectures about behaviours dependent on minds using formal proof schemas that transcended craft skills and eliminated a role for implicit additional assumptions in mapping theory to application. All a theorist in this instance had to do was to imagine a conjecture about the possible behaviours of a system. They could then use logic to determine whether that outcome followed from the syndetic model. If the assumptions held, then the conclusion followed and you could even write QED at the end of a proof. Further, in principle you should not need a tame professor of computer science in your cupboard because proof checkers automated in software could in principle do the work. At the end of the 1990s I was of the view that this should form an important future methodological trajectory for cognitive science. Sadly, the odds of that trajectory being taken up more widely are less than minimal and it currently remains a blind alley. Psychologists and neuroscientists have neither the background interest nor the skills to pursue this form of methodology and the computer scientists who do are not going to accrue brownie points in their own community of practice by focussing on human rather than computer systems.

Although the bulk of the work on ICS through the 1980s and 1990s was supported by grants for research on HCI. I was still very much a part of the APU where I was integrated with basic as well as applied science. This enabled me to think about how the emerging modelling framework could be used in and supported by more basic laboratory paradigms. While I was not ‘a committed member’ of specific basic research communities I was thinking and researching basic paradigms in memory, language, attention, thinking and emotion. I was also in daily contact with researchers who were in those communities. Periodically opportunities arose for some of these strands to appear in written form and examples include working memory [22] [29], thinking [39] and attention. The work on attention once again engaged with formal methods in computer science and we were able to apply process algebra to produce a series of implementations of a glance look model of attentional control based on exchanges among the two meaning systems of ICS and the Body-state subsystem [33] [41] [42] [43]. The actual models were not wholly faithful to the internal organisation of processes within ICS subsystems, but they did make use of a delay line architecture which is my preferred conceptualisation of how a neural implementation of ICS would work [38]. This modelling was important because it showed that distributed control could be modelled at a more abstract level than connectionist networks. Most importantly, it was a very clear demonstration that the dynamic control and co-ordination of meaning systems could be accounted for without recourse to a homunculus or a central executive.




Figure 4. The nine Subsystem architecture after the Body-State subsystem was added in [6]. This shows a specific configuration with buffering (focal awareness) on a stream of information in the Propositional Image.

Our empirical work on HCI has not been covered because here the intention has been to show that the development of the Interacting Cognitive Subsystems, and the techniques for modelling it, were a fulfilment of the objectives outlined earlier. These were the development a model of broad scope and applicability by continually folding it back onto an ever-increasing range of practical topics whilst maintaining consistency with evidence from basic research and seeking to make any craft knowledge explicit in one way or another. The first major opportunity to extend applications of ICS arose in the late 1980's. John Teasdale moved from Oxford to the APU to join the Cognition and Emotion research group. He and I had a number of discussions about two meaning systems and the part they might play in depression. This led to further discussion, not only of depression, but also how ICS might help address evidence from basic research in mood and memory. The work on HCI meant that I had thought about the issue of how exchanges between the two meaning systems in ICS might go wrong – such as interlock (depression) and runaway states (mania, anxiety). I had also thought about and made notes on, how conditions such as variants of amnesia and frontal syndromes could result from compromised components of the two meaning systems. As a result of these initial exchanges we worked closely together and first produced a journal paper addressing issues about how cognitive therapy worked and applied the analysis to depression and mood congruent memory [6]. This was followed two years later by the book, *Cognition Affect and Change* [8] for which 95% of the credit should go to John. Both the journal paper and the book blended arguments from basic laboratory research in mood and memory with concerns with applications of theory in clinical psychology. The earliest version of ICS had eight subsystems and the ninth subsystem, the Body-state subsystem, was added so we could enrich the theory's approach to emotion. The augmentation of ICS was

not straightforward. In the initial forays more than one extra box was added for other senses (taste and smell) as well as a body states and somatic and visceral outputs. After a lot of messing about, it seemed sensible to have just one subsystem which relied on a single coding space that encompassed the body envelope with taste and smell thought of as “stuff” that primarily impacted on the body envelope whereas auditory and visual inputs were of significance distal to the body envelope.

Following on from the publication of the book John and I remained in constant dialogue (and still do) but he focussed on the application of the model to mindfulness and to the development of mindfulness based cognitive therapy [9] [24] [49] [50]. I was keen to widen its application and developed, with Fraser Watts, an account of sleep onset and insomnia [17]; with Rebecca Park an account of Anorexia Nervosa [48] alongside implications for its treatment [53]; accounts of the schizophrenic spectrum and autism [31], bipolar disorder [36] along with further theoretical thoughts on depression [44]. Jon May also used ICS to debate accounts of craving [37] and awareness [30]. Several empirical projects looked at bipolar disorder, borderline personality disorder, and developmental trajectories but with mixed success. Other empirical projects had more success in supporting or developing ideas about depressive ideation and these empirical projects are summarised in the Unit’s quinquennial progress reports. All of this activity was consistent with the longer-term goal of continually folding theory development on to an expanding range of applications and laboratory phenomena. It was not always met with easy pathways to publication. Over time clinical psychology developed a number of approaches to what became called “transdiagnostic” approaches, but ICS was not integrated into that community of practice. After I retired, Felicity Cowdrey and Claire Lomax, who I had collaborated with on a number of empirical ICS studies of bipolar disorder, did use ICS as a basis for discussing the potential of a unified form of basic theory for furthering the transdiagnostic perspective [54].

In the late 1990’s the Unit changed direction, re-orienting away from applied psychology and toward brain science with a consequential loss of opportunities to explore applications other than those closely coupled to medicine. In the transition, I had a managerial role to support the transition in between the directorship of Alan Baddeley and the drive into neuroscience led by the new director William Marslen-Wilson. I was not personally comfortable with this change of direction and when the transition was complete, my 5-year “tour of duty” as assistant director came to an end and I took a short sabbatical. In a bit of a rerun of earlier trials and tribulations over my research future I had another encounter with lived experience of anxiety and depression. I was being enticed away by job offers in applied psychology elsewhere. These were not really what I wanted to do, and my family was firmly rooted in Cambridge. I used my “time out” to explore options of how I could adapt to the new climate, I worked hard to develop a specification of how ICS might map onto neural systems [38]. This too was a bit of a blind alley because I failed to entice an EEG collaborator to pursue research on this class of architecture. It was however the important thinking that underlay the development of the first process algebra model of the glance look model of attention to meanings [33]. After my leave I returned to the Unit and refocussed my research programme on topics in cognition and emotion, creating some links into neuroscience while still seeking ways to explore diversity of applications for ICS. The two non-medical “applied” research strands that emerged, and that I found ways to pursue, were work in cognitive archaeology and in the creative arts.

During my sabbatical leave I revisited the idea that all subsystems had the same internal architecture because, over the course of evolution and rather like cell division, a grandmother

subsystem had subdivided and enabled specialist processing and, eventually more advanced cognitive capabilities. I was fed up with people telling me ICS was unnecessarily complicated and was looking for another line of argument to support the macro-theoretic approach. I initially engaged in a dialogue with an old friend, Dick Byrne, who had switched from research on basic cognitive psychology to become a distinguished and influential primatologist. He had done some brilliant work on great ape eating behaviour and programme level imitation with the “Gorillas in the Mist” of Diane Fossey fame. We discussed how a sequence of architectures with 4,5,6,7,8 and 9 subsystems could help understand cognition and affect in mammals with advanced cognitive capabilities (great apes, cetaceans, crows, hominins) and those without (cats, rats, horses). Dick invited me to join an EU funded interdisciplinary research group on precursors to human culture and we spent three months developing these ideas at Collegium Budapest in Hungary while pursuing wider debates with an archaeologist, a social anthropologist, a comparative psychologist, another primatologist, and a cetacean specialist (dolphins). After that period, I once again worked with the computer scientist David Duke who was the prime mover in developing a key argument about how one subsystem could subdivide its functionality to develop a daughter subsystem of identical structure *and where the resulting architecture had enhanced or augmented functionality*. The result of this work was eventually published in 2007 in the journal *Cognition and Emotion* [40]. This was a somewhat unusual outlet since most researchers in comparative biology, comparative psychology or evolutionary psychology were unlikely to encounter it. It was however the journal in which the ICS paper with Teasdale was published and so in that context might at least be reviewed! The paper sought to lay out a full sequence of cognitive evolution from four subsystems to nine. This contrasted markedly with most approaches that focussed on the emergence of specific magic ingredients like speech, language, ‘symbolic representation’ or theory of mind, it also relied on a single *explicitly specified mechanism* for the differentiation of mental capabilities.

This paper did attract some attention in a small community of practice now called “cognitive archaeology” and led to several presentations at conferences and to three publications. These discussed in more detail, and with different emphases, the properties of the sequence from 4 to 9 subsystems [45] [55] [59]. This work again led to reciprocal argument for science benefiting application and vice versa. Show an archaeologist an artifact in the physical record like pollen in a grave, scratch marks on a stone or pigment on a shell and they will immediately engage in heated debates about whether the hominins who did that had symbolic mental representation and/or indulged in ritual practices. These debates often relied, at best on ambiguous formulations, or on arguments about special magic ingredients that felt more like folk psychology than the real McCoy. One of the ICS key papers in this domain examined how theory might more adequately scaffold inferences drawn from artifacts found in the archaeological record [55]. In particular, it developed the line that each architecture across the sequence from four to nine effectively had a glass ceiling for cognitive capability, and for example, an eight-subsystem architecture could create complex tools but could not develop deep conceptual abstractions. Likewise, our own nine subsystem architecture has limitations – we cannot “see over” implicational models and that renders such meanings and our experiences of them fundamentally ineffable.

Our work was designed to show that cognitive theory could and should usefully bound and guide archaeological inferences – yet another practical application of the theory. The inverse of application informing basic science from an applied problem was developed in the most recent paper in that domain [59]. This explored a simple conjecture advanced in our first paper [45] that the kinds of differentiation in behaviours afforded by tool use could have been

instrumental in the differentiation of meanings. A sequence of thought experiments was topicalised on plausible actions with tools and examined what that implied for semantic roles originally formalised in linguistic case grammars. Slow differentiation of classes of actions and verbalisations about them, could well have created the enabling conditions for the emergence of the ninth propositional subsystem together with implicational meanings and the capability in the central engine to reflect upon them.

It will be of no surprise that this evolutionary argument also informs the emergence of spiritual intelligence. Indeed, Fraser Watts and I discussed this very aspect many years ago and I even contributed a paper for an ISRE conference he organised in 2005. The emergence of the ninth subsystem and with that implicational representation and the central engine of cognition is the enabling requirement for experiencing, alongside reasoning with and about, 4th order abstractions¹. A sense of spiritual meaning is implicational (4th order abstraction). The capability to contemplate such experiences relies on the central engine of cognition with its reciprocal exchanges that generate propositions from implicational meanings and synthesise new implicational meanings and schematic models from propositions. It is through such processing of schematic models relating to spirituality that contemplates can come to know those models and realise aspects of their meaning in propositional and verbal form. Of course, since implicational meanings are an intersection of sensory information and propositional information the experiencing and use of spiritually loaded models of the self, the world and others is also infused with affect and a whole raft of multimodal attributes.

A key feature of the ICS analysis that John and I worked out for the Colorado symposium in 2012 was that an individual's journey into mindfulness could be thought of as a process where cognitive control, with typical day to day thinking habits with their intricate twists and turns, slowly modified to approximate more closely to how a basic mammal governed its allocation of attention over time – looking at, and experiencing, implicational meanings most prominently with occasional glances at the state of other sources of information. In a basic mammal, those sources are, of course, visual, auditory and body-state. Glancing at them has the effect of enriching content in the current landscape present in the moment of the implicational image and with that preserving its contact with self in that body state and environs. The important contrast with those like us with nine subsystems is that the glance element encompasses four sources of information – the three sensory ones and the all-important meaning state made available via the current propositional scenery. In my later treatment of the attentional score idea, I pointed out that basic mammals are necessarily mindful but also that mindful states in a nine-subsystem architecture are augmented by, and infused with, meanings in abstract forms. Thus, while we can think of animals being mindful in one sense, they cannot be *intentionally mindful* – maintaining a focus on attributes of meaning over time in a way that creates conditions for inspecting those attributes contemplatively in diverse ways that may support the development of new spiritually oriented schematic models or the modification of well-established ones.

The kinds of analysis we developed for specifying attributes of cognitive activity for novice, intermediate and expert users of computers provide a template for one way of rendering explicit attributes that would need to be considered in the development of cognitive models of spiritual contemplative practices [61]. The development of such attributes and the rules governing when they apply, require generating attentional scores based upon empirical

¹ Sensory subsystems create first order abstractions; the MPL and spatial-praxic subsystems 2nd order ones; propositions are 3rd order and implicational meanings have 4th order ones.

evidence for how contemplative practices are experienced and expressed. Obviously, novice intermediate, expert and black belt contemplatives would need separate characterisations. Also routes into contemplative practices and their outcomes may vary significantly across a range of well-studied individual difference variables such as verbaliser/visualisers; empathisers/systematisers; extraversion, agreeableness, openness, conscientiousness, and neuroticism (the big five personality traits); not to mention all the research that is indicative of a link between spiritual and mystical experience on the one hand and schizotypy on the other. All of these factors seem highly likely to modulate form and content in attentional scores and create different conditions for what a spiritual assistant should be doing. The SI project cannot engage with a massive empirical programme to sort all that out, but thought experiments, the literature and limited experience sampling could provide important clues for how to proceed. Such thinking would seem to me to be important in the context of a project seeking to make ideas about spirituality computationally explicit. To support the argument for the values of computational models, it is the capability and power of abstraction that counts more than the detail of implementation.

The final arena for the application of ideas based on ICS occurred in the context of a number of interdisciplinary projects involving the arts. We had already capitalised on intersections with the visual arts in our work on HCI, drawing extensively on cinematography [14] [35]. In the late 1990s and early 2000s in exploring a new future research path, I entered into discussions with a number of groups engaged more with performance – music/music therapy, drama and dance. In all of these cases multimodality is present and the topic of a role for implicational meanings came up in across all of these. Dancer Kenneth Tharp and composer Simon Redfern who made up “Artyfartyarts” at the time were in residence to Cambridge with a group of dancers throughout the summer of 2000 engaged with a funded a programme of “conversations by the river.” They invited me to join in. We had lots of interesting discussions and they introduced me to theatre director Dick McCaw and choreographer Kim Brandstrup. In the ensuing year or so I took part in a number of events involving them. These encounters were a strange combination of excitement and frustration. They were talking shops in which I could contribute but did not actually “go anywhere.” Then, in the summer of 2003, arts researcher Scott deLahunta and choreographer Wayne McGregor appeared on the lawn outside my office having a discussion with my colleague Tony Marcel. Since Tony knew of my previous conversations, he invited me to join in and that led to a decade and more of fascinating sci-art research that was very productive in terms of both personal reward – I felt valued by then – and a goodly output of papers. It also culminated in a major exhibition at the Wellcome Trust Gallery in London in the autumn of 2013². This was designed to enhance public understanding of what Sci-Art research could achieve and exposed the work to thousands of people.

This long running collaboration was unique in many ways. First, our research was actually embedded in the day-to-day work of an elite dance company for a decade. In marked contrast to the talking shops I had engaged with previously, we were able to do empirical work to study how the dancers and choreographer expressed their creative skills with both elite dancers and students at the Trinity Laban Conservatoire and elsewhere. The success of the process relied not just on having an open-minded choreographer who was keen on cognitive science, it depended on the presence of a co-researcher, Scott deLahunta, who was a master at bridging the languages of science with those of the world or arts. He was also committed to

² <https://waynemcgregor.com/research/thinking-with-the-body-wellcome-exhibition/>

maintain and develop reciprocal understandings throughout a long running project. Having studied how dancers worked in a studio context, we created a programme of research to establish how insights from cognitive science could augment creative practices in the studio. Scott, Wayne and the dancers were attracted to ICS partly because of its use of implicational meanings and the simple fact that it had a body state subsystem. ICS also enabled discussion of alternate forms of mental imagery techniques often called upon by a range of contemporary choreographers including Wayne McGregor. After a raft of initial papers, we set out to develop an ICS-based intervention to augment the creative skills of elite dancers which came to be called “Choreographic Thinking Tools” [51]. These were subsequently iteratively developed in a three-year project funded by the Paul Hamlyn Foundation to create an educational resource for schools which made use of a simplified version of the intersecting loops of the four central ICS subsystems to encourage differentiation of mental skills in creating movement material³. There were many papers on this topic, but a few key papers offer a sense of the overall project [56], as well as how it was related, and extended to meanings in music [52][57]. A follow-on project directed by my long-term colleague Jon May went on to develop the same choreographic thinking tools for tertiary education and to measure the extent to which such tools enhance creativity in measurable ways [60]. Of most relevance to the spiritual intelligence project is how the idea of an attentional score inter-relates the work on evolution, and imagery in arts studios and mental health clinics [58].

As with other domains in which I was able to apply ICS my own focus on the reciprocity of science and application played through. The development of the model was once again folded back to a new and very different domain from those studied previously *and from which novel technical developments of theory emerged*. As in earlier examples, the work led to augmentation of practice – in this case lectures, workshop material for elite dancers and educational package that supported teaching of creative skills in dance both in schools and workshops for more advanced students and practitioners. The attentional score idea was a technical insight for the development of basic theory that actually came in a dance studio during the creative process for a major new piece of choreography – ATOMOS – that premiered in London at Sadler’s Wells Theatre in 2013.

Concluding observation

Those looking from the outside at the trajectory of ICS research over the last four decades might recall a line from a song “You’re everywhere and nowhere baby, that’s where you’re at” and, in a way, it is an accurate description of what many commentators and critics had to say about ICS often in conjunction with recommendations to focus more specifically on laboratory paradigms with a less “ambitious” scientific appetite. I experienced a few periods of low mood because of such critics and the fact that, when seeking to forge connections within a raft of basic research communities, I was made to feel what might best be described as “unwelcome.” A community of practice specifically devoted to ICS research never developed. Rather, the research relied on a small number of enduring collaborators and intellectual fellow travellers – most notably John Morton, John Long, Jon May, David Duke, Howard Bowman, Su Li, Sophie Scott, John Teasdale, Dick Byrne and Scott deLahunta. The line just quoted is from “Hi Ho Silver lining” sung by Jeff Beck in 1967 when I was an undergraduate studying psychology in London and which I danced to many times most likely in “enhanced” experiential states that I can now no longer recall either the substance of

³ Wayne McGregor, Philip Barnard, Scott deLahunta, Jasmine Wilson and Ellie Douglas-Allan, *Mind and Movement: Choreographic Thinking Tools*. (London: Wayne McGregor/Random Dance, 2013)

those experiences or the substances that helped give rise to them. When I retired in 2011 the Unit organised a meeting in which my key collaborators all talked about the projects summarised above and the “Silver Lining” for me was the realisation that, although not that many people adopted ICS, the long running programme of work still stood as a robust candidate example of how it is possible, with sustained effort over time, to develop a model of broad scope and applicability while also exposing key elements of the craft skills that underlie the development and application of specifically psychological theory.

Listing of Key papers on the ICS framework and its applications with abstracts

This listing below includes papers on Interacting Cognitive Subsystems (ICS) that are mostly oriented towards theory and modelling rather than direct empirical tests – though several papers included in the list also report data. The actual texts for many of the papers listed here are available in electronic form within an ICS folder indexed by date. They are mostly papers that include Barnard as an author but there are some that were produced without my direct involvement, by my close associates Jon May, Giorgio Faconti, John Teasdale, Su Li, Howard Bowman and David Duke. Papers are listed in date order and cover basic laboratory phenomena, cognition and emotion, human-computer interaction, clinical psychology, evolutionary psychology/cognitive archaeology as well as interdisciplinary sci-art work addressing cinematography, art, choreography, music and drama.

[1] Barnard, P.J. (1985) Interacting cognitive subsystems: A psycholinguistic approach to short-term memory. In A. Ellis (Ed.), *Progress in the Psychology of Language, Vol. 2* (pp.197-258). London: Lawrence Erlbaum Associates.

Abstract:

A distributed architecture for human cognition is proposed in which functionally independent subsystems perform specific processing operations and interact with each other over a data network. All subsystems have the same internal organisation of constituent resources that support the representation, storage and recoding of information. Subsystems operate in different domains of processing but are governed by a common set of processing principles. Whilst the framework provides for a broadly based constituent analysis of cognitive resources, this chapter focuses on STM phenomena. Relevant processes are motivated by the requirements of language understanding and production. Standard or “modal” strategies are defined for the use of these processes in short term serial recall tasks. The chapter seeks to provide internally consistent accounts of a range of robust STM phenomena. These include phenomena associated with modality of presentation, word length, articulatory suppression, phonological similarity, suffixes, grouping and very rapid presentation.

Comment – This is the original formulation of ICS actually written in draft form in 1979 as a journal submission. Circulated in that form around MRC applied psychology unit staff and talks given at a number of places. Andy Ellis heard about it and asked me to rewrite it as a book chapter.

[2] Barnard, P.J. (1987). Cognitive resources and the learning of human-computer dialogs. In J.M. Carroll (Ed.), *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction* (pp.112-158). Cambridge, MA: MIT Press.

Abstract:

This chapter focuses upon the development of analytic and theoretical ideas that are applicable to problems in human-computer interaction. The framework of Interacting Cognitive Subsystems (Barnard, 1985) is used to decompose the representational and processing resources of cognition. This decomposition supports "Cognitive task analysis" through which user performance can be related to the functioning of resources. Underlying principles are captured in terms of approximate relationships between four concepts: process configurations; procedural knowledge; the contents of memory records; and the dynamic control of these resources. Example analyses and principles are presented for three sets of experimental evidence. The chapter raises the prospect of formalising these principles so they can be embodied in an expert system which would build "cognitive task models" for predicting user behaviour.

Comment: Cognitive task models describe the mental activity associated with the execution of tasks. These descriptions can be related in a rule base to descriptions of overt behaviour.

[3] Barnard, P., Wilson, M. and MacLean, A. Approximate modelling of cognitive activity: towards an expert design aid. In CHI + GI '87 Human Factors in Computing Systems and Graphics Interface (Toronto 5th-9th April, 1987). New York: ACM, 21-26.

Abstract

Constructs from theoretical psychology can be used to decompose the representational and processing resources of cognition. The decomposition supports "cognitive task analysis" through which user performance can be related to the functioning of resources. Such functional relationships have been formalised and embodied in an expert system. This builds approximate models which describe cognitive activity associated with the execution of dialogue tasks. Attributes of these "cognitive task models" can be used to predict likely properties of user performance.

[4] Barnard, P., Wilson, M. & MacLean, A. (1988) Approximate modelling of cognitive activity with an expert system: A theory-based strategy for developing an interactive design tool. *The Computer Journal*, 31, 445-456.

Abstract:

This paper outlines an approach to 'user modelling'. In the approach, constructs from theoretical psychology, in this case Interacting Cognitive Subsystems, are used to decompose the representational and processing resources of human cognition. This decomposition supports a form of 'cognitive task analysis' through which user performance can be related to the underlying functioning of their cognitive mechanism. Such functional relationships have been formalised and embodied in an expert system. This builds approximate models which describe cognitive activity associated with the execution of dialogue exchanges in human-computer interactions. Attributes of these 'cognitive task models' are used to derive likely properties of user performance. This paper describes two examples of working knowledge bases and discusses their properties.

[5] Barnard, P. (1991). The contributions of applied cognitive psychology to the study of human-computer interaction. In B. Shackel & S. Richardson (Eds.), *Human Factors for Informatics Usability* (pp. 151-182). Cambridge: Cambridge University Press. (See also 2206)

Abstract

This chapter addresses two questions: What progress is being made within applied cognitive psychology towards contributing a principled understanding of the phenomena of system use? and What are the prospects for the systematic application of that principled understanding? Alternative "visions" of an applied science are examined and an organising schema for the role of empirical and conceptual methods is outlined. The achievements of applied research are then reviewed under four headings: the achievements of exploratory empiricism; the achievements of analytic approaches; the achievements of experimental approaches; and theoretical synthesis as a means to pragmatic tools. Having outlined the nature of many of the research activities, the realities of those achievements will be examined from both a scientific and applications perspective. The chapter concludes the assessment by considering a sample of issues that need to be addressed in order to strengthen future contributions.

[6] Barnard, P.J. & Teasdale, J.D. (1991). Interacting cognitive subsystems: A systemic approach to cognitive-affective interaction and change, *Cognition and Emotion*, 5, 1-39.

Abstract:

Interacting Cognitive Subsystems (ICS) is a comprehensive systemic model of the organisation and function of the resources underlying human cognition. We use ICS to provide a conceptual framework for understanding normal and dysfunctional cognitive-affective relationships, and their modification. ICS proposes nine interacting cognitive subsystems, each specialised for handling a specific type of information. We describe the operations of ICS and its account of emotion development and production. ICS emphasises the importance, as part of the total cognitive configuration producing emotion, of a schematic synthetic level of processing that integrates both propositional meaning and direct sensory contributions. Processing at this level corresponds, subjectively, to holistic "sense" or "feeling" rather than to thoughts or images. We contrast ICS with the model underlying cognitive therapy, and illustrate application of the ICS framework to the maintenance of depression and to mood congruent memory.

[7] Barnard, P.J. (1991). Bridging between basic theories and the artifacts of human-computer interaction. In J.M. Carroll (Ed.), *Designing Interaction: Psychology at the Human-Computer Interface* (pp. 103-127). Cambridge: Cambridge University Press.

Abstract:

This book chapter discusses both the past "performance" of basic theory in Human-Computer Interaction and the future prospects for improving its applicability. It is argued that a key role needs to be played by "bridging representations". These are the articulation of key assumptions made in establishing appropriate scientific principles ("discovery representations") and in applying them (Applications representations). The potential value of such constructs is illustrated by reference to several extant and developing forms of HCI theory.

[8] Teasdale, J.D. & Barnard, P.J. (1993). *Affect, Cognition and Change*. Hove: Lawrence Erlbaum Associates.

Abstract:

Moods powerfully bias what we think, remember and perceive. This book traces the development of recent experimental, clinical, and theoretical approaches to understanding mood-dependent cognition, and describes a radically new conceptual framework - Interacting Cognitive Subsystems (ICS). ICS distinguishes two levels of meaning, one specific, the other holistic and integrative, and proposes that only holistic meanings are directly linked to affect. In this way ICS provides better explanations for patterns of experimental findings, for contrasts between "hot" and "cold" cognition, "knowing with the head" versus "knowing with the heart", and for central aspects of psychotherapeutic change. ICS is illustrated by application to experimental findings of mood effects on memory and judgment, to understanding negative thought production and the maintenance of depression, and to cognitive, behavioural and experiential forms of psychological treatment.

[9] Teasdale, J.D. (1993). Emotion and two kinds of meaning: Cognitive therapy and applied cognitive science. *Behaviour Research and Therapy*, 31, 339-354.

Abstract:

The clinical cognitive approach assumes emotional reactions are mediated through the meanings given to events. Cognitive therapy aims to change emotion by changing specific meanings, evaluating the truth value of particular beliefs. Bower's associative network theory of cognition and emotion also primarily concerns specific meanings. This focus on specific meanings causes problems, e.g. the contrasts between "intellectual" and "emotional" belief, between "cold" and "hot" cognition. The Interacting Cognitive Subsystems (ICS) approach distinguishes specific and more holistic, intuitive, levels of meaning. In contrast to alternative approaches, ICS suggests holistic level meanings are of primary importance in producing emotion. The ICS approach to meaning is described and its implications for understanding and treating emotional disorders discussed, together with relevant empirical findings. ICS suggests a therapeutic focus on holistic rather than specific meanings, a role for "non-evidential" interventions, such as guided imagery, and a rational basis for certain experiential therapies.

[10] May, J., Barnard, P.J. & Blandford, A. (1993). Using structural descriptions of interfaces to automate the modelling of user cognition. *User Modelling and User Adapted Interaction*, 3, 27-64.

Abstract:

Our approach to user modelling in human-computer interaction is to build approximate descriptions of the cognition underlying task performance. The technique requires several sets of production rules. One set maps from a real-world description of an interface design to an internal theoretical description. Other rules elaborate the theoretical description, while further rules map from the theoretical description to properties of user behaviour. This paper is concerned primarily with the first type of rule, for mapping from interface descriptions to theoretical description of cognitive activity. We show how structural descriptions of interface designs can be used to model user tasks, visual interface objects and screen layouts. An expert system implementation of the technique has been developed.

[11] Barnard, P.J. & May, J. (1993). Cognitive modelling for user requirements. In P.F. Byerley, P.J. Barnard & J. May (Eds.), *Computers, Communication and Usability: Design Issues, Research and Methods for Integrated Services* (pp. 101-145). Amsterdam: Elsevier Science Publishers, B.V.

Abstract:

After a brief introductory section on an applied science strategy, this book chapter summarises the Interacting Cognitive Subsystems architecture for human perception, cognition and action. It then goes on to show how such an architecture can be used to derive approximate models of human cognitive activity with appropriate form and content for supporting practical decision making in a design context. The core constructs for designing an expert system to model cognitive activity are described in some detail. A simple, but practical, design example is used to illustrate the potential value of the approach.

[12] Barnard, P. & May, J. (1995). Interactions with advanced graphical interfaces and their deployment of latent human knowledge. In Paterno, F. (Ed.) *The Design, Specification and Verification of Interactive Systems* Berlin: Springer Verlag, pp.15-49

Abstract:

Advanced graphical interfaces are increasingly dynamic, multimodal and involve multithreaded dialogues. This paper provides a theoretical perspective that can support an analysis of the issues involved in their use; the Interacting Cognitive Subsystems (ICS) framework. This framework is used to examine alternative ways in which information from different data streams can be blended with in perception, thought and control of action. The potential applicability of the core constructs to the interface design is considered. The paper concludes by outlining a specific strategy for bringing this form of understanding into closer harmony with the formal methods community in computer science.

[13] May, J. & Barnard, P. (1995). The case for supportive evaluation during design. *Interacting with Computers*, 7, 115-143 Year of publication:

Abstract:

The relevance of HCI theory to industry is being questioned, and the emphasis is shifting away from providing generalised support to systematic evaluation methods, typified by Cognitive Walkthroughs (CW). The evidence suggests that CW has not proved as effective as hoped. We examine this evidence, and argue that the problem lies not with CW or its underlying theory in particular, but with its limited scope and in the increasing dissociation of an evaluation method from its theoretical foundation. Evaluation methods retaining a theoretical element would provide the necessary conceptual support to enable designers to identify, comprehend and resolve usability problems, and would also be less limited than dissociated evaluation methods in their breadth and depth of application. We present an ambitious vision, based upon the Interacting Cognitive Subsystems model, of a 'supportive evaluation' tool, and outline Cognitive Task Analysis (CTA), the methodology upon which a proof-of-concept tool has been based. To illustrate how CTA supports the identification and resolution of usability problems, we describe three brief design scenarios, and discuss the role of cognitive modelling in the context of design.

[14] May, J. & Barnard, P. (1995). Cinematography and interface design In K. Nordby, P.H. Helmersen, D.J. Gilmore & S.A. Arnesen (Eds), *Human-Computer Interaction, Interact '95*, pp. 26-30 London: Chapman & Hall

Abstract:

Interface designers are increasingly relying on craft-based approaches to compensate for a perceived lack of relevant theory. One such source is cinematography, where film makers succeed in helping viewers follow the narrative across cuts which change the information on the screen. Cinematography has evolved over the last century, and its rules of thumb cannot be applied directly to interface design. We analyse film makers' techniques with a cognitive theory (ICS) and show that they work by preserving thematic continuity across cuts. Expressing this theoretically allows us to extrapolate away from film, applying it to screen changes in interface design.

[15] May, J., Scott, S. & Barnard, P. (1995) Eurographics Tutorial Notes Series Geneva: EACG.

(See update from 1997)

Abstract:

This Guide is intended to help people who design computer displays to use psychological principles to choose the visual appearance of computer interface objects, their arrangement on the display, and their dynamic behaviour. It will introduce you to some psychological ideas about perception - the process by which people see objects in the world, recognise them and search between them. You don't need to be a psychologist to read this Guide - we've tried to avoid using psychological jargon - but when you have read it, you should be able to use these psychological ideas to analyse your display designs. The techniques this Guide teaches you will let you decide how difficult it will be for people to group objects together, to tell objects apart, to search for objects, and to switch their attention from one part of the display to another. The Guide is organised into several sections. Each section introduces you to some ideas about perception, with some examples, and shows you how these ideas can be seen to affect the usability of display designs. The sections build on each other, introducing the simpler ideas first and the more complicated ideas later, and so this isn't a book that you can 'dip into', like a collection of guidelines might be. You have to read it through section by section - but when you have done that, we hope that you'll have learnt enough to put your new skills into practice, without needing to keep the Guide by your side. An Exercise Companion, containing answers to the exercises within the guide, is included as an appendix.

[16] Duke, D.J., Barnard, P.J., Duce, D.A. & May, J. (1995). Systematic development of the human interface. In: *APSEC '95: Second Asia-Pacific Software Engineering Conference*, (IEE) Computer Society Press, pp.313-321.

Abstract:

The problem of developing software to meet precise specifications has led to the development of mathematical notations for expressing and reasoning about the behaviour of a required or extant system. In this paper we describe a different use of formal models: as tools for gathering and consolidating requirements on interaction between engineered systems and their users. This change in focus reflects the growing use of sophisticated interactive technology in domains, such as medicine, where human comfort or safety is an issue. Not only must software systems function correctly, but the demands that the interface places on users of those systems need to be understood. This problem cannot be addressed by formal models in isolation. Instead, we describe an approach that uses formal models of human information processing, in this case Interacting Cognitive Subsystems, to augment models of

system functions. As a result it becomes possible, at an early stage in system design, to consider the role of human cognition in the correct behaviour of the system.

[17] Watts, F. and Barnard P.J. (1995). A Cognitive Model of Sleep Onset and Insomnia. Working paper posted on Researchgate.
https://www.researchgate.net/publication/321997165_A_Cognitive_Model_of_Sleep_Onset_and_Insomnia DOI 10.13140/RG.2.2.19238.63040

Abstract:

A cognitive model of sleep onset was developed in 1994/5 by Watts and Barnard, using the Interacting Cognitive Subsystems framework (ICS). The model was presented by Barnard at the Annual Conference of the British Psychological Society in 1996. This working paper describes that model in detail. In a context where there has been a growing acknowledgement of the complexity of the onset process, ICS is well-suited to capturing the far-reaching cognitive changes that take place, and to integrate experiential and tacit aspects of cognition. During the process of sleep onset, it is argued that the overall pattern of mental processing activity alters from one in which the underlying information processing mechanism is configured for action in the world, to a state in which processes are configured for rest and the maintenance of low-level vigilance to significant changes in sensation and proprioception. Whilst there may be many different specific routes to the latter state, the core theory of sleep onset holds that there are three cognitive way points shared by all routes to deep sleep. Systematic passage through these way points is perturbed in insomnia. Such perturbations are related both to patterns of recurrent mentation and to an association of that mentation with affect. Implications are discussed for the classification of insomnia, and for the formulation of cognitive processes involved in the treatment of insomnia.

[18] Barnard, Philip J. (1996). Cognitive theory and the design of advanced interfaces. In L K Yong, L Herman, Y K Leung & J. Moyes (eds) *Human Factors of IT: Enhancing Productivity and Quality of Life* Proceedings of the APCHI'96 Conference Singapore, Information Technological Institute, pp.46-63 Abstract:

Abstract:

This paper briefly reviews the problems of developing cognitive theory that can be beneficially applied to the use of advanced interfaces. A sketch is provided of a form of theory that attempts to overcome the problems that have been identified. Three explicit strategies are then outlined for developing technical means of applying theory-based ideas in different design contexts.

[19] May, J. & Barnard, P. (1997). Modelling multimodal interaction: A theory-based technique for design analysis and support. In S. Howard, J. Hammond & C. Lindgaard (Eds), *Human-Computer Interaction: Interact '97*. London: Chapman & Hall, pp.667-668.

This is a newer version of an earlier tutorial [15] by May Scott and Barnard (1995)

Abstract:

This tutorial introduces participants to a way of thinking about how users mentally represent different facets of their interactions with advanced interfaces involving multimodal information. The approach is theory-based. It assumes that human information processing deals with nine different types of mental representation. General rules of combination and decomposition of constituents, basic units and the superordinate organisation of all types of

representation will be presented. The first part of the tutorial introduces the basic theoretical ideas. The second part links this basic theory into the domain of HCI and illustrates how core concepts can be used to analyse the course of interactions with static visual displays, dynamic visual displays and multimodal displays. In the final part a handbook will be introduced describing a simple technique and an associated notation for representing the static and dynamic properties of visual information, task structures and acoustic information.

[20] Faconti, G. & Massink, M, (1997). Syndetic Approach to Referring Phenomena in Multimodal Interaction. In Proceedings of Referring Phenomena '97. Referring Phenomena in a Multimedia Context and their Computational Treatment. Madrid, Spain — July 11 - 11, 1997. Association for Computational Linguistics Stroudsburg, PA, USA. Pages 83-93.

Abstract:

User interfaces of many application systems have begun to include multiple devices which can be used together to input single expressions. Such interfaces (and even the whole application systems) are widely labelled multi-modal, since they use different types of communication channels to acquire information.

[21] Duke, D. J., Barnard, P. J., Duce, D. A. & May, J. (1998). Syndetic Modelling. *Human Computer Interaction*, 13, issue 4, 337-393.

Abstract:

Syndesis n. (pl. ~ es). [mod, L, f. Gk SYNdesis} binding together (sundeo bind together)]. — The Concise Oxford Dictionary, Seventh Edition, 1986.

User and system models are typically viewed as independent representations that provide complementary insights into aspects of human-computer interaction. Within system development it is usual to see the two activities as separate, or at best loosely coupled, with either the design artefact or some third ‘mediating’ expression providing the context in which the results of modelling can be related. This paper proposes that formal system models can be combined directly with a representation of human cognition to yield an integrated view of human-system interaction: a syndetic model. Aspects of systems that affect usability can then be described and understood in terms of the conjoint behaviour of user and computer. This paper introduces and discusses, in syndetic terms, two scenarios with markedly different properties. We show how syndesis can provide a formal foundation for reasoning about interaction.

[22] Barnard, P.J. (1999). Interacting Cognitive Subsystems: modelling working memory phenomena within a multi-processor architecture. In A. Miyake & P. Shah (Eds), *Models of Working Memory: Mechanisms of active maintenance and executive control*. New York: Cambridge University Press. Chapter 9, pp. 298-339.

Abstract:

This chapter outlines Interacting Cognitive Subsystems and its approach to accounting for performance in working memory tasks. The five central features of the theory are:

1. The cognitive mechanisms underlying working memory performance involve a multiplicity of processes and types of mental representation.

2. The detailed properties of performance depend upon the configuration of specific process needed to accomplish the task and the specific types of memory records they access and use in executing the task.
 3. There are no specific capacity limitations on what is stored at any particular level of mental representations. Capacity limitation can arise out of restrictions on the inter-functioning of processes within a wider system.
 4. The accessing and use of memory records requires the generation or revival of a description of the content to be accessed. This can also functionally constrain performance.
 5. There is no unified "central executive" component, central executive functions are themselves accomplished by processing interactions among subsystems.
- 8 These points are illustrated for a range of working memory tasks and are used as a basis for answering eight questions about working memory used to compare this approach with nine other theoretical approaches presented in the book.

[23] Barnard, P. & May, J. (1999). Representing cognitive activity in complex tasks. *Human Computer Interaction*, 14(1/2), 93-158.

Abstract:

To bridge the gap between theory and application, modelling needs to satisfy a requirement for broad scope. Interacting Cognitive Subsystems (ICS) is proposed as a theoretical basis for developing practically oriented representations, and two approaches based upon this theory are described. First, Cognitive Task Models provide a relatively complete representation of the cognitive activity required of a user in the course of an interaction. Second, diagrammatic notations can provide support in small scale problem identification and resolution. With generic form, these notations can be applied across tasks, visual interface and sound interface issues, and can handle static and dynamic situations. While Cognitive Task Modelling can be implemented in a productionrule expert system (ICSpert) and so does not require detailed modelling knowledge on the part of the analyst, the diagrammatic notations do require some theoretical knowledge. Both techniques have been used to represent problems from experimental situations, core HCI scenarios, and 'real world' design projects. They share breadth of scope and abstraction, and their parent theory supports transfer of knowledge across domains of application and from older to newer technologies, and supports feedback between the domain of application and the domain of theory.

[24] Teasdale, J.D. (1999). Metacognition, mindfulness, and the modification of mood disorders. *Clinical Psychology and Psychotherapy*, 6, 146-155.

Abstract:

A distinction is made between metacognitive knowledge (knowing that thoughts are not necessarily always accurate) and metacognitive insight (experiencing thoughts as events in the field of awareness, rather than as direct readouts on reality). This distinction, and its relevance to preventing relapse and recurrence in depression, is examined within the Interacting Cognitive Subsystems (ICS) theoretical framework. This analysis suggests, as an alternative to cognitive therapy with its focus on changing the content of depression-related thought, the strategy of changing the configuration, or mode, within which depression-related thoughts and feelings are processed i.e. changing one's relationship to inner experience.

Specifically, facilitating a metacognitive insight mode, in which thoughts are experienced simply as events in the mind, offers an alternative preventative strategy. Mindfulness training teaches skills to enter this mode, and forms a central component of Attentional Control (Mindfulness) Training, a novel, cost-efficient group preventative programme, for which there is encouraging evidence of effectiveness.

[25] Barnard, P & May J. (2000). Towards a theory-based form of cognitive task analysis of broad scope and applicability. *Cognitive Task Analysis*. Schraagen, J.M.C., Chipman, S.F., & Shalin, V.L. (Eds.) Chapter 10. 147-163. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Abstract:

This chapter outlines the general case for developing theory-based forms of cognitive task analysis of broad scope and applicability, and illustrates one means of realising this objective. A particular cognitive architecture is described, Interacting Cognitive Subsystems, together with an outline characterisation of the mental representations assumed to underpin performance in complex cognitive tasks. This theoretical specification is used to derive a form of approximate model that describe abstract properties of cognitive activity in terms of four components: (1) the configurations of processes involved in a task, (2) the capabilities of individual processes for transforming one form of mental representation to another, (3) the memory records laid down, accessed and used in task performance and (4) a characterisation of the overall dynamic control and co-ordination of the complete mental mechanism. Three routes for applying this type of modelling framework are outlined. The overall characteristics of the approach are illustrated using the most informal of the three methodologies. These illustrations were selected to cover action sequences, and the use of complex graphical interfaces to computer systems. Other potential uses of the method are briefly referenced and wider considerations discussed.

[26] Barnard, P.J., May, J., Duke, D.J. & Duce, D.A. (2000). Systems Interactions and Macrotheory. *Transactions On Computer Human Interaction* 7, 222-262.

Abstract:

This paper proposes a general four component theoretical framework for modelling the behaviour of complex systems. The approach is illustrated here for modelling the behaviour of a modular mental architecture (Interacting Cognitive Subsystems), for individuals interacting with everyday artifacts, and for interactions in groups and larger organisations. Within the framework principles concerning the behaviour of any class of generic "interactor" (anything that interacts in a system), is in principle mathematically formalisable using recent "formal methods" developed in computer science. Such formalisations can be used to "prove" that a claim follows from a theoretical formulation without requiring a running simulation.

[27] Barnard, P.J., May, J. Duke, D. & Duce, D. (2001). Macrotheory for Systems of Interactors. In J.M. Carroll (Ed.) *Human-Computer Interaction in the New Millenium*. Reading MA: Boston: Addison-Wesley Chapter 2, 31-52.

This is a just a minimally edited reprint of Systems Interactions and Macrotheory from 2000

Abstract:

This chapter is a shortened version of a longer journal article. It proposes a general four component theoretical framework for modelling the behaviour of complex systems. The approach is illustrated here for modelling the behaviour of a modular mental architecture

(Interacting Cognitive Subsystems), for individuals interacting with everyday artifacts, and for interactions in groups and larger organisations. Within the framework principles concerning the behaviour of any class of generic "interactor" (anything that interacts in a system), is in principle mathematically formalisable using recent "formal methods" developed in computer science. Such formalisations can be used to "prove" that a claim follows from a theoretical formulation without requiring a running simulation.

[28] Duke, D.J., Duce, D.A., Barnard, P.J & May, J. (2001). Human-Computer Protocols. In: C. Stephandis (Ed), Universal Access in HCI, Volume 3 of the proceedings of HCI International 2001, (pp.296-300). Hillsdale, N.J: Lawrence Erlbaum Associates.

Abstract:

We have recently developed an approach to modelling interaction that encompasses not just the device, but also aspects of a cognitive model. This integrated framework, called syndetic modelling, has been used to reason about the interplay between cognitive and computational resources deployed within an interaction. Here, a new view on this integrated framework is described. We consider the interaction between user and device as forming a hierarchy of protocols, covering different levels of abstraction over the information exchanged. As the protocol is layered, we can discuss interaction in terms of different levels of granularity, better accommodating the representation of new technologies such as vision tracking and speech which can be considered as "continuous" at some levels. We conclude by summarising classes of mathematical representation that can be utilised to represent and reason within such a model.

[29] May, J. (2001). Specifying the central executive may require complexity. In: *Working memory in perspective*. J. Andrade (ed.), Chapter 12, Hove: Psychology Press, 261-277.

Abstract:

The central executive (CE) component of the Baddeley & Hitch (1974) working memory (WM) model was initially intended to avoid the need for the model to deal with phenomena that went beyond the scope of short term memory problems. The application of the model beyond laboratory tasks has inevitably brought more and more of these 'complex' aspects of task performance into play. While the general conception of the CE as an attentional organiser or contention scheduler has allowed some of these aspects to be dealt with, there remains no detailed account of how the CE is organised, nor how it functions, and more importantly, how it might fail to function. In this chapter, May argues that the problem lies in the absence of a clear distinction in the WM model between processing and storage resources, and in the lack of detail about how the CE communicates with the slave subsystems. An alternative approach is exemplified by Barnard (1985; 1999). The emphasis of his Interacting Cognitive Subsystems (ICS) model is upon the flow of mental representations between different levels of representation, and in the competition for processing resources within each cognitive subsystem rather than between them. The chapter explores the implications of this latter approach.

[30] May, J. (2002). An information processing view of fringe consciousness. *PSYCHE: Vol 9 Issue 12.*, PSYCHE: <http://psyche.cs.monash.edu.au/>

Abstract:

In posing the sense of 'Rightness' as a quality-of-processing measure, Mangan runs the risk of a homuncular argument, since some process needs to observe Rightness, as well as the

sensory qualia. Interacting Cognitive Subsystems (ICS) is an information processing account of cognitive activity that is concordant with Mangan's arguments, but which avoids the need for any supervisory system or central executive. The approach models thought as the flow of information between nine different levels of mental representation, and includes a distinction between an unselective diffuse awareness of all active levels of representation, and a selective focal awareness of a single topic of processing. A distinction is introduced between two non-sensory representations: propositional and implicational meaning. While the propositional representations can be easily verbalised, the sensory and implicational representations can only be verbalised via propositional representations. All representations are accessible, although implications and sensory representations are harder to express verbally. As a principled model, ICS can be mapped into anatomical and neural models, supporting argumentation about physical pathways in the brain and functional pathways in the mind.

[31] Barnard, P.J. (2003). Asynchrony, implicational meaning and the experience of self in schizophrenia. In A. David & T. Kircher (Eds.). *The Self in Neuroscience and Psychiatry*, Cambridge, Cambridge University Press, Chapter 6, pp.121-146.

Abstract:

This chapter presents an account of the core signs and symptoms of schizophrenia based upon the Interacting Cognitive Subsystems mental architecture. Four sources of variation are assumed to underlie information processing activity: variation in self-representation, variation in modes of processing, variation in rates of change in the content of mental images; and variation in the synchronisation of exchanges between subsystems. The account proposes that processing exchanges between subsystems specialised to handle propositional and implicational meaning become asynchronous. At short feedback delays, the constituents of implicational representations become intermingled, providing conditions for thought disorder. At longer feedback delays, stream separation would occur leading to conditions allowing the formation of abnormal implicational models and attributions of their origins. Variation in actual symptom expression is then related to variations in modes of processing, rates of change in image content and self-representations

[32] May, J., & Barnard, P. (2003). Cognitive Task Analysis in Interacting Cognitive Subsystems In Diaper, D. & Stanton, N, (Eds.) *The Handbook of Task Analysis for HCI*. Hillsdale: Lawrence Erlbaum Associates, 291-325

Abstract:

Cognitive Task Analysis (CTA) techniques seek to model the mental activity of a task operator. With the rise of computing artefacts, the focus of CTA has changed from supporting the tutoring of operators, to modelling knowledge application, to modelling cognitive processes. Descendants of knowledge based approaches include GOMS, and produce quantitative temporal behavioural predictions for well-defined interfaces. The increasing pace of design, and the dominance of small design teams has led to a demand for more flexible techniques. This chapter describes a particular approach to CTA using a cognitive theory called Interacting Cognitive Subsystems (ICS). A CTA in ICS requires a prior task analysis to have been conducted, but the analyst then identifies the configuration of cognitive processes necessary to transform information about the task, through the phases of goal formation, action specification and action execution, for novices, occasional (normal) and expert operators. The availability of procedural knowledge, experiential and abstracted memories influence the ease of processing, and the scope a design offers for their

development informs ease of learning and skill acquisition. The location of a particular form of buffered processing predicts subjective awareness of different aspects of the task, and of task complexity. Two notations supporting analysis are described. The close coupling of the analytic approach and the underlying theory enables a CTA in ICS to provide supportive evaluation, allowing iterative redesign. It is also allowing further research linking ICS to formal models of systems analysis (Syndetics) and to other methods of TA, namely TKS, to extend both techniques to collaborative and multiple task performance.

[33] Barnard, P. J. & Bowman, H. (2003). Rendering information processing models of cognition and affect computationally explicit: Distributed executive control and the deployment of attention. *Cognitive Science Quarterly* 3(3), 297-328.

Abstract:

In this paper we illustrate the potential of process algebra to implement modular mental architectures of wide scope in which control is distributed rather than centralised. Drawing on the Interacting Cognitive Subsystems (ICS) mental architecture, we present an implemented model of the attentional blink effect. The model relies on process exchanges between propositional meaning and a more abstract, implicational level of meaning, at which affect is represented and experienced. We also discuss how the proposed mechanism of buffer movement can, in the context of the ICS architecture, be extended to account for effects of emotional stimuli and brain damage.

[34] Duke, D., Barnard, P., Halper, N. & Mellin, M. (2003). Rendering and Affect, *Computer Graphics Forum*, 22(3), 359-368.

Abstract:

Previous studies at the intersection between rendering and psychology have concentrated on issues such as realism and acuity. Although such results have been useful in informing development of realistic rendering techniques, studies have shown that the interpretation of images is influenced by factors that have little to do with realism. In this paper, we summarize a series of experiments, the most of which are reported in a separate paper, that investigate affective (emotive) qualities of images. These demonstrate significant effects that can be utilized within interactive graphics, particularly via non-photorealistic rendering (NPR). We explain how the interpretation of these results requires a high-level model of cognitive information processing, and use such a model to account for recent empirical results on rendering and judgement.

[35] May, J., Dean, M. & Barnard, P. (2003). Using Film Cutting Techniques in Interface Design. *Human Computer Interaction*, 18(4), 325-372

Abstract:

It has been suggested that computer interfaces could be made more usable if their designers made use of cinematography techniques, which have evolved to guide the viewer through a narrative despite frequent discontinuities in the presented scene (i.e., cuts between shots). Because of differences between the domains of film and interface design, it is not straightforward to understand how such techniques can be transferred. May & Barnard (1995) argued that a psychological model of watching film could support such a transference. We present an extended account of this model, which allows us to identify the practice of collocation of objects of interest in the same screen position before and after a cut. To verify that film makers do in fact use such techniques successfully, eye movements were measured

while participants watched a commercially released motion picture in its entirety, in its original theatrical format. For each of ten classes of cut, predictions were made about the use of collocation. Peaks in eye movements between 160 and 280 milliseconds after the cut were detected for six of the ten classes, and results were broadly in line with collocation predictions, with two exceptions. It is concluded that film makers do successfully use collocation when cutting in and out from a detail, following the motion of an actor or object, and in showing the result of an action. The results are used to make concrete recommendations for interface designers from the theoretical analysis of film comprehension.

[36] Barnard, P. (2004). Bridging between basic theory and clinical practice. *Behaviour Research and Therapy*, 42(9), 977-1000.

Abstract:

This paper articulates and discusses the parts played by different processes and representations in the overall conduct of applied clinical science. It distinguishes two sorts of representation, theories in the science base and bridging representations needed to map from real world behaviour to basic theory and from theory back to the real world. It is then argued that macro-theories of the "normal" human mental architecture could help synthesise basic theoretical accounts of diverse psychopathologies, without recourse to special purpose clinical-cognitive theories of particular psychopathologies or even specific symptoms. Using the Interacting Cognitive Subsystems model (Teasdale & Barnard, 1993), some specific macro-theoretic variables are identified. Concrete illustrations are given of how the essence of quite complex basic theory can be translated into a simpler representational format to help clinicians conceptualise a psychopathological state and pinpoint relevant variables that might be changed by therapeutic interventions. Some suggestions are also offered about how the inevitable problem of complexity in multiple component theories might be directly confronted.

[37] May, J., Andrade, J. Panabokke, N. & Kavanagh, D. (2004). Images of desire: cognitive models of craving. *Memory*, 12(4), 447-461.

Abstract:

Cognitive modelling of phenomena in clinical practice allows the operationalisation of otherwise diffuse descriptive terms such as craving or flashbacks. This supports the empirical investigation of the clinical phenomena and the development of targeted treatment interventions. This paper focuses on the cognitive processes underpinning craving, which is recognised as a motivating experience in substance dependence. We use a high-level cognitive architecture, Interacting Cognitive Subsystems (ICS), to compare two theories of craving: Tiffany's theory, centred on the control of automated action schemata, and our own Elaborated Intrusion theory of craving. Data from a questionnaire study of the subjective aspects of everyday desires experienced by a large non-clinical population are presented. Both the data and the high-level modelling support the central claim of the Elaborated Intrusion theory that imagery is a key element of craving, providing the subjective experience and mediating much of the associated disruption of concurrent cognition.

[38] Barnard Philip (2004). Mapping Neural Architecture to mental architecture and mental architecture to behavioural architecture. Talk at the Cognition and Brain Science Unit Cambridge.

This talk describes a candidate neural architecture for ICS based on an array of delay lines. It shows how such an architecture could support four core functions necessary for computational cognition: REplication; COmposition; DEcomposition and Selection (RECODES) which together enabled information to be organised stored and recoded for transmission from one subsystem to another. It shows how delay lines of this form could generate EEG patterns that match those observed in both visual and auditory tasks and attentional blink research.

[39] May, J. & Barnard, P. (2006). Reasoning with complementary pathways, not competing processes. In Two Minds Conference, Cambridge, July 2006. Conference Paper available on Researchgate. <https://www.researchgate.net/publication/311562772>.

Abstract:

The core of our argument is that the human mental architecture is composed of nine subsystems of equal status that interact as parts of a coherent overall system, and therefore one mind. Two of these subsystems represent qualitatively different types of meaning, one propositional in nature and the other a more abstract holistic representation, called implicational meaning (Barnard & Teasdale, 1991). Implicational meaning integrates over sensory, conceptual and bodily inputs and so captures affective states. Two aspects of this model relate to the claims of dual process theory. First, the two semantic subsystems stand in rather different relationships to the seven others. Implicational meanings receive direct and therefore fast inputs from visual, acoustic and body state subsystems while the construction and use of propositional meanings relies on inputs derived from longer, and therefore slower, processing routes. Interactions between these two meaning systems are argued to form the central engine of human ideation (Teasdale & Barnard, 1993). The second aspect of the model is that in these interactions processing activity can reflect properties of the two representations to differing degrees as a function of the mode in which meaning is processed. The mode of processing also directly relates to conscious experience. When implicational meanings dominate processing activity over time the same kinds of properties as are proposed for System 1 would be emphasised but when propositional meanings dominate the characteristic properties of System 2 would be more in evidence.

[40] Barnard, P.J., Duke, D.J., Byrne, R.W. & Davidson, I. (2007). Differentiation in cognitive and emotional meanings: an evolutionary analysis. *Cognition and Emotion*, 21(6), 1155-1183

Abstract:

It is often argued that human emotions, and the cognitions that accompany them, involve refinements of, and extensions to, more basic functionality shared with other species. Such refinements may rely on common or on distinct processes and representations. Multi-level theories of cognition and affect make distinctions between qualitatively different types of representations often dealing with bodily, affective and cognitive attributes of self-related meanings. This paper will adopt a particular multilevel perspective on mental architecture and show how a mechanism of subsystem differentiation could have allowed an evolutionarily "old" role for emotion in the control of action to have altered into one more closely coupled to meaning systems. We conclude by outlining some illustrative consequences of our analysis that might usefully be addressed in research in comparative psychology, cognitive archaeology, and in laboratory research on memory for emotional material.

[41] Su, L., Bowman, H. & Barnard, P. (2007). Attentional Capture by Meaning, a Multi-level Modelling Study. *Proceedings of 29th Annual Meeting of the Cognitive Science Society Aug 1st-4th*, Nashville Tennessee, 2007, 1521-1526.

Abstract:

We present a computational study of attentional capture by meaning, based on Barnard et al's key-distractor attentional blink task. We highlight a sequence of models, from an abstract black-box to a structurally detailed white-box model. Each of these models reproduces the major findings from the key-distractor blink task. We argue that such multi-level modelling gives greater confidence in the theoretical position encapsulated by these models.

[42] Su, L., Bowman, H. & Barnard, P. J. (2008). Performance of reactive interfaces in stimulus rich environments, applying formal methods and cognitive frameworks In: the 2nd International Workshop on Formal Methods for Interactive Systems FMIS2007 (held in conjunction with HCI2007), *Electronic Notes in Theoretical Computer Science*, 208, 95-111

Abstract:

Previous research has developed a formal methods-based (cognitive-level) model of the Interacting Cognitive Subsystems central engine, with which we have simulated attentional capture in the context of Barnard's key-distractor Attentional Blink task. This model captures core aspects of the allocation of human attention over time and as such should be applicable across a range of practical settings when human attentional limitations come into play. Thus, we have used this model to evaluate the performance trade-offs that would arise from varying key parameters in Stimulus Rich Reactive Interfaces. A strength of formal methods is that they are abstract and thus, the resulting specifications of the operator are general purpose, ensuring that our findings are broadly applicable.

[43] Su, L., Bowman, H. & Barnard, P. J. (2009). Process Algebraic Modelling of Attentional Capture and Human Electrophysiology in *Interactive Systems Formal Aspect of Computing*, 21(6), 513-539.

Abstract:

Previous research has developed a formal methods-based (cognitive-level) model of the Interacting Cognitive Subsystems central engine, with which we have simulated attentional capture in the context of Barnard's key-distractor Attentional Blink task. This model captures core aspects of the allocation of human attention over time and as such should be applicable across a range of practical settings when human attentional limitations come into play. In addition, this model simulates human electrophysiological data, such as electroencephalogram recordings, which can be compared to real electrophysiological data recorded from human participants. We have used this model to evaluate the performance trade-offs that would arise from varying key parameters and applying either a constructive or a reactive approach to improving interactive systems in a stimulus rich environment. A strength of formal methods is that they are abstract and the resulting specifications of the operator are general purpose, ensuring that our findings are broadly applicable. Thus, we argue that new modelling techniques from computer science can also be employed in computational modelling of the mind. These would complement existing techniques, being specifically targeted at psychological level modelling, in which it is advantageous to directly represent the distribution of control.

[44] Barnard, P.J. (2009). Depression and attention to two kinds of meaning: A cognitive perspective. *Psychoanalytic Psychotherapy*, 23(3), 248-262

Abstract:

The complexity of a mental disorder such as depression is such that a way of interlinking the neural, mental and interpersonal levels is needed. This paper proposes that a theoretical framework which distinguishes, and relates, macrotheory and micro-theory at these levels can serve this purpose. The 'Interacting Cognitive Subsystems' approach to mental architecture is used to show how, via the detailed specification of mental processes and representations, a macrotheory of mental architecture contributes to our understanding of depressed states. In the account advanced by Teasdale and Barnard depressed states are seen as being maintained by an abnormal version of a dynamic dialogue between two qualitatively distinct types of meaning: one is referentially specific, propositional meaning, the other consists of holistic schemata rich in latent content and is called implicational meaning. In depressed states with ruminative and avoidant thought patterns, the mental function of attention is seen as being directed preferentially at propositional meanings. There is a corresponding neglect of attention to implicational meanings. The paper concludes with a brief discussion of how this approach can address transdiagnostic issues and how it may suggest new strategies for therapeutic interventions

[45] Barnard, P. (2010). From Executive Mechanisms Underlying Perception and Action to the Parallel Processing of Meaning. *Current Anthropology*, 51:S1, S39S54.

Abstract :

The dominant conceptualization of working memory distinguishes mechanisms that handle auditory-verbal and visuospatial representations from central executive resources that control and guide them. A straightforward case can be made that executive mechanisms evolved initially in the service of directing attention to salient environmental stimuli or events and selecting adaptive actions under the guidance of affective markers. In this paper, “working-memory capacity” is viewed as an emergent property of interactions between specialist subsystems with no homunculus-like executive. Mental capability could well have advanced via the differentiation of a single multimodal subsystem into additional new specialist subsystems that process not just verbal and spatial representations but also subsystems specialized to manipulate different kinds of meaning. The resulting overall mental architecture would devolve control of action and speech to peripheral mechanisms while allowing central subsystems to focus attention and decision making on meaning. According to this hypothesis, increased mental capability is dually based on the development of more abstract representations and on the observation that the more subsystems there are, the more the mind can do at one and the same time: only the most advanced mental architecture can control walking, talking, and thinking at one and the same time.

[46] Su, L., Barnard, P.J. and Bowman, H. (2010). On the Fringe of Awareness: The Glance-Look Model of Attention-Emotion Interactions In: K. Diamantaras, W. Duch, L.S. Iiadis (Eds): *ICANN 2010, Part III, Lecture Notes in Computer Science*, 6354, 504-509 Abstract:

In previous work, we have developed a “Glance-Look” model, which has replicated a broad profile of data on the semantic Attentional Blink (AB) task and characterized how attention deployment is modulated by emotion. The model relies on a distinction between two levels of meaning: implicational and propositional, which are supported by two corresponding mental subsystems. The somatic contribution of emotional effects is modeled by an additional

bodystate subsystem. The interaction among these three subsystems enables attention to oscillate between them. Using this model, we have predicted the pattern of conscious perception during the AB and the changes of awareness when emotional or other task irrelevant processing occurs. We provide a specific account of the interaction between attention, emotion and consciousness. In particular, the dynamics of two modes of attending to meaning (implicational being more distributed and propositional being evaluative and specific) give rise to fringe awareness.

[47] Su, L., Bowman, H. & Barnard, P. (2011). Glancing and then looking: on the role of body, affect and meaning in cognitive control. *Frontiers in Psychology*, 2:348

Abstract:

In humans, there is a trade-off between the need to respond optimally to the salient environmental stimuli and the need to meet our long-term goals. This implies that a system of salience sensitive control exists, which trades taskdirected processing off against monitoring and responding to potentially high salience stimuli that are irrelevant to the current task. Much cognitive control research has attempted to understand these mechanisms using non-affective stimuli. However, recent research has emphasised the importance of emotions, which are a major factor in the prioritisation of competing stimuli and in directing attention. While relatively mature theories of cognitive control exist for non-affective settings, exactly how emotions modulate cognitive processes is less well understood. The Attentional Blink (AB) task is a useful experimental paradigm to reveal the dynamics of both cognitive and affective control in humans. Hence, we have developed the glance-look model, which has replicated a broad profile of data on the semantic AB task and characterized how attentional deployment is modulated by emotion. Taking inspiration from Barnard,s Interacting Cognitive Subsystems, the model relies on a distinction between two levels of meaning: implicational and propositional, which are supported by two corresponding mental subsystems: the glance and the look respectively. In our model, these two subsystems reflect the central engine of cognitive control and executive function. In particular, the interaction within the central engine dynamically establishes a task filter for salient stimuli using a neurobiologically inspired learning mechanism. In addition, the somatic contribution of emotional effects is modelled by a body-state subsystem. We argue that stimulus-driven interaction among these three subsystems governs the movement of control between them. The model also predicts attenuation effects and fringe awareness during the AB. doi: 10.3389/fpsyg.2011.00348

[48] Park, R. J., Dunn, B.D., and Barnard, P.J. (2011). Schematic Models and Modes of Mind in Anorexia Nervosa I: A Novel Process Account *International Journal of Cognitive Therapy*, 4(4), 415-437.

Abstract:

A deeper understanding of the cognitive-affective mechanisms maintaining Anorexia Nervosa (AN) is required to develop more effective interventions. Clinical challenges posed by AN are reviewed and a novel model of AN is offered to account for these phenomena, framed within an established cognitive architecture (Interacting Cognitive Subsystems). It is proposed that AN is maintained by oscillations between two extreme, yet mutually reinforcing, states of mind. In ‘doing’ mode there is a focus on specific conceptual meanings about the control of eating, shape and weight, with a neglect of broader emotional meaning and bodily states associated with starvation. When control cannot be maintained, individuals move into ‘mindless bodily emoting’ mode. Here attention flips between aversive bodily

sensations and emotional beliefs, resulting in feeling out of control, afraid and self-disgusted. Novel implications for course, recovery and treatment of AN following directly from this theoretical analysis are discussed in a separate sequel paper.

[49] John D. Teasdale & Michael Chaskalson (Kulananda). (2011). How does mindfulness transform suffering? I: the nature and origins of dukkha. *Contemporary Buddhism*, 12, 89-102.

Abstract:

This, the first of two linked papers, presents the Buddha's analysis of the nature and origins of dukkha (suffering) as a basis for understanding the ways in which mindfulness can transform suffering. The First and Second of the Buddha's Four Noble Truths are presented in a way that has proved helpful to teachers of mindfulness-based applications. These Truths offer a framework of understanding that can guide the application of mindfulness to stress and emotional disorders, while stressing the continuity and inevitability of the experience of dukkha in clients, teachers, and those primarily seeking a new way of being. The crucial involvement of self-view and identification with experience are emphasized.

[50] John D. Teasdale & Michael Chaskalson (Kulananda). (2011). How does mindfulness transform suffering? II: the transformation of dukkha. *Contemporary Buddhism*, 12, 103-124.

Abstract:

Mindfulness transforms suffering through changes in what the mind is processing, changes in how the mind is processing it, and changes in the view of what is being processed. The 'bearing in mind' aspect of mindfulness is important in understanding these changes, and is discussed in terms of working memory. The Interacting Cognitive Subsystems perspective recognizes two kinds of meaning, one explicit and specific, the other implicit and holistic. We suggest that mindfulness is a configuration of mind in which working memory for holistic implicit meanings plays a central role. It is here that the processing and view of experience are transformed by the creation of new patterns of implicit meaning. This analysis is applied to mindfulness practice, mindfulness as a way of being, the training of instructors and the use of mindfulness with respect to different aspirations.

[51] deLahunta, S. Clarke., G. & Barnard, P. (2012). A conversation about Choreographic Thinking Tools. *Journal of Dance and Somatic Practices*, 3(1-2),

Abstract:

This paper summarises and discusses interdisciplinary responses to three questions (1) How might we develop new ways of augmenting creativity in movement generation? 2) How can we better connect intellect, imagination and the physical body and enrich their relationship? and 3) Can a scientific understanding of the organisation of the mind provide clues and ideas that can be put into practice and how can somatic approaches contribute? The paper illustrates how the Interacting Cognitive Subsystems approach to embodied cognition and affect can be translated into heuristics to enrich practitioner's skills in thinking about movement innovation.

[52] Barnard, P. J. (2012). What do we mean by the meanings of music? *Empirical Musicology Reviews*, 7(1-2), 69-80

Abstract:

Drawing upon a recent review of the topic by Cross and Tolbert (2009), this paper briefly illustrates the diversity of theories concerning the nature of meanings in music and the challenges that need to be resolved to advance the field. A scheme for layered macro- and micro-theories for neural, mental and behavioural systems is outlined to facilitate the development of a systematic and coherent body of theory. The core of the paper charts the evolutionary origins of a specific macro-theory of the organisation all the components of the human mind. This “mental architecture,” known as Interacting Cognitive Subsystems (Barnard, 1985), incorporates not one, but two qualitatively distinct forms of meaning. Propositional meanings represent referentially specific ideas while implicational meanings encapsulate a yet more abstract and holistic form of meaning that blends conceptual material with the products of immediate distal and bodily sensations. While both forms of meaning interact in the interpretation and expression of musical meaning, such meanings are argued to be primarily implicational in nature. The paper concludes with a short discussion of how this approach might usefully be applied in the development of more precise specification of what music might mean in its various facets.

[53] Park, R. J., Dunn, B.D., and Barnard, P. J. (2012). Schematic Models and Modes of Mind in Anorexia Nervosa II: Implications for Treatment and Course. *International Journal of Cognitive Therapy* 5(1), 86–98.

Abstract:

The ICS account of Anorexia nervosa emphasises how attention to different aspects of experience combined with the impact of cognitive, emotional and bodily inputs can lead to a novel understanding of AN maintenance and routes for recovery. The hallmark of a useful theory is that it should be capable not only of generating new ways of explaining clinical course and understanding phenomenology, but should also generate detailed ideas about how to improve clinical practice and develop new treatments. ICS predicts that cultivating a ‘being embodied’ mode of mind, involving a shift in quality of attention to emotional and somatic experience, as a common route to recovery. In this sequel paper, we elaborate on processes involved in the course of AN, and theoretical and clinical implications of the account. Therapeutic strategies to remodel relationships between thoughts, emotions and bodily states in AN following directly from this theoretical analysis are then discussed.

[54] Cowdrey, F.A., Lomax, C., Gregory, J.D., Barnard, P.J. (2016). Could a unified theory of cognition and emotion further the transdiagnostic perspective? A critical analysis using interacting cognitive subsystems as a case study. *Psychopathology Review*, 4/ 3, 2017, P377-399.

<https://journals.sagepub.com/doi/10.5127/pr.044714>

Abstract:

There is evidence that common processes underlie psychological disorders transdiagnostically. A challenge for the transdiagnostic movement is accounting for such processes theoretically. Theories of psychological disorders are traditionally restricted in scope, often explaining specific aspects of a disorder. The alternative to such ‘micro-theories’ is developing frameworks which explain general human cognition, so called ‘macro-theories’, and applying these systematically to clinical phenomena. Interacting Cognitive Subsystems (ICS) [Teasdale, J.D., & Barnard, P.J. (1993). *Affect, cognition and change: Remodelling depressive thought*, Lawrence Erlbaum Associates, Hove] is a macrotheory which aims to explain aspects of information processing. The aim of this review is to

examine whether ICS provides a useful platform for understanding common processes which maintain psychological disorders. The core principles of ICS are explained and theoretical papers adopting ICS to explain a particular psychological disorder or symptom are considered. Dysfunctional schematic mental models, reciprocal interactions between emotional and intellectual beliefs, as well as attention and memory processes, are identified as being important to the maintenance of psychological disorders. Concrete examples of how such variables can be translated into novel therapeutic strategies are given. The review concludes that unified theories of cognition and emotion have the potential to drive forward developments in transdiagnostic thinking, research and treatment.

[55] Barnard, P.J., Davidson, I. and Byrne, R.W. (2016). Toward a richer theoretical scaffolding for interpreting archaeological evidence concerning cognitive evolution. In T. Wynn and F. Coolidge (eds), *Cognitive models in Palaeolithic archaeology*, pp.45-67. Oxford, UK: Oxford University Press.

Abstract:

The discovery of new evidence in the archaeological record naturally invites interpretation and deeper intellectual probing. What novel attributes of brains, minds or behavioural and cultural lives were necessary for such finds to have come into existence: Symbolic thought? Planning? Language? Enhanced memory? The very fragmentary nature of material evidence and the diversity of interpretations usually makes it difficult to arrive at meaningful consensus. Yet specific changes in patterns of archaeological evidence are set in a much longer arc of evolutionary development from the core mental capabilities of most mammals, through mammals showing evidence of advanced cognition such as apes and extinct hominins, until we reach the mental capabilities that we ourselves are known to possess. Theories of how mental capabilities developed across that full arc of evolution can now be specified in some detail. Aspects of a theory of the full arc can be formally tested on existing species while inferred properties of extinct hominins, assumed to have existed within that arc, can be matched to the more fragmentary archaeological evidence. Rather than interpret a piece of evidence, patterns of evidence should fit into the bigger theoretical jigsaw puzzle in a particular way. This strategy of constraining interpretive arguments around theoretical scaffolding is illustrated using a model of cognitive evolution. This model holds that simple mammalian minds are composed of just four subsystems and that five further subsystems evolved to support our own fully modern cognition. Theoretically derived properties of minds can then be directly mapped onto wider patterns of material evidence concerning evolutionary changes across multiple classes of behaviours - in this case illustrated by reference to nutrition, tool use and medicinal support for health.

[56] Barnard P J and deLahunta, S. (2017). Mapping the audit traces of interdisciplinary collaboration: bridging and blending between choreography and cognitive science. *Interdisciplinary Science Reviews* 42(4):359-380 DOI 10.1080/03080188.2017.1381226

Abstract:

Two long-term sci-art research projects are described and positioned in the broader conceptual landscape of interdisciplinary collaboration. Both projects were aimed at understanding and augmenting choreographic decision-making and both were grounded in research conducted within a leading contemporary dance company. In each case, the work drew upon methods and theory from the cognitive sciences, and both had a direct impact on the way in which the company made new work. In the synthesis presented here the concept of an audit trace is introduced. Audit traces identify how specific classes of knowledge are used

and transformed not only within the arts or sciences but also when arts practice is informed by science or when arts practice informs science.

[57] Barnard, P. J. and deLahunta, S. (2018). Intersecting shapes in music and in dance. In: *Music and Shape*. Editors: D. Leech-Wilkinson and H. Prior. Oxford: Oxford University Press. pp. 328-350.

Abstract:

The idea of shape figures widely in discourse about both dance and music. This chapter focuses on how shared understandings might facilitate developments of theory and practice. It first examines the nature of ambiguities that underlie references to the idea of shape. Those ambiguities create problems when we come to address the issue of how music influences the way dance is created and performed. The chapter draws on two analytic lenses explored over ten years of cumulative, interdisciplinary collaboration within R-Research, a team working alongside the contemporary dance company Wayne McGregor | Random Dance. These two lenses were extended in that context to help locate issues, clarify problems and situate what we can learn from choreographic practice and empirical studies of dance. The first lens is a framework for describing what goes on in the making of an artwork or design processes in general. The second lens is that of mental architecture applied here to examine how the multiple components of the human mind work together in creative and performance contexts. Each of these can provide some insight into the multiple facets of choreomusical relationships and, in doing so can offer some modest augmentations to choreographic practice.

[58] Barnard, P.J. (2019). Paying Attention to Meanings in the Psychological Sciences and the Performing Arts. In: *Performing Psychologies: Imagination, Creativity and Dramas of the Mind*. N. Shaughnessy, & P. Barnard (eds). London: Methuen. pp. 41-66.

Abstract:

This chapter outlines the Interacting Cognitive Subsystems framework for readers with a background in the humanities. Readers are initially invited to think about how a basic mammal attends and acts in response to simple multimodal states of the world before the approach is extended to habitual creative thinking in the human mind. The main body of the chapter strand describes the idea of “an attentional score” and how this relates to multimodally derived forms of meaning. This concept arose after more than a decade of collaboration between cognitive scientists and a prominent London-based contemporary dance company. The chapter describes the development of the idea applied to skills in dance creation and performance. Once established, the idea was readily extensible back into the clinical domain to which the Interacting Cognitive Subsystems theory had already been extensively applied. It furnishes a clear case study of how interdisciplinary synergies and benefits can emerge out of extended collaborative research between creative and scientific processes. An attentional score can be used to help think both about studio-based creativity in the performing arts as well as scaffolding accounts of how psychologists work with meanings and attention in various mental health conditions.

[59] Barnard, P.J. (2019). *Sticks, Stones and the Origins of Sapience*. In: *Squeezing Minds from Stones: Cognitive Archaeology and the Evolution of the Human Mind*, Frederick L. Coolidge and Karenleigh A. Overmann (eds.). New York: Oxford University Press, 102-127.

Abstract:

While sticks and stones have broken countless bones and helped provision thousands of generations of hominins, patterns underlying tool making and use may have had profounder consequences. This chapter explores the conjecture that tool use helped lay the foundations of key properties of modern minds: our propositional meaning system; wisdom and intuitions about meanings with their ineffable qualities and links to emotion; and our ability to walk, talk and think about meanings at the same time. We need to react to similar things with similar thoughts and behaviours (generalisation) while reacting to different things with different thoughts and behaviours. Differentiation within the behavioural systems of our precursor species (actions and vocalisations within their physical and social worlds) must have advanced in tandem with differentiation of their mental and neural systems. Tool use clearly contributed to that differentiation. Such differentiation creates new challenges for grasping what mental states underpinning perception, the control of vocal and physical actions, and bodily reactions all have in common. The emergence of two meaning systems in a specific architectural arrangement (Barnard & Teasdale 1991) is one plausible evolutionary response to those challenges that can account for how we think about meaningful abstractions, innovate and multitask.

[60] May, J., Redding, E., Whatley, S., Łuczniak, K., Clements, L., Weber, R., Sikorski, J. and Reed, S. (2020). Enhancing creativity by training metacognitive skills in mental imagery. *Thinking Skills and Creativity* 38, 100739. <https://doi.org/10.1016/j.tsc.2020.100739>

Abstract:

In a longitudinal study, 240 undergraduate dance students were recruited to assess the effectiveness of a series of workshops designed to develop metacognitive skills in use of mental imagery to support choreographic creativity. The workshops were based upon a theoretical model of mental representations and cognition. The students also completed a creativity test before the workshops, and a newly designed test of flexible thinking before and after the workshops, and a year later. Five forms of the flexible thinking test were created to allow for repeated administration over time, and the forms were shown to be equivalent and to correlate with the creativity test. Students who had taken part in the imagery workshops showed a greater improvement in flexible thinking a year after the training, compared to the scores of students who had not received the training. Evaluations of choreographic assessments by the students' teachers were rated for positive and negative mentions of imagery and creativity, and the control group scored higher than the imagery group on use of imagery immediately after the training, but lower than the imagery group on both creativity and use of imagery four months after the workshop. The findings provide some support for the idea that domain-specific creativity can be enhanced through developing skills in the use of mental imagery to produce novel ideas, and that this also improves domain-general flexible thinking.

[61] Barnard, P. (2023). Paying Attention to Spiritual Meanings: a manifesto for the cognitive modelling of contemplative practices. *Deliverable for “Understanding Spiritual Intelligence: Psychological, Theological, and Computational Approaches,” a project funded by the Templeton World Charity Foundation.* <https://www.issr.org.uk/wp-content/uploads/2023/05/Paying-Attention-to-Spiritual-MeaningsPJB2023.pdf>

Abstract

A range of spiritual contemplative practices and clinical interventions for mental health concerns share practical mental tactics for guiding how to attend and think about our moment-to-moment experiences. Examples include: decluttering mental activity, focussing on the body, avoiding repeatedly pursuing old lines of thought and facing up to or welcoming difficult ideas. Despite these commonalities, the literatures on contemplation and clinical interventions use very different languages and concepts to describe, instruct and explain what is happening in our minds during these practices or interventions. There is now significant interest in attempts to build bridges between these very different communities of practice and develop accounts within shared, scientifically grounded, theoretical formulations. A good example is John Teasdale’s use of the *Interacting Cognitive Subsystems* model of the architecture of the human mind to describe “what happens” in mindfulness (2023). Much needs to be accomplished if we are to develop a body of more formal, and potentially computational realisable theory. This report pursues that agenda by examining how we can specify aspects, not just of the balance between attention to holistic/intuitive thinking and conceptual thinking, but also of how attention is distributed *over time* and to the many *qualitatively different mental subsystem* involved in wider sensory, bodily, imaginal, semantic and multimodal mental processes. The concept of an *attentional score* is outlined and applied to some illustrations of contemplative practices. The report concludes with an outline of possibilities for the computational realisation of this class of cognitive model.