Biologically Inspired Mechanisms for Processing Sensor Rich Environments

Charles J. Hannon

Computer Science Department
Texas Christian University
Fort Worth, Texas 76129
C.Hannon@tcu.edu

Abstract

Biological organisms use a combination of attention and arousal to control the amount of sensory data being sent to their central nervous system's higher order processing. We propose using what neuroscience has learned about these natural mechanisms to construct a biologically inspired model of how sensor input can be filtered and fused in a large-scale DAI system. We then show how this model can be implemented using our Goal Mind agent environment and discuss how such an implementation can be tested.

Introduction

Biological organisms rely on a large number of simple sensors to provide constant input about the changing world around them. To timely process these sensor rich environments, the central nervous system of these organisms must use both an attention and arousal mechanism to both filter and fuse the resulting data into a workable set of critical input. A number of large-scale DAI applications, such as smart home projects and battle management systems for the 21st century battlefield, already exhibit sensor collections that could gain from a biologically inspired model that improves sensor fusion and input filtering. In this paper, we will discuss the neuroscience background for such a model, a resulting theoretical model drawn from this neuroscience research and how such a model can be implemented using our Goal Mind environment (formally called AMEBA) [2].

In this discussion, we will limit the type of sensors being used in a DAI application to those similar in complexity to a biological system. While this, at first, might seem to limit the resulting mechanisms' general application, in reality, it does not. Almost all complex sensors are made up of a collection of simple sensors ganged together for the ease of use in their primary domain of operation. For example, a digital camera internally relies on a grid of simple light sensors which work together to produce the resulting captured image. This is not unlike the human eye which internally consists of a set of bar, edge and hue detectors that decompose the image being projected on the back of the eye into a collection of simple sensor inputs. What we perceive as the image being generated by our eyes is actually constructed in the visual cortex of our brain from not only the output of these detectors but also what our brain expects the image to look like. Most vision systems being used in robotics and other AI related fields end up processing camera input using the same type of detectors that biological systems use, so the proposed mechanisms can be used with even these types of complex sensors, as long as, access to the internal processing of the sensor can be gained.

The Attention and Arousal Mechanism

In biological systems, the attention mechanism is primarily used to block low-value sensory input so that the Higher Order Processing (HOP) resources are not overwhelmed by sensory data. The arousal mechanism provides a way for attention to be switched from one set of inputs to another. A breakdown in either the attention or arousal mechanism can have serious consequences, causing such conditions as Attention Deficit Disorder and even Autism. However, there is some debate as to the neurological processes behind attention and arousal, and whether these processes differ based on the modality of the sensory input. A detailed discussion of both the psychological and neurological basis for attention and arousal theories can be found in [1, 4, 5, 7]. Here, we will only provide a brief overview which is based primarily on these sources.

Psychologically-based theories for both mechanisms tend to put more emphasis on the effect of modality than do neurologically-based theories. This is not really surprising since an experimental psychologist must use some modality to conduct her research and this testing method will somewhat color the results. On the other hand, cognitive neuroscientists tend to propose the same underlying mechanism (such as lateral inhibition for attention or the gated dipole for arousal) regardless of modality. Since our intention is only to use these theories as an analogy for a DAI system, it would seem safe to assume that we can ignore the effect of modality and propose a unified model.

Early work on attention (by Broadbent in the late 1950's) focused on the idea that sensory input was filtered before being analyzed. This view has become know as the early selection theory. In 1963, Deutsch and Deutsch suggested that the filtering occurs after analysis but before response processing. This view is now known as the late selection theory. However, both of these theories have conceptual problems. The early selection theory does not provide a suitable mechanism for attention switching while the late selection theory implies that we can simultaneously use cues from all sensory inputs in our internal processing of the attended input and this simply is not true.

In 1969, Treisman suggested that unattended sensory input is not completely blocked, but simple attenuated. This theory has been shown to both fit experimental data and connectionist explanatory mechanisms. It provides room for the arousal mechanism by providing a low level of off-target sensor data that can eventually force attention switching.

Another aspect of attention is the number of sensory input falling within the attention spotlight. Normal human subjects can focus their attention on few or many sensory input (increase or decrease the spotlight beam) depending on the requirements of the task at hand and move the attention spotlight through the sensor domain as required to complete the task. While this phenomenon has been primarily demonstrated through experiments with visual input, the ability clearly crosses modality boundaries. Stein and others have demonstrated that animals attend better to coherent multisensory input than any single sensory input [8]. Murphy presents an overview of how biologically-based sensor fusion has been defined and used in AI systems [6].

A Model of Attention and Arousal

Based on the psychological and neurological research into the attention and arousal mechanism, we propose a computational model that attempts to emulate how these mechanisms work. Our model is based on a hierarchical collection of three types of modules, each containing a filter and fuser component. Each of these modules supports a piece of the overall attention and arousal mechanism of the resulting system.

As shown in Figure 1, the Time Filter/Fuser supports a single sensor input. When no inhibiting signal (I_s) has been applied, the input is simply passed through the filter/fuser. When an inhibiting signal has been applied, the input is blocked and the arousal mechanism is activated. Based of the threshold provided by T_τ, a differential greater than plus or minus the threshold between the sensory input at time t and time t+τ will cause the signal Δ_τ to fire.

As shown in Figure 2, the Space Filter/Fuser supports any number of sensory input of the same modality. In addition to allowing a sensor input to pass through the

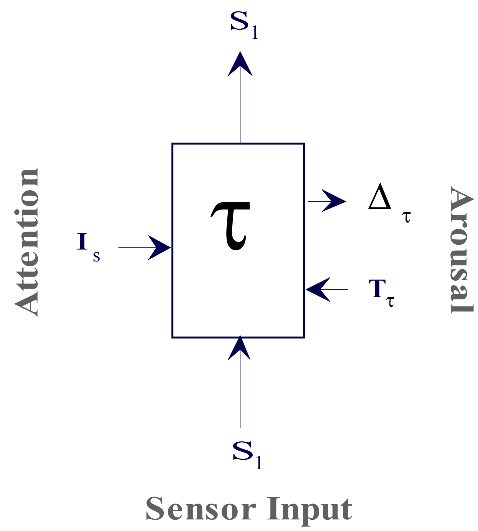

Figure 1. The Time Filter/Fuser

filter/fuser, the attention mechanism generates a set of combined sensory output such as the average sensor value and the maximum and minimum sensory value. The inhibiting of any of these output can be controlled using the inhibiting signal (I_s). When inhibiting signals have been applied, the arousal mechanism is activated. Based on the threshold provided by T_σ, a differential greater than the plus or minus the threshold between any of the sensory input will cause the signal Δ_σ to fire.

As shown in Figure 3, the Mixed Modality Filter/Fuser supports any number of sensory input of any number of modality. In addition to allowing a sensor input to pass

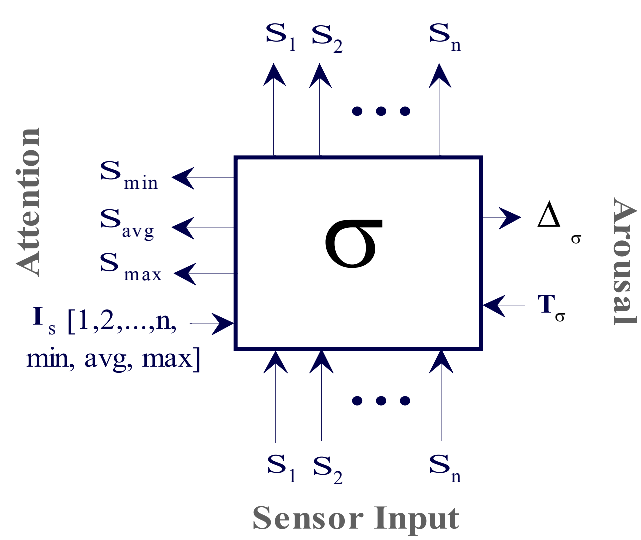

Figure 2. The Space Filter/Fuser

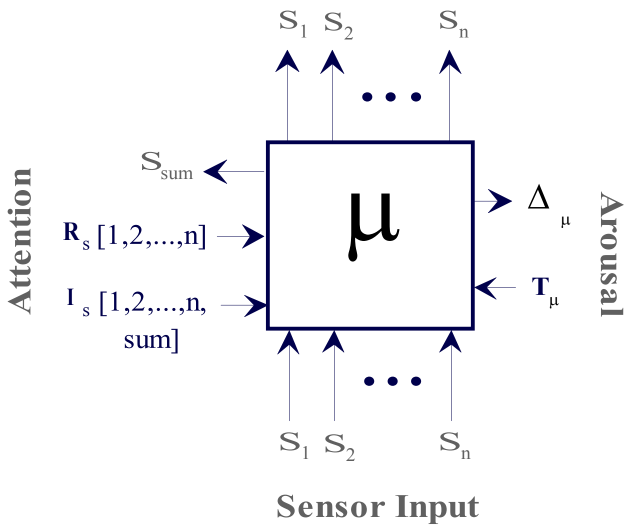

Figure 3. The Mixed Modality Filter/Fuser

through the filter/fuser, the attention mechanism generates a summation sensory value which reflects how the different input modalities are related. The inhibiting of any output can be controlled using the inhibiting signal (I_s). In addition, Bower's taxonomy is supported by allowing the input to be recalibrated using the R_s signal [6]. When inhibiting signals have been applied, the arousal mechanism is activated. Based on the threshold provided by T_μ, a differential greater than plus or minus the threshold between any of the sensory modalities will cause the signal Δ_μ to fire.

These filter/fuser components are designed to be used as modular components in an architecture like Goal Mind which provides the maximum flexibility in the way they are connected. To support a flexible implementation, any sensory input should be allowed to run to any number of different filter/fuser components and the filter/fuser components should be allowed to be stacked on top of each other. After a brief overview of the current Goal Mind system, we will address both the Goal Mind implementation and how it will be tested.

The Goal Mind Architecture

The Goal Mind system, part of our Gold Seekers toolset, is the next generation of our Adaptive Modeling environment for Explanatory Based Agents (AMEBA) architecture [2]. While attempting to productize AMEBA for academic use, it became apparent that some early design decisions made the support mechanism and GUI control of this system too inflexible for general use. Based on AMEBA, the Gold Seekers project is designed to produce a set of powerful free tools for both general distributed application and multi-agent intelligent system design and implementation. The first of these tools is Alchemy [3]. Alchemy is a general distributed processing environment which supports: 1) the asynchronous processing model needed by our cognitive-based approach, 2) a GUI-driven dynamic generation, operation and testing environment, and 3) a multi-level security facility for safe operation over the Internet or other public networks. The second tool is Goal Mind, a redesign of AMEBA to run on top of Alchemy. While we now have a working copy of Alchemy and Gold Mind, we are still a number of months away from being able to release them for general use.

Goal Mind, like AMEBA, attempts to capture the explanatory force of a connectionist neural model while allowing the use of the better-understood representation and reasoning methods of symbolic AI. From a system perspective, it attempts to provide processor transparency within a parallel system and a flexible method of process and knowledge management. The key element that supports these requirements is the etheron process framework that allows agent components to be built from Alchemy processing nodes. An etheron provides a container for an instance of any inference or routing mechanism needed by the system. Once contained, the etheron supports the mechanism with, 1) a standard way to load and store knowledge, 2) interfaces to a set of predefined management tools and 3) a generalized set of communication channels for talking with other etherons.

Goal Mind models draw their explanatory depth from the environment's ability to support hierarchical cognitive processing. Using adaptive distributed processing and generalized inter-process communication, cognitive functions can be modeled at different levels of abstractions without changing the logical relationship between these functions. Thus, a function like the conceptual reasoning about the world and self can be simulated with a reasoning and knowledge storage system which has far less capacity than that of a real human. This allows us to preserve the overall model's explanatory depth, as long as we preserve explanatory relationships between cognitive components. To ensure that we preserve these relationships, our modeling research is driven by both the evidence from experimental psychology regarding the architecture of the mind and the evidence from neuro-physiology regarding the architecture of the brain.

The Goal Mind Implementation

Like AMEBA, Goal Mind supports a set of Representation and Inference Mechanisms (RIMs), Coded Response Mechanisms (CRMs) and InterFace Nodes (IFNs). The sensor filtering and fusion approach being proposed here represents a shift in our research from developing applications where the biologically-inspired sensors are simulated by the IFNs and used to study human cognition, to applications where the IFNs serve as interfaces to real sensors. While this shift represents a new effort to support real-world missions for our tools, both the existing AMEBA system and new Goal Mind system have always been designed to support this role. Therefore, the use of the new attention and arousal model to support large-scale sensor arrays does not radically change the design or operation of a Goal Mind multi-agent model.

Figure 4 presents a high-level view of one agent in our current test model. In this agent, the IFNs can either emulate sensor input or condition real sensor input used by the agent. The filter/fuser components are implemented as CRMs. The Attention Control and Arousal Control components are implemented as RIMs using our SKIPS knowledge base engine to support rule-based control of these systems. The Sensor Processing RIM is for the most part any existing SKIPS components which has been used in previous AMEBA test models.

At start-up, all threshold values are set by the Arousal Control component by sending threshold messages to each of the filter/fuser components. We currently do not adjust the thresholds during testing, but adding this capability would require only a simple modification to the Arousal Control rule-base. For most purposes, the fact that any filter/fuser can output any input generates a great deal of redundant messages so we currently inhibit all individual input and only turn on the ones from the Time Filter/Fuser when a threshold is reached. While this approach seems to work, we are still researching the best way of handling the routing of input to the rest of the system.

Our research approach has been to incrementally increase both the number of sensors and reasoning complexity of our test model over time. In the existing test model, the sensor input is still relatively basic. The model divides sensor input into five rooms, with two temperature and two light sensors per room. In the rule-bases, the second sensor per modality per room is viewed as redundant to the first and is used to provide a certain level of fault tolerances regarding the sensor data coming from each room. This model allows us to study the agent's ability to reason about the cause of temporal, spatial and dual-modality temperature and light relationships from the environment. The end result of this iterative modeling approach will be a model that greatly increases the complexity of the sensor input, and thus, the level of understanding the model has about its environment.

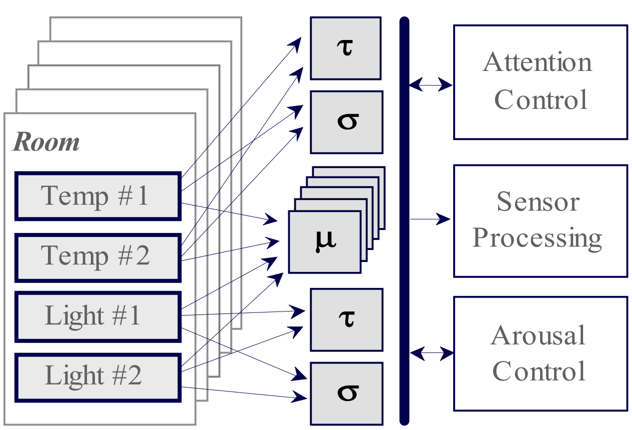

Figure 4. Implementation of the Current Test Model

Discussion of Goal and Testing Criteria

The goal of this research is to show that our attention and arousal model improves both the overall performance of an agent or multi-agent system in a sensor-rich environment and an agent's ability to detect and respond to critical input that is currently not in its attention spotlight. There is clearly a limit to the amount of sensory input a centralized reasoner can timely process. Further, the amount of useful information you can gain by distributing the sensory input across a decentralized set of reasoners is limited since it is difficult to gain complete insight from multi-modal sensor fusion unless a single processing point in your system can actually see all of the raw input needed for the task at hand and this 'spotlight' of needed input is dependent on the state of the world and the agent.

In early tests with high rates of sensor updates with the current model, this model demonstrates the level of improvement we would expect from a trade-off in the overall increase of the model's communication and processing latency (resulting from the addition of filter/fuser and control components) being offset by the decrease in the amount of processing needed by the sensor processing HOP to process each sensor update cycle. These results also logically follow from the model's design since it takes more work (measured in number of rule firings) for the sensor processing HOP to process a useless input than it takes for a filter/fuser to remove it from consideration under the intelligent control of the much slower cycling attention and arousal mechanism. Gold Mind's processor thread design of also helps the overall performance of the model since components at rest (i.e., not processing data) tend to consume very little processing bandwidth. However, it is still too early to completely quantify this improvement in raw processing time since the model is still quite small.

We have also run tests to determine the current models speedup when the components were distributed across our Beowulf-like cluster in different ways. So far we are seeing speedup results that are characteristic of other Alchemy (and AMEBA) models, such as those reported in [2] and [3]. As with these other models, we see fairly good speedup (on the order of 0.5n) up to about 8 processors, but beyond this point the speedup tends to fall off rather quickly. Our analysis of these results indicate that this fall-off is a result of the current model being too small to computationally stress more than 8 processors.

While we have good reason to believe that in an even larger model (with many more sensors and sensor types) our attention and arousal model will do an even better job of improving the overall processing time needed for the larger number of input, the model's time performance is not the only expected improvement in a larger test model. As we know from biologic examples, at some point a centralized processor without a front-end filter/fuser mechanism will simply lose the ability to provide any useful evaluation of the input. We are currently designing a larger test model to determine at what point this loss of ability occurs within the Gold Mind framework. This model will use the dynamic nature of the Gold Mind tool to: 1) allow the filter/fuser mechanism to be turned off and on, and 2) allow the sensor processing HOP's workload to be distributed across a number of cookie-cutter components that each process only a part of the input.

Using this larger model, we should be able to provide some further understanding regarding the point at which a centralized reasoner using an attention and arousal mechanism begins to significantly outperform either a centralized approach without a front-end filter/fuser mechanism or a decentralized approach that subdivides the same set of sensor input between different processes. Based on initial studies, this critical processing point is expected to occur when the input rate reaches several hundred sensor input per time slice, but the actual number could vary greatly from this early estimation.

Future Work

Once we have completed testing with the our simulated and tightly controlled real input test models, the next step is to apply the attention and arousal model against a real-world environment and see how well it performs. The smart home environment is ideally suited for this real-world test since agents with attention and arousal mechanisms should be able to perform a number of tasks better than agents without them. For example, one of the tasks our test agents have been given is to detect possible fire conditions without overreacting to abnormalities in both the sensor input and the sensors themselves. As anyone who has been exposed to fire alarm systems that are prone to false alarms knows, humans will eventually stop believing any reports of fire by these systems. A smart home agent that demonstrates a human-like ability to reason about possible fire conditions may be able to respond quicker to real fires and better suppress false alarms. The improvement in reasoning should be greatly aided by the ability to better fuse and filter sensor input since the agent is thus able to consider a larger set of input evidence.

Conclusion

While still in the early stages of our research, we are already seeing signs that the attention and arousal model proposed in this paper will improve large-scale sensor processing. It is also beginning to shed some light on the actual cognitive processes on which it is based. This dual use of Goal Mind has always been one of the most exciting features of this research. With the addition of an explanatory attention and arousal model, we believe that the Gold Seekers project has taken yet another small step toward a unified understanding of cognition.

References

1. Anderson, J. R., 1995. *Cognitive Psychology and its Implications*. New York: W. H. Freeman and Company.
2. Hannon, C. and D. J. Cook, 2001. Developing a Tool for Unified Cognitive Modeling using a Model of Learning and Understanding in Young Children. *The International Journal of Artificial Intelligence Tools*.
3. Hannon, C., 2002. Alchemy: A Geographically Distributed Processing Environment. *Submitted to PDPS-2002*.
4. Levine, D. S. 1991. *Introduction to Neural and Cognitive Modeling*. Hillsdale, NJ: Lawrence Erlbaum Associates.
5. Martindale, C. 1991. *Cognitive Psychology, A Neural-Network Approach*. Brooks/Cole Publishing Co., Pacific Grove, CA.
6. Murphy, R.R. 1996. Biological and cognitive foundations of intelligent sensor fusion. *IEEE Transactions of Systems, Man and Cybernetics – Part A: System and Humans*, 26(1):42-51.
7. Posner, M. I., 1995. Attention in Cognitive Neuroscience: An Overview. In Gazzaniga, M. S. (Editor-in-Chief) *The Cognitive Neurosciences*. MIT Press: Cambridge 615-623.
8. Stein, B. E., M. T. Wallace and M. A. Meredith, 1995. Neural Mechanisms Mediating Attention and Orientation to Multisensory Cues. In Gazzaniga, M. S. (Editor-in-Chief) *The Cognitive Neurosciences*. MIT Press: Cambridge 683-702.