Modelling Perceptual Phenomena using Temporal Abstraction Networks

Neil Madden and Brian Logan

University of Nottingham
School of Computer Science and Information Technology
Nottingham, UK
{nem,bsl}@cs.nott.ac.uk

Abstract

We present Temporal Abstraction Networks, a novel cognitive architecture which can be used to model a variety of perceptual phenomena. The architecture is based on processes operating on collections of time-limited buffers in a parallel model of cognition and draws on aspects of the Multiple Drafts theory [Dennett and Kinsbourne, 1992]. We briefly describe the architecture and show how it can be used to model two relevant experiments from the literature: colour phi [Kolers and von Grünau, 1976], and the cutaneous "rabbit" [Geldard and Sherrick, 1972].

1 Introduction

Modelling cognitive phenomena in which the time of perception plays a role is an important challenge for cognitive science. A perceptual event is by its very nature transient. In order to reason about perceptual events we either have to extract information from them as they occur, or try to recreate details from causal evidence after the fact. A number of experimental studies, e.g., [Kolers and von Grünau, 1976; Geldard and Sherrick, 1972] have suggested that interpretations of events can override direct sensory evidence. For some sequences of perceptual events of short duration, the interpretation of individual events in the sequence depends on the characteristics of the sequence as a whole. This 'backwards referral in time', in which later events influence the perception of earlier events, is difficult to account for within a serial model of cognition without incorporating implausible delays (basically delaying sensory experience "until all the data are in").

Dennett and Kinsbourne [1991; 1992] have proposed the *Multiple Drafts* theory as a way of modelling such cognitive processes. The Multiple Drafts theory is based on a parallel, distributed view of cognition, in which large numbers of processes work independently on multiple interpretations of data simultaneously. These are the multiple *drafts*. Eventually a single draft may become dominant, but no draft is ever entirely safe from revision.

In this paper we present Temporal Abstraction Networks, a cognitive architecture for perceptual processing which draws on aspects of the Multiple Drafts theory. A Temporal Abstraction Network (TAN) consists of a network of inference processes working in parallel on collections of data over different temporal intervals (see Figure 2 for an example). TANs can be used to model a variety of perceptual phenomena which present difficulties for more conventional serial models of cognition, such as ACT-R [Anderson and Lebiere, 1998] or Soar [Newell, 1990; 1992]. As an illustration, we show how TANs can be used to model two perceptual phenomena that have been claimed to cause problems for serial models of cognition [Dennett, 1991]; namely, colour phi [Kolers and von Grünau, 1976] and the cutaneous 'rabbit' [Geldard and Sherrick, 1972; Geldard, 1977]. The architecture and models have been implemented using the SIM_AGENT toolkit [Sloman and Poli, 1996].

The remainder of this paper is organised as follows. In the next section we give a brief overview of the Multiple Drafts theory, focusing on its implications for modelling perceptual phenomena in reactive agents. In section 3 we introduce the time-limited buffers and buses which form the key components of the Temporal Abstraction Network architecture and describe how these can be combined to give models of reactive perceptual processing in simple agents. In section 4 we present models of Kolers and von Grünau's 'colour phi', and Geldard and Sherrick's 'cutaneous rabbit' experiments. In section 5 we briefly discuss related work before considering the implications of our approach for the Multiple Drafts theory in section 6.

2 The Multiple Drafts Theory

The Multiple Drafts theory of consciousness proposed by Dennett and Kinsbourne [1991; 1992] is an attempt to explain general cognitive processes, and how they can give rise to consciousness, without appealing to a *Cartesian theatre* — a central process where everything "comes together". Instead, the Multiple Drafts theory posits a highly parallel view of cognition where processing and interpretation are carried out in a distributed manner. Different interpretations constitute the multiple *drafts* which compete for (temporary) dominance. A draft which survives long enough can become relatively uncontested, and so become the dominant interpretation of events. However, no draft is ever entirely safe from further revision or reinterpretation.

The advantages of a parallel architecture over more traditional serial theories of cognition are most apparent when time is taken into account. An agent situated in an environment must act in a timely fashion in order to respond to events occurring in the world. However, as events continue to be perceived during the selection and performance of an action, there is often a need to revise the interpretation(s) of events to take account of new information. Rather than commit to a single interpretation of an event, and then later possibly have to backtrack or revise the interpretation in the face of new evidence, it is quicker to keep track of multiple possible interpretations of an event (multiple drafts) in parallel and simply drop those which are no longer supported by evidence. It is also preferable to be able to make decisions without delaying processing until all possible data has arrived, but without having to commit to a single interpretation of events too early. By allowing multiple drafts to exist simultaneously we can select actions based on the current dominant interpretation, while still allowing future information to revise or create new drafts which may then influence future action selection. A serial model of cognition commits us to either delaying processing or facing possible costly back-tracking (the two alternatives are characterised by Dennett as “Stalinesque” and “Orwellian” revisionism, respectively. [Dennett, 1991] p.116–24).

The cognitive architecture presented in this paper draws on several aspects of the Multiple Drafts theory. The architecture allows the formation of networks of parallel processes, with no single central process in ultimate control. Different processes work on (potentially) conflicting interpretations of events, and these drafts may persist for different lengths of time, depending on whether they are considered useful by other processes. No direct revision of drafts occurs; instead new interpretations are generated which outlast the obsolete drafts.

In this paper we focus on reactive models of perception, and do not attempt to model higher deliberative processes, or to demonstrate how reports of perceptions are assembled. In addition, we do not model long-term memory or persistence of drafts beyond the short time-scales of immediate perception.

3 The Temporal Abstraction Network Architecture

The Temporal Abstraction Network architecture consists of a set of processes that make inferences based on *symbolic* data within a certain temporal window. The processes are connected via a bus architecture, allowing the conclusions drawn by one process to form the inputs to other processes (including themselves), see Figure 1.

3.1 Time-limited Buffers

Each inference process has an *input buffer* with specified capacity and duration. The *capacity* of a buffer is the number of items that may be present in the buffer at any given time. The buffer's *duration* is the maximum length of time elements can remain in the buffer before they are forgotten. The duration and capacity of a buffer are independent of each other, e.g., a buffer may have large capacity but short duration or small capacity but longer duration. New inputs are added to the buffer in first in first out fashion — items arriving at a full buffer cause the oldest items in the buffer to be overwritten.

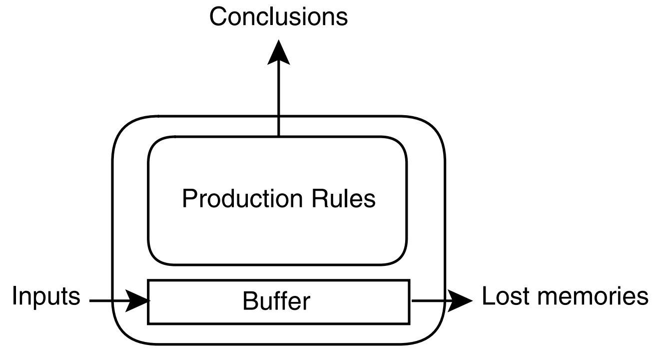

Figure 1: Conceptual model of an inference process.

Each process can draw conclusions based on the current contents of its input buffer, using an inference procedure as detailed below. Each inference process may have a different buffer duration, allowing different process to draw conclusions based on events occurring over different lengths of time. At any given point, the contents of the buffers constitute the entire state of the perceptual system.

3.2 Inference

The system as a whole runs in cycles. Each inference process contains an *inference engine*, a set of production rules that are used to spot patterns in input data, and draw conclusions based on these patterns. At each cycle each inference process's production rules are matched against the contents of the input buffer and a single rule fired. The inference engines all operate at the same rate, and the production rules have a chance to fire at each cycle regardless of whether any new inputs have arrived at the corresponding buffer. Each rule can generate a single output which is automatically transferred to the output bus of the process (see below).

More precisely, at each cycle each process:

1. Removes expired items from the buffer. An expired item is one which has been present on the buffer for longer than the buffer duration.
2. Adds any new items that have arrived since the last cycle, over-writing the ‘oldest’ items if the total number of items exceeds the capacity of the buffer.
3. Matches rules against current contents of the buffer.
4. Selects a single rule and runs it (if possible).
5. Transfers the output of the rule to the output bus.

Inference processes can be categorised into three main types based on the kinds of inferences they perform: processes that filter information; those that apply transformations; and those that integrate features of many items in a buffer and produce conclusions based on the properties of the composite grouping. An individual inference process may perform some combination of the above three categories of abstraction.

3.3 Buses and Information Propagation

The production rules within each process are used solely to draw conclusions based on the current state of the input buffer — they cannot alter the buffer in any way. Instead, output from the rules are passed on to a *bus*, that transmits them to the input buffers of other processes. In this way, data can be abstracted as it progresses through the network of connected processes, with different abstractions persisting for different lengths of time. We envisage that processes further up a chain (further from the initial percepts) would have buffers spanning a longer duration of time than the lower level processes, allowing the system as a whole to remember more abstract conclusions, while most of the details are forgotten.

To use the terminology of the Multiple Draft theory, the output of a process would represent a *draft* interpretation of events, and there may be multiple simultaneous drafts existing in multiple buffers at any given time. In addition, the bus connection architecture means that a single conclusion from a low-level process can be delivered to multiple higher-level processes, allowing for multiple drafts to be formed based on the same data, potentially producing different conclusions or interpretations. The many-to-many connections of the buses mean that the architecture is not limited to a rigid hierarchy with all conclusions eventually converging on a central process. There may be *local bottlenecks* where some combination of outputs are brought together for integration, but this does not imply that further integration must occur: the flow of information may diverge as well as converge.

The survival of a particular draft depends solely on its relevance to higher-level processes; drafts that are irrelevant to higher-level processes (in other words, do not match any patterns in the rules) will simply expire from the buffers or be overwritten by more recent drafts.

The architecture assumes that buses propagate information instantaneously—that no time is spent transferring information from one process to another, even over several buses. In practice, this means that an output from one process on a particular cycle will be present on the buffers of connected processes by the start of the next cycle. In the current implementation conclusions are transferred immediately (i.e., during the current cycle) but only made visible to rules at the start of the next cycle—this means order of execution of the processes is not important.

3.4 Feedback Loops and Alarms

While each process cannot write to its own buffer directly, it can do so in a round-about way, by making use of a feedback loop. Each process is connected to two buses; one for input, another for output. However, any bus can be connected to an arbitrary number of other processes, and other buses. This allows the output bus of a process to be connected to the input bus, creating a feedback loop. Feedback loops are useful for a variety of reasons, the most important of which is keeping track of information over a longer period of time than the input buffer allows. By using a feedback loop, a process can periodically (e.g. every cycle) send itself a message keeping track of important information, for instance keeping a running total of events that have occurred.

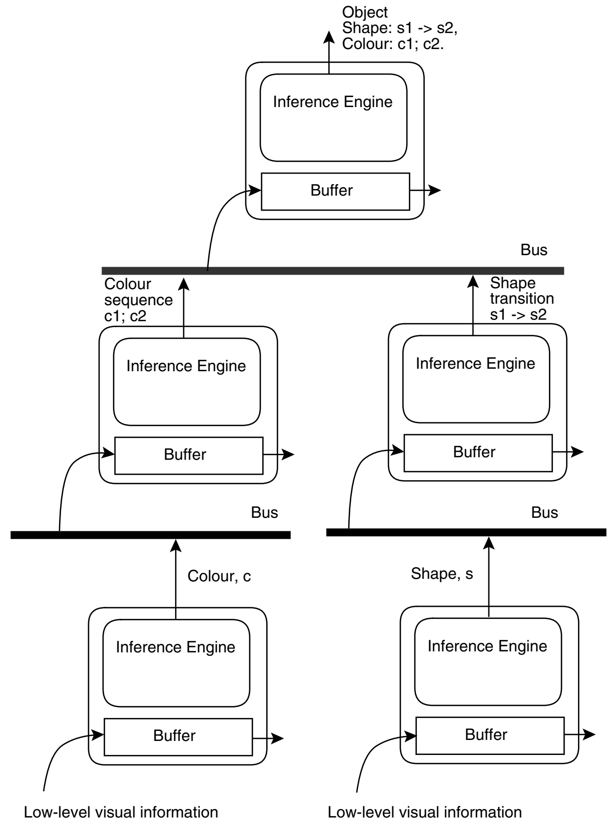

Figure 2: Network for the colour phi experiment

In addition to feedback loops, buses can be connected together in other ways, in order to provide different functions. One such function is to provide an alarm mechanism [Norman and Long, 1996; Sloman, 1998], by which particularly important events detected at a low level could bypass intermediate levels of processing and be passed directly to higher level processes for immediate action. This can be accomplished by connecting the low level bus to the high level bus via a process that acts as a filter, to determine which events require immediate attention (for instance, acting as an attention filter [Logan, 2000]).

4 Models

We have developed Temporal Abstraction Network models of a number of perceptual phenomena. In this section we briefly describe two such models: the colour phi phenomenon, and the cutaneous "rabbit", and show how the TAN architecture was used to model them. The models were implemented using the SIM_AGENT toolkit [Sloman and Poli, 1996].

4.1 Colour Phi Phenomenon

Kolers and von Grünau's [1976] colour phi phenomenon demonstrates an interesting aspect of how visual stimuli are perceived over time. In the experiments subjects were briefly shown a coloured shape at a certain position. The shape then disappeared and was swiftly replaced by another shape of a different colour in a different position. The stimuli were typically presented for 150ms with a 50ms gap between them. A number of variations on the experiment were performed, with some keeping the colour constant and changing the shape, some changing the colour but not the shape, and others changing both. Subjects were asked to indicate how the colour and shape changed. The results show that subjects perceive shape as changing continuously between the two points, whereas colour is perceived as changing abruptly somewhere between the two points. Moreover, the colour "filled-in" the intermediate shapes.

The model that we have created discriminates colour (a property of surfaces) separately from shape (a property of edges) and also discriminates changes in colour separately from changes in shape. The colour processes interpret a change in colour as an abrupt change (or rather, make no interpretation at all of changes in colour), while shape processes perceive changes as continuous (a single object changing shape). We believe this is a reasonable dichotomy, as objects in the natural world often vary in apparent shape over time (e.g. as a result of relative motion), but rarely change in hue. This is also consistent with findings that colour plays only a small role in perception of movement and other properties of objects [SFN, 2005]. Further processes then try to integrate this change data into a conclusion of the form "shape s_1 changing to s_2 AND colour c_1 followed by c_2". No "filling-in" of intermediate stages is performed at all by the perceptual processes. Again, this is reasonable, as we have reached a level of abstraction at which the agent reasons about *behaviour* of objects over time, rather than the properties of individual perceptions of the objects. This is a more useful level of abstraction for making predictions about the future behaviour of objects. By identifying objects and making predictions about their movements, an agent can anticipate and plan ahead rather than simply responding to events as they occur. The use of short-duration buffers allows the agent to aggregate information from several events in succession and spot overall trends in data, while the bus architecture means that the use of these buffers does not preclude the propagation of such information to other processes for immediate action. Once an overall trend is spotted, it is not necessary to recall the original data for revision, or to spend time "filling-in" the missing pieces. Instead, it is sufficient to generate a conclusion describing the overall trend. It is only when explicitly asked to describe the events (when asked to form a report) that an attempt is triggered to recreate the (phantom) intermediate stages.

Figure 2 shows the network of inference processes used to model this experiment. The two lowest-level processes extract shape and colour information from low-level visual sensory data. The details of how these processes operate is dependent on the underlying representations of visual data, and are omitted. The next layer of processes are also split into a colour/shape divide. Each process in this layer has a longer duration buffer than the lower level processes, and aggregates information about colours/shapes over that time period. The buffers can vary in duration independently of each other, but the results from colour phi suggest that a duration of at least 50ms is required for both. The outputs of these intermediate levels express a *change* in the underlying property. This change is encoded as a start and end state together with an indication of the sort of change (continuous transition or abrupt change). The final process at the top of the figure integrates shape and colour information to describe the recent changes in properties of an object. This information can then be used by further processes to make predictions and decide upon appropriate actions. The intermediate processing and raw sensory data are not typically made available to other processes (although they could be), and so this information disappears when it expires from the short-duration buffers in the low-level processes. Thus, when a higher-level process wishes to generate a report of the episode, it only has the high-level interpretation of a continuously changing shape and an abrupt colour change.

4.2 Cutaneous "Rabbit"

Geldard and Sherrick's cutaneous "rabbit" experiments [1972; 1977] offer evidence for another perceptual phenomenon called *sensory saltation*. In the experiments a series of short 'taps' (of about 2ms duration) were delivered to the arm of a subject. The taps are delivered with intervals of between 0–500ms. In the original experiment the taps were delivered in sequences to different locations on the arm — for instance, five taps at the wrist, followed by five between wrist and elbow, and then five more at the elbow. Subjects reported that the taps had been more or less evenly spaced along the arm — as if a little rabbit was hopping up the arm. This effect is illustrated in Figure 3. Variation in the interval between taps (inter-stimulus interval, ISI) causes differences in the effect felt. If the ISI exceeds approximately 200 ms then the effect is not felt; the taps are felt at their correct locations. With an ISI of 20 ms or less the number of taps felt becomes illusory, for example, 15 taps may be perceived as just 6. Inter-stimulus intervals between these extremes cause variations in the apparent spacing and intensity of the displaced taps, but an overall sensation that the taps were more or less evenly spaced between the location of the first tap and that of the last.

A model of this experiment is shown in Figure 4. The model uses a feedback loop to aggregate information about individual taps into information about a sequence of taps. As in the colour phi model, this aggregation allows the agent to reason about and predict the actions of an object over time, rather than being concerned with the details of individual perceptions. The lowest level process in the figure simply processes low-level information to determine the presence of a single tap. This process has a buffer duration of 20 ms and a capacity of just a single element (in this case, the 'element' is actually a collection of low-level data). If more than one tap occurs within this time-frame, then the newer tap simply overwrites any previous tap. This is consistent with the experimental results that taps occurring within 20 ms of each other are merged in this way with the location of the newer tap dominating.

The next layer of processing consists of a single process with a buffer of duration 200 ms and a capacity of 2 elements. When a tap arrives at an empty buffer (which can happen at most once every 20 ms) a new aggregate conclusion is gener-

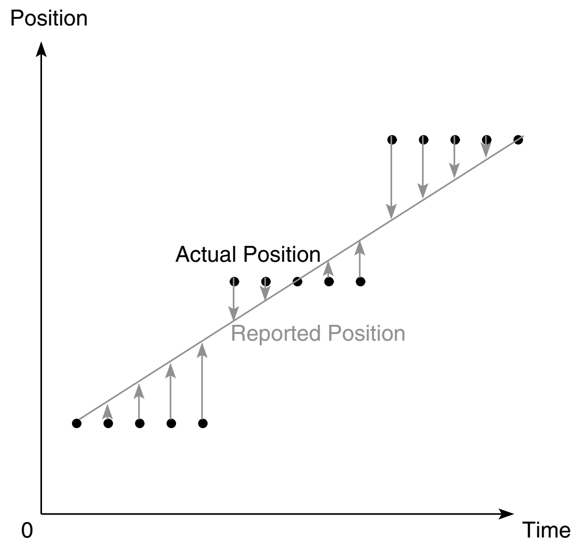

Figure 3: Actual vs. reported position of a series of taps. A train of taps is delivered with regular time interval to three positions on the arm. The subject reports that the locations of the taps are evenly distributed.

ated, taking the position of the tap as the start and end position of the 'run', and initialising the count of taps in this run to 1. This conclusion is passed to the output bus, where it is transmitted to other processes, but also, via a feedback loop, back to the input bus of the intermediate process. If a subsequent tap arrives before this aggregate fact expires from the buffer (i.e., within 200 ms) it overwrites the data about the previous tap (as this is older than the aggregate conclusion) and a new conclusion is formed which adjusts the end point of the run to the new tap position and increases the tap count by 1. The buffer duration ensures that any gap of 200 ms or more causes the previous 'run' to be forgotten, and thus any subsequent tap will be perceived as the start of a new run, which is consistent with the experimental data. It is important to note that although the buffer has a duration of 200 ms it is not necessary to delay conclusions for 200 ms. Instead, the process produces a conclusion whenever a new tap is felt (at most, once every 20 ms), and these conclusions can be acted upon immediately. For instance, Dennett [1991] suggests that a subject would be able to press a button after perceiving two taps at the wrist, but then still later report that the taps were evenly spaced. The process at the top-right of Figure 4 is waiting to do just that: it looks for a sequence of two taps at the wrist and then initiates the button press. The process on the left, however, is more conservative: it waits for a tap sequence to end before drawing a conclusion. (Detecting the end of a tap sequence can be done by employing a two-element capacity buffer and comparing the start position of sequential tap-run inputs). It is the output from this process that is eventually used to generate a report of the experience.

The model presented above suffices to explain the basic data of the cutaneous 'rabbit' experiments. Later work by Geldard [1977] appears appear to show that individual tap timings are preserved while the locations are displaced in in-

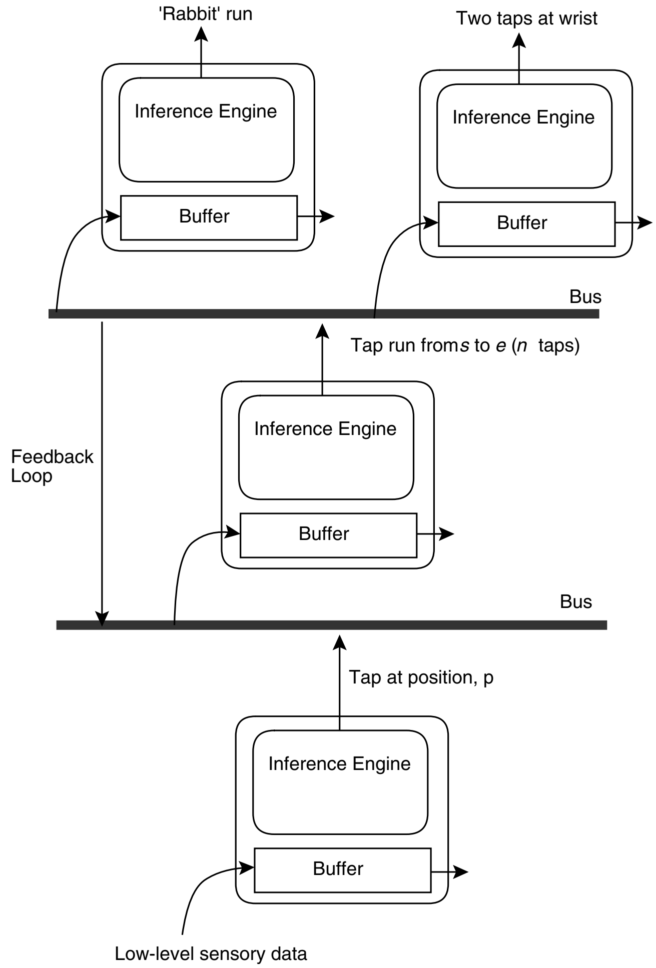

Figure 4: Network for cutaneous "rabbit" experiment.

verse relation to the interval between the tap and a subsequent offset tap. An initial model of these results can be created by converting the count of taps in a run into a list of timings of individual taps. However, it is not clear if the results of these later experiments are entirely consistent with the original findings in [Geldard and Sherrick, 1972]. The initial findings that the train of taps is felt more or less evenly spaced would produce a situation as shown in Figure 3 where each tap location is adjusted towards an average. In the later experimental results, the adjustment is always towards a subsequent tap, which would suggest a report of bunching of taps towards the end of the train. While it is not clear at this stage which model is to be preferred, the architecture developed is capable of accommodating either.

5 Related Work

There has been considerable work on formal representations of time both by philosophers and Artificial Intelligence researchers (e.g., see [Allen, 1991] for a survey). However, there has been relatively little work on cognitive models of how humans represent and reason about events occurring over time at the reactive level. Most of the popular cognitive architectures such as Soar [Newell, 1990] and ACT-R [Anderson and Lebiere, 1998] concentrate on a primarily serial model of cognition which is less suited for modelling the sorts of experiments discussed in this paper. There has been some work done on modelling temporal perception in ACT-R [Taatgen *et al.*, 2004]. However this work concentrates on the perception of the passage of time itself, and uses a Temporal Module that can keep track of a single timer that can be used to measure and reproduce the interval between two events. In contrast, our work concentrates on the perception of phenomena in which the timing and sequence of events plays a crucial role. That is, we are studying the effects of time of perception rather than perception of time. In our architecture there are no explicit timers, or measurements of the passing of time; instead, the duration of buffers and the cycle time of inference processes can be used to make judgements about the temporal relationships between events.

The Multiple Drafts theory offers an alternative view of a highly parallel system with less of an overarching structure to cognition. However, to the best of our knowledge, there is no implemented cognitive architecture which captures all the features of the Multiple Drafts theory. Some parallel models of cognition come close to the "feel" of the Multiple Drafts theory. For example, the CopyCat architecture [Mitchell, 1993] is based on a Pandemonium-style collection of competing and cooperating stochastic processes. However, in contrast to the Multiple Drafts theory (and our architecture) which allows for both parallel processing and multiple simultaneous drafts, the CopyCat model employs many parallel processes working on a single solution (a single draft). The EPIC architecture [Kieras and Meyer, 1997] is also parallel, with multiple production rules executing simultaneously. However, unlike the Multiple Drafts theory, there is an explicit executive process in EPIC responsible for conflict resolution, resource arbitration and other *meta-management* tasks. The Global Workspace Theory [Baars, 1997] proposes a parallel, distributed architecture, but in contrast to the Multiple Drafts theory it explicitly advances the notion of a central workspace as a mechanism for sharing information and coordinating the parallel processes.

We are not aware of any previous attempts to model perceptual phenomena such as colour phi or saltation using these architectures. The architecture we have presented is capable of using separate processes to achieve the sort of integration found in the CopyCat, EPIC and Global Workspace models locally without requiring a single, global process, although it would be possible to construct a "central workspace" process within the architecture.

6 Discussion

The Multiple Drafts theory argues for replacing a central executive process in the brain with a parallel, decentralised view of processing. We agree with this view. A number of experiments, including the two discussed in this paper, suggest that a parallel view of human cognition is to be preferred, at least at the level of reactive perception. However, abandoning a *single* central executive process does not mean that information cannot be brought together *locally* for integration. The TAN architecture we have presented in this paper is a middle-road between serial central processing architectures and parallel Pandemonium architectures. The architecture is based on parallel networks of simple processes drawing conclusions based on individual snapshots of events occurring within a short time-span, connected via buses and feedback loops. TANs are capable of local (serial) integration while maintaining multiple simultaneous drafts: information flow can diverge as easily as converge. This conclusion is in contrast to that of Dennett and Kinsbourne who suggest that the only alternative to the Cartesian theatre is a strictly parallel architecture, where local integration is replaced by a more chaotic Pandemonium approach (e.g., [Kinsbourne, 1994] pp. 1324). The architecture and models we have presented allow for both separate analysis of aspects of a stimulus and local integration, without appealing to a central executive process where everything comes together.

The Temporal Abstraction Network models that we have developed demonstrate the ability of the architecture to model a variety of perceptual phenomena at the reactive level in which time and the temporal ordering of events plays a key role. In future work we plan to concentrate on extending the architecture to account for action selection, as well as expanding on the details of how reports are generated. One interesting area for future research will be to look at Libet's controversial experimental results [Libet, 1985] on voluntary action.

References

[Allen, 1991] James F. Allen. Time and time again: The many ways to represent time. *International Journal of Intelligent Systems*, 6(4):341–355, July 1991.

[Anderson and Lebiere, 1998] John R. Anderson and Christian Lebiere. *The Atomic Components of Thought*. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 1998.

[Baars, 1997] Bernard J. Baars. In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, 4(4):292–309, 1997.

[Dennett and Kinsbourne, 1992] Daniel C. Dennett and Marcel Kinsbourne. Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences*, 15(2):183–247, 1992.

[Dennett, 1991] Daniel C. Dennett. *Consciousness Explained*. Penguin, London, 1991.

[Geldard and Sherrick, 1972] F. A. Geldard and C. E. Sherrick. The cutaneous 'rabbit': A perceptual illusion. *Science*, 178:178–179, 1972.

[Geldard, 1977] F. A. Geldard. Cutaneous stimulis, vibratory and saltatory. *Journal of Investigative Dermatology*, 69:83–87, 1977.

[Kieras and Meyer, 1997] David E. Kieras and David E. Meyer. An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, 12(4):391–438, 1997.

[Kinsbourne, 1994] Marcel Kinsbourne. Models of consciousness: Serial or parallel in the brain? In M. S. Gazzaniga, editor, *The Cognitive Neurosciences*, pages 1321–1330. MIT Press, Cambridge, MA, 1994.

[Kolers and von Grünau, 1976] P. A. Kolers and M. von Grünau. Shape and color in apparent motion. *Vision Research*, 16:329–335, 1976.

[Libet, 1985] Benjamin Libet. Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8:529–566, 1985.

[Logan, 2000] Brian Logan. A design study for the attention filter penetration architecture. In *Proceedings of the AISB'00 Symposium on How to design a functioning Mind*, pages 94–101, Birmingham, UK, April 2000. AISB, AISB.

[Mitchell, 1993] Melanie Mitchell. *Analogy-Making as Perception: a computer model*. MIT Press, Cambridge, MA, 1993.

[Newell, 1990] Allen Newell. *Unified Theories of Cognition*. Harvard University Press, 1990.

[Newell, 1992] Allen Newell. Précis of "Unified theories of cognition". *Behavioral and Brain Sciences*, 15:425–492, 1992.

[Norman and Long, 1996] Timothy J. Norman and Derek Long. Alarms: An implementation of motivated agency. In Mike Wooldridge, Joerg P. Müller, and Milind Tambe, editors, *Intelligent Agents II — Agent Theories, Architectures, and Languages (ATAL-95)*, volume 1037 of *LNAI*, pages 219–234. Springer-Verlag, 1996.

[SFN, 2005] Brain facts: a primer on the brain and nervous system, 2005.

[Sloman and Poli, 1996] Aaron Sloman and Riccardo Poli. SIM_AGENT: A toolkit for exploring agent designs. In Mike Wooldridge, Joerg Mueller, and Milind Tambe, editors, *Intelligent Agents Vol II (ATAL-95)*, pages 392–407. Springer-Verlag, 1996.

[Sloman, 1998] Aaron Sloman. Damasio, Descartes, alarms and meta-management. In *Proceedings IEEE International Conference Systems, Man, and Cybernetics (SMC 98)*, pages 2652–2657, Los Alamitos, California, 1998. IEEE, IEEE Computer Society Press.

[Taatgen *et al.*, 2004] Niels Taatgen, Hedderick van Rijn, and John Anderson. Time perception: Beyond simple interval estimation. In *Proceedings of the sixth International Conference on Cognitive Modeling*, pages 296–301, Pittsburgh, PA, 2004.