MEMORY ASSOCIATION MACHINE: AN ACCOUNT OF THE REALIZATION AND INTERPRETATION OF AN AUTONOMOUS RESPONSIVE SITE-SPECIFIC ARTWORK.

by

Benjamin David Robert Bogart

Hons. B.F.A. Ryerson University, 2003

a Thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Science
in the School
of
Interactive Arts and Technology

SIMON FRASER UNIVERSITY
Summer 2008

APPROVAL

Name: Benjamin David Robert Bogart

Degree: Master of Science

Title of Thesis: Memory Association Machine: An Account of the Realization and Interpretation of an Autonomous Responsive Site-Specific Artwork.

Examining Committee: Dr. Tom Calvert
Chair

_____

_____
Prof. Thecla Schiphorst, Senior Supervisor

_____
Prof. Maia Engeli, Supervisor

_____
Dr. Philippe Pasquier, Supervisor

_____
Dr. Maria Lantin, External Examiner,
Director of Intersections Digital Studio,
Emily Carr Institute

Date Approved: _____

Abstract

This thesis is an account of the realization and interpretation of the autonomous responsive electronic media artwork "Memory Association Machine" (MAM). Realization and interpretation are components of the creative process that braids conceptual, site-specific, electronic media art and artificial intelligence practises. The meaning of MAM is dependent on its unique location in space and time. MAM relates *itself* to its context using three primary processes: *perception*, the *integration* of sense data into a field of experience, and the *free-association* through that field. MAM perceives through a video camera, integrates using a Kohonen Self-Organizing Map, and free-associates through an implementation of Liane M. Gabora's model of memory and creativity. These processes are as important as MAM's physical appearance, are composed of computational elements, and allow the system to respond to context autonomously.

Keywords: Memory; Machine Creativity; Situated; Self-Organized Map; Responsiveness; Free-association

Subject Terms: Computer art – Canada; Art and technology; Installations (Art); Artificial Intelligence; Creative ability in technology

For Dad.

"An image is a reference to some aspect of the world which contains within its own structure and in terms of its own structure a reference to the act of cognition which generated it. It must say, not that the world is like this, but that it was recognized to have been like this by the image-maker, who leaves behind this record: not of the world, but of the act."

— What is an Image, Cohen (1979)

Acknowledgements

I would like to thank my committee, Thecla Schiphorst, Philippe Pasquier and Maia Engeli, for providing invaluable support and critique. In particular Thecla's efforts have greatly improved the quality of this text.

I would like to thank my peers who have engaged in much discussion about the project. In particular Greg Corness has been a constant influence. Steven Barnes has inspired much discussion around neurophysiology in the development of the free-association system. Aaron Levisohn and Marek Hatala have provided much feedback and many questions concerning the project.

I would like to thank Gordon Pritchard and the SFU Surrey research support staff, without whose efforts the on-campus installations of MAM would not have been possible. I would like to thank Joyce Black and Tiffany Powell for their administrative support. I would like to thank Joanne Curry for her interest in the project and for setting up the demonstration of MAM for Bing Thom.

This work would not have been thinkable without the inspiration of David Rokeby.

Contents

Approval ii

Abstract iii

Dedication iv

Quotation v

Acknowledgements vi

Contents vii

List of Figures ix

Preface x

1 Introduction 1

2 Artistic Practise 5

2.1 Conceptual Art 5

2.2 Site-Specific Art 6

2.3 Electronic Media Art 8

2.4 The Artist / Author 11

2.5 Artificial Intelligence 12

3 Theories of Creativity 15

3.1 Dimensions of Creativity 16

3.2 The Cognitive Mechanisms Underlying the Creative Process 17

3.2.1 Sparse, Distributed, and Content Addressable Memory 18

3.2.2 The Structure of Memory 21

3.2.3 The Creative Thought Process . . . 22
3.2.4 Remembering as Recreating and Recreating as Remembering. . . . 22
3.2.5 "Honing" The Thought Process . . . 23
3.3 Intersections . . . 24

4 Creative Machines **26**
4.1 "AARON" . . . 26
4.2 "Pockets Full of Memories" . . . 28
4.3 "The Giver of Names" & "n-cha(n)t" . . . 29

5 Memory Association Machine **32**
5.1 The System . . . 32
5.1.1 Kohonen Self-Organizing Maps . . . 35
5.1.2 Memory System . . . 38
5.1.3 Free-Association System . . . 44
5.2 Future Development . . . 47
5.3 MAM and Creativity . . . 52
5.3.1 Evaluating MAM using Boden's Framework . . . 52

6 Growing Form from Context **54**
6.1 Realization→Interpretation . . . 55
6.2 Artifacts-as-Processes . . . 57
6.3 FLOSS & Collaboration . . . 59
6.4 Meaning . . . 61
6.4.1 The Perturbation of the Viewer . . . 61
6.4.2 Software as Artist Statement . . . 62

7 Conclusion **64**

Bibliography **66**

A Implementation **72**

B Journal **73**

C Log of Practise **93**

D SOM Test Data **148**

List of Figures

3.1 Memory is both *sparse* and *distributed.* 19
3.2 Types of Memory Distribution 19
3.3 How Memories are Stored 21
3.4 The Bridging Concept 23

5.1 Installation of MAM for the 2007 Pure Data Conference 33
5.2 The Architecture of "Memory Association Machine" 33
5.3 A SOM trained on data containing two independent clusters. 36
5.4 A comparison of feature-map representations. 39
5.5 The representation of the field of experience. 40
5.6 Effect of Gaussian alpha channels in the representation of the field of experience. . . 40
5.7 Comparison between MAM and canonical SOM training processes. 42
5.8 Comparison of feature-maps trained using the MAM and canonical methods. 43
5.9 Comparison of feature-map structure. 43
5.10 Comparison between code-books and MAM visualization. 44
5.11 The propagation of free-association signals. 45
5.12 Illustration of Free-Association 45
5.13 Illustration of the proposed SOM training method. 49
5.14 Comparison of feature-maps resulting from various training methods. 50

6.1 The braided practise and its major components. 54
6.2 Realization→Interpretation 56
6.3 The nested realization→interpretation loop. 58
6.4 The realization→interpretation loop of the embodied artifact-as-process. 58
6.5 The components of meaning in an embodied artifact-as-process. 62

Preface

Here I am, face to face with the sublime, completely in awe of the truth before me. What is there to say? Once one is looking right into the face of the sublime what is the purpose in saying anything? What can a statement hold that has value with this before you? What is the meaning of art in this context? It is the drive to understand. Art is a way of understanding, a way of playing with those signs that make up our reality. Art is playing with ideas, making mistakes, creating processes, learning and becoming. I have realized I want to make art that lives. To make art that has an existence, a meaning, beyond me—the creator. To create work that continues to evolve and is connected to the play of signs from other systems of knowledge. For the piece to live it has to have meaning, a place beyond the intentions of the artist. It is not just a manifestation of ideas but a living entity. It cannot be summed up or reduced to a "concept". It is so attached to those systems around it that it would not exist without them.

Excerpt from "Art in the Face of the Sublime" Bogart (2001)

Chapter 1

Introduction

As a youth growing up in Vancouver, BC, I had always felt that my interests were in some space between creativity, art and technology. In my mind there was a dichotomy between art on one hand and science and technology on the other. I had taken a single high-school art class and knew I was not interested in expressing myself through art. In 1998 I enrolled in the Emily Carr Institute of Art and Design[1] and completed the foundation year. During that year my work shifted from tactile to conceptual. I was already more concerned with *doing* than in expressing or representing.

I recall one project where the purpose was to create a precise reconstruction of some detail of the building. I had gotten the mix of the plaster wrong—leading to a brittle and unsuccessful reproduction. The following project was to take that previous artifact and transform it in some way. Due to my distaste for the object, the idea of destroying it came naturally. Somehow a pile of plaster dust and flecks of paint did not strike me as an art object; art is, after all, about the creation of objects. I encased a portion of the plaster dust and paint chips in a small acrylic cube. The object became secondary to the process. The artifact was just a physical trace—a proof of the destruction of an object that no longer existed.

The preface is an edited portion of a paper I wrote in my second year at Ryerson University[2] (Bogart, 2001). The assignment was to articulate what drove me to making art; a question that I had not given much thought to at that point. Soon after moving to Toronto it became clear that my place between art, science and technology was not unique. There is a historical precedent for the combination of art and contemporary technology.[3] My first exposure to this world was my attendance at the Subtle Technologies conference in 2000,[4] a conference where artists, scientists and

[1] `http://www.eciad.ca/`

[2] `http://www.ryerson.ca`

[3] Any tool could be considered a technology. Here contemporary technology refers to modern technology such as electronics, computers, robotics, etc.

[4] `http://www.subtletechnologies.com/`

technologists come together. This led me to the Interaccess Electronic Media Arts Centre,[5] a 25 year old artist-run centre focusing specifically on art practises making use of contemporary technology. It was there that I found my artistic community—firmly rooted in the artworks and approaches of artists such as David Rokeby and Norman White.

Since then it has become clear that my interest in art is less a drive to express myself than a way of understanding the world—a method of asking questions, and imagining the possible. This thesis reflects my place between the worlds of science, technology and art. The text is an account, a study of, my creative process in the construction of the artwork "Memory Association Machine" (MAM). What MAM *does* is of central importance in this study. MAM is not just an object, but also a process. It is a way of reconsidering artistic authorship, the art object and its context. This reconsideration is more important than my own personal expression.

The initial drive for the project was to create an autonomous system whose form[6] is as independent of my intention as possible. I sought to create an artwork whose form is influenced as much by *itself* as by me. MAM cannot be reduced to its appearance. The appearance is the perceptible trace of what the system is *doing*. I refer to this trace as the *external properties* of the system.

Through the process it has become clear that the more I attempt to remove my control, over the system's external properties, the more mechanisms I implement. As more mechanisms are implemented more of my intention is encoded. The result is a paradox where the only way of reducing my influence is by more deeply integrating my intentions in the system. In order to completely remove my intentions, I would have to be causally disconnected from it. This would result in nothing at all—not an object nor a process. In order to lower my influence, I invite the system, and its physical context, to engage in the process.

Site-specific art is a movement where meaning is situated in, and dependent on, the artifact's context. That context may be physical, political or temporal. The artwork is connected to the world around it. This movement sets the precedent for the collaboration between author and context in MAM. Context is defined as those visual properties, perceivable by MAM, that make its place in space and time unique.

The site-specific artist chooses a site and then develops the artwork in relation to that site. The location is often a public place that invites the community to experience, and form a relationship with, the artwork. An issue with some early minimalist site-specific artwork is the artist having a differing conception of the site than the community who shares it. "Titled Arc", by Richard Serra, was removed from its site due to public pressure. Since site-specific artwork is often a physical artifact, it remains static even as the physical, political and temporal site changes.

MAM is an enacted process. The purpose of the process is simply to *create* a relationship with its context. It changes its external properties in response to changes in its context. The artistic

[5] `http://www.interaccess.org`

[6] Form here is referring to the shape of the system, not the process itself, but the output of the process.

process, of the creation of MAM, is a fusion of attributes of conceptual, site-specific and electronic media art practises.

MAM is an autonomous responsive site-specific installation. The system is autonomous in that it operates without the need for an interactor. The system is responsive as it is continuously collecting sensor impressions[7] of its context. MAM is site-specific as the system grows its own external properties from its context. The mechanism that relates MAM to its context is the artist's encoded intention. This mechanism is based on three parallel processes: *perception*, the *integration* of experience,[8] and the *free-association* through the field of remembered experience (memory). These processes allow the system to consider elements of its context in relation to one and other and originate associations between those elements.

As this research is contextualized within an artistic practise I am concerned with the qualitative aspects of the computational process. The research is grounded in a "reflective practise" (Schön, 1983) that is dependent on "tacit knowledge" (Polanyi, 1983). This research is a braided practise where artistic techniques are combined with artificial intelligence (AI) in order to implement a model of memory and creativity developed in cognitive science. A formalization of the cycle of practise has formed through the process of constructing MAM, and influenced by the philosophy of Merleau-Ponty (1968). This cycle is composed of two operations: the physical realization of the artist's action and the immediate interpretation of the results of that realization. The artifact is continuously being altered, refined and reconsidered.

MAM is implemented through a computational process. The computational process encodes my artistic intention in software. The software development is centrally concerned with the realization of the concept. A secondary concern is the stability of the system for long-term exhibition. MAM is written in the Pure Data (Puckette et al., 1996) visual programming system.[9] A significant character of this research is that all software used in the development of the project is FLOSS:[10] a term that encapsulates both the "Free Software"[11] (FL) and "Open Source Software"[12] (OSS) movements. FLOSS is used for the operating system that runs MAM, as well as the writing and statistical tools used in this research.

This research project contributes to electronic media art practise through a novel combination of AI and a cognitive model of creativity. The logging of the creative process through the creation of MAM adds a deep account of my experience working with AI techniques in an artistic frame. The

[7] A sensor impression is data, from the physical world, which is processed and perceived. In the case of MAM sensor impressions are visual images.

[8] An experience is the result of perception considered in light of memory.

[9] For more information on Pure Data, as it pertains to MAM's development, see Appendix A.

[10] There is one unfortunate exception to this statement. The graphics driver used in the installation of MAM is a proprietary driver provided by Nvidia®.

[11] `http://www.fsf.org/`

[12] `http://opensource.org/`

account serves as an example of what is technically and artistically possible in a FLOSS oriented artistic practise, as well as a case study in the creation of large projects in Pure Data.

The next chapters discuss the related literature. Chapter 2 contains a detailed description of the artistic practises used in the research, including electronic media art (Wilson, 2002; Manovich, 2001), site-specific art (Kwon, 2004), and their fusion in this practise. Chapter 3 discusses theories of creativity. Particular attention is spent on individual creativity (Boden, 2004; Gabora, 2002b). Chapter 4 is a short survey of creative machines (Cohen, 1979, 1995; Rokeby, 1990, 2001; Legrady and Honkela, 2002). Chapter 5 contains a detailed description of the implementation of MAM and its relation to the ideas introduced in Chapter 3. Chapter 6 is a description of the artistic process resulting in MAM and a discussion of the formalization that has resulted from that development. The research is concluded in Chapter 7. Additionally the thesis includes both the journal that contains my reflections through the progress of MAM (Appendix B) and the subversion log that contains my day to day notes through the process of research (Appendix C).

Chapter 2

Artistic Practise

A number of contemporary artistic practises are concerned with a deconstruction of the relationship between the artwork and the author. This research braids the practises of conceptual, site-specific, and electronic media art. The essence of conceptual art is that ideas themselves can be art. The art object is a trace of the true artwork—the concept. In site-specific art the artwork's meaning is dependent on the context of the work.[1] Such artworks are often installed in a public setting and refer to elements of that site. Conceptual and site-specific art are embedded in the contemporary consideration of art practise. The essence of electronic media art, also known as "new media"[2] (Manovich, 2001), and "information art" (Wilson, 2002), is that the material of the artwork is, or depends on, a technological apparatus.

As in any categorization of culture these practises have blurry boundaries. Wilson (2002) states that "[a]rtists resist categorization. Artworks are typically multi-layered, addressing many themes simultaneously. Many artists purposely try to confound preexisting categories". The purpose of this chapter is to communicate the foundation of this artistic enquiry. Artistic practises are complex and evolve organically. I consider the essential ideas at the centre of each of these movements in order to clarify their influence on my research. Additionally interconnections or overlaps between the categories are discussed to provide breadth. The following sections describe the central artistic practises to illustrate the shifting role of the author, viewer and context.

2.1 Conceptual Art

> Ideas can be works of art; they are in a chain of development that may eventually find some form. All ideas need not be made physical. (LeWitt, 1999)

[1] Kaye (2000) provides a background on performative site-specific practises.

[2] The author rejects this term due to its alignment with commercial design.

In conceptual art the artistic idea is the artwork. The status of art object is rejected for artistic purpose: "Conceptual art was conceived as a democratic means of making art ideas cheap and accessible by replacing the conventional 'precious object' with 'worthless' and/or ephemeral mediums such as typed sheets, xeroxes, snapshots, booklets, streetworks" (Lippard, 2001). The trace of conceptual art is still collected and fetishized. The object has survived "...the most consequential assault on the status of that object: its visuality, its commodity status, and its form of distribution" (Buchloh, 1990).

In 2002 I saw a collection of the conceptual artworks of Yoko Ono at the Art Gallery of Ontario[3] (AGO). A number of the pieces were intended to be interactive. The concept required the viewer to take some action in order to "complete the work of art" (Kaji-O'Grady, 2002). For example "Ceiling Painting (YES Painting)" (Yoko Ono, 1966), consists of a ladder, a sheet of paper with the word "YES"[4] typed on it, and a magnifying glass hanging from the ceiling on a metal chain. The word "YES" can only be read if the viewer climbs the ladder and views the text using the magnifying glass. When the work was originally exhibited the viewer was expected to climb the ladder. In the AGO exhibition, the ladder was placed on a plinth. This sent the signal that this was an art *object*, and not an artwork to be *experienced*. Even though "Ceiling Painting (YES Painting)" is regarded as a piece of conceptual and interactive art, contemporary exhibitions of the work elevate the object above the concept—at least to the viewer who is not already familiar with the work. To a familiar viewer the object is a trace of the concept—without the experience of interaction. In this case, the ideals of conceptual art have not broken free of the museum's hold on the object.

"Ceiling Painting (YES Painting)" requires the viewer to "complete" the work through their interaction with its material. This dependence on interaction shows a fundamental shift in the relationship between the author, the object and the viewer. Conceptual artists have "reflected upon the construction and the role (or the death) of the author just as much as they redefined the conditions of receivership and the role of the spectator" (Buchloh, 1990).

2.2 Site-Specific Art

Site-specific artwork gives "...itself up to its environmental context, being formally determined or directed by it" (Kwon, 2004). Where conceptual art can be considered a rejection of the object, site-specific art can be considered a rejection of the gallery context. Traditionally the artist chooses a site and the artwork is meant to be meaningful in relation to that particular site. Serra describes the importance of the site in relation to the artwork: "to remove the work is to destroy the work" (Kwon, 2004). These artworks are often static and their relationship to site is constructed by the artist.

[3] `http://www.ago.net/`

[4] I assume the word was typed in upper case, due to the title of the work, but I cannot be sure as I have never climbed the ladder.

Although the context changes the artifact does not change in response. In some cases the artwork becomes an impediment to the community, as in the case of Richard Serra's "Titled Arc"—installed in the Federal Plaza in New York City in 1981. The work was removed in 1989 due to community complaints. Such cases have brought together community and site-specific art practises that result in artworks created in collaboration with the community sharing its site.

In 1963 Hans Haacke created "Condensation Cube". The work consists of a transparent acrylic cube, water and the "climatological conditions of the environment". As the cube is lit by sunlight the temperature inside the cube increases. The humidity inside the cube and the temperature difference between the inner and outer environment causes the water to condense inside the cube. The artwork depends on its physical environment to determine its form. The state of the water is a central component of the artwork. It is an object, but an object that changes its form depending on its surroundings. The work can be placed in any context and its form is still dependent on the state of that environment. Haacke considers the relationship between the object and its environment:

> A "sculpture" that physically reacts to its environment is no longer regarded as an object. The range of outside factors affecting it, as well as its own radius of action, reaches beyond the space it materially occupies. It thus merges with the environment in a relationship that is better understood as a "system" of interdependent processes... A system is not imagined, it is real. (Lippard, 1997)

From this statement it could be induced that the object (the cube itself) becomes a component of a different artwork as its environment changes from gallery to collection. The object does not change but the "system of interdependent processes" does. Haacke considers the artistic concept as not just "imagined", but as implemented in the system of processes, and therefore real.

There are two different approaches to site-specific artwork. The first, of that "Condensation Cube" is an example, results in an object whose external properties change in response to the environment. The object becomes integrated into its environment, to varying degrees, as a participant. The environment inside "Condensation Cube" is influenced more by its environment than visa versa. The second approach results in static artifacts that are constructed by the artist to have a specific relationship with their environment. In order to relate to the environment the artwork references aspects of its context. The creation of the artifact often nurtures or inspires a dialogue between the context, the artist and the viewer. These works usually concentrate on the social, historical and cultural aspects of context. In this approach public works[5] are produced in a collaboration between the artist and the community. The resulting artwork contains symbols, references and the stories of the people who share the site.

[5]Some site-specific community-based artworks are produced in reference to some event. They are meant to bind the contemporary community to the events of the past. A powerful example of such work is the "Hamburg Anti-Fascist Memorial" conceptualized by artists Jochen Gerz and Esther Shalev-Gerz. `http://www.rrz.uni-hamburg.de/rz3a035/antifascist.html`

2.3 Electronic Media Art

As my experience in a community of electronic media artists has grown so has my conception of art and my understanding of its history. A definition of the term 'electronic media art' will be constructed through an analysis of the various labels that refer to largely the same artistic approach: new media, information arts, and electronic media. This analysis contextualizes my artistic practise in the approximate forty year history of artists working with contemporary technology. The term "media" is highly present in these labels. The Oxford English Dictionary provides two definitions pertaining to this discussion. Firstly the definition of "medium" is as follows:

> c. Any of the varieties of painting or drawing as determined by the material or technique used. Hence more widely: any raw material or mode of expression used in an artistic or creative activity. (Oxford English Dictionary Online, 2008)

Medium is used to describe the material or "mode of expression" of artistic activity. Secondly the dictionary provides the following definition of "media":

> 2. Computing. A physical object (as a disk, tape cartridge, etc.) used for the storage of data. Cf. MEDIUM n. 4e. (Oxford English Dictionary Online, 2008)

The media is considered a physical object created for the express purpose of storing computational data. Through the lens of these two definitions let us define "new media" as novel material used for artistic expression and/or the storage of computational data. Manovich (2001) makes a distinction between artistic artifacts that make use of new media for "exhibition and distribution" and those that result from a production process dependent on new media. Examples of artifacts using new media for exhibition and distribution are films and texts that are distributed on-line. These artifacts do not depend on computational media to exist. A film on celluloid is almost identical to its copy distributed on DVD.[6] A text printed on a press resembles the text seen on a computer display. These are examples where new technology is used to simply redistribute old media, rather than invent a "new media". According to Manovich the creation of a new media would require a greater use of technology in the production of artifacts, rather than in their exhibition and distribution. Manovich proposes five "Principals of New Media". The terms are Manovich's, but the descriptions are my interpretation of those terms:

1. Numerical Representation: all artifacts (text, sounds, images, moving images, etc.) manipulated by a computer system, regardless of origin, must be encoded numerically (usually digitally). The result is that artifacts can be constructed or manipulated *computationally*.[7]

[6] The purist may argue the celluloid version is far superior in quality and resolution.

[7] Digital data does not imply "numerical" representation. It is simply a pattern of data that is encoded, and interpreted, as numerical. Perhaps a better term would be *electronic* representation where any media can be encoded in electrical patterns.

2. Modularity: every artifact may include other artifacts. Components are easily exchanged between artifacts and can be manipulated in context or independently of their parent structures.[8]

3. Automation: once computational manipulators and generators are encoded they can run without the need for human intervention.

4. Variability: the material is mutable. An artifact can be changed in place—replacing the original. Any number of copies of an artifact are possible—each identical to the original.[9]

5. Trans-coding: material can be converted from one medium (e.g. an image) to another medium (e.g. a sound). This is a result of the numerical representation of new media.[10]

These principals form a sketch of new media. In relation to the broader conception of electronic media the core features are:

- Trans-coding, which enables representation
- A formal language with which to describe the trans-coding process (e.g. mathematics or computation)
- Automation, where the mechanism is able to execute processes

The labels "information arts" and "electronic media" are used to broaden our conception of electronic media art practise. I counter the label "new media" for two reasons: firstly the label's dependence on the term "new" is problematic as technology is continuously changing. A case in point is the characterization of digital networks as "New Media" even though they have been in use since 1973.[11] Secondly the label has been widely accepted by commercial design, which has left little room for the conception of new media as fine art.

The label "Information Arts" has been put forth by Wilson (2002). "Information Arts" is more appropriately concerned with the properties of the artwork being created with technology, than being concerned with the notion of newness. Wilson (2002) organizes artistic sub-practises into categories based on "scientific disciplines and areas of technology". Those categories that pertain to this research are:

- Algorithms, Mathematics and Artificial Life[12]

[8] This may be a global property of data encoded in digital form, which begs the question: Do all digitally encoded artifacts exhibit this "modularity"?

[9] This may be a global property of data encoded in digital form.

[10] This term could supersede "numerical representation" as a numerical, or electronic, representation must always result from encoding or trans-coding. Without trans-coding numerical representation would not be possible.

[11] The inception of ARPANET.

[12] Wilson (2002), Chapter 4

- Digital Information Systems and Computer Media[13]

"Information Arts" considers the relationship between art research, science and technology. The artworks listed in the book "Information Arts" Wilson (2002) run a wide gamut but can all be considered art-research at the intersection of science and technology. The label covers such a wide range of approaches that it extends beyond the scope of this research. I consider *processing* (or trans-coding) more appropriate than "information" as the central principal of electronic media arts.[14]

The third label "Electronic Media Art" is not associated with a particular researcher/theorist. It is largely synonymous with "computer art" and "digital art". Artists such as Normal White have constructed artworks composed of electronic components since the 1960s. Artworks considered under this label are largely electronic. As the digital computer became more common many artists shifted their interest in analog electronics to the digital computer. There are artists that work exclusively in electronics to produce complex artworks—rejecting the use of computers. Electronic media art is the practise of using electronics and computers for artistic purposes. A number of arts organizations, which were conceived before the term "new media" was popularized, contain the word electronic in their names. Notable examples include: Ars Electronica (1979 to present), The Interaccess Electronic Media Arts Centre (1983 to present), The International Symposium on Electronic Art (1988 to present), and the Dutch Electronic Art Festival (1994 to present). A key importance of this label is its reference to the history of electronic media art.

Now that we have a sense of the various characteristics of electronic media art I will discuss two subsets of the practise: interactive electronic media and electronic media installation art. Yoko Ono's "Ceiling Painting (YES Painting)" highlights the history of interactive art not enabled by electronic media technologies. This history includes conceptual and performance art, where the viewer is required to participate in order to complete the work (Alberro, 2003), and kinetic sculpture, where the artwork is a mechanical system that is physically manipulated by the viewer (Popper, 1968). In interactive electronic media art the behaviour of the viewer is sensed by the system and results in a change of its external properties. The work is a "transforming mirror" (Rokeby, 1995), where the viewer recognizes their behaviour as it is reflected back at them through its external properties.

Installation art represents another shift in the artist's conception of the art object. The artifact is expanded into an entire environment. The artist transforms the exhibition space into an artwork that encompasses the viewer. In interactive electronic media installation art the environment is the surface of interaction. Interface is often provided by environmental sensors[15] and projections in place of standard interfaces, such as mice, keyboards and monitors.

---

[13] Wilson (2002), Chapter 7

[14] This consideration highlights my interest in the process (trans-coding) over the artifact (representation / information).

[15] Sensors that measure the position, temperature, weight, image, and action of the viewers/interactors.

The following features of electronic media art are based on Manovich's principals of new media, combined with attributes of Wilson's categorization of information arts and a history of electronics in art practise:

- The artwork is intrinsically dependent on its technology. It is not old media simply displayed and distributed using new[16] technology.
- The central principals of the practise are *trans-coding*—which implies the representation of, and formal operation (processing) on, those representations—and *automation*. Both of these principals can be reduced to the phrase: *Autonomous Processing*.
- The art practise follows from the history of artists using contemporary technology.
- The practise has a strong interdisciplinary connection to both science and technology.
- The primary material of the artwork is a computational process running on contemporary technology.[17]
- In interactive electronic media art installation standard computer interfaces are replaced by environments and the viewer is expected to participate.

2.4 The Artist / Author

Historically, the artist is "imagined as an isolated figure of exceptional creative powers who suffers for his art" (Barker et al., 1999). There is a mythology surrounding the "creative genius" and the artifacts the, predominantly male, artist creates. The artifact is a record of genius and often collected and fetishized. Conceptual art is a movement away from the object itself, towards the ideas that give the artwork purpose. It is then the artist's role to generate ideas, which are communicated through, or implemented in, the artifact. The result is an emphasis on the process over the object. The artist is primarily concerned with ideas. Physical materials are used as a means to express, explore or manifest, those ideas. Site-specific art is a rejection of the, traditionally contrived, context of the art gallery. The artist is concerned with the relationship between the artifact, or concept, and its context. The physical, environment, social, political, and historical contexts are all potentially integrated into the creative process and reflected in the artifact or concept. The artist is working within a broader context, rather than in isolation. This is particularly true of community based site-specific art practises. The electronic media artist follows from this historical arc. The importance of the process and concept in conceptual art is a strong precedence for electronic media art where

[16] In this case "new" refers to technology created after the inception of the media itself. For example a digital projector is "new" technology when considering the film media presented on it.

[17] e.g. A network, digital computer, microprocessor or electronic circuit.

the artist encodes concepts directly in computational processes. As the artist is working with the products of engineers (technologies) she is engaged in a collaborative process where the artwork is a result of the interaction between artistic intention and technology. The role of the artist has shifted through the practise of electronic media art:

- The artist works in the realm of ideas as much as the realm of materials. In some cases there is no physical material at all.
- Artworks are often systems that are difficult to collect, purchase and conserve.
- The artist actively seeks to reduce her control over the artwork in order to encourage the participation of the technology, the viewer, the community or the context.
- The practise is inherently collaborative due to its dependence on technology.

A natural consequence of the shift from physical artifact to process (Possiant, 2007)[18] is the artist's engagement with other influences. This shift is highly related to the "birth of the reader" (Barthes, 1977) where the reader is considered as important as the author in the interpretation of a work. In interactive electronic media art the viewer is often the centre of attention. The artwork seeks to engage in a dialogue with the viewer. The emphasis on the reader/viewer is stated in the context of conceptual art: "The artist may not necessarily understand his own art. His perception is neither better nor worse than that of others" (LeWitt, 1999).

The preceding section is an analysis of the intersections between the artistic practises of conceptual, site-specific, electronic media, installation and interactive art. These practises have reconsidered and reshaped the relationship between the creator and the artifact, in order to integrate the context and the viewer in the process.

2.5 Artificial Intelligence

Artificial intelligence (AI) is included in this section because of its intersection with electronic media art practise. In this braided practise AI is considered a thread of equal weight. Artistic practises are expected to shift the use and consideration of AI methods and AI techniques are expected to transform the artistic practise.

What is artificial intelligence? Barr and Feigenbaum (1981) provide a general definition of AI as a "part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior—understanding language, learning, reasoning, solving problems and so on." Stephen Wilson considers the relationship between AI and art:

[18] Possiant speaks of the shift from material to interface, but the very idea of an interface implies a process. The shift to interface requires a shift to process.

> Artificial intelligence is one of these fields of inquiry that reaches beyond its technical boundaries. At its root it is an investigation into the nature of being human, the nature of intelligence, the limits of machines, and our limits as artifact makers.[19] I felt that, in spite of falling in and out of public favor, it was one of the grand intellectual undertakings of our times and that the arts ought to address the questions, challenges, and opportunities it generated. (Wilson, 1995)

AI is primarily concerned with developing techniques that endow machines with intelligent behaviour. Specifically AI tends towards the construction of systems that exhibit artificial reasoning for the purpose of problem solving.[20] My use of AI is to move forward electronic media art in the construction of artworks that are meant to relate themselves to their context. As suggested by Wilson's quote, AI, as a set of ideas and techniques, can certainly be expanded beyond the discipline's normal limits. A prime example is Agre's critique of AI that aims to demonstrate "...a positive method of inquiry that maintains a dialogue between the philosophical and technical dimensions of AI research" (Agre, 1995). One of the contributions of this project is to provide an example of how AI can be considered in light of artistic practise. AI is bound to the preceding discussion of artistic practise in the following ways:

- AI is dependent on the computational technologies used to implement its techniques.
- AI practise constructs systems that autonomously process data in order to accomplish their assigned task. This task could certainly be an artistic one, which places AI techniques near the centre of the principals of electronic media art—automation and processing.
- AI has proven to be a practise that is explored by electronic media artists.[21]
- AI is inherently about ideas, rather than the physical media in which they are implemented. This feature is highly congruent with conceptual art.
- AI systems are meant to operate in a context. A "problem domain". This relationship to context could be considered in the light of site-specific art practises.
- AI is, by definition, invested in a shift of the role of the author (programmer) as the artifact (system) is expected to behave not only autonomously but also intelligently.

In this chapter I have discussed a number of artistic disciplines, and considered AI in relation to the key attributes of electronic media art practise. The essence of conceptual art could be considered

[19] Wilson's characterization of creators as "artifact makers" indicates that even in the technological arts that the object can be considered the central focus—over the process the artifact implements.

[20] For discussion on the metaphors used in AI, specifically "problem solving", see Agre (1997).

[21] A selection of artworks are discussed in Chapter 4.

a rejection of the notion that the artwork can be reduced to the physical object that manifests its concept. Conceptual art emphasizes the artistic idea over the artifact. Site-specific art is centrally concerned with the relationship between the artwork and its historical, social, physical and political context. Electronic media art is concerned with the use of technology for the purpose of artistic enquiry. The essence of the practise is the autonomous processing of computational representations where the intentions of the artist are implemented in a computational form. The dominant uses of technology are often subverted and critiqued through the practise. Electronic media art, as it follows from these other artistic disciplines, is largely in line with the features of AI research. This research is a braiding of these different approaches into a single unified practise.

Chapter 3

Theories of Creativity

The artistic concept of "Memory Association Machine" (MAM) is a machine that relates itself to its context—where its external properties are formed through that relationship. The machine must *create* its relationship to its context. MAM is a machine that is intended to act *creatively*. This chapter will discuss two conceptions of creativity that have been used in the development of MAM. The Oxford English Dictionary Online (2008) defines "creativity" as:

1. Having the quality of creating, given to creating; of or pertaining to creation; originative.
2. Spec. of literature and art, thus also of a writer or artist: inventive (cf. INVENTION 3b), imaginative; exhibiting imagination as well as intellect, and thus differentiated from the merely critical, 'academic', journalistic, professional, mechanical, etc., in literary or artistic production. So creative writing, such writing; also freq. in the U.S. as a course of study.

The first definition points to the importance of originality in relation to creativity. In the case of MAM the external properties of the system should be an original result of its attempt to form a relationship with its context. The second definition requires more than originality, including imagination and intellect. Is it possible for a machine to be autonomously creative? Boden (2004) argues that machines can be *"considered"* creative in the same way that machines can be thought of as intelligent according to the "Turing Test" (Turing, 2004). This research is concerned with originality and not imagination nor intellect. Boden (1998) states that "[c]reativity is a fundamental feature of human intelligence...". This chapter will discuss creativity as an integral part of the human mind and not just of intelligence.

The central sources of theory on creativity concerning this research are the work of Margaret A. Boden (Boden, 1994, 1998, 2004) and Liane Gabora (Gabora, 2000, 2002a,b). Creativity will be considered an individual, rather than social, operation.[1] MAM is conceptualized as an individual

[1]For a social model of creativity see Csíkszentmihályi (1988).

entity rather than a model of social interaction.[2] Boden's characterization of creativity is used as an evaluative tool for MAM as well as a selection of creative machines discussed in Chapter 4. Gabora provides a model of creativity that is central to MAM's relationship to context.

3.1 Dimensions of Creativity

Margaret A. Boden's conception of creativity is the primary method of characterizing the potential creativity apparent in MAM and other creative machines. Boden (2004) defines creativity as "...the ability to come up with ideas or artefacts that are *new*, *surprising and valuable*" (Boden's italics). I refer strictly to artifacts rather than ideas in this discussion. Boden describes two types of novelty (newness): P-novelty and H-novelty. P-novelty is when the artifact is *new* for the creator, but not new when considering the history of artifacts.[3] H-novelty is attained when the result of P-novelty proves to be new in consideration of the history of artifacts. An artifact can be *surprising* in that its qualities are somehow unexpected in the context of its construction. Boden (1998) describes *valuable* as "interesting, useful, beautiful...". Boden's ellipses imply that the criteria for determining value is not global but domain dependant. I interpret the three classes of creativity specified by Boden (2004) to refer to artifacts rather than ideas:[4]

- Combinational creativity results from linking together existing artifacts, or their properties, in the production of an original artifact. These combinations should be "improbable".[5] For example an artist attaches a handle, associated with one class of vase, to a vase of another class.

- Exploratory creativity is accomplished by moving through a structured space of possible results. A space of possibilities could be the space that all vases could occupy. Exploring this space would be exploring all the options of what a vase could be.

- Transformational creativity is the alteration of the structured space of possible results. The new vase would expand the space by changing what is considered a vase. For example a vase meant for a zero gravity environment may not be initially recognized as a vase—lacking the opening at the top and flat bottom.

Boden's key terms (P-novelty, H-novelty, new, surprising and valuable) all require an external evaluation of the artifact. The criteria for considering an artifact as surprising and valuable is socially

[2] A social model of creativity, implemented as a creative machine, is discussed by Saunders and Gero (2001a,b, 2002).

[3] The "history of artifacts" is the collection of all known artifacts created by any practitioner.

[4] Boden (2004) tends to use "ideas" as atoms of creativity. The following descriptions are my interpretation of the analog of each point when considering artifacts rather than ideas.

[5] Evaluating the artifact as "improbable" is analogous to considering the artifact as surprising.

constructed and domain specific. With all this concentration on evaluation, what is the initial seed that originates ideas and artifacts that can then be evaluated?

Creativity can be considered a two step process. Some *originator*, the kernel of creativity, creates a *new* item. At this point the item is only new in terms of the mechanism that originated it, and may not even be new for the creator. This item then goes through a process of evaluation that filters all but the most new, surprising and valuable items. Boden's argument can be summed up in one statement: an artifact can only be considered "creative" if it has been successfully *evaluated* as such.[6] These two steps are *both* required for a creative result. Emphasizing one over the other results in a partial model of creativity. The reduction of creativity to evaluation would be a significant error. The result of the first step, in isolation, may not originate something *highly* surprising or valuable but certainly *could* originate something *new*. If we were to execute only the second step, evaluation, then *nothing* would originate at all. Reducing creativity to a two-step process, as I have done here, is an oversimplification. There are multiple iterations of origination and evaluation. This would explain the difficulty in locating the "seed" of a creative process in its nest of recursive evaluations.

Boden offers a definition of what is creative—an idea or artifact, that is new, surprising and valuable. This definition is acceptable when the entire creative process is collapsed into a final result, but what of the whole creative process? Failures and dead-ends are part of the creative process, and contribute to the result, but would not be considered creative in Boden's conception. The notion that creativity permeates a process, through constant refinement, is closer to the conception of creativity proposed by Gabora (2002b). The focus of this research is not on the evaluation of artifacts but the exploration of a computational implementation of a creative process.

3.2 The Cognitive Mechanisms Underlying the Creative Process

The design of the processes that relate MAM to its context are directly inspired by "The Cognitive Mechanisms Underlying the Creative Process" (Gabora, 2002b). It was her talk entitled "Dawn of the Creative Mind: The Origin & Evolution of Innovative Ideas" at the School of Interactive Arts and Technology[7] that initiated the transformation of a vague idea into a cohesive artistic and technical possibility. The model is highly suited to computational implementation. This section begins with my interpretation of her model of human creativity, and traces the ideas through a selection of her publications. The selections are made to contextualize the research in terms of Gabora's ideas. This interpretation is developed through the practise of developing MAM.

[6] For Boden creativity is the result of the two-step process, so before the evaluation the item should not be considered creative at all, but simply as an unclassified response.

[7] The talk occurred on February 21st 2007, at Simon Fraser University, Surrey.

3.2.1 Sparse, Distributed, and Content Addressable Memory

The key to Gabora's theory of creativity is a cognitive model of human memory. This model is characterized by three features: memory is sparse, distributed (but constrained), and content addressable. Consider a space of all possible stimuli. This is the stimulus space. The space occupies the same number of dimensions as the stimuli. A dimension is an independent axis on that the stimuli can vary. A single number, of any scale, is one dimensional. The number can only vary by moving up or down. For example, a digital sound file is multidimensional. The original sound is broken up into a number of samples. Each sample is an amplitude at a particular moment in time. A single sample is one dimensional; it can only vary in amplitude (up or down, loud or quiet). The digital sound is an array of these samples—organized in time. To obtain the number of dimensions of the sound the number of dimensions of each sample is multiplied by the number of samples. A one second sound file, at CD quality, has 44100 samples, and therefore can vary on 44100 different dimensions.

A stimulus space designed for these sounds would have 44100 dimensions. Each sound would be a single point located in that space. The space implies all *possible* variations of the sound file, as every variation would be in a different location. A practical stimulus space would occupy a very large number of dimensions.[8] Human experience could vary in so many ways that storing all possible experience would be impossible. The two primary features of memory are that it is both *sparse* and *distributed*. Memory is sparse because the number of stimuli that can be stored is significantly less than the number that can be perceived. Figure 3.1 on page 19 (a) represents all possible stimuli implied by a discrete stimulus space.[9]

The grey area represents the discrete stimulus space. The unfilled circles represent the location of each possible stimulus. As it would be impossible to store all of the possible stimuli that *could* exist, memories are separated from one and other. The possible memory locations in Figure 3.1 (b), shown by unfilled circles, are sparsely distributed over the space of all possible stimuli. The unfilled circles in Figure 3.1 (b) represent memory locations that have not been associated with stimulus. The filled circles represent memory locations where stimuli is stored. The memory space is the space of possible memories where stimulus can be stored. Similar stimuli are associated with nearby memory locations.

In non-distributed memory each location in the memory space stores one particular stimulus. Figure 3.2 on page 19 (a) represents non-distributed memory. Each stimulus (a hatched circle) is stored in a single memory location. The location of a memory is addressed by its position in the stimulus space. In fully distributed memory each stimulus is stored equally across all memory locations. Figure 3.2 (b) illustrates the stimulus occupying every memory location. In non-distributed

[8] Imagine on how many dimensions even a banal experience, like riding the bus to work, could vary?

[9] A continuous, rather than discrete, stimulus space implies an infinite number of possible stimuli.

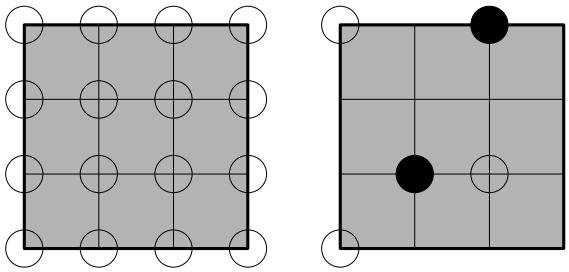

Figure 3.1: Memory is both *sparse* and *distributed.* The unfilled circles in (a) represent all possible stimuli that can occupy this space. The unfilled circles in (b) are *potential* memory locations. Filled circles in (b) are memory locations where stimuli is stored.

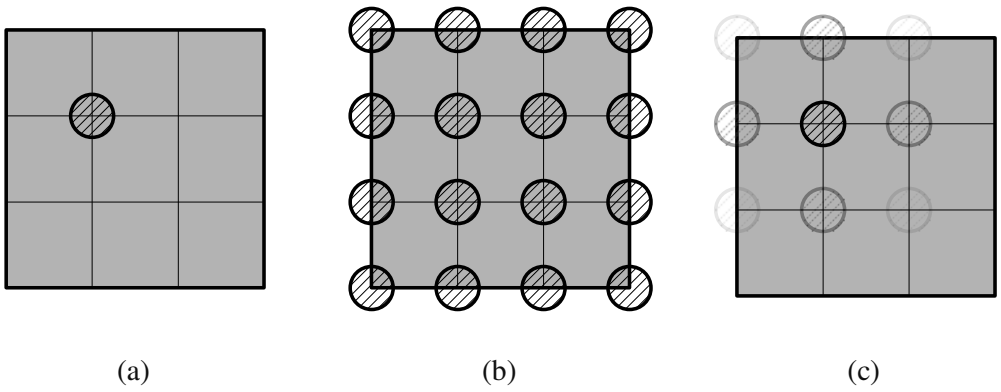

Figure 3.2: Types of Memory Distribution. Non-distributed memory (a) stores each stimulus, the hatched circle, in a single location. Fully distributed memory (b) stores each stimulus in every location. Distributed, but constrained, memory (c) stores each stimulus, to a degree, in a limited number of locations—shown here as transparency.

memory each location can be addressed. Each address refers to a single memory location. In fully distributed memory each stimulus is located in every location. Addresses are irrelevant because every address refers to the same stimuli.

According to Gabora neither of these models is feasible for human memory. The second primary feature of memory is that it is "distributed (but constrained)" (Gabora, 2000). A stimulus is stored in multiple, but not all, memory locations. The stimulus is stored in each location to a degree. Figure 3.2 (c) represents the storage of a particular stimulus in multiple locations of distributed, but constrained, memory.[10] Figure 3.2 (c) pictures nine memory locations. The stimulus is stored in the centre location to the greatest degree, represented as the transparency of the hatched circle.

A third feature of memory is that it is content addressable. In content addressable memory there is a "[s]ystematic relationship between the content of an experience... and the memory locations where it gets stored..." (Gabora, 2002b). A search within, non content addressable memory, requires every location to be examined for stimulus with the desired properties.[11]

Imagine a memory space, that is not distributed nor content addressable, as a set of lockers. The goal is: locate the locker (the memory location) containing a pair of red running shoes (the stimulus). The number of the locker (the address) is not correlated with the content of the locker. To find the red shoes we must examine the content of every locker. Content addressable memory adds two features to the lockers. Firstly they would have frosted glass doors that correlate each locker's content with its outward appearance (address).[12] The content addressable nature of the memory allows the search to proceed without the requirement of examining the content of each location. Without opening a single locker we can scan over their doors in search of red. The second feature is that the lockers reorganize themselves based on their content. All lockers containing shoes would be located near one and other.

What happens if the shoes we are searching for are actually orange and not red? We would search lockers containing shoes but not find any coloured red. If the query does not *exactly* match the address nothing is found—even if there is a correlation between the address and the content. Searching in a fully distributed memory would be pointless, as all lockers contain the same items.

This combination of distributed, but constrained, and content addressable memory is the foundation of Gabora's theory of creativity. In this type of memory *multiple* lockers can store the *same* pair of shoes to varying degrees. These lockers have doors that illuminate when they contain an item similar to the one being searched for. In our search for red shoes we would see an area of illuminated lockers. One particular locker would be brighter than the others. After opening the illuminated locker we would discover a pair of *orange* shoes. The memory illuminates the lockers

[10] This illustration does not represent the sparse nature of memory.

[11] Computer memory is usually not content addressable. In order to find a file on the system, every memory location must be searched until the file is found.

[12] The doors are frosted because they do not show the exact content of the locker but give an indication as to its contents.

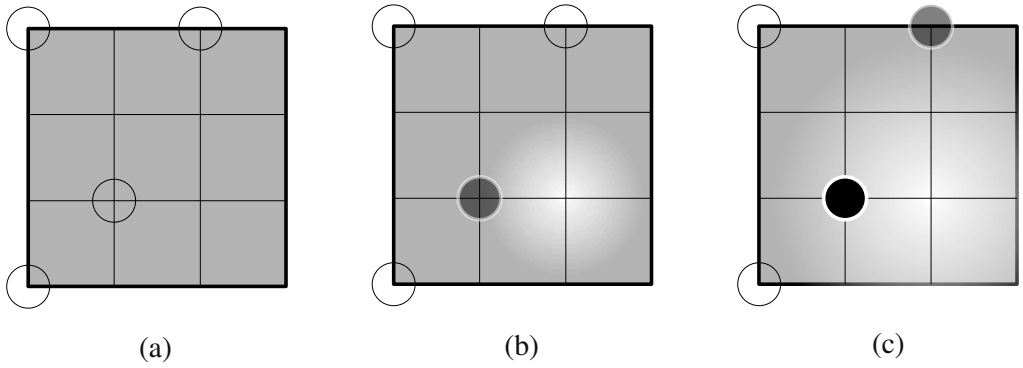

Figure 3.3: The memory space before any stimulus has been stored is pictured in (a). Stimulus activates one location using a narrow activation function (b). Stimulus activates two locations using a wide activation function (b).

based on the similarity between their content and the query. Even if we were to open one of the slightly illuminated lockers we would still find the orange shoes. The ability for the memory to retrieve items similar to those requested is the central basis of Gabora's conception of creativity.

3.2.2 The Structure of Memory

Figure 3.3 on page 21 (a) shows a memory space in that no stimulus has been stored. The space of all possible stimuli (the grey box) contains a number of locations where memories could be stored (the unfilled circles). In Figure 3.3 (b) a new stimulus activates a location in stimulus space—represented as a white gradient. The activation centre does not overlap with any memory location but its outer margin does. This results in the partial activation of that memory location (represented as the transparency of the filled circle). The degree of activation is inversely proportional to the distance between the initial activation and the memory location. In Figure 3.3 (c) two memories are associated with the stimulus due to a wider activation function. The nearest location is activated to a greater extent.

The content of memory locations are not overwritten. The stimulus is appended to the content of each activated memory location. Each location stores a combination of stimuli—each to a varying degree. This overlapping of stimuli results in a memory structure where memories are linked by similarity.

3.2.3 The Creative Thought Process

How does this memory structure explain creative thought? As memory is distributed there is no single location where a stimulus is located. All stored memories are a combination of other similar memories. Memory can never retrieve an exact copy of a previous experience. Every experience has been coloured by, and remembered in context of, other experiences. The structure of memory is effected by experiences that have occurred since that memory's initial storage. According to Gabora (2000) this memory structure allows the construction of concepts (abstractions of many experiences) as clusters of related memories. The concept of "bicycle" is an activation of memories associated with stimuli that share bicycle features. The linking of memories into networks of concepts is labelled a "world-view" in Gabora's vocabulary. Gabora (2000) considers this construction of concepts as a "creative act". A concept is an activation that binds a set of properties from various stimuli. This process is intrinsic to how humans process information. We are being creative in every aspect of our lives by integrating new experiences in the web of remembered experience.

In "Toward a Theory of Creative Inklings" Gabora (2000) describes a creative "inkling" as a "perturbation" of the world-view. A new experience links two islands of, until that moment, unrelated stimuli. The concept transforms those islands of thought. In this sense creativity is a new way of seeing, or understanding, the world. During every moment of our lives we're reconsidering and reintegrating our experiences and changing the world, both internal and external, through our behaviour. Gabora (2000) considers the world-view as in a state of potentiality. New concepts are always possible. It is very difficult to predict how a new experience will change the world-view. The result of the perturbation is a collapse of potentiality into a concrete (re)experience of memory.[13] The wider the activation function, used to store stimuli, the more memories would be evoked as a result of the perturbation.

What is the mental state analog of these mechanisms? The wider the activation function the more distributed memories would be. The result is that "streams of thought" (Gabora, 2000) are longer and tend to be more abstract. Abstraction occurs when an activation traverses a highly distributed memory—moving through diverse stimuli. The larger the gaps between islands of memories the more abstract the thought process. There is experimental evidence that mental states such as defocused attention (Dewing and Battye, 1971; Dykes and McGhie, 1976; Mendelsohn, 1976), and high-sensitivity (Martindale and Armstrong, 1974; Martindale, 1977), are correlated with highly creative individuals. These mental states may reflect a wide activation function.

3.2.4 Remembering as Recreating and Recreating as Remembering.

Gabora (2002a) explains some subtlety in the model discussed above. As a set of memories are

[13] A mathematical model of this notion of potentiality and collapse, in relation to human creativity, is available in Gabora and Aerts (2002).

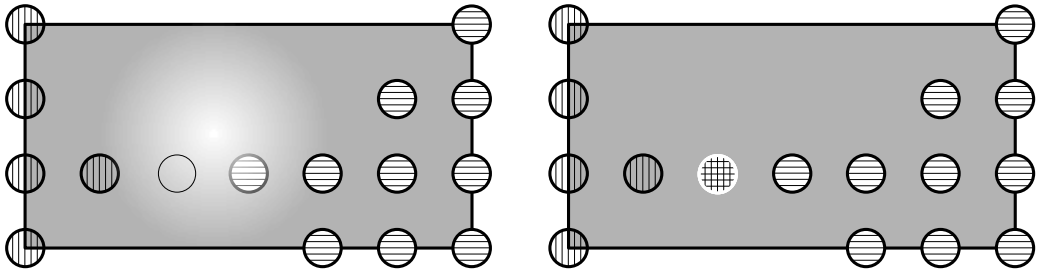

(a) (b)

Figure 3.4: The Bridging Concept. A memory space is occupied by two clusters of experience (a). The hatching indicates the cluster of each memory. Memories at the margin of both clusters, and an available memory location, are activated by a wide function. The activation combines the stimulus from both clusters and stores the result in the unassociated memory (b).

activated through the process of reminding the resulting blended stimulus is stored back into memory. If this activation occurs in a region where two clusters of stimulus are separated, a wide activation can link those clusters. Such an activation is pictured in Figure 3.4 on page 23 (a). The activation function is wide enough to activate memories at the margins of both clusters. The stimulus resulting from this activation is a hybrid of the stimuli stored in both clusters. Its storage in memory links the two clusters as illustrated in Figure 3.4 (b). This feedback loop constructs broad concepts through the creation of memories that bridge independent regions of experience. Concepts in Gabora (2002a) are considered the same as in Gabora (2002b) except for an explicit mention that memories are stored and retrieved from memory simultaneously. According to Gabora (2002a) the essence of creativity is not in the items of experience, memories, but in the network of concepts that organize them into a coherent world-view.

3.2.5 "Honing" The Thought Process

In "Cognitive Mechanisms Underlying the Creative Process" (Gabora, 2002b) discusses how associative and analytic modes of thought could be considered in light of her model. Individual memory locations contain specific stimuli. These memories are combined[14] to create concepts, and recall memories, at various levels of abstraction. Gabora (2002a) introduces the notion that the activation function could be controlled intentionally. A wide activation function is linked with associative thinking. A narrow function is linked with analytical thinking.

A typical creative process could be initiated with brain storming. A wide activation function is used to explore many possible ideas. As the creative process continues the activation function

[14]How memories are combined is a function of the character of the activation function.

narrows, leading to more refined and analytical ideas. By controlling the size and shape of the activation function the creator can shift between associative and analytical modes.[15] Following brain storming the artifacts are refined using a more and more analytical thought process.

In Gabora's model creativity is intrinsic to how the human mind makes sense of the world. The features of memory that make this model possible are:

- Memory is *sparse*. Many more stimuli are possible then there are locations to store them.
- Memory is *distributed* (but constrained). Stimuli are not stored in independent locations but are overlapped to a degree. Memories are linked in a structure.
- Memory is *content addressable*. There is a systematic relationship between the address (location of the memory) and its content. All memories need not be searched to find an item with the desired features.

These features allow the thought process to occur. A stimulus activates a certain area of memory. That activation stores the stimulus in nearby locations. Simultaneously memories are stimulated and the experience of their content is evoked. Memories contain a blend of stimuli that are interpreted in the context of current experience. Each activation sets off a cascade of activations through the memory. This propagation of activation is a thought process. The thought process both retrieves experiences and builds upon the overlapping structure of memory. Regions of overlapping activation become abstract concepts in that they encompass the properties of multiple stimuli. The creative process mirrors this cognitive thought process where the creator expands or contracts their activation function to attain associative or analytic modes of thinking. Experimental evidence shows that creativity may be correlated with broad activation of the brain (high-sensitivity[16]) and associative thinking (de-focused attention[17]).

3.3 Intersections

This section provides a reflection on the intersections between Gabora and Boden's ideas in context of the development of MAM. Boden argues machines could be considered creative. Gabora proposes a model of memory in which creativity is a central component and is suited to computational implementation. Gabora's conception of the relationship between memory and experience is highly complimentary to the requirements of MAM.

Boden may not consider Gabora's conception of the construction of memory as creative as memories are integrated and considered in terms of similarity, not newness, surprise nor value. The

[15]The "variable focus" of the creative individual is supported by the characterization of dualities in the creative mind apparent in Barron (1963).

[16]Martindale and Armstrong (1974); Martindale (1977)

[17]Dewing and Battye (1971); Dykes and McGhie (1976); Mendelsohn (1976)

essential process on which Gabora's model of memory is based, constrained and distributed memory, is in essence the *combination* of experiences into a structure. Many of these combinations would not be surprising. The combinations are only evaluated by being compared to other memories.

There is a strong link between Boden's "exploratory" and "combinational" creativity and Gabora's conception of the "thought process". The thought process is a trajectory through memory. The *combinations* of the properties of stored stimuli are *explored*. This trajectory propagates through the memory space activating and constructing new relationships between elements of experience.

The conception of "transformational" creativity can be applied to Gabora's model. As the trajectory propagates through memory the construction of concepts change the space of memories. A thought process can bridge two isolated islands of experience fusing them together to form a new territory that would not have been possible before.

Because memories are placed in memory according to their content, and overlapped to a degree dependent on the context, the placement of memories is an evaluation of stimuli. Creativity is not a two step process but a continuous contextual evaluation. Boden's argument is that the result of the whole process must be "new, surprising and valuable" to qualify as creative. Low level cognitive mechanisms are constantly generating novelty that are evaluated at many levels. It is only the highest level that results in Boden's conception of creativity. Gabora's conception is that creativity permeates the process of constructing a world-view from memories, and memories from stimuli.

The relationship between the context and the artifact is the central theme in the development of MAM. This relationship falls directly in line with Gabora's conception of the relationship between agent, memory, and environment. The world-view is not a perfect reflection of the environment, nor random, but a function of the cognitive processes of perception and memory.

Chapter 4

Creative Machines

This chapter contextualizes MAM in relation to selected artistic projects that, in my consideration, exhibit creative behaviour. This section is focused specifically on projects, related to machine creativity, initiated by artists. These projects highlight the historical trajectory in which MAM is placed. The works of Cohen (1995, 1979), Legrady and Honkela (2002) and Rokeby (1990, 2001) will be described, interpreted and categorized in terms of Boden's framework. Cohen's project "AARON" is discussed because of its seminal position and its place at the intersection of computation, cognition and art. Legrady's installation "Pockets Full of Memories" is included because it applies an AI approach similar to that used in MAM. Rokeby's "The Giver of Names" and "n-cha(n)t" are discussed as MAM follows a similar creative trajectory.

4.1 "AARON"

Amongst the most notable examples of creative machinery are the AARON programs (Cohen, 1995, 1979) initiated in 1973 and continuing to the present. AARON is a set of programs that draw, and more recently paint, in a style distinct from much computer graphics. When I first began working on MAM I had only a cursory understanding of Cohen's work. There is a significant overlap between this research and the research resulting in AARON. This overlap is due to similar artistic drives behind both projects.

AARON is a line of artistic enquiry. The project began with a "...desire to understand more about the nature of art-making processes than the making of art itself allows..." (Cohen, 1979). Two keywords in this quote emphasize the connection between my work and Cohen's—*understand* and *processes*. The project emphasizes the system's process of creation over its results—the drawings. The purpose of this emphasis is to examine the property of free hand drawing that allows a set of marks, on a page, to evoke objects in the physical world. Cohen rejects the labels reference,

symbol and representation, in preference to his label "standing-for-ness". The essence of his enquiry is an examination of meaning itself. Cohen makes two key points in the relationship between the artist and viewer in relation to meaning. Firstly Cohen places the majority of the responsibility of meaning-making on the viewer: "I am proposing that the intended meanings of the maker play only a relatively small part in the sense of meaningfulness" (Cohen, 1979). Secondly Cohen states "...the minimum condition for generating a sense of meaningfulness...[does not]...include the assumption of an intent to communicate..." (Cohen, 1979). The structure of the cognitive processes used in the construction and reading[1] of images result in the "meaningfulness" of the image: "...the exercise of an appropriate set of these cognitive processes would itself be sufficient to generate a sense of meaningfulness" (Cohen, 1979). Those processes can be computationally executed—as they are in AARON.

As a collection of programs, AARON can "create" in a number of different styles. Each style uses a different variant of AARON that implements a different set of rules. Examples of these variants are "abstract AARON", which creates abstract landscapes, "acrobat AARON" which creates acrobatic figures, and "jungle AARON" which creates scenes of figures in a complex jungle ground that evoke Gauguin. AARON programs contain sets of rules that encode specific compositional and stylistic laws specified by Cohen.[2] Each component of the composition—figure, ground, and object—is each an instantiation of the model those rules encode.

The choices AARON makes are generated by a weighted random number generator and constrained by rules (Cohen, 1979). The results of these choices are applied to create paintings that, in a recent version, are physically produced by a painting machine. AARON has no sensory system. It receives no feedback from the results of its actions on the canvas. The system contains an internal representation that effects the placement, pose and arrangement of items in the picture plane. The internal model[3] is realized in that perfect theoretical vision regardless of the properties of the physical artifact. The only feedback between the physical artifact and AARON is through Cohen himself.

AARON could only be considered creative in a *symbiotic* relationship with its creator.[4] Early "paintings" were drawn by AARON but painted by Cohen.[5] Cohen does not consider AARON an artist. AARON's artwork is the result of a collaboration between Cohen and AARON.

Cohen states that a software system is a natural approach to art-making because artistic composition is rule-based. I agree that graphic composition, in a particular style of painting, is rule-based.

---

[1] Cohen considers the construction and reading of images based on the same set of cognitive processes.

[2] These laws result in the cognitive processes that allow AARON to create meaningful images. Cohen did not intend to model his own painting style but drawing and painting in general. Viewers of AARON's work repeatedly mention the similarity between the styles of Cohen and AARON.

[3] AARON's rules could be considered a model of creative intention.

[4] Every "creative" machine depends on its creator to be creative.

[5] Cohen has been known to paint a drawing differently than specified by AARON.

It does not follow that all aspects of artistic creation are. AARON exhibits combinational creativity in that it combines and arranges different types of objects, figures, and ground in a composition. The elements of a AARON painting are not selected from a database of pre-constructed elements, but are constructed on the fly.[6] AARON exhibits exploratory creativity. The drawing of a figure is an exploration of the space of possibilities within the constraints of the model. The space is searched by constrained randomness. There is no correlation between how figure was drawn last time and how it is drawn this time.

Boden considers AARON's paintings surprising. The paintings are highly constrained by AARON's rules. The paintings resulting from a particular program are very similar. As a collection of programs AARON creates surprisingly convincing[7] line drawings. Surprise, in Boden's conception of creativity, is considered in light of the domain. AARON does not exhibit transformational creativity as the space of possibilities is a fixed model that can only be changed by Cohen.

The value of the AARON project, in relation to MAM, is Cohen's conception of the relation between intention, cognition and the meaningfulness of the artifact. Thirty years later, I follow Cohen in considering art as a method of enquiry, not simply expression. Art is centrally concerned with meaning, standing-for-ness, and the relationship between the viewer and the artist. Standing-for-ness is how an artifact can evoke meaning in the viewer. The creation of standing-for-ness does not require the intention to communicate. The cognitive process of making something meaningful is at the centre of Cohen's enquiry. For Cohen an image is a trace of the cognitive process that makes the world meaningful.[8]

4.2 "Pockets Full of Memories"

George Legrady's "Pockets Full of Memories" (PFOM) (Legrady and Honkela, 2002) was made possible by a commission from the Centre Pompidou Museum of Modern Art in 2001. The project was revisited and exhibited in the Dutch Electronic Arts Festival in Rotterdam, Netherlands in February 2003, and then at Ars Electronica in September 2003. PFOM is included because what it *does*, organize a continuously changing database of content from the world, is echoed in MAM's process.

PFOM is an interactive installation that constructs a digital archive of viewer participation. The project was produced in a highly collaborative[9] context. The system uses a Self-Organized Map[10] (SOM) to visualize and organize a database of continuously changing content—provided by the audience. PFOM is presented using two parallel approaches. The first is an interactive installation

[6] Every time AARON draws a ball it is a different ball.

[7] By "convincing" I mean the viewer is convinced that the creator could certainly be a human artist.

[8] This statement refers to the quote on page v.

[9] The project was produced by a large international team. An account of the production process is discussed in Steinheider and Legrady (2004).

[10] The Self-Organized Map is a central artificial intelligence technique used in MAM—discussed in Chapter 5.

that allows the audience to contribute content. The second is an on-line version where users are able to contribute meta-data, to the existing content, over the Internet.

The installation consists of a large projection and kiosks with flat-bed scanners. The audience is encouraged to scan an image of an artifact in their possession. The kiosk prompts the participant to answer questions regarding the meaning of the artifact. The answers are stored in a database and bound, as meta-data, to their corresponding images. The meta-data is fed to the SOM that plots corresponding images in the projection. Artifacts attached to similar meta-data are plotted closer together than artifacts with dissimilar meta-data.[11]

The database is being studied by an independent research project led by Dr. Steinheider. The research is concerned with the "semantic specificities" (Legrady and Honkela, 2002) of the archive. The goal of the research is the production of a "...cultural overview of how the members of this particular audience have described themselves through their object and description choices" (Legrady and Honkela, 2002). The database is new, surprising and valuable enough to warrant direct study.

PFOM's organization of contributed items, based on similarity, exhibits combinational creativity. The visualization is a combination of a selection of items from the database. This organization exhibits exploratory creativity as the space of the database is explored by the visualization. PFOM encourages creativity in the audience as is illustrated by the myriad of body parts contributed as artifacts. The emphasis in PFOM is on the interaction between the audience and the database content, rather than the creative properties of the method of content organization.

4.3 "The Giver of Names" & "n-cha(n)t"

David Rokeby is a great influence on my work. Where the AARON project is linked conceptually to the development of MAM, and the technical use of a SOM connects PFOM to MAM, the work of David Rokeby is foundational to my work in both technical and conceptual realms. The projects discussed in this section were driven by Rokeby's desire to remove himself from the external properties of the artworks.[12] This desire manifests itself in my work as MAM is meant to relate itself to its context.

David Rokeby's work could be divided into two broad and partially overlapping categories. The first is the interactive installation work—for example "The Very Nervous System". The second is populated by works that tend to be generative and autonomous. They are concerned with machine perception, knowledge, and language. The second category contains artworks that are most related to MAM. This section will discuss "The Giver of Names" and "n-cha(n)t". These works are creative[13]

[11] An image of the SOM representation of the database is available at `http://legrady.mat.ucsb.edu/chron_map.html`

[12] D. Rokeby, personal communication, April 9, 2008

[13] The characterization of these systems as creative is the interpretation of the author. The artists intention was not centrally the construction of creative machines.

in that they construct free writing passages initiated by their physical context.

The first exhibition of "The Giver of Names" (Rokeby, 1990) was in 1998 in Guelph, Ontario.[14] The system perceives through a video camera pointed at a pedestal. The floor around the pedestal is scattered with children's toys the audience is encouraged to place in the camera's view. "The Giver of Names" attempts to give names to the objects it sees. The system creates a free writing passage,[15] written in proper grammatical structure, inspired by those objects. Their colour and shape are associated with concepts in the system's knowledge-base. The knowledge-base is based on WordNet[16] and is expanded with information returned from a "reading" system that extracts relations between words in texts fed to it. This disembodied knowledge is linked to the machine's sensory impression of the physical world—resulting in poetic passages. The movement through the knowledge-base is a response to the viewers' action—not agency in the system.

First installed in 2001 in Banff, Canada, "n-cha(n)t" (Rokeby, 2001) builds on the themes of language and interpretation implemented in "The Giver of Names". The project is composed of a network of identical interconnected units. Each unit is able to hear[17] and speak[18] poetic passages and has a copy of a knowledge-base similar to that used in "The Giver of Names". The hearing mechanism of each unit attempts to translate microphone input into text.[19] The translated text selects the "object of interest"[20] and stimulates the knowledge-base—resulting in a free writing passage. The passage is spoken through each unit's voice synthesis mechanism. The hearing apparatus is a directional microphone that picks up sounds in close proximity. Each unit's microphone is tuned to ignore the sounds from other units. The units communicate their object of interest to one and other over an Ethernet network. In the absence of external stimuli the units tend to chant a mix of phrases and sentences in synchrony. The chant is an emergent result of the constancy of the perceived environment and the knowledge-base. When one unit's microphone picks up a sound it results in change of that unit's object of interest. The inconsistency of the object of interest between units results in a disturbance of the chant. The chant can also be disturbed by "slips" of timing between the units.[21] Once the system has been disturbed it will eventually return to convergence—where all units chant in synchrony.

"The Giver of names" and "n-cha(n)t" are attached to the physical world through sensors that allow them to respond to their context. Their knowledge-base is implanted by the artist. It is not

[14] A prototype of "The Giver of Names" (Rokeby, 1990) was exhibited in Toronto, Canada in 1997.

[15] Free writing is a method where the author's process of associative thinking is recorded.

[16] WordNet is an electronic lexicon of the English language: `http://wordnet.princeton.edu/`

[17] "Hearing" is accomplished using a microphone and voice recognition

[18] "Speaking" is possible through speakers and voice synthesis.

[19] The process of translation is an "attempt" because it is often inaccurate.

[20] The "object of interest" is the concept in the knowledge-base that is currently being stimulated. The phrase is how Rokeby described the location in the knowledge-base in personal correspondence.

[21] "Slips" occur because the timing mechanism of each unit is not synchronized to other units. This results in some temporal drift between units.

constructed from an analysis of sensor impressions. The systems exhibit combinatorial creativity by combining words from the knowledge-base to form free writing passages. Exploratory creativity is exhibited in the systems as they explore the structure of their knowledge-bases. Judging these systems as transformationally creative is difficult due to the question of what constitutes a structured space of possible results. The system's vocabulary, and therefore the space of words, is fixed. The associations between words in the knowledge-base are in flux as they are effected by sensory experience. The result of these effects remain in the system for a short period and dissipate over time. In order to exhibit transformational creativity the systems need to change their spaces of possibility. For each word selected a space is created for words that may follow it. This space is transformed by both the words that preceded it and by the sensory impressions of the machine. The grammatical structure of the poetic passages is fixed.

Cohen's work highlights the history of artists working with technology in a drive for enquiry, rather than expression. At the centre of this enquiry is a question of meaning or standing-forness. PFOM and MAM both use an AI technique to organize a constantly changing collection of content. MAM is directly in line with the work of Rokeby. The external properties of MAM, "The Giver of Names" and "n-cha(n)t" are causally dependant on their physical context. All three works implement computational processes of *origination* in relation to that context. Rokeby and myself have embarked on an impossible process to remove ourselves from the external properties of our artworks. The projects discussed in this chapter have been selected to highlight key points in the research leading to MAM. These projects are driven by a question of meaning. For Cohen an image is a trace of the cognitive process of understanding the world. For Rokeby and myself the construction of systems that originate entails a drive to reduce our influence over their external properties.

Chapter 5

Memory Association Machine

MAM is an artwork whose external properties are the result of a collaboration between two major influences: the artist's intention and the system's context. These influences are bound together through the enacted process that relates the artwork to its context. MAM's enacted process explores its context, remembers and integrates its sensory impressions into a field of experience, and free-associates through that field. These three processes occur in parallel and are visualized on each of the artwork's three displays. These processes are initiated by a computer controlled camera that randomly pans, tilts, and zooms to explore the installation's visual context. A photograph of MAM,[1] as installed for the 2007 Pure Data Conference, is pictured in Figure 5.1 on page 33. The camera's sensor impressions are fed into the left display and the system. The system's integration of remembered sensor impressions is visualized on the centre display. Both the centre and right displays present the process of free-association. The centre display visualizes the activation of memory locations while the right display presents a cinematic montage that reflects the sequence of activations. MAM relates to its context by transforming its sensor impressions into an *original* representation.

Section 5.1 describes the mechanisms that relate MAM to its context. The section discusses the Self-Organizing Map, and MAM's Memory and Free-Association systems. Future development of MAM's software is discussed in Section 5.2. The chapter concludes with a discussion of Boden's framework in relation to MAM's process.

5.1 The System

MAM is composed of two primary subsystems: the "Memory System" and the "Free-Association System". These systems correspond to the processes of memory integration and free-association and

[1] This version of the project was exhibited under MAM's former title "Self-Other Organizing Structure #1" (SOOS1) (Bogart, 2007).

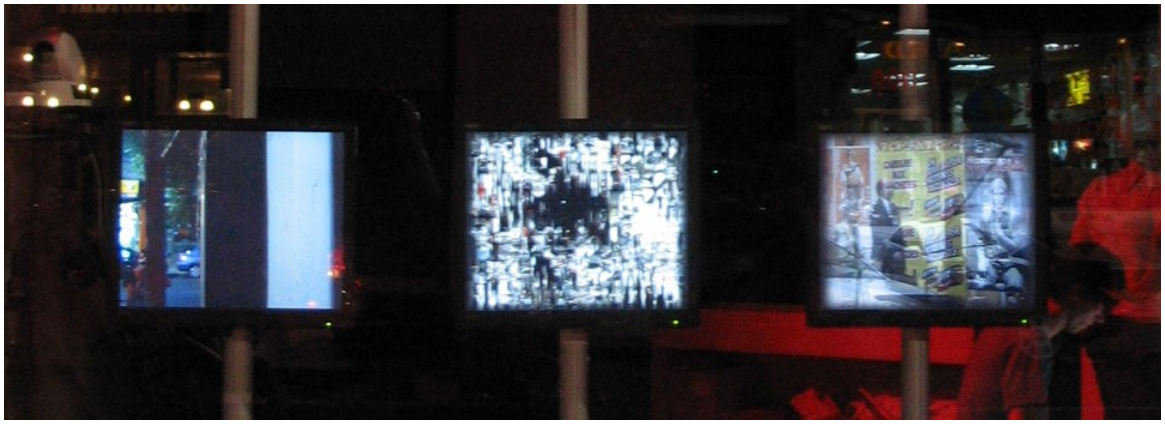

Figure 5.1: Installation of MAM for the 2007 Pure Data Conference. The left display shows the current stimulus—a feed directly from the camera in the upper left. The centre display is a visualization of the field of experience. The right display is a cinematic montage representation of the process of free-association.

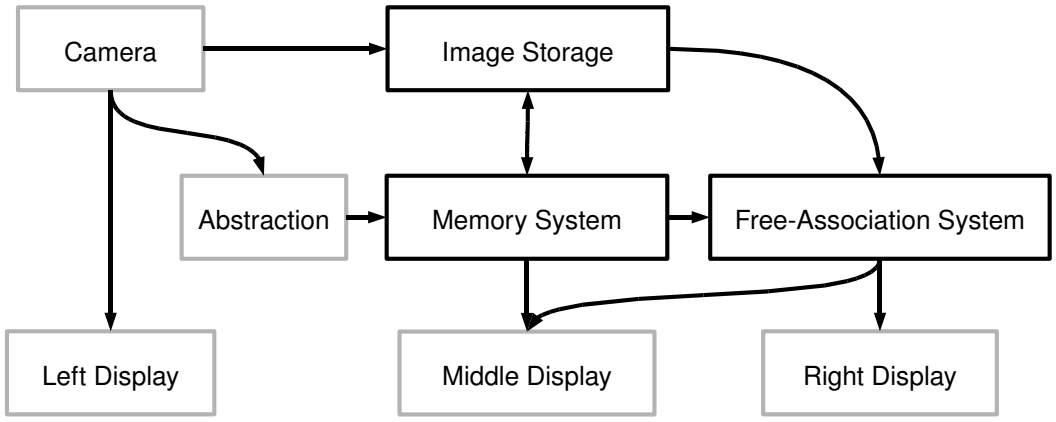

Figure 5.2: The Architecture of "Memory Association Machine"

are independent networks of numerous identical units. In addition to these systems is a storage mechanism that archives the sensor impressions (images) collected by the camera. Figure 5.2 on page 33 shows an overview of the system's major components. The camera image is sent to three locations: the left display, the abstraction mechanism and the image storage system. The abstraction is fed into the Memory System. The Memory System and the Free-Association System both access image storage.

MAM has three states: waking, dreaming and suspended. The selection of state is determined by a schedule that reflects the activity of the installation location.[2] The system is initiated in a waking state in which the system explores its context through the camera. Every twelve seconds a pan/tilt/zoom triplet is randomly generated[3] and sent to the camera via an RS-232 port. A single frame is grabbed for each fixation[4] and a unit of the free-association system is activated. This activation is a perturbation of the memory field and results in a cascade of activations of the system's previous experience.

In the dreaming state the camera ceases to explore—sensor impressions are not collected. Every twenty-five seconds a time-seeded random number generator activates a unit in the Free-Association System. In this state the system is isolated[5] from external stimulus—activations are not initiated by the context. As these activations propagate through the memory field they call up sensor impressions from previous experience.

In the suspended state the displays are put into DPMS[6] sleep and the Free-Association System no longer propagates activations. This state reduces wear on the camera and displays during times when there is both little light and activity in the system's context.

The memory and free-association networks interact in negotiation with the system's context to generate the system's external properties. The Memory System stores and integrates MAM's sensory experience. Central to the Memory System is a Kohonen Self-Organizing Map (SOM) (Kohonen, 2001). The SOM is an unsupervised artificial neural network that acts as an arbitrary pattern classifier. Patterns are organized into categories based on their similarity to other patterns. In MAM the patterns presented to the SOM are the sensor impressions captured by the camera. Each SOM category is a location in MAM's memory field. The SOM locates sensor impressions in the memory field by comparing new sensor impressions with remembered ones. The details of the SOM will be discussed in Subsection 5.1.1.

The Free-Association System is a network of simple units that are independent of the SOM. This network allows activations to propagate between units and is similar to a cellular automata.

[2]The selection of state is hard-coded to the time of day.

[3]The triplets are constrained within a range. Additionally certain combinations of triplets are ignored so that sensitive areas viewable by the camera are not explored.

[4]A fixation is the moment the camera has reached the requested pan/tilt/zoom position.

[5]The system is isolated from its context as the camera is not collecting sensor impressions.

[6]VESA Display Power Management Signalling

The activation of units selects sensor impressions from the system's memory. The choices of what impressions to select are not predetermined but a result of the system's negotiation with its context. The Free-Association System is based on the notion of the creative thought process as introduced by Gabora (2002b).

5.1.1 Kohonen Self-Organizing Maps

A Self-Organizing Map (SOM),[7] is an unsupervised artificial neural network (ANN) designed for classification. The SOM implementation used in MAM is provided by the ann_som Pure Data external (Zmölnig, 2001). The ANN is an AI approach inspired by neurophysiology. An unsupervised ANN is able to classify inputs without the benefit of any information provided to it.[8] These networks restructure themselves in response to the input patterns presented during training. ANNs are characterized by being composed of numerous simple components, inspired by neurons, which are massively interconnected.[9] In mathematical terms the SOM is a non-linear projection of a high-dimensional data-space onto a low dimensional "feature-map" that preserves topology.[10] A SOM is able to categorize an arbitrary input pattern, with a finite number of dimensions, into a finite and fixed[11] number of categories. The number of categories and dimensions are specified before training proceeds. The SOM is a projection as it maps values from input space onto values in output space. The input space is roughly analogous to the stimulus space discussed previously. Every possible input is a point in input space. Points in output space are the categories in which input patterns are classified. The output space is analogous to the memory space. Each point in output space is a memory location. The SOM projection is a mechanism that determines in which memory location (category) a particular stimulus (input pattern) should be located. Memory locations (categories) could be organized differently than the stimuli (input patterns) and remain content-addressable. The memory locations of a SOM reorganize themselves in order to best represent the topology of the stimuli. In Gabora's model, memory locations have a fixed position in stimulus space.

The SOM consists of a network of units,[12] each corresponding to a category. Units are usually arranged in a 2D Euclidean lattice that reflects the output space. Figure 5.3 on page 36 (c) shows a 12x12 unit SOM. The training data is a collection of one dimensional numbers ranging from 0 to 1. Figure 5.3 (a) is a visualization of the training data, sorted by value, represented using a grey-scale

[7] Also known as a Kohonen network.

[8] A supervised ANN learns by example. The correct answer is required for the network to learn.

[9] For a survey of ANNs see Medler (1998).

[10] The Oxford English Dictionary Online (2008) defines topology as "[t]he way in which constituent parts are interrelated or arranged". The term concerns structure.

[11] "Adaptive Resonance Theory" (Carpenter and Grossberg, 1994), "Incremental Grid Growing" (Blackmore and Miikkulainen, 1993), and "Growing Cell Structures" (Fritzke, 1991) networks could allow for the addition of a new category in response to an input pattern.

[12] Units are also known as nodes or neurons.

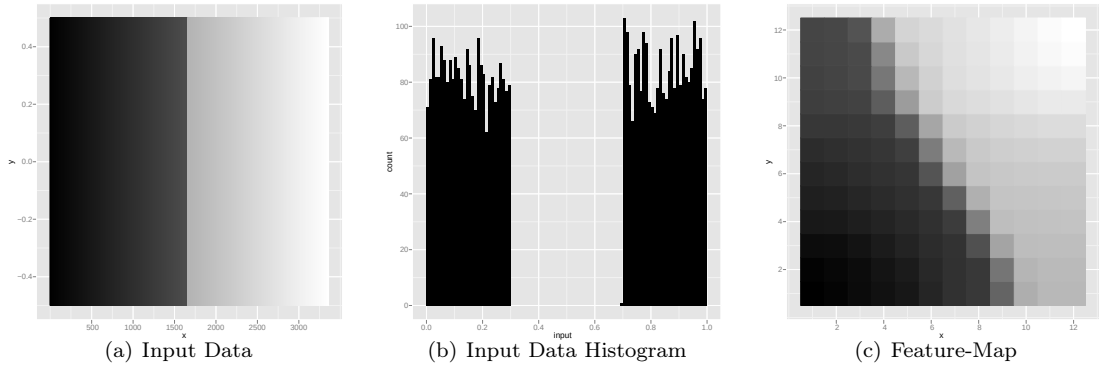

(a) Input Data (b) Input Data Histogram (c) Feature-Map

Figure 5.3: The relation between input patterns and the resulting feature-map. Figure (a) represents the sorted values of the input patterns presented to the SOM. Figure (b) shows the distribution of input values that shows two clusters. This structure is reflected in the trained SOM (c) where two regions of similar input values are apparent.

where 0 is black and 1 is white. The training data is not uniformly distributed. The distribution is depicted in the histogram in Figure 5.3 (b). The histogram shows the training data is grouped into two independent clusters. One cluster is composed of darker shades ranging from approximately 0 to 0.3. The other cluster is made up of lighter shades ranging from approximately 0.7 to 1. The training data contains no values between approximately 0.3 and 0.7. Through training each of the input patterns is associated with a particular category. Each of the units in Figure 5.3 (c) is shaded in the value it has been associated with. Once a SOM has been trained the result is a feature-map that represents the structure of the input patterns. The feature-map resulting from training the SOM (Figure 5.3) shows two regions corresponding to the dark and light clusters of the input patterns.

The input space is defined as a set of real numbers ($\Re^n$) with n dimensions. Points from the input space (the space of all possible inputs) are mapped to points in the output space (the position of all the units in the map). The SOM projects the values from Figure 5.3 (a) into the output space—Figure 5.3 (c). Input patterns are defined as $x = \{\xi^1, \xi^2, \ldots, \xi^n\} \in \Re^n$ where x is a single pattern. The pattern is composed of the magnitudes of each of its dimensions (ξ) in the set of all possible inputs. For example, if x is a sound file each of its samples are the ξ components. If x is one of MAM's sensor impressions, ξ is each of the colour components for each pixel.

The set of all input patterns, Figure 5.3 (a), is defined x_i^j; i is the index of the input pattern, where $1 \leq i \leq n_p$ and n_p is the number of patterns; j is the index of each dimension, where $1 \leq j \leq n$. Similar inputs are associated with similar categories. Inputs are considered similar according to the sum of the distances between their vector components, $dist(x_1, x_2) = \sum_{j=1}^{n} (x_1^j - x_2^j)^2$. The smaller the distance the more similar the inputs. This measure is based on the pairwise comparison

of corresponding components of input patterns.[13] Categories are considered similar based on their Euclidean distance in output space. Dissimilar inputs end up being associated with nearby categories when two or more highly dissimilar clusters of points are presented to the network. This is the case in Figure 5.3. The clusters compete for categories in the finite SOM—causing folds. Folds occur when the Euclidean distance between units in output space is not proportional to the distance between the patterns they are associated with.[14] Figure 5.3 (c) shows the input patterns closest to black and white are associated with units the greatest distance apart.

Each unit has a sensor for each dimension of input, and a code-book vector—also known as a weight vector—which contains the same number of elements as the unit has sensors. The code-book vector is defined as $m_i = \{\mu_i^1, \mu_i^2, \ldots, \mu_i^n\}$ where i is the index of the unit and μ is the component corresponding to each dimension. Each unit has a "neighbourhood" (N_i) defined as the set of units within a certain radius (in output space) of the unit i. During training the values of the code-book vectors approach the values of the input patterns. Once training is complete the network's structure (code-books) matches the structure of the input patterns—preserving topology. The training procedure, as implemented in MAM, is as follows:

1. Set the dimensions of the code-books of all units to 0.[15]
   $m_i = 0$

2. Present the network with an input pattern.[16]

3. Select the unit that contains the code-book with the smallest Euclidean distance from the input pattern. This unit is the "Best Matching Unit" or BMU. The BMU most closely resembles the most recent input pattern.[17]
   $c = argmin_i(\sum_{j=1}^{n}(x_i^j - m_i^j)^2)$, where c is the index of the BMU.

4. For each unit in the neighbourhood of the BMU, add to the code-book vector, $m_i(t)$, the difference between the input pattern and the code-book, $x(t) - m(t)_i$, multiplied by the learning rate, $\alpha(t)$, and the neighbourhood function, $\theta(v, t)$. The learning and neighbourhood functions will be discussed in Subsection 5.1.2.

[13] This measure of similarity does not consider two inputs, with identical blocks of values occupying different dimensions, as similar.

[14] Folds in feature-maps can be visualized using the U-Matrix (Ultsch and Siemon, 1989; Kraaijveld, 1992; Ultsch, 1993) method. For more information on SOM visualization methods for highlighting folds see Kaski et al. (2000); Merkl and Rauber (1997); Rauber (1996); Tasdemir and Merenyi (2006); Vesanto (1999); Ultsch (1993, 1999, 2003b,a). Blackmore and Miikkulainen (1993); Blackmore (1995); Fritzke (1991); Yin (2002) propose alternatives to the SOM for cluster visualization.

[15] Often code-books are set to random values. According to Kohonen this was initially done to demonstrate how robust the SOM is—even when used with arbitrary initial conditions. Training can be accelerated by calculating initial code-books based on the training data. See Kohonen (2001) Section 3.7.

[16] Give the units something to compare their code-books to.

[17] During the first iteration, due to the ann_som implementation, if the code-books of all units are equal the unit with the largest index is chosen as the BMU.

$m_i(t+1) = m_i(t)+\theta(v,t)\alpha(t)(x(t)-m_i(t))$, where t is the current time step, and v the distance, in output space, between the BMU and the unit i.

5. Repeat from 2.

In a canonical SOM the amount that the code-books are changed (the learning rate), and the size of the neighbourhoods (the neighbourhood function), both decrease monotonically over time.[18] Training is complete when the mean of the difference between all input patterns and all units' code-books, $mean_i(\sum_{j=1}^{n}(x_i^j - m_i^j)^2)$, ceases to decrease. Training typically takes from hundreds to thousands of iterations. Training duration depends on the number of sensors and the size of the network.

5.1.2 Memory System

As MAM explores its context, the SOM integrates sensor impressions into a field of experience. The field of experience is MAM's equivalent of Gabora's "world-view". The field is a lattice representing the SOM's organization of the system's remembered sensor impressions. The sensor impressions are fed into the computer as a full-frame, 640x480 pixel, 30fps video stream. At twelve second intervals, the 12x12 unit SOM is fed with a 100x75 pixel RGBA sub-sampled frame of the video stream. The raw values for each pixel, represented as four RGBA[19] floating point values, correspond to the 30,000 sensors of each SOM unit. The SOM is used to index the storage of images, rather than store the images itself. The BMU for each sensor impression is used to specify its storage index. The storage area has the same number of locations as the SOM has units. Each SOM unit is associated with its corresponding storage location. Figure 5.4 on page 39 shows two representations of a trained SOM. The SOM was trained over 5024 iterations with the 400 100x75 pixel RGBA images included in Appendix D. The SOM was trained using linearly decreasing learning and neighbourhood functions. The learning rate ranged from 1 to 0 and the neighbourhood size from 8.5 to 0. The "instar" neighbourhood function was used: the learning rate was applied by a factor inversely proportional to the linear distance from each neighbour to the BMU. Figure 5.4 (a) is a visualization of the code-books. In this visualization the values of each code-book component represents the RGBA channels of a 100x75 pixel image. The visualization shows the internal structure of the SOM network. Figure 5.4 (b) illustrates the structure of the feature map by representing each unit with its corresponding stored sensor impression. The latter is referred to as the MAM visualization. As sensor impressions are associated with the units' with the most similar code-books, during training, the code-book visualization resembles the MAM visualization. If more than one sensor impression is associated

[18] This is not the case in the MAM implementation—discussed in Subsection 5.1.2.

[19] The alpha value is passed to the SOM even though it is opaque for all pixels in all frames. This is due to the high CPU usage of extracting the alpha values from the 30,000 element message in Pure Data. A solution would be writing a Pure Data external to remove every 4^{th} element of a message.

with the same category, the most recent sensor impression replaces the previous one. In this method of storage the resulting visualization is a combination of multiple iterations of training as a single BMU is generated for each iteration.[20] The visualization reasonably reflects the structure of the underlying code-books. In installation the mean time a sensor impression is held in the system is approximately two hours, with a range of approximately thirty-three hours.[21] The memory for unusual impressions is longer, in duration, than for common impressions as common stimulus is seen more often and therefore replaced more often. Sensor impressions are stored at their full resolution—not the sub-sampled resolution fed to the SOM.

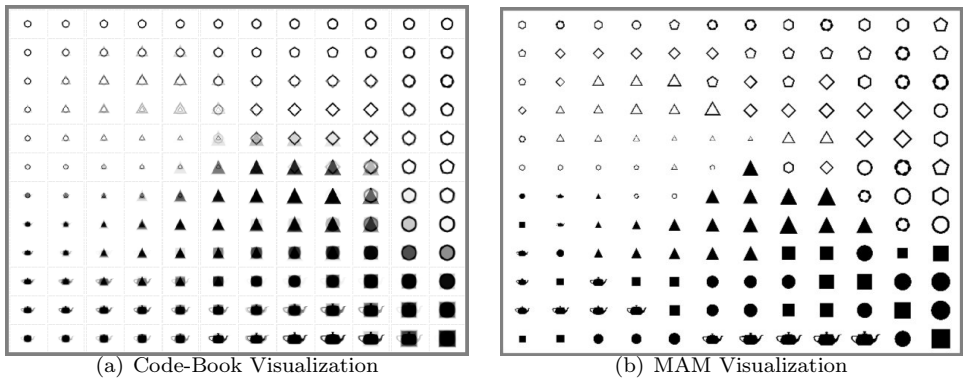

(a) Code-Book Visualization (b) MAM Visualization

Figure 5.4: A comparison of feature-map representations. Figure (a) is a direct visualization of the code-books. In figure (b) each unit is represented by its associated input pattern.

The memory system is represented as a Euclidean lattice as pictured in Figure 5.5 on page 40. The display area is divided by the number of units in order to uniformly allocate display space. Units are represented as scaled circles with Gaussian alpha channels. The "feathered" edge allows the impressions to blend with nearby neighbours.[22] As a result the structure of the impressions takes precedence over the structure of the lattice. Only in regions associated with impressions with little spacial variation[23] does the lattice of circles become visible.[24]

MAM is intended to be in constant negotiation with its context. The SOM must be trained continuously to incorporate new experiences in its structure. The canonical SOM is trained on a finite set of data and reaches convergence when the network reflects the topology of the input patterns. In MAM the network is constantly converging, but is not meant to, and cannot, reach convergence. Allowing the SOM to converge implies that its process of relating to its context could

[20]The visualization in Figure 5.4 on page 39 (b) is based on the BMUs generated over the iterations of 1088 to 4997—out of 5024.

[21]These values include the time images are kept during the dreaming and suspended states.

[22]A additional effect of the lack of hard contrasting edges is the reduction of 'burning' of the LCD display.

[23]For example those impressions that are out of focus.

[24]This is somewhat rare and always limited to a regional cluster—due to the SOM organization.

Figure 5.5: The representation of the field of experience.

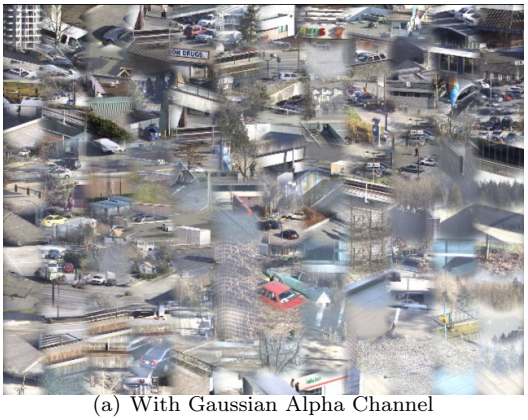

(a) With Gaussian Alpha Channel

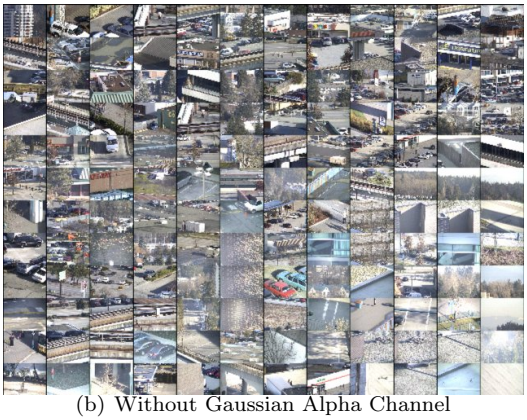

(b) Without Gaussian Alpha Channel

Figure 5.6: This figure compares feathered, and scaled, (a) to un-feathered (b) edges in unit visualization. The result is the content of the memories takes precedence over their Euclidean arrangement.

be complete. As the context changes so should the SOM structure. Continuous training is enabled by using cyclic learning and neighbourhood functions. These functions control the rate at which the self-organization evolves and is refined. They are driven by the cosine equation defined as $\frac{cos(t/100)+1}{2}$—where t is an integer counter that ranges from 628 ($\approx 200\pi$) to 0. The range ensures the resulting cosine function loops seamlessly. The result of a neighbourhood size that ranges from 1 to 0 (inclusive) is that the code-books of few units are updated for each iteration. The initial choice of this range was based on the assumption that ann_som was normalizing the neighbourhood to the size of the network—which has proven to not be the case. The function is iterated each twelve seconds and results in a period of approximately two hours. The cyclical functions allow the SOM to respond to a continuous flow of new input patterns and integrate them into a constantly reorganizing field. A comparison between the training progress of the canonical SOM and MAM's implementation is pictured in Figure 5.7 on page 42.

Code-books are constantly refined through the training process. This makes the initial code-books increasingly insignificant over time.[25] As the training process is continuous the SOM is replacing structure—created as a result of previous experience—with structure that reflects current experience. The network oscillates between durations of large and small change.[26] A byproduct of this training method is that the MAM visualization takes considerably longer to populate than the canonical method. Figures 5.7 (a) and (b) present the training process using the canonical method. Figure 5.7 (a) represents the code-books directly and (b) is the corresponding MAM visualization. Units in Figure 5.7 (d), trained with the MAM method, are clearly associated with stimulus more slowly than units trained with the canonical method Figure 5.7 (b).

The relatively small number of units in the SOM—combined with the complexity of sensor impressions—results in a memory field that is often interpreted, by viewers, as unorganized. Figure 5.8 on page 43 compares the feature-maps resulting from the canonical (a) and MAM (b) training methods. The MAM method results in a more complex, but not unorganized, feature-map. In Figure 5.9 on page 43 (a and c) units are associated with a uniformly distributed random selection of images.[27] In (b) and (d) units are associated with images using a SOM trained using the MAM method.

The resolution of the data fed into the SOM was determined through trial and error based on system performance. The discrepancy between the resolution of the images, used in the memory representation (320x240 pixels), and the resolution the SOM is presented with (100x75 pixels) results in an incongruity between what is visible to the SOM and what is visible to the viewer. This likely contributes to the interpretation of the memory field as unorganized. The extreme difference

[25] In the case of MAM, which trains over hundreds of thousands of iterations, the initial code-books are of little consequence.

[26] Large changes result from a large neighbourhood and learning rate—small changes from a small neighbourhood and learning rate.

[27] Images are selected from a finite set of test data.

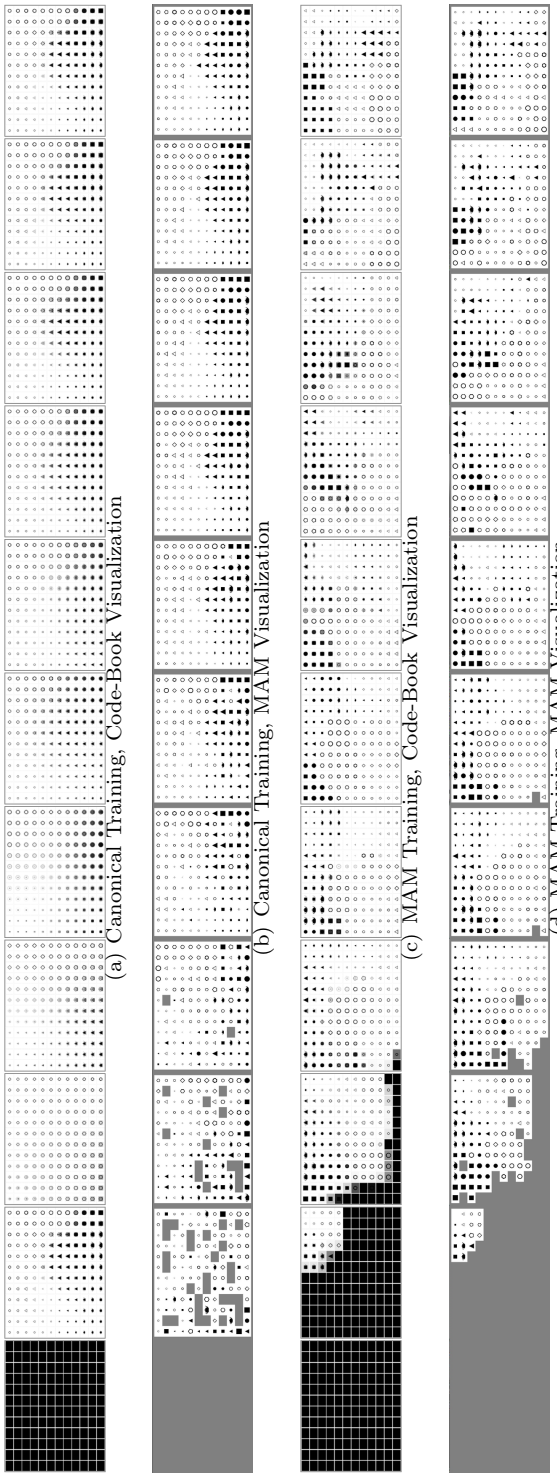

Figure 5.7: Comparison between MAM and canonical SOM training processes. Each 500th iteration of the training process is pictured. Figures (a) and (b) show the progress of canonical SOM training. Figures (c) and (d) track the MAM training method. Figures (a) and (c) are representations of the code-books, while (b) and (d) are MAM visualizations.

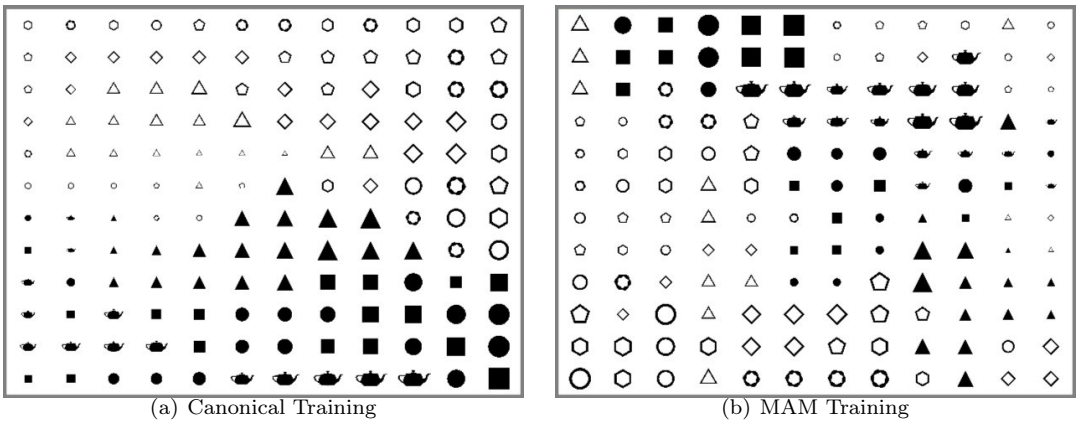

(a) Canonical Training (b) MAM Training

Figure 5.8: Comparison between feature-maps trained using the canonical (a) and MAM method (b). Both figures are presented using the MAM visualization.

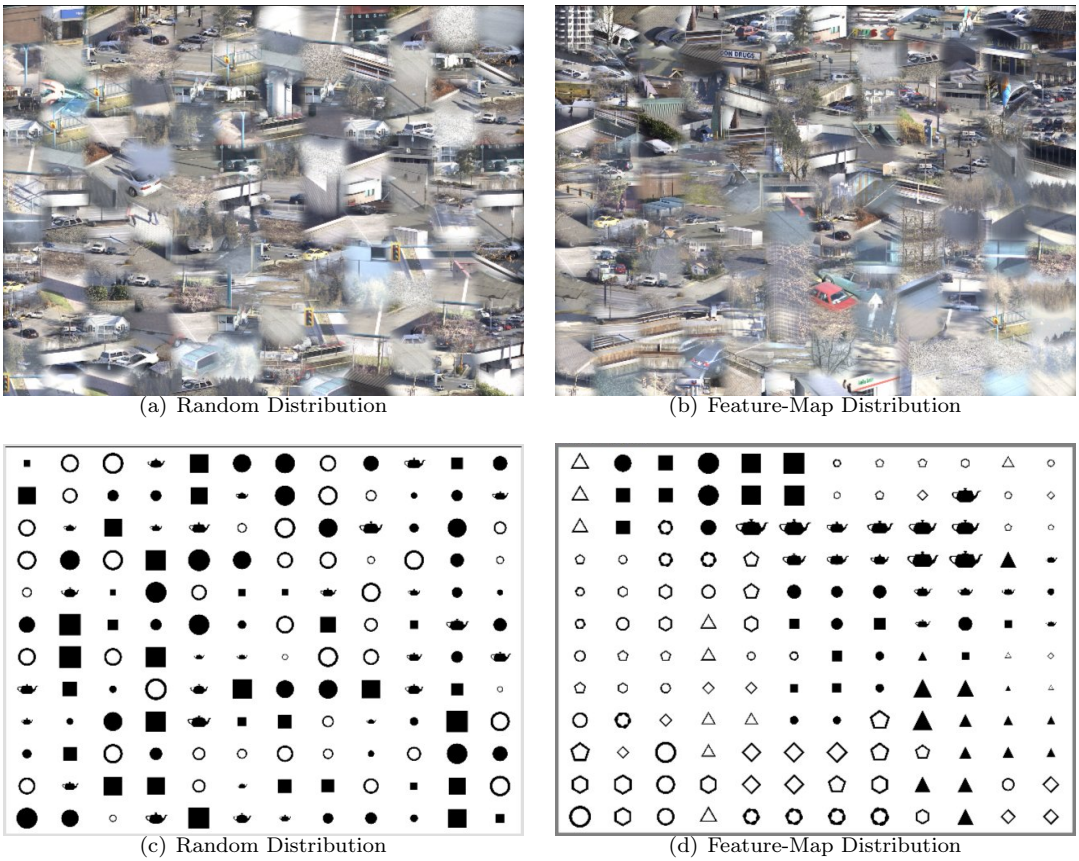

(a) Random Distribution (b) Feature-Map Distribution

(c) Random Distribution (d) Feature-Map Distribution

Figure 5.9: Comparison of the random association of sensor images to units (a and c) and feature-map trained using the MAM method (b and d).

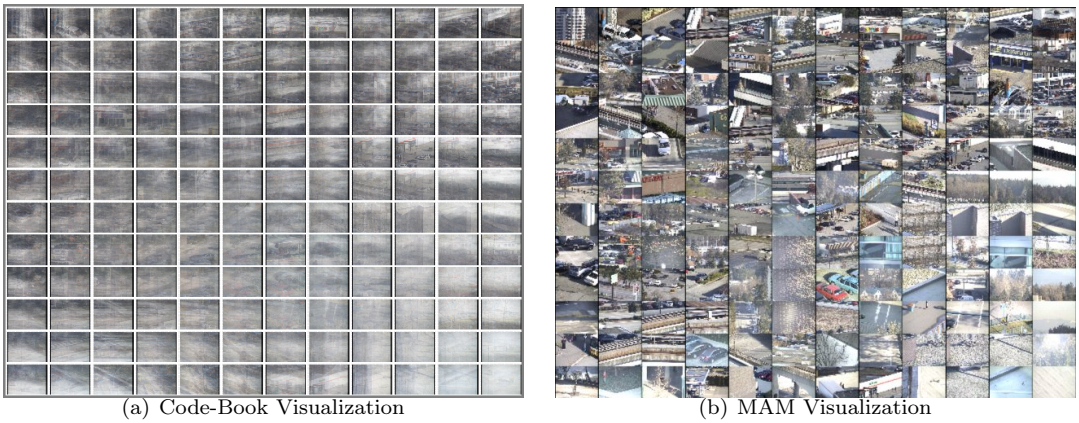

(a) Code-Book Visualization (b) MAM Visualization

Figure 5.10: Comparison between code-books (a) and MAM visualization (b) of SOM trained on a finite set of sensor impressions using the MAM method.

between the dimensionality of the input data and the feature-map results in a feature-map that is highly folded. The memory field takes on a quality very different than that of evenly distributed randomness—as pictured in Figure 5.9.

Gabora's conception of memory structure is highly analogous to the operation of the SOM. Each SOM unit is similar to a memory location. Input patterns are transformed and added to units' code-books. The neighbourhood function distributes the patterns over a number of units. This reflects two features in Gabora's conception: each memory location stores a combination of stimuli and a stimulus is associated with multiple memory locations. Both of these features are visible in Figure 5.10 on page 44. The complexity of sensor impressions is reflected in the complexity of the code-book structure. Each unit in Figure 5.10 (a) is a combination of many sensor impressions.[28]

5.1.3 Free-Association System

The free-association is a traversal through MAM's field of experience. The Free-Association System is made up of a network of identical units in the same arrangement as the SOM—a 12x12 unit Euclidean lattice. Figure 5.11 on page 45 (a) shows the arrangement of a small 5x5 lattice of free-association units. Each free-association unit corresponds to a stored memory. The purpose of the free-association network is to propagate activations to select stored memories.

For each camera fixation the SOM selects the memory most similar to the current stimulus. The BMU is then passed to the Free-Association System, which activates the corresponding free-association unit. This is pictured in Figure 5.11 (a) where unit 17 is activated. The activation then

[28]If the SOM was used directly for storage, rather to index storage, independent input patterns would not be retrievable. MAM would acquire a more ephemeral character if SOM code-books were used in place of the stored sensor impressions.

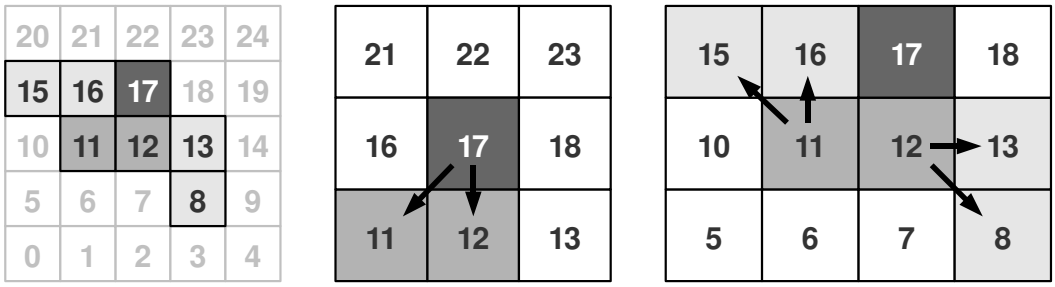

(a) (b) (c)

Figure 5.11: The propagation of free-association signals. Figure (a) shows a small 5x5 lattice of free-association units. Unit 17 is activated by external stimulus. Unit 17 passes a degraded signal to units 11 and 12 (b). The signal is further degraded and passed to units 15, 16, 8 and 13 (c).

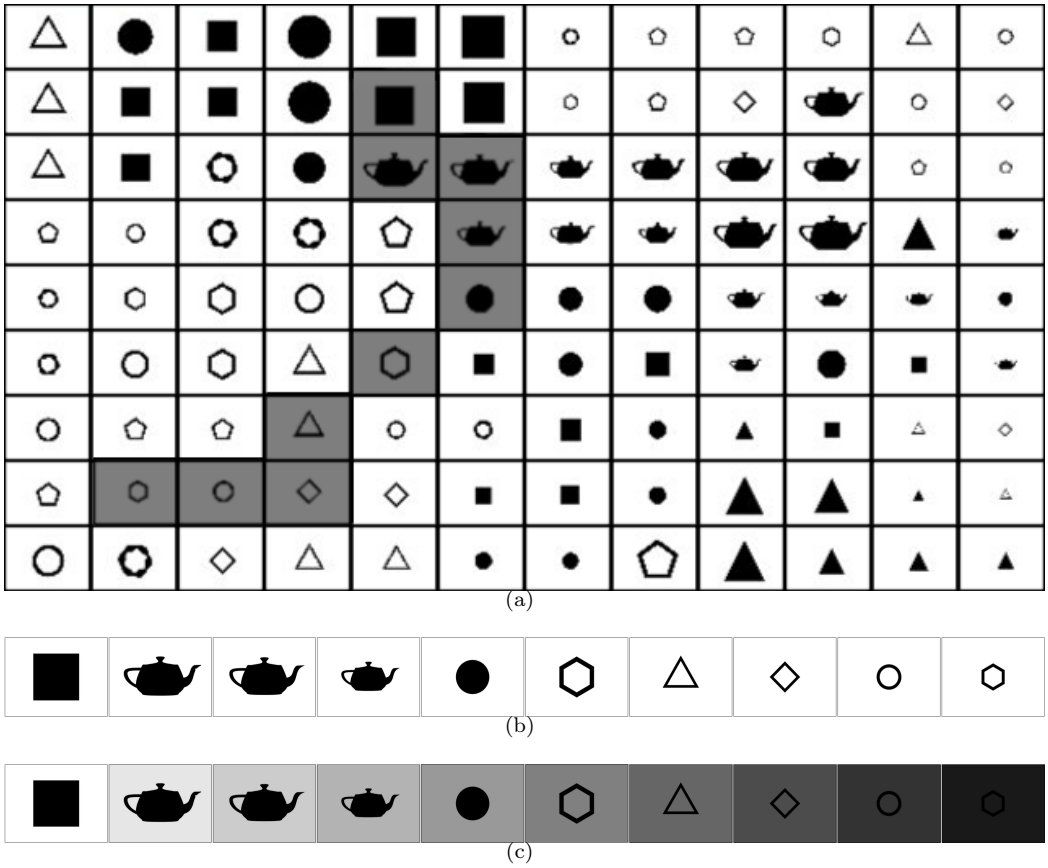

(a)

(b)

(c)

Figure 5.12: Illustration of the free-association process. Figure (a) is a portion of a feature-map trained using the MAM training method. Figure (b) shows the sequence of memories selected by the free-association. The signal degradation is represented as the darkness of each memory in figure (c).

propagates through the network—selecting memories to varying degrees. Once activated each unit generates two random numbers between 0 and 7—which correspond to the 8 directions in which a signal can be propagated. Once unit 17 is activated two random numbers are generated, 0 and 1, which correspond to the lower left and lower middle neighbours. As pictured in Figure 5.11 (b) the signal is degraded and passed to units 11 and 12. For each transfer of the signal it is decreased by 20%—10% during the sleeping state. The degree of activation falls off proportional to the distance between the initial activation and each unit receiving it. Once units 11 and 12 are activated they each generate two random numbers. Unit 11 generates 5 and 6, which activates units 15 and 16. Unit 12 generates 2 and 3, which activates units 8 and 13, pictured in Figure 5.11 (c). These units are degraded by an additional 20%. The degree of activation is illustrated by the darkness of the units in Figure 5.11.

Figure 5.12 on page 45 (a) represents the free-association units with their associated sensor impressions. Signals in Figure 5.12 (a) are propagated in one direction in order to illustrate the quality of free-associations. The darker units show a trajectory through the field. The free-association sequence is visualized as a cinematic montage on the right display. Figure 5.12 (b) shows the sequence of activated memories. Figure 5.12 (c) shows the same sequence including the degree of activation where darker images are less activated than bright images.[29] Signals are propagated between free-association units using a custom message containing the following fields:

- The (x,y) position of the sending unit.
- The (x,y) position of the destination unit
- The value of the signal.

The use of (x,y) pairs, rather than unique indices, was chosen to facilitate the calculation the immediate neighbours of the propagating unit. Messages are processed by the units whose (x,y) indices match the destination field.

Signals are temporally delayed as they propagate between units. For each direction the signal is delayed for a time-seeded random duration—ranging between 500 and 1000 milliseconds. The delay reduces the number of activations occurring nearly simultaneously as they occlude one and other in the montage. Once a unit has been activated it becomes inhibited. For a duration of 2 seconds it will not propagate any signals. The inhibition and directional control of propagation are required to keep the system from over-stimulating itself. Early implementations simply used up all the resources of the hardware only moments after the initial stimulation.

The cinematic montage can display four sensor impressions simultaneously on four independent layers. As the free-association signal propagates through the network the units' IDs—the ID in this case is the unique index, not a pair of (x,y) indices—and degrees of activation are fed into

[29] The signal decay used in the illustration is 10%.

the montage mechanism. The mechanism is composed of four FIFOs—each corresponding a layer. Activations are stored in the FIFOs in the order of activation and emptied one item at a time. For each item popped[30] the system retrieves the sensor impression corresponding to that unit ID, fades it in, delays for a duration, and then fades it out. The sensor impression is visible for a duration proportional to the unit's degree of activation ($duration(activation) = 1000 \times activation + 500$) where the resulting duration is in milliseconds. The opacity of each impression follows a similar relation ($opacity(activation) = .8 \times activation$). The result is a montage of cascading sensor impressions—starting with similar and brighter images that are visible for a longer duration. As the activation decays dissimilar impressions become darker and are visible for shorter periods.

The free-association considered in Gabora's theory is intentionally tuned to a position on the continuum between analytical and associative modes. The analytical mode is accomplished with a narrow activation function—which only activates memories highly similar to the initial activation. Associative thinking is the result of a wide activation function that activates memories at the fringes of similarity. MAM uses an activation function fixed in size and shape. The combination of the complexity of the feature-map structure and the size of the activation function leads to highly complex free-associations. Gabora's conception of an activation that bridges two islands of experience is apparent in Figure 5.12 where the propagation crosses from a region of filled shapes to a region of unfilled shapes. The relation between the first memory (a large filled square) and the last memory (a small unfilled hexagon) is highly abstract. This complexity in MAM is analogous to the kind of highly associative thinking that is the basis of highly creative thought in Gabora's conception.

5.2 Future Development

This section outlines a selection of future plans for the software development of MAM. These particular changes were chosen because of their potential positive impact on the project. Both the production journal (Appendix B) and subversion log (Appendix C) contain all development ideas that have occurred throughout the process. The changes are ordered in terms of their impact on the software. Short-term plans are listed first and followed by longer term developments. The system currently implements a novel method of integrating and free-associating through its remembered experience. MAM is stable and has been running, at the time of writing, for over 135 days and nearly 500,000 iterations. I have been awarded a Canada Council Production Grant to continue working on the project under the title "Dreaming Machine".

[30] "Popped" indicates that an item is removed from a buffer.

Improved Training Method

The current learning and neighbourhood functions were intuitively chosen. The resulting feature-map is highly folded. The most obvious way to reduce folds is to increase the number of units in the SOM. The 12x12 SOM was chosen to align with the performance constraints of the development hardware. A version of MAM with a 20x20 unit SOM has been exhibited but did not show any appreciable reduction in folds. Performance is the major factor limiting the size of the SOM.

The degree to which the feature-map is folded is due to the small neighbourhood size. The neighbourhood oscillates between a radius of 1 and 0—corresponding to a total of 9 units in the largest neighbourhood—which slows down training significantly. The use of a cosine equation to drive both the learning and neighbourhood functions is appropriate. The smoothness of the equation results in the code-books being changed in small incremental steps. The training method includes variation in the size of the neighbourhood that contributes to the self-organization. The amplitude of the neighbourhood function should be increased in order to decrease the time required for all units to be associated with sensor impressions, and increase the quality of the self-organization.[31] The MAM training method uses a hard-limit for the neighbourhood function. All code-books in the neighbourhood are changed to the same degree. A linear fall off neighbourhood function would change the code-books proportional to their distance from the BMU. During training new stimuli would have a greater effect at the centre of the neighbourhood (the BMU) than at the edges.

A larger linear-fall off neighbourhood function is proposed as an improved training method. Figure 5.13 on page 49 presents the training processes of the MAM method (a) the proposed method (b) and the canonical method (c). All three cases represent each 500th of a 5024 iteration training session. A neighbourhood size ranging from 4 to 0 is used in the proposed method. In figures (b) and (c) units are associated with input patterns more quickly than in (a).

Figure 5.14 on page 50 is a comparison between feature-maps. The feature-map resulting from the proposed method is less folded than the MAM method and more closely resembles the canonical method. To associate SOM units with sensor impressions as quickly as possible, a neighbourhood that encompasses the whole network should initiate training. Early in training the neighbourhood should shrink to a more appropriate size.

Improvement to ann_som external

The ann_som external (Zmölnig, 2001) does not include a function for generating random code-books. It is not practical to generate them from Pure Data as it would require sending the ann_som external over 4 million random values.[32] A function for generating random code-books will be

[31] At 12 seconds per iteration it takes as much as 8 hours for every unit in the field of experience to be associated with stimulus.

[32] 30,000 sensors multiplied by 144 units.

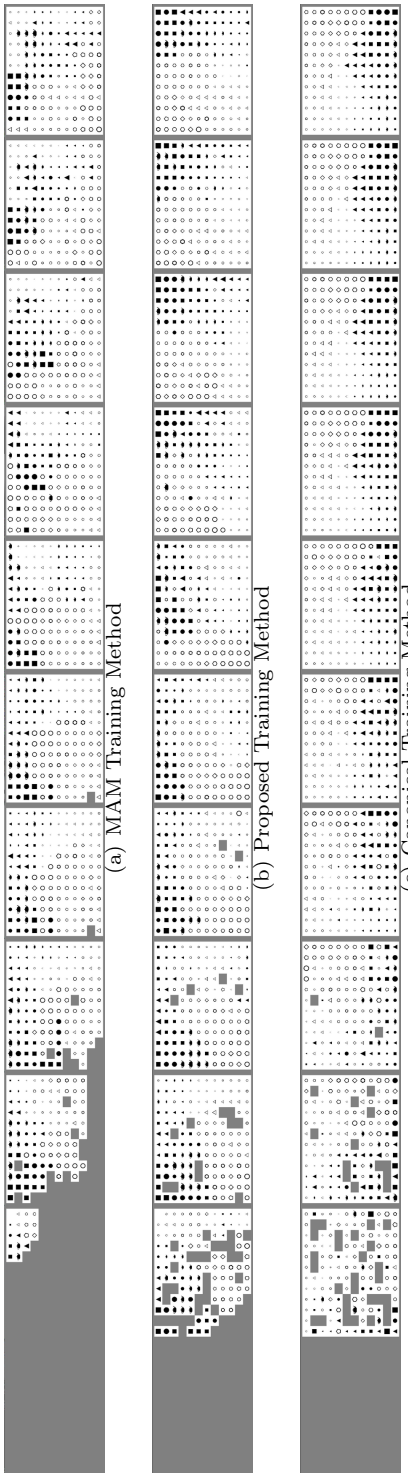

Figure 5.13: The SOM training progress using the MAM method (a) and the proposed method (b). Figure (c) uses canonical SOM training for comparison. All figures are represented using the MAM visualization method.

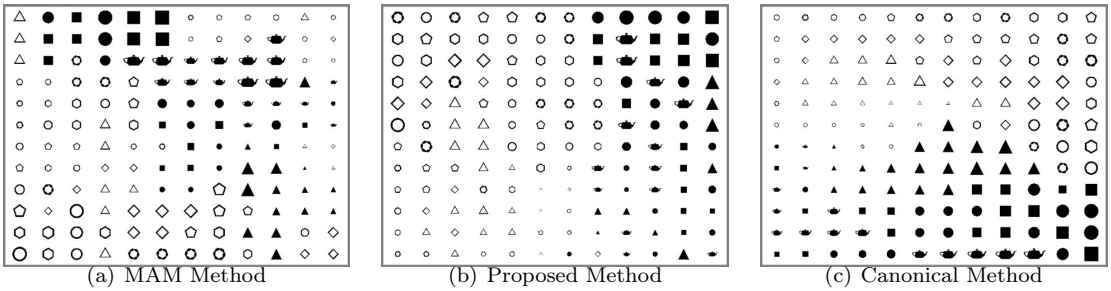

(a) MAM Method (b) Proposed Method (c) Canonical Method

Figure 5.14: Comparison between feature-maps created using the MAM training method (a), the proposed method (b) and the canonical method (c). All figures are represented using the MAM visualization method.

implemented.

Integration of memory and free-association systems

There is no synchronization of the free-association visualization between the middle and right displays. This is due to the montage mechanism being independent of the free-association units. An ideal solution would be to integrate the two systems. Each free-association unit would determine when its associated sensor impression would be visible in the montage. The timing of the montage could be bound directly to the timing of the free-association signals.[33]

Reduction of the use of random variables

Random variables are used to control camera direction and provide variation and constraint to the free-associations. The use of randomness makes the camera's exploration of context arbitrary. Randomness was chosen to avoid the explicit implementation of a model of intention. To resolve this arbitrary element the camera's direction should be a result of the free-association process. The system's exploration of context would be a result of its process of relating to context. It is not clear how this could be accomplished at this time. An alternative, likely to be implemented in "Dreaming Machine", is a simple stimulus-response model of attention.[34] The camera could look in a direction to locate the source of a sound heard from a particular direction.[35]

[33] An implementation has been attempted but required too many resources on the development machine. This is due to the increase from 4 to 144 (the number of units) montage layers.

[34] "Dreaming Machine", as proposed, uses two sensor modalities—light and sound. The system would integrate and free-associate through both visual and sonic fields of experience. It is currently unclear how the two sensory modalities should be related.

[35] This would require a system to locate sounds in space. This system would also result in a field of experience that contains only sound emitting, or reflecting, objects. It would also a require an installation where the system is not separated from its context—for example by a window.

The propagation of free-association signals requires variation for two reasons. It reduces the number of directions in which the signal is passed and keeps the montage mechanism from presenting occluding sensor impressions. An alternative to these random variables is to use the structure of the SOM to constrain and add variation to free-association trajectories. Every unit in the SOM is separated by a fixed distance in output space. It is this distance that is used to degrade the free-association signal. Each unit also occupies a position in input space that reflects its similarity to its neighbours. Under this proposal the free-association signal would degrade proportionally to the distance between sender and receiver units in *input space*. The amount of degradation would be based on the degree of similarity between sender and receiver. Free-associations within clusters of similar memories would last longer and tend to terminate once they near folds. A reinforcement model could be implemented to increase the signal, in some conditions, in order to enable the trajectory to traverse folds. This model could make the connection between two units stronger the more often those units pass signals.[36] The stronger a connection the less the signal would degrade as it is passed through a reinforced connection. The input space distance could also provide variation in the timing of propagation.[37]

MAM as self-sufficient installation

A long-term goal for MAM is to bring the system closer to being appropriate for long term public exhibition. The obsolescence and fragility of computers makes permanent electronic media art installations challenging. A permanent installation should be able to operate for 10 years without maintenance. This likely means a solid-state system with no moving parts, sealed for weather, and highly heat tolerant. Ideally the system would not rely on external infrastructure and generate its own power. These constraints would clearly change MAM but fulfil its goal as an example of the potential of electronic media as permanent public art.

Images vs sub-symbolic micro-features

A significant difference between MAM's memory system and Gabora's conception is that according to Gabora the values stored in memory are micro-features of stimulus. An entire image could certainly not be considered a micro-feature. An interesting avenue of future development would be a version of MAM that stores micro-features. Components of stimulus, rather than entire sensor impressions, would be retrieved from memory. These components would be combined to form an experience. At this time it is unclear how MAM could combine these micro-features into an intelligible structure.

---

[36] This is inspired by the concept of "long-term potentiation" in neurophysiology.

[37] At the time of writing ann_som does not allow easy access to code-book values. A new mechanism would have to be implemented to access code-books from within Pure Data.

5.3 MAM and Creativity

MAM is not a system that creates an artwork, but a system that expresses creativity, as framed by Gabora's theory. MAM is creative in its relationship to context. Context is narrowly defined as that which the machine is able to sense—visual images. MAM relates to its site by relating the current sensor impression to elements of its memory. These relationships are manifested in MAM's external properties. The mechanism that enables this is provided by the artist, but it is system that determines the relations between its own memories. The relation between elements of its experience are a result of its unique integration. They cannot be determined in advance or predicted. Gabora's theory is not bound to any particular creative endeavour but conceives of creativity as central in the construction and recollection of memories. MAM should not be evaluated for creativity based on its external properties but on the process that relates it to its context.

5.3.1 Evaluating MAM using Boden's Framework

The aspect of *newness* is the focus of this research above *surprise* and *value*. Newness comes from MAM's ability to be different for each new context. The results of its embodied negotiation create a unique, and therefore new, reflection on its context. As the external properties shift in response to a changing environment the system constantly generates novel relations between elements of context. The diversity and complexity of the real-world environment should guarantee that the system never receives an identical stimulus twice. The primary value of the project is in the process that makes its relation to context possible.

Combinatorial creativity is inevitable in a connectionist network that supports learning. This is because the shift of the units' code-books change the topology of the network—combining the features of input patterns in various ways. Exploratory creativity is also present in these systems as, through the learning process, the network explores the space of the input data. In order for a connectionist network to exhibit transformational creativity it would have to change the space of possibilities.

The combination of a SOM and model of free-association endow MAM with exploratory creativity because the free-association traverses its memory. The finite memory serves as the space of possibilities.[38] Elements of memory are related to one and other through free-association. The memory system is a constantly shifting field of experience as the system continues to integrate new sensor impressions. MAM exhibits transformational creativity through its ability to add to, and remove from, its space of possibilities over time. The space is being transformed through its process of self-organization.[39] Two identical memory traversals, occurring at sufficiently different times, would

[38] As the SOM is a 12x12 grid of memories. It has a fixed space of possibilities.

[39] A more fundamental transformation would be MAM changing the number of memories in the field.

select entirely different sets of memories from experience.[40]

Boden's conception of creativity is largely dependent on the evaluation of ideas or artifacts. This conception is at odds with this research. MAM is not a machine designed to be creative in its external properties. There is no mechanism that serves the role of evaluator for those properties. The process of integration and free-association exhibits all three of Boden's classes of creativity. MAM exhibits combinational creativity as the Free-Association System combines elements from memory. Exploratory creativity is apparent as MAM explores both its context, using its sensors, and its field of experience through the Free-Association System. MAM's field of experience is its structured space of possibilities. The system's ability to reorganize that field, in response to context, is a transformation of that space.

[40]Due to the free-association signal being propagated in random directions any particular traversal is unlikely to repeat itself.

Chapter 6

Growing Form from Context

> . . . [W]hen professionals write or speak about their own crisis of confidence, they tend to focus on the mismatch of traditional patterns of practice and knowledge to features of the practice situation—complexity, uncertainty, instability, uniqueness, and value conflict—of whose importance they are becoming increasingly aware. (Schön, 1983)

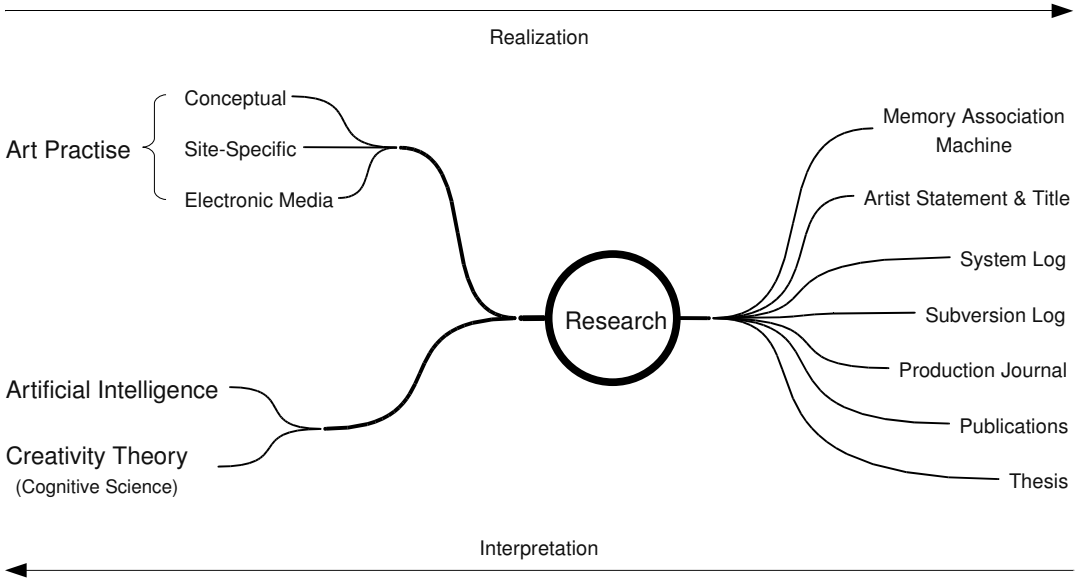

Figure 6.1: The braided practise and its major components. Research is the connective tissue that relates the contextual, theoretical and conceptual ideas that have informed it (on the left) to the physical traces of the research (on the right).

MAM is a manifestation of a process that actualizes the potentiality of context. MAM's physical context nurtures a form in the shape of its external properties. This actualization of potential is also reflected in the research. The context of ideas, pictured on the left of Figure 6.1 on page 54, nourish the knowledge resulting from research. The physical traces of that knowledge, the actualization, are pictured on the right of Figure 6.1.

The formalization of the creative process of MAM, based on the realization and interpretation of the artifact, is discussed in Section 6.1. The consideration of computational process as artistic material is discussed in Section 6.2. The collaborative aspects of this research, described in Section 6.3, are: the creation of a machine that can operate in a collaborative context and the use of free and open source software (FLOSS). Section 6.4 is a discussion of meaning in relation to MAM and the viewer.

The research process has two levels: the level of concrete production that has lead to the artwork, and the formalization of that process. The first level is specific, concrete and personal, and tied directly to MAM's production. The second level is abstract: it could be applied to other creative productions and formalizes what has occurred through development. I have worked between these two levels, working in production, and at the same time observing my own process. These two levels are not separate. The validity of the formalization is in the practise itself.[1] As MAM developed it became clear that the system was dependent upon intuitive choices. This research is deeply invested in a "reflective practise" (Schön, 1983) where "tacit knowledge" (Polanyi, 1983) is of significant value.

This chapter describes the nature of the practise through the actualization of MAM. A summary of previous chapters is provided to ground the discussion: Chapter 2 discusses relationships between conceptual, site-specific and electronic media art, and how those combine in this artistic practise. Boden and Gabora's theories of creativity are discussed in Chapter 3. The implementation of this theory, using AI techniques, is described in Chapter 5. These threads of enquiry, listed on the left side of Figure 6.1, combine to historically situate the research, providing inspiration, concepts and context. The actualization of the research results in the publications, journals and logs, that document the work, are listed on the right side of Figure 6.1. The formalization of practise has been constantly refined to reflect my experience. The formalization divides creative practise into two operations: realization and interpretation. Realization is the path of intention from the artist to the world, whereas interpretation is the path of causation from the world back to the artist.

6.1 Realization→Interpretation

The relation between the formalization's iterative operations are pictured in Figure 6.2 on page 56. Realization is the manifestation of abstract concepts in the construction of artifacts. Interpretation is where concepts are constructed, revised and reconsidered in the light of what has been accomplished

[1] For a discussion on confidence in practise based research see Candlin (2000).

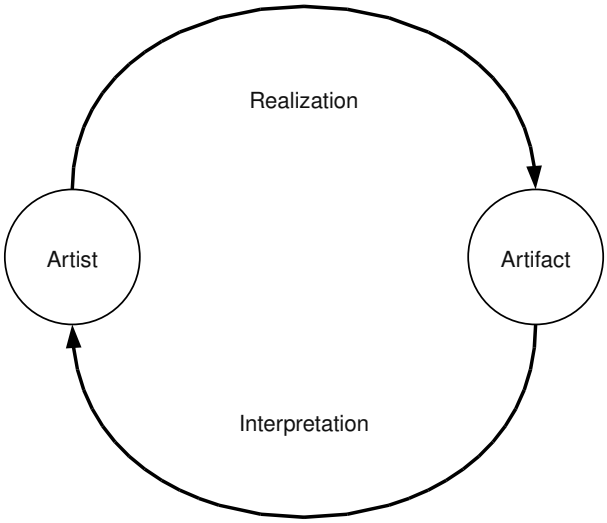

Figure 6.2: The artist/agent realizes the artifact, which is then interpreted.

in realization. The process is driven forward by the tension between what is expected and what is accomplished.

This push and pull of practise reflects the process of phenomenological embodiment conceptualized by Merleau-Ponty (1968) as the "intertwining" that binds the mind-body. The structure of realization→interpretation resembles Agre's notion of a "dynamic" that "pertains to a recurring causal chain whose path passes back and forth between components of the individual's machinery and objects in the world" (Agre, 1997). The formalization has been developed through the creative practise in an artistic context, but could hold for other creative practises such as design. Gero (2002) considers design as an "activity during which the designers perform actions in order to change the environment." Gero's conception of creativity in design process is aligned with this formalization: "By observing and interpreting the results of their actions, [designers] then decide on new actions to be executed on the environment" (Gero, 2002).

I consider artistic material as the structure in which the artist does their work of creation. In sculpture the material is physical, the clay or wood in which the artist's vision is manifest. The artistic material of MAM is the software that encodes my intentions for its computational process. Both the software, and its executed process, are interpreted. One component is built, tested for its desired function, and then components are attached and tested again.[2] Every change of software can

[2]This incremental methodology is similar to that proposed by Brooks (1992).

have a large effect on the system's function.[3] This incremental development follows the structure of realization→interpretation as the programmer's intention is compared to the results of every software change.[4]

The artist is then working with two different materials: the material of the software and the material of its executed process. When MAM's external properties are observed they stimulate the artist's memory and experience as considered by Gabora's model of creativity. Interpretation is a placement of the stimulus in the context of the experience of the viewer.[5] The interpretation may uncover some new insight that shifts the development of the project.

Figure 6.2 does not include a representation of the goals or purpose of creative process, as they often evolve through practise. Every interpretation of the artifact causes a "perturbation" (Gabora, 2000) of the world-view. In MAM's process the perturbation is caused by the physical context of the work. The system is only able to effect the context in a limited way. In the artist's creative process the source of perturbation is the result of the artist's action. This is a major difference between the creative process of the artist and that of MAM.[6]

The next realization is made in a slightly different context than the last. The creative process can move from one expected outcome to something entirely different. This flexibility allows creative process to find new goals through practise. Any stimulus from the environment can inspire an interpretation that leads to a new creative process. This interpretation becomes a discovery that seeds a new creation.

Figure 6.2 is a simplified representation. A single artifact may contain other artifacts, which in turn may contain additional components, as pictured in Figure 6.3 on page 58. MAM is composed of thousands of small computational components related to one and other at different levels. A creative process need not contain a single artist, as is the case in collaboration.[7] Each component is embedded in an independent realization→interpretation loop. The time it takes to complete a loop is entirely dependent on the nature of the creation.

6.2 Artifacts-as-Processes

> The age-old functions of art—to provide communication, experience, insight, or entertaining distraction—remain what they have always been: a matter of human beings working individually or collectively to create something that did not exist previously and is the

[3] Function is referring to what the software does, which may not be what it was intended to do.

[4] The creation of software is analogous to the creation of a physical artifact. By no means is computation virtual. It is physically manifested as electrical patterns.

[5] The artist is also considered a "viewer".

[6] If the external properties of MAM are considered its artistic artifact then its creative process could be more closely aligned with the artist's by focusing the attention of the camera on its own displays.

[7] A single artifact may also contain a number of components at the same level—not depicted in this figure.

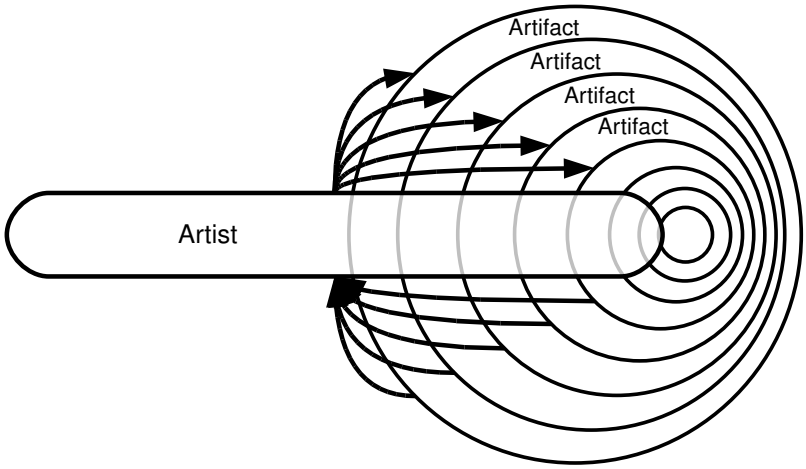

Figure 6.3: The nested realization→interpretation loop.

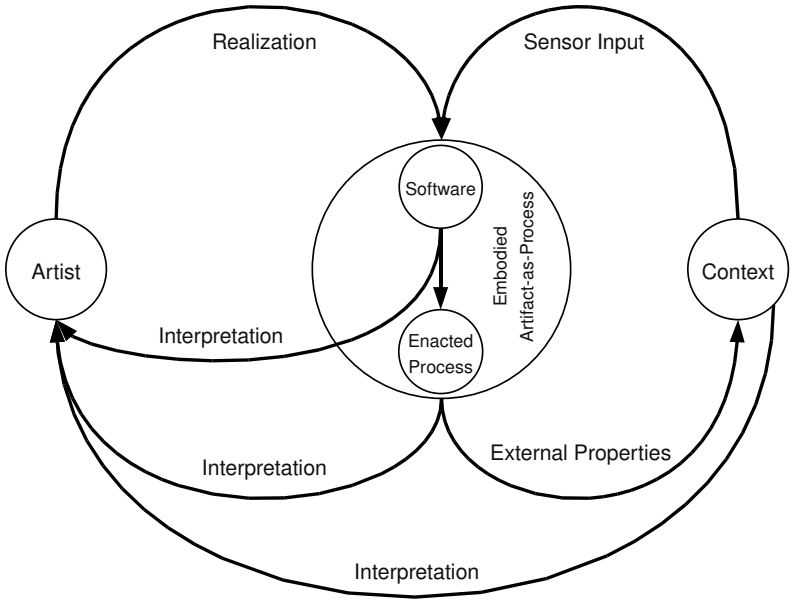

Figure 6.4: The realization→interpretation loop of the embodied artifact-as-process.

> product of imagination meeting material (even the "immaterial" of electronic of means) to be made manifest as process or form. (Drucker, 1997)

An artifact-as-process is an artifact composed of computational processes. They are snapshots of the process of *their* creation and are enacted processes in *themselves*. The artifact-as-process can not be reduced to software code. Software is instructions for a process, rather than an enacted process itself. Embodied artifacts-as-processes are systems that change their *own* external properties to relate to, through the reflection and interpretation of, their embodied context. The context can impact the artwork (through sensors), and the artwork can impact the context (through its external properties).

Figure 6.4 on page 58 shows the causal influences on the embodied artifact-as-process. These influences are organized in terms of their relationships of realization and interpretation. At the centre of the diagram the software is executed, and processes sensor input, resulting in the enacted process. Three elements in the production of an embodied artifact-as-process are interpreted; the software, the external properties and the context:

- The software is directly interpreted as it is being written.
- The enacted process is interpreted, in light of the installed context,[8] through its manifestation in the external properties.

The behaviour of the system (its external properties) determines the future developments of the software. Monterege (1989) considers the importance of evaluating a system in context, rather than relying on isolated tests:

> The real demonstration of the functionality of a system does not lie in a test. This demonstration is achieved by observing the behavior of the system and its impact on its intended environment.

6.3 FLOSS & Collaboration

MAM is a collaboration: a site of the interaction of influences. The essence of collaboration permeates the artistic process. MAM's production is effected by the following influences: the software, the physical context, the hardware, the operating system, art history and cognitive science. MAM reflects this dependence on its technological and social context by being physically responsive to its environment.

For a time I considered my interest in the system acting beyond my intentions as a removal of myself—a removal of my intention from the system. In this attempt I was instituting software

[8] An interpretation of context is required in order to determine the qualities of the system's relationship to its context.

mechanisms that more deeply ingrained my intention in MAM. This removal was an attempt to remove my "hand" from the work.[9] My intention shifted from the design of the system's external properties to the design of the interface between the context and the system. This effort moved to a different level, rather than removed, my influence over the system.

MAM is produced in a FLOSS context. With a single exception[10] the installation[11] uses entirely FLOSS. The whole research project, from writing this thesis, to designing its figures, to generating images of SOM code-books, use only FLOSS.[12] As much as possible I use older hardware so that my work can serve as an example for those with limited means of production. The Elmo pan/tilt/zoom camera[13] used in MAM requires no license to access its serial control protocol.[14] The code used to control the camera has been released under the General Public License[15] (GPL) so that it can be easily used by other practitioners working in Pure Data.

The use of FLOSS is not simply a practical choice, but deeply embedded in the whole production of MAM. While MAM itself can be considered a collaboration between myself and its context, it is also a collaboration between the many developers who release the software I used to create MAM. This type of collaboration is present in all electronic media art because the artist works in a technical context that is dependent on the work of technologists, engineers and software developers. This project is embedded within these power relations by the choice of software. Pure Data exists because the author, Miller Puckette, wanted to create a platform for computer music that would be as widely available as possible.[16] The choice of what software, and hardware, is used in a project situates that production in an ideological context. At its most basic level MAM is about sharing influence and control in a community—not about the sole genius of a single author.

---

[9] The "artist's hand" refers to the individual character or style of the artist present in their work.

[10] MAM uses a proprietary graphics driver provided by Nvidia™.

[11] The long-term testing installation machine (sr-00150) runs on an AMD Athlon™64 2500+ with an Nvidia™FX5600 graphics card, running Debian Etch (`http://www.us.debian.org/`).

[12] The thesis has been written in LyX (`http://www.lyx.org/`), using Ubuntu (`http://www.ubuntu.com/`) on a eight year old AMD Duron™800Mhz PC. All figures have been designed in OpenOffice Draw (`http://www.openoffice.org/`) and visualizations of code-book data have been produced in R (R Development Core Team, 2007) using the ggplot2 graphics package (`http://had.co.nz/ggplot2/`). Image manipulation has been done in GIMP (`http://www.gimp.org/`) and ImageMagick (`http://www.imagemagick.org/`). The analysis of SOM code-books and MAM logs have also been accomplished in R running on a three year old AMD Athlon™64 3200+ running Debian (`http://www.us.debian.org/`).

[13] PTC-100 (`http://www.elmousa.com/presentation/index.html`)

[14] Companies like Canon™require a developer to agree to a license before accessing the protocol used to control their cameras.

[15] GPL: `http://www.fsf.org/licensing/licenses/`

[16] Pure Data was written from scratch by Miller Puckette after his first computer music system (MAX) became a commercial product.

6.4 Meaning

This section discusses the consideration of meaning as conceived through MAM's creation. Meaning is considered in light of the context of Gabora's theories discussed in Chapter 3. The consideration is composed of two parts: the first, discussed in Subsection 6.4.1, is consideration of the interpretation of artifacts in terms of Gabora's model. This consideration is aligned with Cohen's conception of standing-for-ness—where meaning is constructed largely by the viewer. This section is not a semiotic or semantic analysis of meaning in MAM, but a description of how this project has lead me to think about artistic meaning. Subsection 6.4.2 discusses the notion of software as artist statement and a structure to guide the interpretation of artifacts-as-processes.

6.4.1 The Perturbation of the Viewer

The meaning of an artwork, for the viewer, concerns the relationship between the whole of the viewer's experience and the artifact. The artwork becomes meaningful as the perception of it, and its properties, are situated in the world-view of the viewer. A few points follow that map this conception of meaning to my experience of interpreting artistic artifacts:

- The experience of an artifact stimulates a memory of previous experience—linking the artifact to the history of remembered artifacts. The artwork is made meaningful through its location within the world-view. The world-view is enriched and *perturbed* through the inclusion of the new stimulus.
- The experience of an artifact goes beyond being situated in the world-view of the viewer; it inspires the viewer to act. As the stimulus *perturbs* the world-view activation expands through memory. This activation has the potential to bridge conceptions of experience—leading to a new way of looking at the artifact, and the world.

In contemporary artistic practise the artist statement and title, the texts that accompany the artwork, are significant in the process of interpretation.[17] These texts and the artwork are engaged in a complex relation that transforms the meaning of both. The reading of the artwork is coloured by the text, and the reading of the text coloured by the artwork. The text and artwork are two physically independent artifacts that are highly correlated in terms of meaning-making. A tension is almost inevitable in the integration of the meaning of both artifacts. This tension "... is resolved by a reworking of the visual configuration and the meaning of the title so that some kind of correspondence or 'fit' is established between the two" (Franklin, 1988). The "reworking of the visual configuration" is the viewer reconsidering the image in relation to the title. The title perturbs the world-view and allows the image to be seen differently. Shapes that were initially ambiguous become clear. Franklin

[17] For a discussion on the importance of titles in the interpretation of painting see Heffernan (1985).

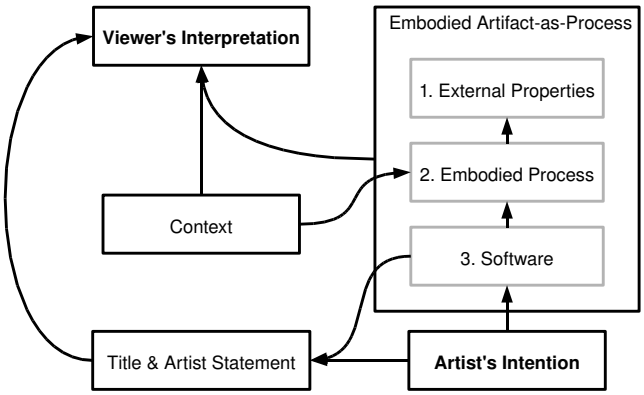

Figure 6.5: The components of meaning in an embodied artifact-as-process.

is concerned centrally with the title in relation to the meaning of the artwork. I would argue that the artist statement functions in largely the same fashion. It is not the title alone, but all text accompanying the work that binds the meaning of the artifacts.[18]

6.4.2 Software as Artist Statement

A valuable source, in the interpretation of artistic intention, is the software of the embodied artifact-as-process. The software is often invisible to the viewer and yet has a significant impact on what the artwork is meant to do, and does. The role of the text in conceptual art is as a gateway between the physical material and the concept of the work. The embodied artifact-as-process contains a number of components that each contribute to the reception of the work—as pictured in Figure 6.5 on page 62. Each of these components is considered a layer dependent on the components beneath it. The first (topmost) layer is the external properties of the artifact—at a particular moment in time. This layer is the primary site for reception, as it is the most accessible to the viewer.[19] The external properties of the artwork are dependent on the second layer—the enacted process. This layer is more difficult to interpret as it is manifested in the changes, over time, of the external properties. The enacted process is available to the viewer through an invested interpretation of the artwork over

[18]Every aspect of the sensed environment effects the world-view and has the potential to shift the interpretation of the artwork.

[19]By "accessible" I mean that any attributes of this layer are directly and immediately perceivable by the viewer.

time.[20] The enacted process is dependent on the context and the third layer—the software that encodes the artists intention. The software is not visible to the viewer directly as it is abstracted by the layers above it. The software is similar to the artistic idea in conceptual art, which is not directly accessible, but apparent through its realization in the material of the artwork and the text.

The embodied artifact-as-process cannot be appropriately interpreted in isolation, as it is dependent on its context. The viewer's interpretation is dependent on the artwork, the text that accompanies it, and the context in which it is installed. The text provides a primer for reading the artist's intention through the external properties of the work. The context is of central importance in the interpretation of MAM, whose external properties must be considered in relation to that context. The causal dependence of the external properties on the context would only be accessible by observing the context.

The process of the development of MAM is a reflective practise deeply rooted in tacit knowledge. The research methodology is a braided practise that is contextualized within the foundations of art-practise, a cognitive theory of creativity and AI techniques. The research outcomes are the MAM system, the journals documenting the process, this thesis and other publications. A formalization of this practise is composed of two processes, realization and interpretation, which allows creative flexibility and a rigorous equilibrium between MAM's implementation and its artistic concept. This research project is intrinsically collaborative, from its braiding of multiple disciplines to MAM's technological implementation, which is situated in the FLOSS movement. The research has resulted in a consideration of meaning that is manifested in the structure of components that effect the reading of MAM.

[20]The interpretation is invested because it requires patience and time to uncover the operation of the system through its external properties.

Chapter 7

Conclusion

My interests in art are concerned centrally with *process*. Art is about *action*, an experiment, an attempt to *do something*. It is not simply representation. The artifact is documentation of the process. MAM is a snapshot of that project—a fusion of artistic practises and computational techniques informed by cognitive science.

The artistic practises of site-specific, conceptual and electronic media art are highly complimentary. Each of these practises is engaged in questions of the art object, the author and the material. Site-specific art rejects the gallery in order to place art in the community. The artist is engaged in a collaboration between their artistic ideas and the context in which the artwork exists. Conceptual art moved the artist from an engagement with materials to an engagement with ideas. The artifact is a manifestation of the concept. The artistic idea is central. At its essence electronic media art is an artistic exploration of the creative potential of both the construction, and use of, contemporary technology. The artistic practise resulting in MAM inherits the importance of context, in the meaning of the artwork, from site-specific art. The combination of electronic media art and site-specific art leads to a practise where the artifact interactively responds to its context. The importance of the artistic idea in conceptual art manifests itself in this practise through the consideration of software as the implementation of the artistic idea.

MAM is considered a creative machine, but not a machine meant to create an artistic product. Creativity has a central role in the way MAM relates to its context. Boden considers creativity largely from an evaluation stand-point. MAM has no creative product that can be evaluated in terms of being new, surprising and valuable. MAM exhibits all three types of creativity discussed by Boden, combinational, explorational, and transformational. In counter point to Boden's conception of creativity, as a process depending on evaluation, Gabora considers creativity as intrinsic to the way the human mind processes and remembers stimulus. The essence of this model of creativity is based on two attributes of human memory: memory is considered both sparsely distributed and content-addressable. This structure of memory allows stimulus to be stored in a way that integrates

new experiences in terms of the whole of the remembered experience. The integration of stimulus both changes the memory, and changes how that memory is recalled.

MAM relates to its context using three primary processes that are executed in parallel. Firstly MAM perceives its visual context through a camera. These sensor impressions are integrated into a constantly shifting field of experience. MAM free-associates through these experiences by tracing the effect of the current stimulus's perturbation of that field. Each of these processes is visualized on a corresponding display of a triptych. The combination of the SOM and model of free-association is a novel method of organizing and relating the constantly changing external context.

The process of constructing MAM is centrally an artistic practise. Through practise a structure of creative process has evolved. This formalization describes both my own creative process, and the future of what MAM's process could be. In the formalization, creativity is considered intrinsically embodied and cyclical. The creation of an artifact is made up of two operations. Realization is the path of intention from the artist to the material of the artifact. The results of this change are interpreted in the second operation, where the artifact impacts the artist. The rigour of this research project is a result of the reflective practise itself. The practise is invested in the use of tacit knowledge and involves my continuous questioning of what MAM should do and how it should do it.

MAM is an illustration of the potential at the centre of this research. MAM is one of many possible realizations of the intersections between art-practises AI techniques, and theories of creativity, embodiment and memory. This thesis is as much a part of the artwork as MAM is part of this thesis.

Bibliography

Agre, P. (1997). *Computation and Human Experience*. Cambridge University Press.

Agre, P. E. (1995). The soul gained and lost: artificial intelligence as a philosophical project. *Stanford Humanities Review*, 4:1–19.

Alberro, A. (2003). *Conceptual Art and the Politics of Publicity*. MIT Press.

Barker, E., Webb, N., and Woods, K. (1999). *The Changing Status of the Artist*. Yale University Press.

Barr, A. and Feigenbaum, E., editors (1981). *The Handbook of Artificial Intelligence*, volume 1. Morgan Kaufmann.

Barron, F. (1963). *Creativity and psychological health: origins of personal vitality and creative freedom*. New York: Van Nostrand.

Barthes, R. (1977). The death of the author. *Image, Music, Text*, pages 142–148.

Blackmore, J. and Miikkulainen, R. (1993). Incremental grid growing: encoding high-dimensional structure into a two-dimensional feature map. In *Neural Networks, 1993., IEEE International Conference on*, pages 450–455 vol.1.

Blackmore, J. M. (1995). *Visualizing high-dimensional structure with the incremental grid growing neural network*. PhD thesis, University of Texas at Austin.

Boden, M. A. (1994). *Dimensions of Creativity*. MIT Press.

Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103:347–356.

Boden, M. A. (2004). *The Creative Mind: Myths and Mechanisms*. Routledge; London, 2nd edition.

Bogart, B. D. R. (2001). Art in the face of the sublime.

Bogart, B. D. R. (2007). Self-other organizing structure 1: Seizures, blindness & short-term memory. In Brouse, A., editor, *Proceedings of the Second International Pure Data Conference, Montreal; August 21-26, 2007*. Pure Data Conference 2007.

Bogart, B. D. R. (2009). *The Handbook of Research on Computational Arts and Creative Informatics*, chapter Memory Association Machine. IGI Global.

Brooks, R. A. (1992). Intelligence without representation. *Foundations of Artificial Intelligence*, 47:139–159.

Buchloh, B. H. D. (1990). Conceptual art 1962-1969: From the aesthetic of administration to the critique of institutions. *October*, 55:105–143.

Candlin, F. (2000). A proper anxiety? practice-based phds and academic unease. *Working Papers in Art and Design*, 1.

Carpenter, G. A. and Grossberg, S. (1994). *Adaptive Resonance Theory*. Boston University, Center for Adaptive Systems and Dept. of Cognitive and Neural Systems.

Cohen, H. (1979). What is an image? In *Proceedings of IJCAI*.

Cohen, H. (1995). The further exploits of aaron, painter. *Stanford Humanities Review*, 4:141–158.

CollabNet (2000). Subversion.

Csíkszentmihályi, M. (1988). Society, culture, and person: A systems view of creativity. *The nature of creativity*, pages 325–339.

Dewing, K. and Battye, G. (1971). Attention deployment and nonverbal fluency. *Journal of Personality and Social Psychology*, 17(2):p214 – 218.

Drucker, J. (1997). Digital reflections: The dialogue of art and technology. *Art Journal*, 56(3):2.

Dykes, M. and McGhie, A. (1976). A comparative study of attentional strategies of schizophrenic and highly creative normal subjects. *The British Journal of Psychiatry*, 128:50.

Franklin, M. B. (1988). "museum of the mind": An inquiry into the titling of artworks. *Metaphor and Symbol*, 3:157 – 174.

Fritzke, B. (1991). Unsupervised clustering with growing cell structures. In *Neural Networks, 1991., IJCNN-91-Seattle International Joint Conference on*, volume ii, pages 531–536 vol.2.

Gabora, L. (2002a). *Creative Evolutionary Systems*, chapter The beer can theory of creativity, pages 147–161. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Gabora, L. and Aerts, D. (2002). Contextualizing concepts using a mathematical generalization of the quantum formalism. *Journal of Experimental & Theoretical Artificial Intelligence*, 14(4):327.

Gabora, L. M. (2000). *Toward a Theory of Creative Inklings*, pages 159–164. Intellect Press, Oxford England.

Gabora, L. M. (2002b). Cognitive mechanisms underlying the creative process. In T. Hewett and T. Kavanagh, editors, *Proceedings of the Fourth International Conference on Creativity and Cognition*, pages 126–133.

Gero, J. S. (2002). Computational models of creative designing based on situated cognition. In *Proceedings of the 4th conference on Creativity & cognition*, pages 3–10, Loughborough, UK. ACM.

Heffernan, J. A. W. (1985). Resemblance, signification, and metaphor in the visual arts. *The Journal of Aesthetics and Art Criticism*, 44(2):167–180.

Kaji-O'Grady, S. (2002). National identity at arakawa & gins' site of reversible destiny—yoro, japan. *name Fabrications: The Journal of the Society of Architectural Historians, Australia and New Zealand*, 12:19–34.

Kaski, S., Venna, J., and Kohonen, T. (2000). Coloring that reveals cluster structures in multivariate data. *Australian Journal of Intelligent Information Processing Systems*, 6:82–88.

Kaye, N. (2000). *Site-specific Art: Performance, Place and Documentation*. Routledge.

Kohonen, T. (2001). *Self-Organizing Maps*. Springer.

Kraaijveld, M. (1992). A non-linear projection method based on kohonen's topology preserving maps. *Pattern Recognition, 1992. Vol.II. Conference B: Pattern Recognition Methodology and Systems, Proceedings., 11th IAPR International Conference on*, pages 41–45.

Kwon, M. (2004). *One Place After Another: Site-Specific Art and Locational Identity*. MIT Press.

Legrady, G. and Honkela, T. (2002). Pockets full of memories: an interactive museum installation. *Visual Communication*, 1(2):163–169.

LeWitt, S. (1999). Sentences on conceptual art. *Alberro and Stimson*.

Lippard, L. (2001). *Hot Potatoes: Art and Politics in 1980 (1981)*, pages 107–118. Blackwell Publishers Ltd.

Lippard, L. R. (1997). *Six Years: The Dematerialization of the Art Object from 1966 to 1972...* University of California Press.

Manovich, L. (2001). *The Language of New Media*. MIT Press.

Mark Danks, Günter Geiger, Johannes M. Zmölnig, Chris Clepper, and James Tittle II (1995-). Graphics environment for multimedia.

Martindale, C. (1977). Creativity, consciousness, and cortical arousal. *Journal of Altered States of Consciousness*, 3:69–87.

Martindale, C. and Armstrong, J. (1974). The relationship of creativity to cortical activation and its operant control. *Journal of Genetic Psychology*, 124:311–320.

Medler, D. A. (1998). A brief history of connectionism. *Neural Computing Surveys*, 1:61–101.

Mendelsohn, G. A. (1976). Associative and attentional processes in creative performance. *Journal of Personality*, 44(2):341.

Merkl, D. and Rauber, A. (1997). *Alternative ways for cluster visualization in self-organizing maps*, pages 106–111. Helsinki University of Technology, Neural Networks Research Centre, Espoo, Finland.

Merleau-Ponty, M. (1968). *The Visible and the Invisible, trans.* Evanston: Northwestern University Press.

Monterege, J. T. (1989). The creation of digital consciousness. *SIGART Bull.*, 109:30–33.

Oxford English Dictionary Online (Accessed 2008). Oxford english dictionary online.

Polanyi, M. (1983). *The tacit dimension.* Gloucester, Mass.

Popper, F. (1968). *Origins and Development of Kinetic Art.* New York Graphic Society.

Possiant, L. (2007). *Media Art Histories*, chapter 12, pages 229–250. MIT Press.

Puckette, M. et al. (1996). Pure data.

R Development Core Team (2007). *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Rauber, A. (1996). Cluster visualization in unsupervised neural networks. Master's thesis, Technische Universitit Wien, Austria.

Rokeby, D. (1990). The Giver of Names.

Rokeby, D. (1995). *Transforming mirrors: subjectivity and control in interactive media*, chapter 7, pages 133–158. SUNY Press.

Rokeby, D. (2001). n-Cha(n)t.

Saunders, R. and Gero, J. S. (2001a). Artificial creativity: A synthetic approach to the study of creative behaviour. *Computational and Cognitive Models of Creative Design V, Key Centre of Design Computing and Cognition, University of Sydney, Sydney*, pages 113–139.

Saunders, R. and Gero, J. S. (2001b). The digital clockwork muse: A computational model of aesthetic evolution. *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in Arts and Sciences*, 1:12–21.

Saunders, R. and Gero, J. S. (2002). How to study artificial creativity. *Proceedings of the 4th conference on Creativity & cognition*, pages 80–87.

Scharmer, C. O. (2000). Three gestures of becoming aware. conversation with professor francisco varela.

Schön, D. A. (1983). *The Reflective Practitioner: How Professionals Think in Action*. Basic Books.

Steinheider, B. and Legrady, G. (2004). Interdisciplinary collaboration in digital media arts: A psychological perspective on the production process. *Leonardo*, 37(4):p315 – 321.

Tasdemir, K. and Merenyi, E. (2006). Data topology visualization for the self-organizing map. *European Symposium on Artificial Neural Networks (ESANN'2006), Bruges, Belgium, d-side eds*, pages 277–282.

Turing, A. (2004). *The Essential Turing: Seminal Writings in Computing, Logic, Philosophy, Artificial Intelligence, and Artificial Life, Plus the Secrets of Enigma*, chapter Computing Machinery and Intelligence (1950). Oxford University Press, USA.

Ultsch, A. (1993). Self-organizing neural networks for visualization and classification. *Information and Classification*, pages 307–313.

Ultsch, A. (1999). Data mining and knowledge discovery with emergent self-organizing feature maps for multivariate time series. *Kohonen Maps*, pages 33–46.

Ultsch, A. (2003a). Maps for the visualization of high-dimensional data spaces. *Proc. Workshop on Self organizing Maps*, pages 225–230.

Ultsch, A. (2003b). U*-matrix: a tool to visualize clusters in high dimensional data. *Dept. of Computer Science University of Marburg, Research Report*, 36.

Ultsch, A. and Siemon, H. (1989). Exploratory data analysis: Using kohonen's topology preserving maps. Technical Report 329, University of Dortmund, Germany.

Vesanto, J. (1999). Som-based data visualization methods. *Intelligent Data Analysis*, 3:111–126.

Wilson, S. (1995). Artificial intelligence research as art. *Stanford Electronic Humanities Review*, 4(2).

Wilson, S. (2002). *Information Arts: Intersections of Art, Science and Technology.* MIT Press.

Yin, H. (2002). Visom - a novel method for multivariate data projection and structure visualization. *Neural Networks, IEEE Transactions on*, 13:237–243.

Zmölnig, J. M. (2001). ann_som: Component of the Artificial Neural Network library for Pure Data.

Appendix A

Implementation

MAM is written entirely in the Pure Data visual programming system (Puckette et al., 1996). The visual representation of the system, its external properties, is created using the "Graphics Environment for Multimedia" (GEM) (Mark Danks et al., 995) that provides OpenGL access to Pure Data. Programs written in Pure Data, called "patches", are constructed by visually connecting boxes, called "objects", that represent computational elements. The syntax is visual but functions and operations are selected textually. A simple Pure Data patch that roughly corresponds to the C analogue:

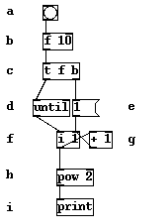


```
float b=10;
for (int f=1; f<=b; f=f+1) {
   printf("%d", pow(f,2));
}
```

Object "a" is a bang. It is an interactive object that initiates the processing of this patch. Object "b" creates a new float with the value of 10. The collection of objects d,e,f and g combine to form a "for" loop. The loop iterates 10 times (the value of object b). For each iteration the object "f" is incremented by 1. This value is squared by object "h" and finally printed out by object "i". At the time a writing the MAM patch is made up of approximately[1] 27,000 such objects. As the project progresses each step of development is managed by the Subversion Version Control System (CollabNet, 2000). A log of changes made to the patch is included in Appendix C.

[1]This estimate is highly inaccurate as the patch is highly nested and many objects are instantiated in multiple locations over multiple levels.

Appendix B

Journal

This appendix is the production log for Memory Association Machine. The text describes the progress of the project at a high-level. A log tracking the day to day practise is included in Appendix C. By in large issues of behaviour and concept are considered, rather than technical details. The following uses the original language and is edited for grammar and structure. In these entries I address myself as both "I" and "you".

April 12th 2007

I will probably be using tomboy to do the real notes. Tomboy needs better support for saving related groups of notes in different places.[1] Why does the SOM return the same BMU for any image[2] for the first 10 or so iterations?[3] Sometimes it looks like a new image should be captured, but the node stays empty, perhaps the node update stuff is still not working quiet right? Pauses between creative actions are a little slow—maybe too much dependence on random variables.[4]

The system is very much not critical about what it is looking at, seems some computer vision to basically evaluate each image is needed.[5] The use of a weighted randomness for the pan/tilt could be

[1] The "real notes" are referring to this journal. Tomboy was not often used as the version included in Debian Etch did not have any remote sync features. The journal ended up being a text file included in the subversion repository.

[2] An "image" is analogous to a "sensor impression".

[3] This was caused by an interaction between small neighbourhood, learning function and that the first BMU, when all units are identical, is always the one with the largest index.

[4] There was initially a significant problem with keeping the activation of nodes, and the montage that was representing it in sync. Over time the "live" activations would happen further away in time than the montage that showed that behaviour.

[5] Using computer vision to give the system some critical awareness of the content of sensor impressions has not been explored. This is due to two issues: firstly that it is unclear *how* the system should be critical. Secondly computer vision used to evaluate the content of the images would constrain the generality of "context" and force the system to observe objects in the context with specific attributes.

nice, though the table would have to be very large. How would the weights be determined, especially considering the interaction between zoom and pan/tilt? This may not be an issue though since the weighted randomness can't be worse than uniformly distributed randomness.

Maybe the BMU stays the same for so long because the node weights are not randomized by default and I should be doing it randomly.[6] It could be the learning rate and neighbourhood size functions that are not working well.[7] I would really like to get going with all three screens as separate panels. I will be hard to arrange those in this space though.[8]

April 19th 2007

I received positive feedback after the discussion around SOOS1[9] in AI with Marek today.[10] Some ideas:

- Trigger camera movement from a lack of activation. A longer activation cascade would allow the camera to settle on one place during the free-association—rather than moving at a fixed interval.
- Use RF context,[11] rather than random variables, for the individuality of each unit—the random variables that allow each free-association unit to behave slightly differently.[12]
- Increase the resolution of the SOM and abstraction of camera image.[13]
- Increase the delay between propagation of the free-association signal.[14]
- Add a scheduling system, that can be controlled via the terminal, so that the camera turns on

[6]Using random weights would certainly have kept the same BMU from being chosen during the initial iterations.

[7]Indeed the intuitive choices for learning and neighbourhood functions have turned out to be problematic and based on false assumptions.

[8]The initial version used a single 1280x1024 19 LCD screen in which each of the three images was presented at 640x480. The installation space referenced is the apex of room 14-800 in the SFU Surrey campus.

[9]The original title of "Memory Association Machine" was "Self-Other Organizing Structure #1". The former title made reference to both the Self-Organizing Map, and the work of Liane Gabora in regards to abiogenisis.

[10]The project was initiated in Marek Hatala's IAT813 Artificial Intelligence for Art and Design.

[11]The initial conception of context included the radio frequency spectrum at the site.

[12]The "individuality" of each unit was how much each unit should degrade the signal and change the timing of the free-association signal propagation.

[13]The resolution of the SOM has been fixed at 12x12 units due to hardware constraints. The installation of the project at the Pure-Data Convention 2007 used 20x20 units on faster hardware. The abstraction (the sub-sampled version of the video stream) was initially 4x3 pixels and later increased to 40x30 pixels for the Pure-Data Convention. The resolution is currently stable at 100x75 pixels.

[14]The first attempts to keep the montage and activation in synchronization was by increasing the speed of popping the montage FIFOs in order to match the speed of the activation. Sensor impressions were difficult to read at this speed and the two systems would still go out of synchronization by the end of each day.

and off at a certain time.[15]

- Add a mode to turn on an "auto" capture of all creative associations.[16]
- Add a log mode to save all images captured and the value of the important variables.[17]
- A possible computer vision criteria for an image could be the range of pixel values. The system could then ignore images with very little variation—for example images of the grey sky or out of focus.
- How to make a connection between two nodes strengthen the more times it is triggered (long-term potentiation)? How could this be done so that the potentiation would drop down to normal level over a longer period (minutes to days)?

April 20th 2007

The buffer that stores the cascade of activation grows like crazy (see rev_41_activation_buffer). I've decreased the rate of the activation (300ms+50-150ms). I've also increased the rate at which the content gets read from the buffer (750ms).[18]

Use this command to plot the buffer size over time:[19]

```
cat soos.log | grep activation-buffer-size | cut -f 2 -d "␣" > data-
    analysis/rev_43_activation_buffer_size.log
```

The new activation system is not working well. The camera is static for too long. Maybe attaching the activation to the camera is not a good idea.[20]

April 21st 2007

I changed the delay of the auto-save image capture from 50ms to 500ms, since it was not capturing much.[21] The schedule now turns off the SOM feeding. I should check on the state of the buffer size

[15]The system now uses a scheduling system to manage the systems state. A number of MAM's variables can be controlled through a local network socket accessed via the terminal program.

[16]A mechanism to capture the unit IDs composing a free-association has not been implemented.

[17]MAM now uses a complex logging system. An image of the field of experience is captured when the learning rate is one or zero.

[18]The "buffer" is the montage FIFO. (The montage only implemented one buffer at this point of development. This describes an attempt to keep the free-association and montage in synchronization.

[19]This command converts the log entries containing the size of the montage FIFO so they can be read and plotted in R (R Development Core Team, 2007).

[20]An early attempt was made to replace the camera's fixed interval with a test of the lack of activation in the free-association network. This implementation failed and I returned to the fixed interval.

[21]I was attempting to capture the field of experience and the montage during a free-association. This proved difficult due to the random delays in the propagation of free-association signals and in the montage buffer.

over time.[22] I need to have the schedule turn off the SOM feeding as well as the camera![23]

The cosine learning rate takes 42min to complete a cycle—does that make sense? Seems like it is probably too slow, at least based on the 10min memory of the system, which is a very inaccurate estimation.[24]

Don't forget to try out a zooming effect so that the current activation is centred on the screen but zoomed in so that the images are more easily seen.[25] Having the frame follow the activation pattern would be really tricky though. I suppose it could be done based on the X/Y ranges of activation of the nodes. Need to have a version that uses multiple GEM windows to do that.[26] That is a good long-term plan since it would make the placement of the windows more flexible. This would mean changing the coordinates of all graphical elements—which are positioned considering the current display layout.

May 3rd 2007

The free-association screen was solid white after the changes from last week.[27] Looking at the buffer showed that the buffers were getting filled up too fast for the system to deal with—these were added to the SVN under rev_50. Note that now that the buffer size can be printed from a socket command (debug 1)[28] the code to trim all by the activation buffer size is as follows:[29]

```
cat soos.log | grep activation-buffer-size | cut -f 3 -d " " > data-
    analysis/rev_50_activation_buffer_size.log
```

Maybe I bit off too much with the proportional delay.[30] How best to slow down the playback rate? Maybe an equation that relates the size of the buffer and the pattern of activation to the determine the delay time? Bigger buffers empty faster than smaller ones? How can I even test this—since I need to wait for some time for images to get collected into the SOM? I will try and

[22] Still dealing with constant montage synchronization problems.

[23] The schedule system is being altered so that the camera can rest overnight. While the camera is static the SOM should no longer be fed sensor impressions otherwise the field of experience fills with slight variations of the same frame.

[24] The time of the cosine cycle is based on the frequency of the equation and the camera interval. Earlier versions had much shorter intervals—as short as 4seconds. This certainly contributed to problems with the montage keeping up with free-associations. The 10min memory (based on camera interval * number of units) is totally inaccurate due to the complexity of when sensor impressions in the SOM are replaced.

[25] An alternative solution was to increase the scale of each unit, during activation, so that its associated sensor impression can be more easily read.

[26] The code to create multiple graphics windows in Gem has not yet been released.

[27] The "free-association screen" is referring to the montage display. When geometries are not textured they appear as solid white in GEM.

[28] The log of can be turned on or off via terminal.

[29] The format of the log has changed, requiring a different command to parse the data.

[30] The proportional delay is the mechanism that makes a sensor impression visible on the screen for a duration proportional to the degree of activation.

speed up the rate the buffer gets emptied to test and see what is happening with the white and the qualitative effect of the activation being proportional to the signal strength.

May 1st 2007

I'm not happy with the way the montage flashes—a smoother more continuous aesthetic would be an improvement.[31] It looks like the image in the buffer is switched while the image is not at 0% opacity. The signal to move the camera does not seem to happen when we have a poor activation—as in no propagation. It has a tendency to stare. Should the montage use cuts, fades or some combination?

The ramp up[32] and down time is 2000ms but a new set of values gets sent each 1000ms![33] That is the cause of the popping. Will try and implement a method where the buffers pop, then wait a delay before popping the next value. The delay should be proportional to the strength of activation—1000ms for up and down * max number of nodes (lets say 15?)[34] = 15000ms. Lets say activation*1000+1000ms for now.[35]

Should add a remote command to allow me to save subsequent frames for high-quality video documentation. How slow would it be to write those 1280x1024 TIFFs? Also don't forget to try the multiple-window implementation of GEM so you can test things like zooming of the memory system. I think the LCD is a little tired.[36]

May 6th 2007

The buffer certainly fills up now.[37] I should add a counter that tells me how many memories are being activated in one cascade. This would be helpful to find a quantitative method to control the rate the buffer gets emptied. I need to find a way to keep the camera from moving during the individual activation that does not lead to a cascade.[38] I'll leave this for another session.

---

[31] This "flashing" was due to the speed at which the montage was required to proceed in order to be synchronous with the activation.

[32] "Ramp" is referring to the fade in and out speed.

[33] The layers change images before the fading animation has time to complete. This causes a sudden, rather than smooth, transition between images.

[34] 15 is a blind guess as to how many memories may be activated for one free-association.

[35] This is the relationship between the degree of activation and the duration the sensor impression is visible in the montage.

[36] Saving high-quality video documentation using pix_write is impractical. Best to record video of the display itself. At this point the display is showing some ghosting from the hard-edge grid of memories in the field of experience.

[37] In working on the aesthetic aspects of the montage the synchronization has been compounded. Aesthetically speaking the montage and free-association are at odds in terms of timing. Shorter fast activations are preferred for the free-association—slow long animations are favoured by the montage.

[38] Since the camera is looking in a new direction at a fixed interval it is quite possible that a cascade of activations will last longer than the interval—leading to a montage that is a combination of both trajectories not separated by any rest space.

July 20th 2007

Since adding the dual screens (1280x1024x2) PD keeps crashing—seemingly right after it comes out of the sleep state in the morning. If this smaller 10x10 configuration does not fix it—and the memory usage looks fine during the crash—then I'll try removing the sleep stuff to see if that makes any difference.[39] I now believe that the buffer going empty very quickly over time is actually the buffer emptying after sleep. That is the buffer always fills up faster than the montage can process it, so it is only at night that things calm down, and the FIFO is emptied.[40] I should check out the numbers for the size of the cycle on Insitu[41] to see if that corresponds to the time the system is on during the day.[42] Why is the buffer stimulated at night though? I also really need to rewrite the part that propagates the signal. It is not working at all right now because of too many activations. Maybe a single thread in a random direction is a proper replacement at this stage.[43]

July 24th 2007

The dual screens really give a nice impression. The resolution of the thumbnails is perhaps a little low, but actually ends up looking pretty nice. Since I'll aim for more nodes I think this is no problem. The memory and CPU usage is very low (around 20% for both) leaving around 200MB of RAM for more nodes and future work (on this machine with only 1GB of RAM). The next version will have some changes to the node propagation so I'm looking forward to seeing that in person. The bigger images make the lower brightness of the screens less of an issue.[44] Perhaps we should have a "smooth"[45] object in the node animation, because it is certainly a little awkward.[46] I will look into how I can draw a line to show the path of propagation. I made the aesthetic decision that it would be better to have fewer, longer cascades than more frequent ones—even if that means having the camera integrate experience at a slower rate. This version is now sampling each six seconds rather than each four seconds.[47]

---

[39] The "sleep stuff" is the mechanism that puts displays into DPMS sleep.

[40] This was the initial inspiration for a dreaming state. The montage would continue once the activation in the free-association has terminated.

[41] "Insitu" is my development machine at home.

[42] At this point the log-file only showed the order of events, rather than the timing of specific events.

[43] This proposal involves removing the code to propagate a signal in a range of directions and have each unit propagate the signal only in one direction as pictured in Figure 5.12 on page 45.

[44] If it's bright enough for the camera to capture nice images it is usually too bright to see the displays. If it is dark enough for the displays then its usually too dark for the camera to capture any readable images. It is important that the camera and displays are both in the same space so that there is no power relationship between the observers and the observed. Ideally all people in the context of the work should be both observed and observing.

[45] The smooth object is a PD abstraction, written by Thomas Grill, that acts as a low-pass filter on a control signal.

[46] The activation animation of a unit is characterized as awkward due to do the linear interpolation used.

[47] The sampling rate (the rate at which an sensor impression is presented to the SOM) is synchronized to the interval between camera movements.

August 5th 2007

For the show in Montréal[48] I'll be going with a 20x20 grid and using circles rather than 4:3 rectangles for the memory map.[49] This will show the SOM structure better and be less strict than rectangles in a Euclidean grid. The activation path stuff is almost certainly not going to work for this show so I'm going to disable it.[50] I need to figure out how to make an SVN release, and an easy way to revert to a previous version, for the final version for the show.[51] I should try using more than 12 sensors per image for tests after the design layout stuff is pinned down.[52]

August 7th 2007

The reason why the cascade of activations dies off so fast is that there are too many walls formed by dead nodes that no longer accept activations.[53] I will do a test will a long inhibition delay and two directions of propagation.

August 13th 2007

In order for SOOS1[54] to be the most visible in Montréal the idea is that the schedule will be made more complex. It will have three, rather than two (sleeping and waking), states. The system will have an additional "dreaming" state where a metro feeds the activation system, but not the ANN itself.[55] It will dream (with longer activations?) on its previous experience. The third state will be a inactive state where the displays and camera sleep. The sunrise and sunset in Montréal (in Montréal time) are: 6 to 20. I suppose the sleeping state could then run from 20 to 24.[56]

---

[48] Pure-Data Convention 2007

[49] The "memory map" is the field of experience.

[50] The idea of drawing a line to show the progression of the free-association has not been revisited.

[51] I was not yet familiar with SVN tagging and branching.

[52] At this point the SOM is still being fed a 4x3 pixel down-sampling of sensor impressions.

[53] The inhibition for this revision was too extreme—causing short free-associations.

[54] The abbreviation used for MAM's original title—"Self-Other Organizing Structure #1"

[55] The purpose of the dreaming state is that the work can still be shown after sunset—when the context is too dark for the camera. In this mode the SOM is bypassed and only the free-association network is active.

[56] The times are specified in the way they are in the schedule system: waking from 6am to 8pm, dreaming from 8pm to 12am, and sleeping outside that range.

August 16th 2007

The system now seems to be working very well. I hope it will also work as well in the Montréal installation (installing on the 20th and 21st of August). I've checked in the log-file for this revision (109) to take a look at buffer sizes over time. Looks like the auto-save is pretty useless, as the montage seems to be black on every single capture. This is not a huge issue, as the montage as a still is largely useless. The memory map still does get captured, but seemingly without visible activation. I'll have to take some manual xwd[57] dumps over network in installation.

When the system goes to sleep it would be nice if, rather than just going black, the displays slowly fade to black before the DPMS command is executed. This more closely relates to the transition between conscious and unconscious states.[58] Another idea is to make the camera blurry when its sleeping since we have no control of it starring off into space while dreaming.[59] Maybe shutting down the camera iris to be really dark could also work. For the data logging I should add a field for the time of each message so the time between messages can be determined, rather than just the order in which they happen.

August 20th 2007

Installation in Montréal. Time to set up the view of the camera. The new ranges for pan/tilt are:[60]

- Pan: 3200 to 14100
- Tilt: 1200 to 3200
- Zoom: 241 to 2481

September 10th 2007

I was finally able to create a high-resolution video (huffyuv AVI) from the captured frames. I will be able to make the final check-in of the branch used for the PDCON[61] installation. Things to do for the next version are making a more in depth logging system that includes the learning and

[57] "xwd" is a *NIX utility for dumping the X display to a file—a screen-grab.

[58] This is based on my own qualitative view of my own transition between sleeping and waking states.

[59] In the dreaming state the camera simply remains looking in the direction in which it was last looking during the waking state. The scan-converter used to convert the NTSC video from the camera to a VGA signal for the display does not have any power management features. Even in the sleeping state, where the system is inactive, the display showing the live camera remains on. In the long term a solution needs to be found for this inefficiency—for example display hardware that can be put into into DPMS sleep, remotely, through serial or Ethernet.

[60] The ranges of these values are specified in the camera's units. They specify what range in which the camera should look, and zoom into.

[61] Pure Data Conference

neighbourhood rates—to see the effect of the learning rates on the BMU.[62] This data could be an interesting source for a visualization of the internal process. SOOS will be installed in the fishbowl[63] this semester. Hopefully it can be a real installation and will work as both a platform for the development of the system and installation aspects.[64] In addition to a time tag for each event the log should include the number and file name of the auto captured frames—to examine the data in relationship to the captured images.

The PDCON07 installation was running for 3 days before it crashed due to the disk filling up. The system managed to capture 16000 frames in that time. The video is not quite what I expected. The flows of the learning function, visible in the test patches, is not apparent.[65] This could be due to the low resolution (20x20 nodes) or perhaps the diversity of visual information much more complex than solid colours. It is also clear that the total memory of the system is actually quite long, as certain images do stay in the buffer for extended periods of time. Perhaps the nodes are being overwritten by images so similar that the eye does not see the difference.[66] I think the latter is unlikely due to the positioning of people and so on. Seems some memories could actually stay in the system for over one day. The new visual design of the memory nodes[67] is really aesthetically interesting, very dynamic and complex, maybe even too dense. I would like to project the memory screen at a very high-resolution to show the depth of the content collected. It seems without the visual scaffold of the Euclidean grid, the eye looks only at the contents of the frames. Each location creates a unique composition of images whose quality tends not to change even with the changing light patterns of the environment.

September 14th 2007

While reading Scharmer (2000) (for Aaron Levisohn's directed reading class) I started to think about the experience of the machine itself. I had a few ideas. One was that it is great that Varela and I agree on this idea of the self existing on multiple levels. I wonder if he would agree with the idea

[62]The concern here is about the relation between the learning rate and progression of the BMUs. It takes some time for all the units in the SOM to be associated with some sensor impressions—which can be problematic for some installations.

[63]The "Fishbowl" is the demo space at Simon Fraser University, Surrey in room 3840. The room has windows looking down onto the "Grand Hall", which is a large space with lots of traffic.

[64]It was difficult to find installation locations, that have enough flexibility to experiment with installation details, and include a view of a complex and dynamic context. My time in 3840 was extremely valuable in this respect.

[65]The test patches are the Pure Data patches that were used to initially explore the SOM external. They had many more units than the current installation, but very few sensors—only three for RGB components. These patches showed a highly complex organic growth of clusters in the data. This is probably not visible in the installation due to a much more folded feature-map.

[66]This observation is due to the fact that I was looking at time-lapse videos of the memory system visualization over time. This was the first time I got a sense of how long sensor impressions stay in the system.

[67]The "new" visual design is the replacement of hard-edge rectangles for larger circles with Gaussian alpha channels.

of the mind being present in all parts of the flesh,[68] but to varying degrees. In terms of the dream system, the installation should decide when it's time to dream. Not really *decide* but have the dream be a result of its process.[69] In order to stimulate a dream the first idea, to come to mind, would be the lack of external stimulation. When the world outside becomes too familiar, too similar to the world in its memory, it could go to sleep. I have no idea what the mechanism would be that would allow the machine to wake up again. What causes a human to wake up? His or her internal cycle? What drives this internal cycle? Does a person wake up from external stimulation only? That seems unlikely—for example when a nightmare wakes someone up. An intense stimulation can break the cycle—internal or external. What is it that drives a dreaming/sleeping cycle?[70]

September 19th 2007

Starting to get the system working again on sr-00150[71] (the SFU Surrey Research installation machine). Things are going OK, but twice now PD refused to open the patch. I've done two fresh checkouts to test this. I'm right now determining the pan/tilt zoom ranges for the camera, but I'm not yet 100% sure of the camera placement due to the amount of space (and lighting considerations) of the other pieces in the room.[72] So I've just talked to the organizers of the publicity event and it seems that the plasma screens that were hopefully going to be rented breaks the budget. I'm back to (probably at best) these 191Ts.[73] I hope the scan-converter that was ordered will actually show up in time for the installation. Ranges (based on the current camera position, with the tripod leaning on the window ledge):

- Pan: 2200 to 16600
- Tilt: 0 to 3700
- Zoom: 241 to 2000

[68] Here "flesh" is being used in reference to Maurice Merleau-Ponty's conception.

[69] The connectionist approach used in this project mean that there is no reasoning or deliberation. The word "decision" is perhaps an inappropriate term. The important point is that, in the long term, all variables in the system should be tied to the systems embodied process and not be random nor arbitrary.

[70] This is the initiation of the ideas behind the dreaming machine project.

[71] The installation in Montréal was run on my own shuttle PC, which was able to easily run a 20x20 unit SOM.

[72] This is the first effort towards installing the project in the "fishbowl".

[73] Samsung 191T 19" LCD Displays.

September 24th 2007

For the fishbowl installation you did not actually put the ranges determined into the system![74] Now fixed in the next commit. The suspend code used for Micro[75] and Insitu does not work on sr-00150.[76] Post-fishbowl to-do:

- Better debugging for R (R Development Core Team, 2007) analysis: Date and Time for all events and all events routed through DEBUG, where the date and time is added for each event.[77]
- Add learning and neighbourhood rates to DEBUG.
- Look through all previous logs and reflect on those efforts. Look for new ideas not yet implemented (like the visualization of the activation path).
- Add a counter so that the auto-saved image files can be located in DEBUG logs.

October 31st 2007

New ranges (based on discussion with Maia Engeli) for camera in "fishbowl" space:

- Pan: 1600-7800
- Tilt: 0-4400
- Zoom: 2000-2489

Seizure data in people, according to Steven Barnes, is periodic![78] Those strange spikes seen in the last few iterations could very well be analogous to seizure activity in human brains. To verify I will try running the system continuously (in wake state) for 24hours a day to see if those spikes still happen. Then also try running the system for 24 hours in the sleeping state to see what happens.

---

[74] The system was running for some time using the ranges from Montréal! Luckily those ranges actually worked quite well for the 3400 installation.

[75] "Micro" is the shuttle PC used to run the installation of the project for the Pure Data Convention 2007.

[76] Even though both machines are AMD64s with Nvidia graphics cards and the same Debian Etch OS, they respond differently to the "xset dpms force suspend" command.

[77] This was an overhaul to the logging system to make it more effective.

[78] Steven Barns has been a great influence and inspiration to the development of this project. We have often discussed potential correlations between the character of "Memory Association Machine" and the human mind. In particular the possibility of the system having a seizure—characterized by an over-activation of neurons.

November 7th 2007

After collecting data (now readable using R in the new logging format) for almost 7 days of only waking state it seems clear that the spikes seen in previous revisions is indeed an artifact of, most likely, sleeping. I would then need to test the dreaming state to see if that is happening, or if I am right to suspect that the spikes are simply an artifact of the transition from one state to another (waking to sleeping?). The new format will tell us for the next run.

I talked with Greg Corness on Monday about the timing. It's hard to check correlation between the live camera image and the free-association because there is little gap between the states with the 10s fixed interval. Perhaps a way to remove the fixed interval for capturing images would be to have each unit check if it has been activated in the last say, 3s, if no units have had any activation then trigger a new camera position.[79]

I should also add the ability to change the schedule remotely. If the system was to run in a dreaming only state it will not have collected any images! Either place-holders need to be used or the system has to be run, for a period, in waking state and then the schedule changed to force it to sleep for the rest of the run.

November 8th 2007

Changed the image sub-sampling so that the SOM units have 480,000 sensors. This seems to work in terms of CPU usage. It is difficult to tell when oprofile[80] seems to be broken. There is a long pause[81] when the image is converted to a list. Since this happens only each 12s, and not much else is really happening, it makes sense that we delay the activation of the network (from the SOM) until after the CPU lag has passed.[82] The use of projectors rather than LCD screens is working VERY well. The room has a great design to show off the project with two, maybe three, big projections. Must see how this will work in the future. People actually stop. The projector is very bright even with the windows open enough for the camera. The screen-saver went off once at the start—disconcerting. I will want to use bigger screens for the installation at the Surrey Art Gallery. I need to call about the feasibility, and brightness, of available projectors. I am looking forward to seeing more than just a bunch of numbers on the image![83]

---

[79] This has not been implemented.

[80] A Linux profiling tool.

[81] The long pause is a CPU spike that halts rendering, leading to a static display for that period.

[82] The premise here is that since the visualization of the memory system is static, when its not activated, a delay could hide the lag by ensuring the image is converted to a list before the network is activated. In practise the appropriate delay has not been determined and there remains a slight pause in rendering during this operation.

[83] During testing the system loads an image of a unique number for each unit into each memory location. That way it is easier to determine the system is working. Without the numbers the system starts off with a entirely black screen that is difficult to differentiate from a screen blanker. I was able to borrow projectors and had the space to myself for exploring installation options. Due to the short time I had the projectors, the system did not have time to

Rather than RF you could use parameters of the images themselves to give the nodes individuality rather than randomness.[84]

November 14th 2007

I saw the space at the Surrey Art Gallery today. It looks like an installation with the camera in the outer hall, and the screen within the theatre, could work. We'll see how it goes. MAM crashed today after 5.648275 days. I could not see any reason for this in the data that was logged. Adding log fields for RAM, CPU, and perhaps disk usage readings, could help with issues like this in the future. I wonder how to get that kind of data into PD. I suspect the cause of the crash was having the kernel updated, which overwrites the grub menu configuration file, which means that "noapm acpi=off" was missing! I know that causes instability problems with the Nvidia driver. So now those kernel flags are added again. Lets see how long it lasts this time.

November 16th 2007

Rather than having the zoom range fixed, giving only closeup or wide views, how about attaching the zoom range to the learning rate—or some other internal rhythm?[85]

November 19th 2007

In order to measure how long the memory is the system should generate events for each time a unit, in the network, is overwritten with a new image from context. This is actually contained in the BMU values—since the winning unit will always save a new image. So the question of analysis would be how to measure the mean distance between repeating BMUs. I currently have no idea how to do this in R.

November 20th 2007

A "video" version of MAM could store each subsequent image, associated with a particular category, rather than the newest image replacing older images. These images could then be played in a loop,

---

associate each unit in the network. The image of the animated numbers projected on the wall of the room was highly compelling.

[84] Another approach to giving the unit individual variation was to use some analysis of the sensor impression each one is associated with. This avenue has not yet been explored.

[85] The quality of the memory system visualization would then move through cycles of differing quality, based on the effect of integrating images with different zoom levels. Since the time it takes for the whole memory system to be updated is quite long, it would often be a mix of different zoom levels. This idea has not yet been explored.

which over time would show the variation of images within one category. This would certainly change the "random" perception of the SOM's organization.[86]

November 26th 2007

MAM will be moving out of 3400 by the end of this week. This afternoon it is being installed at the Surrey Art Gallery for the eMixer event. Unfortunately the last test (with schedule disconnected—not yet verified) lasted just under 2 days.[87] One good thing is that this time it ends hard without any sputtering as it did with revision 157B.[88] The quality of the failure is much more like revision 157, which was (I think) caused by the acpi flags being missing from the kernel boot. Did I do a system update since then?[89] Did I reboot the machine? So it turns out that the kernel on sr-00149[90] is newer (compiled on May 9th), than the one on sr-00150 (compiled on Mar 26th). I'll try and upgrade the kernel to the same one—to see what happens. I realized that the test to see if it was the known kernel bug, that X crashes, not PD, was not tried because I manually killed X. I did not try and kill PD first to see if X would come back.[91] So at this point I unable to know if it was PD or Nvidia that crashed this time. It indeed was not the suspend operation.[92] From this point on the revision names used in data files are no longer based on the revision of soos-parent—since now it is too modularized and does not change much. Now revision names will always refer to the revision of the trunk/ directory![93] SAG ranges:

- Pan: 2700-15500
- Tilt: 0-1300
- Zoom: 0-2489

[86]This is referring to the field of experience as being perceived as random.

[87]The system crashed after two days.

[88]Some revisions contain letters to reference the run. In some cases a crash will happen and the software is run again without making any changes. The new runs, named after the svn revisions, get alphabetic suffixes to differentiate them.

[89]As the menu.1st file is edited by hand to include "acpi=off noapm", once the kernel is updated those flags are removed and have to be manually added.

[90]This is the Surrey Research Machine used for development and testing, not installation.

[91]Initially the crashes caused by missing "acpi=off noapm" flags to the kernel was due to an interaction between the Nvidia proprietary driver and the kernel. After rendering OpenGL for some time X would crash. The test to see if it was PD or X, that has crashed, is to terminate the PD processes and see if X would return. If X has crashed then it would not recover after killing PD.

[92]The code used to put displays into DPMS sleep (suspend) was considered a likely cause of some PD crashes at this time. The suspend code runs commands using the [shell] external.

[93]Soos-parent.pd originally contained all the code for the project, but has become increasingly modularized through the development of the project. Since SVN tracks the changes of the files it was the children of the patch, not soos-parent.pd, that changed most often. All revision number references before this point refer to the revision of the soos-parent.pd file, rather than the directory containing all the project files.

After some more testing it seems that the crash *is* related to the kernel interacting with the Nvidia driver, rather than PD. The last time the no-vmlinus showed up at 97% CPU, and X did not recover.[94] Try removing shell,[95] try again, and hope this kernel lasts long enough. After the SAG[96] installation, try going back to the CK kernel.[97]

November 29th 2007

I need to get ready for the demo for Bing Thom[98] on the 10th of December. I need the installation equipment as well as poster/signage for consent,[99] and to describe the work. Based on the kviz[100] reading, the relationship between the input space (the position of the images in super high dimensional space), the position of the units for categories in the feature map, and even the position of the points in pan, tilt, and zoom space, is a very interesting area for potential visualizations.[101] Perhaps the start of a shape that defines the relationship between these three spaces could the impetus for 3D form in response to context.

---

[94]"no-vmlinus" here is referring to the oprofile flag to not profile kernel functions. The result is that the profiling output shows the CPU usage of all the kernel processes under one heading "no-vmlinus". When the crash was due to an interaction between the kernel and the Nvidia drivers, oprofile showed the kernel using a very high CPU usage while X was frozen. Experience has shown that a very high kernel CPU usage is likely to correspond to a freeze due to kernel and Nvidia driver interaction.

[95]The "shell" PD external is still considered a possible cause of instability.

[96]Surrey Art Gallery

[97]The installation machine (sr-00150) was initially running a performance tuned kernel using the CK patches. At this point the stock Debian Etch kernel replaced it.

[98]This was a demo of the project for Bing Thom. The architecture firm that designed the SFU Surrey Campus.

[99]I was hoping to demo a proper installation. The "consent" signage is simply signage that communicated the fact that images of people would be captured. The project had recently received ethics approval.

[100]IAT814 Knowledge Visualization

[101]The essence of the MAM mechanism could be considered a mapping from multiple spaces to others. The combination and relationship between these spaces are highlighted here as a potential exploration where the system generates a sculptural form, rather than animation, that is the result of its embodied negotiation. When I first started at SFU Surrey I wanted this research project to have a dynamic physical sculptural form. This still remains an interesting direction of future work.

December 3rd 2007

When reading for my review of SOM approaches,[102] I was reminded of the issues of activations wrapping around the edges of the SOM.[103] This has the potential to greatly increase the abstraction[104] of the free-associations. Should give it a try to fix for thesis.[105]

December 5th 2007

When I arrived today to take a look at the system the camera was stuck in the 0,0 position[106] and stayed there for over 12 seconds.[107] The SOM proved that the camera had been starring in the corner of the room for some time.[108] This is most likely a bug in the security ignoring code.[109] The good news is that the system stayed stable while it was running the last week (with [shell] removed).[110] The failure seemed not power related as the computer was still running. I did not need to do an address reset[111] on the camera (which needs to be done when the camera loses power). Will need to add events[112] to log the pan, tilt, zoom position so I can have some idea of what is happening.[113]

December 6th 2007

The installation for Bing Thom will now be in 3950 and 3956. If I change the system so that not one image, but a series of images, are stored by each unit in the network, then the free-association can

[102] For IAT814 I wrote a paper that reviews approaches to the visualization of clusters in Self-Organizing Maps.

[103] A bug in the free-association network allows signals to propagate from the left edge of the network to the right edge of the network, and from the top to the bottom. As the SOM used in the project is not a toroid, but a simple Euclidean plane, the left and right edges are not related in terms of topology and should be hard borders. The fact that the activation can pass over these borders makes the resulting free-association contain a break of topology, where subsequent sensor impressions may bear little to no relation relative to the SOM organization.

[104] For some time the trajectory of the free-association was characterized as a continuum from literal (where the stimulus and result are similar) to abstract (where the stimulus and result are dissimilar).

[105] At the time of writing the cause of this bug has not yet been discovered. A toroidal SOM may be more appropriate for MAM.

[106] Referring to the lowest position index for both pan and tilt, which corresponds to the lower left extreme of the camera's range.

[107] The camera should look in a new direction each twelve seconds if it is functioning properly.

[108] If the camera fails to move but the SOM continues to function, then all the units in the memory system end up being associated with the same image, except for minor differences due to camera noise and lighting changes.

[109] This is the code that restricts the camera from looking at the security desk.

[110] At this point the system appears stable with the shell external removed. The consequence is that the schedule is unable to put the displays to sleep.

[111] For an unknown reason the SFU Elmo PTC-100 cameras do not respond to pan, tilt, zoom serial commands on start-up. A serial command resets the address of the camera (used when there are multiple cameras chained on the same serial port). Only then does the camera respond to serial pan, tilt and zoom commands.

[112] Here "events" are referring to events that are logged while the system is running.

[113] This particular issue has never occurred again—even after longer installations.

be driven directly by the memory units and not through the independent network,[114] which is not based on node state.[115] That way the free-association will not require a buffer and could happen at the rate decided by the memory unit itself. For example the activation would only propagate once all the images in a particular memory unit have been presented. Then more activation patterns could be explored, as the free-association would respond accordingly in all situations.[116]

December 8th 2007

Watching the flow of pan, tilt, zoom random numbers during the latest installation made me wonder about the quality of the randomness. In particular I suspect that trandom[117] has some issues. Added an event for pan/tilt/zoom but capitalized the first letter, which is not the same as other event messages, that should be fixed.[118] I need to do a more in depth analysis of the randomness in MAM. I should be able to do the chaos style time-plotting by using seq() in R to skip odd values, and use those as the Y axis in a plot.[119] This way any periodicity should be visible as multiple clusters.

December 18th 2007

The ptz-ignore[120] system is not working properly. Seems a range of tilt values OR a range of pan values were constrained—meaning a cross shaped area was ignored. This explains why the range seemed so constrained. A fix would be to ignore a set of values only if the pan / tilt values are within both the pan AND tilt range—not OR.[121] I'll have to look at the code again to see how to fix it.

The distributions of the [trandom] and [random] values is certainly uniform. I don't see any difference between the two distributions. I have been unable to find any periodic structure in the

---

[114] The "independent network" is the free-association network.

[115] The state of activation of a unit in the free-association network is independent of the state of the montage system.

[116] This text describes two ideas. One is to solve the problem of synchronization between the montage and free-association network. This could be solved by having the montage not as a separate system, but have the montage happen inside of each free-association unit. Each free-association unit would then contain an OpenGL entity to show both the unit activation, and the montage presentation of the image. This is a solution because the montage would then run at the rate of the activation of nodes in the network. This added flexibility of having the montage driven directly by the free-association network would allow much more flexibility in the exploration of patterns of activation (the rules that define how a unit propagates a signal). The montage animation would take care of itself. The second idea is that the units store more than one image. This is the "video" version of the project, and is proposed for "Dreaming Machine"

[117] "trandom" is a Pure-Data abstraction that provides a time-seeded random number generator.

[118] As the events are spread all over the system, its hard to keep their field names consistent.

[119] Here I am referring to the phase-plot method of showing time-structure in a sequence of values. The resulting plots have not yet shown any periodicity during the existing installations. The description in this entry does not accurately describe the phase-plot method.

[120] The name of the abstraction that keeps the camera from looking in a certain range of directions.

[121] Here "AND" and "OR" are referring to the logical operators && and || .

data. Greg Corness said something about the possibility that multiple [random] objects sometimes produce the same values under high CPU load.

Jan 2nd 2007

In addition to, or perhaps in place of, the current fixed amount a free-association is decayed, as it propagates from one unit to another, the distance between the code-book vectors of the source and destination units could be used to determine the amount the signal is decayed. This way activations would be longer in highly correlated clusters, and activations would be less likely to cross boundaries.[122] Since more abstract activations would require more energy[123] there should be some way to either increase the activation or inject more energy in the associative network.[124]

Today I will take down MAM from 3400. It should contain a good run of data for analysis, including testing the pan/tilt ignore system, and the quality of the learning due to the capture of frames with learning rates near 1 and 0.[125]

After a brief look at the new captures (at 0 and 1 in the learning cycle) its hard to tell which is "unorganized" as after a point they both look highly organized. I should do a test capturing a bunch of images and writing them directly to disk, and using those to test learning functions. Images are never captured in the sleeping state, was that intentional?[126] The log confirms that images are captured at the right points in the learning cycle.

Despite my fears that the initial random variables used in the ann_som object are not time seeded,[127] the BMUs are not identical for different runs, which means either they are time-seeded, or that the self-organizing process causes the tends to individualize very quickly.[128] The BMUs are different even near the start of the process for all runs. I'm not sure if my phase-plane function is plotting correctly. It seems that most of R's internal RNG generators give a very similar phase structure as mine, a hexagonal shape. Most data from SOOS plots show this same structure, except for activations,[129] which gives a very different shape.

---

[122] This is where the idea of using code-book distance for signal decay is initially considered.

[123] Only larger signals would have enough magnitude to survive the large decay at a cluster boundary.

[124] The "associative network" refers to the free-association network.

[125] The system captures a frame of the memory visualization when the learning rate reaches either 0 or 1. This is done to reduce the number of frames captured and compare how organized the feature-map appears after each learning cycle.

[126] As the SOM is not training in sleeping nor dreaming states, the feature-map should not change while the system is in these states. It is unclear why I thought it was unintentional at that moment.

[127] I had assumed that the "init" method for the ann_som external would randomize the code-books. Turns out that "init" without arguments has no effect on code-books. The ann_som external contains no internal function for generating random code-books. It depends on the patch to provide the random values for each unit.

[128] As the code-books are not initially random the different BMUs are due to the SOM training on different sensor data.

[129] Here, "activations" refers to the collection of which units get activated.

I need to figure out a way of exploring the BMU data[130] so as to see the length of time between repeat occurrences of a value. That way I can see how long memories are stored and get an impression of the range of how long memories are kept.

Things I need to check: if the ignore system is causing multiple BMU's to be sent out (one for the attempt and one for the corrected position).[131] I have seen what looks like multiple independent activations in different parts of the map. Also the pan/tilt values recorded by the log for revision 217 are *wrong* because the log was created from the direct ptz-random output and not the ptz-ignore output![132]

Jan 14th 2008

While working on the chapter for the book[133] I was looking again at the memory length.[134] I figured out a way of measuring the distance, in time, between the repetitions of a particular BMU. The results are interesting (some in the SVN log) but a look at the histograms really make me wonder. The histogram (and phase plots) show the distribution of the BMUs is very complex. At first it seemed as though "0" was the most common time lag between repeated BMUs, meaning 12s of memory. A closer inspection shows the data to be even more complex, and there appears to be an auto-correlation. The peak near 0 is actually near 12s, and the auto-correlation is happening at 12s intervals, which is the granularity of how often the BMUs get sent.[135]

To make matters more complex I realized that the BMU name is also used as the variable that drives the dreaming state. In that case the BMU does not overwrite the images in memory, but just activates them. I need to change this to get any good indication of BMUs.[136] I'll use the dream state in the log to compare the distributions of the sleeping BMUs to the waking (actual) ones. I really need to change the dream "BMU" to a different name, and log it differently, to make things clearer in the future.

---

[130] Which BMUs are chosen by the SOM over time.

[131] The ptz-ignore system works by trapping sets of values that do not fit the criteria, and then regenerating a new set of random values.

[132] The log showed the pan/tilt/zoom value attempts, not the results that fit the criteria.

[133] "The Handbook of Research on Computational Arts and Creative Informatics" (Bogart, 2009)

[134] The duration that memories are kept in the system.

[135] In retrospect this structure makes sense. There should not be any memory of 0 seconds, since that would mean no image was stored. Images can only be stored each 12s at the fastest. Additionally the auto-correlation is obviously due to the fact that memory durations can only occur in 12s increments.

[136] A good indication of the memory duration is not effected by the dreaming state.

April 1st 2008

A way of increasing the diversity of visual data, in the SOM, would be to have the range of the camera pan/tilt, and perhaps zoom, change over time. At any one moment in time the SOM would contain a subset of the images from its context. A mechanism moves that subset around the superset. The trajectory of this range would need to be determined. One idea would be to use some statistical abstraction of the free-association. Or perhaps even an additional network of nodes that chooses the trajectory of the subset based on the activation of previous memories. Does mapping the space of memories onto the space of PTZ space make any sense?[137] This approach needs more thought but could very well be a method for driving the camera movement from the free-association itself. How this would effect the timing (make it pointless?) is yet unknown.

[137]The idea here would be that the position of SOM units in output space would be superimposed on the PTZ space. A free-association trajectory would also be a pan/tilt trajectory for the camera. It is unclear what the result would be. This particular mapping seems arbitrary. The SOM space and PTZ space have no relation.

Appendix C

Log of Practise

This is the raw subversion log used through the development of MAM and this thesis. The list of changed files for each entry has not been included. Notes for "SOMLab" are included: A set of patches and R programs for SOM visualization and the analysis of training methods. The content is edited for typing errors and spelling—not grammar.

---

r1 | bbogart | 2007-03-30 17:28:12 -0700 (Fri, 30 Mar 2007) | 2 lines

Initial import of SOOS files

---

r2 | bbogart | 2007-03-30 17:42:14 -0700 (Fri, 30 Mar 2007) | 1 line

cleanup of initial check-in

---

r3 | bbogart | 2007-03-30 17:43:12 -0700 (Fri, 30 Mar 2007) | 1 line

cleanup of initial check-in -- again

---

r4 | bbogart | 2007-03-30 18:17:30 -0700 (Fri, 30 Mar 2007) | 1 line

node-som is now working with simple on/off activation. Whole system activates, but no stack overflows nor huge CPU spikes. Makes use of an arbitrary 100ms closed spigot.

---

r5 | bbogart | 2007-03-30 19:12:54 -0700 (Fri, 30 Mar 2007) | 1 line

non-cascading activation with a fall-off corresponding to unit distance

---

r6 | bbogart | 2007-04-02 19:34:11 -0700 (Mon, 02 Apr 2007) | 1 line

started again from scratch. This time I have a working propagation of a signal, but only to a specific node (based on an equation) but this does propagate through the network, just in the order of the instance numbers.

---

r7 | bbogart | 2007-04-04 15:07:02 -0700 (Wed, 04 Apr 2007) | 5 lines

Added "limit" abstraction
the activation is working in three directions now, any directions should work, but may cause feedback. Will try more and see if signal inhibition is needed.

---

r8 | bbogart | 2007-04-04 15:19:11 -0700 (Wed, 04 Apr 2007) | 2 lines

Four direction propagation, no inhibition

---

r9 | bbogart | 2007-04-04 17:18:37 -0700 (Wed, 04 Apr 2007) | 7 lines

Working propagation w/ inhibition. The system behaves a lot like a cellular automata though, which does not work quite that well. I've also added a random activation/deactivation of the initial signal, which is biased towards the reduction of signals. numbers between 60 and 70 seems to work well for the split (60%-70% of the random numbers decrease the signal strength.

---

r10 | bbogart | 2007-04-05 11:34:26 -0700 (Thu, 05 Apr 2007) | 3 lines

changed the camera range to see about tilt problem.
added activation amount to brightness of free-association

---

r11 | bbogart | 2007-04-07 15:09:52 -0700 (Sat, 07 Apr 2007) | 6 lines

Added a complex version of inhibition based on the re-enforcement of connections in the network. Problem is, it seems the distribution of messages from a particular direction is uniform, so the system evenly inhibits which does not create a useful system. Will add randomness to the choice of direction of propagation, not just the value propagated.

---

r12 | bbogart | 2007-04-07 16:44:35 -0700 (Sat, 07 Apr 2007) | 6 lines

Random directions, sometimes no direction it seems. Also disconnected both normal "global" inhibition and the per-connection inhibition

(the latter was always uniformly distributed). Disconnected the individual propagation factor for each connection also. Why does the signal die down when there is no signal degradation?

---

r13 | bbogart | 2007-04-09 15:56:06 -0700 (Mon, 09 Apr 2007) | 5 lines

ramped down the speed of activation to see if that helps with the CPU spikes, activation behaviour is quite nice with this set of variables. Will see if the camera still tilt upwards when we always send the same tilt level command.

---

r14 | bbogart | 2007-04-10 13:35:24 -0700 (Tue, 10 Apr 2007) | 8 lines

Added masks to the videos to give a less harsh edge.
Slowed down the rate that the activation gets spit out. which gives a much more fluid cinematic feel, it is really starting to look beautiful. Added scaling so that the whole memory is visible, but the activated ones get more opaque and increase in size. Works quite well. I like the aesthetic of this version.

---

r15 | bbogart | 2007-04-11 13:16:02 -0700 (Wed, 11 Apr 2007) | 6 lines

Changed the rand-delay so that it is not possible to have the upper and lower limit be the same number. cleaned up the parent patch. Made installation ready but have not tested it, increased the size of the delay between movements. There is a slight grey border around the free association I don't like.

---

r16 | bbogart | 2007-04-11 14:55:52 -0700 (Wed, 11 Apr 2007) | 2 lines

Unsure what I changed here...

---

r17 | bbogart | 2007-04-11 15:01:38 -0700 (Wed, 11 Apr 2007) | 2 lines

Slight changes, cleanup

---

r18 | bbogart | 2007-04-11 15:02:18 -0700 (Wed, 11 Apr 2007) | 2 lines

more slight changes...

---

r19 | bbogart | 2007-04-11 15:04:42 -0700 (Wed, 11 Apr 2007) | 2 lines

Check in time-randomizer

---

r20 | bbogart | 2007-04-11 15:06:03 -0700 (Wed, 11 Apr 2007) | 3 lines

Almost done rearranging the directory structure, at least this is possible with svn!!!

---

r21 | bbogart | 2007-04-11 15:06:45 -0700 (Wed, 11 Apr 2007) | 2 lines

Moved Image-Processing inside soos!

---

r22 | bbogart | 2007-04-11 15:08:23 -0700 (Wed, 11 Apr 2007) | 2 lines

Removed fractal image file, why did I check that in anyhow?

---

r23 | bbogart | 2007-04-11 15:10:33 -0700 (Wed, 11 Apr 2007) | 2 lines

Added a script to start the installation

---

r24 | bbogart | 2007-04-11 15:21:44 -0700 (Wed, 11 Apr 2007) | 2 lines

whoops, messed up the syntax for the script

---

r25 | bbogart | 2007-04-11 16:15:56 -0700 (Wed, 11 Apr 2007) | 2 lines

Added the switch to turn off screen-saver in blackbox

---

r26 | bbogart | 2007-04-11 16:55:50 -0700 (Wed, 11 Apr 2007) | 4 lines

The camera does not move on its own, but now is triggered by the signal from the memory signal. This is so that the live video image is settled during the time that the free association happens.

---

r27 | bbogart | 2007-04-11 17:01:49 -0700 (Wed, 11 Apr 2007) | 4 lines

The patch does not seem to be integrating each new image into the ANN, is

```
this because the ANN is not providing a new BMU? Or is the node simply
not updating?

------------------------------------------------------------------------

r28 | bbogart | 2007-04-11 17:07:36 -0700 (Wed, 11 Apr 2007) | 3 lines

the node buffer update was going to a cold inlet, why??? changed now.
Added a "debug" message to turn on and off debugging.

------------------------------------------------------------------------

r29 | bbogart | 2007-04-12 10:22:52 -0700 (Thu, 12 Apr 2007) | 6 lines

Almost there! Added a flash to associate the (non-creative) activation
of a node, and attach that to a flash of the video image.
There was no visual signal that showed when the ANN was fired.
Now I'll tweak the size of the memory for the 1280x1024 frame, and add a
image-capture function.

------------------------------------------------------------------------

r30 | bbogart | 2007-04-12 10:56:15 -0700 (Thu, 12 Apr 2007) | 3 lines

Tweaked the size of the memory nodes
added a "save" function to grab a TIFF screen-shot. "grab bang"

------------------------------------------------------------------------

r31 | bbogart | 2007-04-12 11:21:41 -0700 (Thu, 12 Apr 2007) | 3 lines

Delay style inhibition was disabled! Put it back in now, hopefully
properly. Have not watched the effect of it yet.

------------------------------------------------------------------------

r32 | bbogart | 2007-04-12 13:55:56 -0700 (Thu, 12 Apr 2007) | 4 lines

Changed the range of tilting back to 100% range, added a remote-control
mode so that "panonly 1" can be sent from the terminal in order to turn
off tilting for the SFU cameras. The panonly tilt follows the horizon.

------------------------------------------------------------------------

r33 | bbogart | 2007-04-12 14:52:18 -0700 (Thu, 12 Apr 2007) | 3 lines

Added a little file for my rough notes that targeted towards a particular
revision. Perhaps this note will move to tomboy and get removed from svn.

------------------------------------------------------------------------

r34 | bbogart | 2007-04-12 17:06:54 -0700 (Thu, 12 Apr 2007) | 5 lines

increased size of nodes
```

```
added 1s delay (1s) to image-capture/snapshot
made a slightly less signal degradation (0.13)
Made the default image brightness 80%

------------------------------------------------------------------------
r35 | bbogart | 2007-04-17 11:23:35 -0700 (Tue, 17 Apr 2007) | 4 lines

Added Trunk/Branches/Tags, changed pd files to binary with:
svn propset svn:mime-type application/octet-stream


------------------------------------------------------------------------
r36 | bbogart | 2007-04-17 11:24:09 -0700 (Tue, 17 Apr 2007) | 2 lines

Removed SOM and RF as they are only there for historical reasons.

------------------------------------------------------------------------
r37 | bbogart | 2007-04-17 17:17:38 -0700 (Tue, 17 Apr 2007) | 4 lines

Changed the activation pattern so that more memories get activated in a
more interesting way. Not yet running on the soos1 machine. (changed > <
to >= <=)

------------------------------------------------------------------------
r38 | bbogart | 2007-04-17 17:19:15 -0700 (Tue, 17 Apr 2007) | 2 lines

Cleaned up version w/ proper folders

------------------------------------------------------------------------
r39 | bbogart | 2007-04-19 14:52:43 -0700 (Thu, 19 Apr 2007) | 11 lines

* Changed the rate the images of the free association get spit out to
1000ms.
* Changed the signal for the camera to move based on when the free
association is done. (+ 1s)
* now the initial activation in the network does feed into the free
association.
* capture file-name changes per sessions automatically
* autosave 1 will save each free association.
* auto on and off. 6am to 10pm
* log entry

------------------------------------------------------------------------
r40 | bbogart | 2007-04-19 16:35:58 -0700 (Thu, 19 Apr 2007) | 4 lines

Changed the delay after activation has ceased to 0s, it was seizing too
much! Also changed the suspend command, hopefully it will actually work
```

now.

---

r41 | bbogart | 2007-04-19 23:06:49 -0700 (Thu, 19 Apr 2007) | 3 lines

30,000 entries in the free-association buffer. Maybe too slow. Tried subtracting 20% from the signal for each propagation.

---

r42 | bbogart | 2007-04-20 10:58:35 -0700 (Fri, 20 Apr 2007) | 4 lines

This new directory will hold data-plots. The activation buffer is the buffer that holds the results of the activation cascade. Will investigate how fast it climbs in R.

---

r43 | bbogart | 2007-04-20 11:57:59 -0700 (Fri, 20 Apr 2007) | 4 lines

I've decreased the speed of the activation itself (300ms+50-150ms) I've also increased the rate at which the content gets read from the buffer for the free association (750ms).

---

r44 | bbogart | 2007-04-20 12:27:01 -0700 (Fri, 20 Apr 2007) | 2 lines

Buffer is still filling like crazy, did my changes make any difference?

---

r45 | bbogart | 2007-04-20 15:32:45 -0700 (Fri, 20 Apr 2007) | 7 lines

Changed patch to remove all FIFO buffers. Now the activation is slowed (a random # from 500-1000ms for each direction) and gets fed directly into the free association. Then each of the 4 layers has its own FIFO buffer, so it should fill up 4 times less quickly. The activation should not happen at a rate that the free association is able to keep up. The activation buffer log only refers to the buffer of the first layer.

---

r46 | bbogart | 2007-04-20 23:23:25 -0700 (Fri, 20 Apr 2007) | 2 lines

Notes

---

r47 | bbogart | 2007-04-21 13:20:21 -0700 (Sat, 21 Apr 2007) | 4 lines

Changed the delay on the autosave image capture, since it was not capturing much. Changed the autosave delay to 500ms from 50ms. Schedule

```
now turns off the SOM feeding. Made a log entry.

------------------------------------------------------------------------
r48 | bbogart | 2007-04-21 15:20:21 -0700 (Sat, 21 Apr 2007) | 2 lines

Added log-file to see how the activation is looking long-term.

------------------------------------------------------------------------
r49 | bbogart | 2007-04-21 15:49:54 -0700 (Sat, 21 Apr 2007) | 2 lines

Added a few comments about the cosine cycle.

------------------------------------------------------------------------
r50 | bbogart | 2007-04-21 16:22:50 -0700 (Sat, 21 Apr 2007) | 3 lines

More logging, talking about the possible "zoom" effect for the memory
system.

------------------------------------------------------------------------
r51 | bbogart | 2007-04-22 22:18:30 -0700 (Sun, 22 Apr 2007) | 4 lines

Last revision was using the %d argument from makefilename, which does
not work with makesymbol. PD crashed just after the sleep command was
executed, hopefully that typo was the cause.

------------------------------------------------------------------------
r52 | bbogart | 2007-04-22 22:19:12 -0700 (Sun, 22 Apr 2007) | 2 lines

Added another data-set to look at activation/free association over time.

------------------------------------------------------------------------
r53 | bbogart | 2007-04-22 22:23:09 -0700 (Sun, 22 Apr 2007) | 2 lines

Added a "cam" message to turn cam on and off as the schedule does.

------------------------------------------------------------------------
r54 | bbogart | 2007-04-22 22:25:03 -0700 (Sun, 22 Apr 2007) | 2 lines

Turning off camera ("cam 0") also turns off the SOM feed.

------------------------------------------------------------------------
r55 | bbogart | 2007-04-27 20:39:20 -0700 (Fri, 27 Apr 2007) | 6 lines

rerouted the activation buffer size to the debugger.
Shell seems to cause PD to watchdog after sending the suspend command.
The log should be more useful now without the activation stuff in it.
Need to figure this out for the longer term installation, maybe use
```

pyext to execute the command instead.

---

r56 | bbogart | 2007-04-30 11:51:20 -0700 (Mon, 30 Apr 2007) | 12 lines

Changed the metro that controls the FIFO readout to 1000ms, to match the delay of the line's for the fade. I wonder is 1000ms will make the buffers fill up too fast.

will log the activation in this case to see what the buffers are doing. It would be nice to have the rate the free association gets spit out be non-interval. Perhaps the speed would be proportional to the level of activation? The more activated a memory the longer it will stay on screen? Associating each activation with a time stamp and then feeding them out at the same relative time rate (but slowed down) could work also.

---

r57 | bbogart | 2007-05-01 17:30:02 -0700 (Tue, 01 May 2007) | 9 lines

Added some notes, checked the sleep schedule, and it should be working, changed to 9am to 6pm. Changed the free association so that the delay and fade times are proportional to the signal strength, slow at the start and fast at the end. This should probably be reversed so the more abstract associations (further along) get more emphasis than the early (more direct) ones. Activation*1500ms+1000ms = delay and fade (in one direction) time. Would be interesting to have the fade-out and fade-in work independently.

---

r58 | bbogart | 2007-05-01 17:37:06 -0700 (Tue, 01 May 2007) | 3 lines

Now the nodes that get the update signal also send a signal to tell the camera to move, should cut down on the starring.

---

r59 | bbogart | 2007-05-01 17:39:55 -0700 (Tue, 01 May 2007) | 3 lines

Added more notes

---

r60 | bbogart | 2007-05-03 18:15:58 -0700 (Thu, 03 May 2007) | 6 lines

Added notes
tried to get rid of the popping to no avail, perhaps the line is too slow for the activations so they pile up. Now the delay is 1000 to

2500ms inversely proportional to the activation, so initial is fast and abstractions slower and more meditative.

---

r61 | bbogart | 2007-05-03 18:17:52 -0700 (Thu, 03 May 2007) | 2 lines

made the debug squares invisible for the free association popping.

---

r62 | bbogart | 2007-05-06 11:24:13 -0700 (Sun, 06 May 2007) | 14 lines

Added global update signal (rather than using the local one that triggers individual nodes). This is also used in a counter to see how many nodes get activated in each cascade. Camera position is now tied to the 4s sampling of the rendering rate. The camera moves, and then 2s later the data is fed into the SOM. It makes more sense in the long term to have the sampling of the video happen in response to when the camera is in its new position. What effect will this have another than reducing the number of fixed internal components? Each camera settle would then be an activation in the network. That could cause problems with the generative timing of the montage, which is dependant on the rate of the SOM sampling. Or perhaps not since it would just mean that the next montage could run without the current being finished. As long as the buffers don't fill up that is acceptable.

---

r63 | bbogart | 2007-05-06 11:30:31 -0700 (Sun, 06 May 2007) | 8 lines

Crap, I forgot to add the new abstraction "free-association-layer" the last version was doing delays between 1000 and 2500ms. This version is doing 250-1250ms. This is to test if it helps with buffer filling. Using the max and average number of activations in a cascade (now being logged in debug mode) will allow me to figure out a better equation to keep the buffers from filling. Activations are larger than expected though, as many as 35 activations (and that was in a short test).

---

r64 | bbogart | 2007-05-06 12:21:14 -0700 (Sun, 06 May 2007) | 8 lines

Changed the delay function for the montage/free-association from inverse back to being normal. The hunch being that there are less strongly activated nodes than weakly activated nodes. Therefore having weakly activated nodes stay up longer means much much much longer free-associations. See if this works with the same time range as before 250-1250ms. Also we have a rough log of the activations in each cascade to look at.

---

r65 | bbogart | 2007-05-07 12:01:01 -0700 (Mon, 07 May 2007) | 6 lines

First the cascade stops, then the buffer (which fills up to being very large) empties, and then it gets to the bottom I get a bunch of PD watchdog messages. I tried removing -rt from the run script as a starting point to see if that could be the cause. Will shorten the buffer spit out time further now.

---

r66 | bbogart | 2007-05-07 12:02:54 -0700 (Mon, 07 May 2007) | 3 lines

Made signals get spit out faster. If this is not working well, will need to slow down (increase) the delay of activation in the network.

---

r67 | bbogart | 2007-05-29 10:49:34 -0700 (Tue, 29 May 2007) | 5 lines

Parsed the log-file into buffer size and per-cascade logs and checked in. Note that this rev_63 activations per cascade is actually a revision of a previous one, which was also rev63. Will try slowing down both the montage and the propagation rate in the next commit.

---

r68 | bbogart | 2007-05-29 11:43:59 -0700 (Tue, 29 May 2007) | 5 lines

Tired to make the montage and the related signal propagation slower, propagation is much slower (s500+500 to s1000+1000, montage is s500 to s1000) where "s" is the signal at that node at that time. Lets see if that helps with the buffer and make the montage nicer.

---

r69 | bbogart | 2007-05-29 11:46:08 -0700 (Tue, 29 May 2007) | 3 lines

Made the camera zoom in more always, hopefully to increase the variation of the images. From 241 to 1000 out of 2489.

---

r70 | bbogart | 2007-05-30 15:42:19 -0700 (Wed, 30 May 2007) | 4 lines

Change rand-delay from f*1000+1000 to f*1500+1000 as the buffer is still getting filled. Included the log of the number of activations in data-analysis under rev69.

---

r71 | bbogart | 2007-06-07 12:05:32 -0700 (Thu, 07 Jun 2007) | 3 lines

```
Added the following data files. These files correspond to rand-delay
rev79, soos-parent rev69. Seems the buffer does fill up, but it was
less full today than it was yesterday. Will plot in R to see what is
going on.

------------------------------------------------------------------------

r72 | bbogart | 2007-06-07 13:08:29 -0700 (Thu, 07 Jun 2007) | 5 lines

slightly more data from the same revision.
This data shows a clearly periodic behaviour, but the buffer only emptied
once it gets to about 4000 items, and that takes close to 60,000
iterations. Will try and slow down the activation further.

------------------------------------------------------------------------

r73 | bbogart | 2007-06-07 13:12:17 -0700 (Thu, 07 Jun 2007) | 3 lines

Activation delay is now rand*2000+1500. What does this cause for the
buffer pattern? Hopefully the buffer does not fill...

------------------------------------------------------------------------

r74 | bbogart | 2007-06-19 11:25:53 -0700 (Tue, 19 Jun 2007) | 3 lines

new logs for the last run, made no changes, the camera stopped working
and/or someone moved it.

------------------------------------------------------------------------

r75 | bbogart | 2007-07-16 16:01:27 -0700 (Mon, 16 Jul 2007) | 3 lines

Another look at how the buffer is behaving long-term. Still seems to be
huge, should change rand-delay to an extreme amount for the next trial.

------------------------------------------------------------------------

r76 | bbogart | 2007-07-16 22:17:32 -0700 (Mon, 16 Jul 2007) | 2 lines

Added data-file for R analysis

------------------------------------------------------------------------

r77 | bbogart | 2007-07-18 16:10:03 -0700 (Wed, 18 Jul 2007) | 14 lines

Big Changes for Montreal Installation (Pd Con)
Now there are two screens on the 14th floor, one for the memory system,
the other for the free association.

made both pan and tilt now the default since the camera has been
repaired.

dual screen config for pd-con, removed live video feed. Increased
```

```
number of nodes to 16x16

Made rand-delay up to 3000 to 6000ms. That is really long, if that does
not solve the buffer issue then I need a new solution.
```

```
r78 | bbogart | 2007-07-19 12:35:56 -0700 (Thu, 19 Jul 2007) | 21 lines

Previous revision crashed after a little while, looks like it was using
too much memory. Need to look at top carefully next time.

now 12x12, seems 16x16 was using too much ram, need more ram on these
machines.

change node-som to .44 x 0.33 units

loaded the video-mask.tif file into a buffer to save memory. Now uses
256MB less RAM while rendering) leaving 664480KB left for memory system.
7.3728 MB per image = only 90? Lets try it with 144 nodes and see how it
behaves.

Old Method (pix_image * 144):
726372 before PD starts - 506436k after = 219936k
rendering: 320528k

New Method: (pix_buffer)
721096 before PD starts - 676648 after = 44448k
rendering:      56244k
```

```
r79 | bbogart | 2007-07-20 14:11:37 -0700 (Fri, 20 Jul 2007) | 5 lines

Changed the array to 10x10 (but did not change the node size to proper
values) Just to see if its perhaps memory usage that is the problem with
the larger frame-buffer? (double size) Forgot to check the memory usage
during the crash this time...
```

```
r80 | bbogart | 2007-07-20 14:16:47 -0700 (Fri, 20 Jul 2007) | 3 lines

Some notes on the behaviour of today and insight into the reason why the
buffer empties periodically!
```

```
r81 | bbogart | 2007-07-21 09:38:05 -0700 (Sat, 21 Jul 2007) | 9 lines
```

```
Rewrote the whole signal propagation system. Now the signal only
propagates in two directions, and the allowable signals are (actually)
limited to 0-1 whereas before some other numbers were leaking in. Also
the previous version had a bunch of redundant activations I could
not explain, and this version fixes that problem. Lets see how this
now effects the buffer filling.

Now the direction, variation of delay, and variation of degradation
are time-seed random variables.

------------------------------------------------------------------------
r82 | bbogart | 2007-07-21 09:46:11 -0700 (Sat, 21 Jul 2007) | 2 lines

This version has the DPMS stuff removed for testing.

------------------------------------------------------------------------
r83 | bbogart | 2007-07-21 09:49:04 -0700 (Sat, 21 Jul 2007) | 2 lines

Back to 12x12 because that amount is needed in MANY places!!!

------------------------------------------------------------------------
r84 | bbogart | 2007-07-21 10:26:46 -0700 (Sat, 21 Jul 2007) | 3 lines

New version of constructor that tells each abstraction how many there are
(via argument)

------------------------------------------------------------------------
r85 | bbogart | 2007-07-21 10:29:18 -0700 (Sat, 21 Jul 2007) | 2 lines

Fixed bugs in argument order to make constructor work properly.

------------------------------------------------------------------------
r86 | bbogart | 2007-07-21 22:00:28 -0700 (Sat, 21 Jul 2007) | 2 lines

Missing files for the new propagation system.

------------------------------------------------------------------------
r87 | bbogart | 2007-07-21 22:08:08 -0700 (Sat, 21 Jul 2007) | 2 lines

Added patch to make sure Elmo is working.

------------------------------------------------------------------------
r88 | bbogart | 2007-07-21 22:11:16 -0700 (Sat, 21 Jul 2007) | 2 lines

Missing another file for the new free association system.

------------------------------------------------------------------------
```

```
r89 | bbogart | 2007-07-23 22:19:31 -0700 (Mon, 23 Jul 2007) | 3 lines

Added som-constructor, changed node-som to calculate size based on
constructor args. Needed some float() stuff to make expr work as expected.

------------------------------------------------------------------------
r90 | bbogart | 2007-07-24 10:57:44 -0700 (Tue, 24 Jul 2007) | 3 lines

Added logs for the behaviour of the new activation system. It is still
getting filled. Need to increase signal degradation.

------------------------------------------------------------------------
r91 | bbogart | 2007-07-24 12:13:23 -0700 (Tue, 24 Jul 2007) | 6 lines

Added some more data to the R project.
Changed the memory sub-patch into an abstraction, changed "constructor"
to som-constructor, made the som size dynamic (and set to 15x15) and
changed the 4s per capture issue into 6s per capture. Need to tackle the
damn problems with DPMS.

------------------------------------------------------------------------
r92 | bbogart | 2007-07-24 12:58:13 -0700 (Tue, 24 Jul 2007) | 7 lines

Changed the direction of propagation to 1, increased the delay to 10s
between camera movements. fixed the odd flipping problem and "invalid
enumerant" messages. 15x15 grid, decreased signal degradation. Montage
is looking nice. Constrained tilt so that the camera looks mostly down
(for the 14th floor installation). Also at 10s for each SOM feed, a
higher resolution version of the image fed into the some is possible.

------------------------------------------------------------------------
r93 | bbogart | 2007-07-29 10:39:56 -0700 (Sun, 29 Jul 2007) | 3 lines

Added smooth (control signal low-pass) to make the node activation
animation nicer.

------------------------------------------------------------------------
r94 | bbogart | 2007-07-29 11:45:53 -0700 (Sun, 29 Jul 2007) | 5 lines

Added a new line for the bttv card on micro, and added a device message
to the comport object to see the USB serial adaptor. The patch would
need some changes to work on this machine (micro) due to differing
devices.

------------------------------------------------------------------------
r95 | bbogart | 2007-07-30 13:57:54 -0700 (Mon, 30 Jul 2007) | 2 lines
```

```
Fixed calculation for dynamic sizing of nodes.

------------------------------------------------------------------------
r96 | bbogart | 2007-08-05 18:25:55 -0700 (Sun, 05 Aug 2007) | 7 lines

Lots of changes this checkout, this is all the work going into getting
the system working on the new installation machine "micro". fixing small
bugs and testing patches for the extra stuff. Next commits will be
making the circle design default and making sure the camera and serial
devices work on micro.


------------------------------------------------------------------------
r97 | bbogart | 2007-08-07 12:51:56 -0700 (Tue, 07 Aug 2007) | 6 lines

20x20 nodes should work well for the Montreal installation. I'm going to
decrease the size of the circles, and slow down the cascade of
activations as they are happening too fast. Also the cascade of
activations stops for before the signal has degraded, I'm not sure why
this is.

------------------------------------------------------------------------
r98 | bbogart | 2007-08-07 20:06:57 -0700 (Tue, 07 Aug 2007) | 3 lines

Cleaned up the remote-control system and general cleanup to integrate
the circle-node stuff into soos-parent.

------------------------------------------------------------------------
r99 | bbogart | 2007-08-08 09:28:54 -0700 (Wed, 08 Aug 2007) | 4 lines

Should be able to run the new patch in "installation" mode now, will try
a test in blackbox first to see why the CPU usage climbs as much as it
does.

------------------------------------------------------------------------
r100 | bbogart | 2007-08-10 10:10:47 -0700 (Fri, 10 Aug 2007) | 14 lines

This will be the last commit with the circle-node work, as that is now
the current version. Those files for testing that approach will be
removed on the next commit. This configuration is working very well.
Here are the plans for the next commit:

Left & right edges of full-screen are cut off, make test patch.
Try increasing the data fed into the SOM.
Time for schedule seems to have too large a delay (30min or more?)
Some activated nodes at screen border do show on free-association,
double check render priorities.
```

```
The system takes a very long time to fill up 400 nodes, hopefully this
will decrease when sending more data to the SOM. Another option could be
to increase the learning frequency.

------------------------------------------------------------------------

r101 | bbogart | 2007-08-12 10:06:05 -0700 (Sun, 12 Aug 2007) | 2 lines

Added more data from the last tests.

------------------------------------------------------------------------

r102 | bbogart | 2007-08-13 10:46:51 -0700 (Mon, 13 Aug 2007) | 28 lines

Very close to the final version for Montreal. Patch works well, except
the auto-save function does not work (to save images of each cascade).
Here are the notes for this revision:

screen saver is on!!
This is not quite resolved yet, it seems DPMS starts the screen saver!!!
Left edge of full-screen is cut off, make test patch.
        See test-framing-problems. Issue is 1280x1024 is NOT 4:3!!!
        Set new resolution to 1280x960 for each screen
Try increasing the data fed into the SOM.
        (40x30)=4800,(20x15)=1200
Time for schedule seems to have too large a delay (30min or more?)
        Changed sleep time to <= rather than <. See if that works.
Some activated nodes at screen border do show on free-association,
double check render priorities.
        Seems like priority "0" is not the lowest priority, using 1, 2 now.
The system takes a very long time to fill up 400 nodes, hopefully this
will increase when sending more data to the SOM. Another option could
be to increase the learning frequency.
       We'll see on the next test...
Autosave does not work very well, add a function for directly setting
pix_write and also to reset the capture file-name (include time also)
        Left the old autosave stuff, added functions for resetting the
file-name, quality and to send raw messages directly to pix_write
(like 'auto 1' to save each frame)
Black images are stored in the buffers at start-up now.
Added receive for pix_write to control how images are saved.

Will write an email to see what the hours are. Could a quick sleep mode
that free-associates over its memory without the camera moving or the
ANN being fed? This could be very interesting for the night events at
the SAT. Probably
worthwhile.

------------------------------------------------------------------------
```

r103 | bbogart | 2007-08-13 11:35:04 -0700 (Mon, 13 Aug 2007) | 5 lines

Fixed the autosave function for saving the associations by recreating
the trigger-from-ann signal removed previously.
Added BMU values to the DEBUG output so we can see how the ANN behaves
over time.

---

r104 | bbogart | 2007-08-14 17:10:48 -0700 (Tue, 14 Aug 2007) | 13 lines

Changed the montage method, and how the free association follows. Added
a dream state and a third schedule argument for it. In dream mode the
interval between stimulations is a result of the duration of the
proportions themselves. I also used some of the gemhead priority setting
code from the nodes in the montage. I'm not sure if there is a
correlation between the stimulated node and the montage, I'm having
trouble seeing it, but it must be the case due to the same variable
being used for the node and the image buffer. The schedule stuff for
waking and sleeping should work now.

I'll make a tag of this release if a test over the next few days does
not cause any problems I need to immediately fix.

---

r105 | bbogart | 2007-08-14 23:10:34 -0700 (Tue, 14 Aug 2007) | 3 lines

Added raw log for test of soos for Montreal. 104 is broken, dreaming
system stopped associating.

---

r106 | bbogart | 2007-08-15 11:38:28 -0700 (Wed, 15 Aug 2007) | 2 lines

Added images to identify the numbers on each node.

---

r107 | bbogart | 2007-08-15 13:49:23 -0700 (Wed, 15 Aug 2007) | 3 lines

Went back to the old metro system for the Dreaming state. Hopefully this
fixes the situation.

---

r108 | bbogart | 2007-08-15 14:48:16 -0700 (Wed, 15 Aug 2007) | 3 lines

Changed all the PD files to binary.
Hopefully this is why micro will not open my PD files committed on Insitu!

---

r109 | bbogart | 2007-08-16 00:17:56 -0700 (Thu, 16 Aug 2007) | 2 lines

Final changes before leaving for Montreal tomorrow.

---

r110 | bbogart | 2007-08-16 00:20:58 -0700 (Thu, 16 Aug 2007) | 2 lines

Added tag for the final patch before installing in Montreal.

---

r111 | bbogart | 2007-08-16 11:23:32 -0700 (Thu, 16 Aug 2007) | 2 lines

Added log for a data analysis for the last overnight run of the system.

---

r112 | bbogart | 2007-09-04 14:28:55 -0700 (Tue, 04 Sep 2007) | 2 lines

Final version of SOOS1 used in PDCON07 Installation.

---

r113 | bbogart | 2007-09-10 13:13:19 -0700 (Mon, 10 Sep 2007) | 3 lines

Added some notes to the log-file and also added a tag for this revision of the software.

---

r114 | bbogart | 2007-09-14 09:33:08 -0700 (Fri, 14 Sep 2007) | 2 lines

Added more notes from inspiration from "The Three Gestures of Becoming Aware".

---

r115 | bbogart | 2007-09-18 21:01:07 -0700 (Tue, 18 Sep 2007) | 2 lines

Added an abstraction that makes switching from micro to Insitu video input easier.

---

r116 | bbogart | 2007-09-18 21:51:19 -0700 (Tue, 18 Sep 2007) | 2 lines

Forgot to add this abstraction that deals with controlling the camera while in installation in Montreal.

---

r117 | bbogart | 2007-09-18 22:19:59 -0700 (Tue, 18 Sep 2007) | 4 lines

```
Made a few changes to patch to (hopefully) load on sr-00150 now.
Decreased to 12x12 node, and made the v4l deviceline stuff work for
sr-00150 by default.
```

```
r118 | bbogart | 2007-09-18 22:22:26 -0700 (Tue, 18 Sep 2007) | 3 lines

Small changes to see if sr-00150 will now load the patch. No idea
what is causing it to freeze while trying to load.
```

```
r119 | bbogart | 2007-09-18 22:29:14 -0700 (Tue, 18 Sep 2007) | 2 lines

Disconnected devicename message for USB serial device (on micro)
```

```
r120 | bbogart | 2007-09-18 22:37:15 -0700 (Tue, 18 Sep 2007) | 4 lines

renamed pix_video- (since sr-00150 was not loading the patch) to pix_video.
sr-00150 is now loading the patch, I just checked out a fresh copy. I need
to figure out why this happens (and so rarely).
```

```
r121 | bbogart | 2007-09-19 10:12:13 -0700 (Wed, 19 Sep 2007) | 2 lines

Slight changes to test-camera for sr-00150 ptz range tuning.
```

```
r122 | bbogart | 2007-09-24 10:28:57 -0700 (Mon, 24 Sep 2007) | 10 lines

Changed schedule from sub-patch to abstraction. This version worked on
micro (for a few days) but seems to crash on Insitu after about 10 shell
commands. Next commit will revert to the suspend method used in rev88
(back on he 14th floor) which worked for very long periods on this
machine (sr-00150). Will also make a tag for this micro vs sr-00150 schedules.
That may not be the best way to do it, since schedule has much more than
sleeping stuff in it. I don't have time to rewrite that patch to wrap
the shell stuff in a suspend abstraction. Scratch that, I should just
wrap suspend and make a new abstraction.
```

```
r123 | bbogart | 2007-09-24 10:40:22 -0700 (Mon, 24 Sep 2007) | 4 lines

Suspend wrapper so that it will be easier to switch modes for sr-00150 and
micro. Maybe in the future these should be branches? How to merge
changes between them would be painful.
```

---

r124 | bbogart | 2007-09-24 10:42:31 -0700 (Mon, 24 Sep 2007) | 6 lines

Minor changes to get patches working on sr-00150 for the Demo day on the 26th. Seems to crash after 5 cycles (10 shell commands), the NV driver uses a lot of CPU usage, but that could be a result of watchdog. Will merge in changes from sr-00150 suspend stuff when running on the 14th floor.

---

r125 | bbogart | 2007-09-24 10:49:36 -0700 (Mon, 24 Sep 2007) | 5 lines

Changed schedule to use the new "suspend" wrapper to suspend the X display. Next step is to add that abstraction to tags and make a revision based on the r88 method of suspending (which worked and was well tested on sr-00150)

---

r126 | bbogart | 2007-09-24 10:57:21 -0700 (Mon, 24 Sep 2007) | 2 lines

This should tidy up all the changes to fix suspend on sr-00150.

---

r127 | bbogart | 2007-09-24 10:58:20 -0700 (Mon, 24 Sep 2007) | 3 lines

Renamed the suspend files with the proper pd extension. I need to make sure these are considered binary in svn.

---

r128 | bbogart | 2007-09-24 11:06:52 -0700 (Mon, 24 Sep 2007) | 3 lines

These are all the files that were not marked as binaries in svn, which could cause merging problems in the future.

---

r129 | bbogart | 2007-09-24 11:20:05 -0700 (Mon, 24 Sep 2007) | 2 lines

Put new pan/tilt/zoom ranges into "move-camera" and added to notes.

---

r130 | bbogart | 2007-09-24 11:22:31 -0700 (Mon, 24 Sep 2007) | 4 lines

Removed arbitrary delay for the image auto-saver. Now activations should not be visible in the saved images, and the pause when saving should be less obvious.

---

r131 | bbogart | 2007-09-24 11:25:02 -0700 (Mon, 24 Sep 2007) | 2 lines

Added the 10 day testing log for Sept 26th Demo day.

---

r132 | bbogart | 2007-10-11 13:14:59 -0700 (Thu, 11 Oct 2007) | 3 lines

added log-file that shows development of system over SFU demo day installation.

---

r133 | bbogart | 2007-10-12 10:42:45 -0700 (Fri, 12 Oct 2007) | 3 lines

Added a branch for the very old experiments with the SOM. Contains a few changes from the original "rev1" files.

---

r134 | bbogart | 2007-10-13 10:09:17 -0700 (Sat, 13 Oct 2007) | 5 lines

This version of the patch allows much more controlled data to be presented to the network. In this case the probability of RGB values are specified to see how regions of similar inputs, and boundaries are created in the resulting SOM feature maps.

---

r135 | bbogart | 2007-10-17 12:04:25 -0700 (Wed, 17 Oct 2007) | 2 lines

Added BMU log for latest installation run (for SFU demo day)

---

r136 | bbogart | 2007-10-17 13:03:30 -0700 (Wed, 17 Oct 2007) | 2 lines

Added file that stores BMU for a particular log-file.

---

r137 | bbogart | 2007-10-17 14:39:18 -0700 (Wed, 17 Oct 2007) | 2 lines

Added BMU_124 to R

---

r138 | bbogart | 2007-10-17 14:49:44 -0700 (Wed, 17 Oct 2007) | 2 lines

Added extra files for BMU progressions for different revisions.

---

r139 | bbogart | 2007-10-22 15:36:17 -0700 (Mon, 22 Oct 2007) | 3 lines

Updates to R data file, did research into how to do time-tagged irregular time series data in R.

---

r140 | bbogart | 2007-10-26 12:25:50 -0700 (Fri, 26 Oct 2007) | 7 lines

Initial draft of the time-tagging part of the new logger system (to facilitate loading log data into R. The use of PD symbols may cause long-term problems, make sure they are only generated when debugging. Using the default POSIX format: "%Y-%m-%d %H:%M:%S" expect the centre space is an _ so: "%Y-%m-%d_%H:%M:%S". Next steps to move all the events into a centralized system.

---

r141 | bbogart | 2007-10-26 16:11:39 -0700 (Fri, 26 Oct 2007) | 6 lines

Changing infrastructure to the new event logging system. finished the logger and the message format. Changed only the BMU, image capture, and schedule and ANN learning and neighbourhood events thus far.

Am I missing anything that needs to be logged?

---

r142 | bbogart | 2007-10-31 11:14:09 -0700 (Wed, 31 Oct 2007) | 3 lines

Added directories for thesis in SVN, included OpenOffice spreadsheet of thesis development plan (for Maia).

---

r143 | bbogart | 2007-10-31 16:29:14 -0700 (Wed, 31 Oct 2007) | 7 lines

Added notes as per discussion with Steven Barns about activation spikes.
Changed schedule as per his suggestion to make it run in "waking" state for 24 hours/day continuously (no sleeping).
Changed range of camera to see largely closeups (Maia Engeli's suggestion) this will also increase the diversity of data presented to the network, maybe more nodes would be needed.

---

r144 | bbogart | 2007-11-05 12:13:29 -0800 (Mon, 05 Nov 2007) | 18 lines

Added header file to make columns easier to read in R.
The format for creating the final log file (suitable for R) from the raw soos.log is as follows:

```
cat log_header.txt > soos.log &&  cat soos.log.raw | grep LOG | cut -d "
" -f 2- >> soos.data
```

Then loading into R:

```
data =
read.table(file="soos.data",colClasses=c("POSIXct","factor","factor","numeric"))
```

The best way to interpret the time in the proper class is still in development. It will be something like:

```
data[,1] = timeDate(data[,1])
```

---

r145 | bbogart | 2007-11-05 15:57:57 -0800 (Mon, 05 Nov 2007) | 3 lines

Added file showing shortcuts for the R commands to load the time series data from soos (using the new data format)

---

r146 | bbogart | 2007-11-05 16:23:59 -0800 (Mon, 05 Nov 2007) | 4 lines

Added a few more comments, got plotting to work for the time-series, can even create a zoo object, but the plotting needs a lot of work, need to check out plotting methods.

---

r147 | bbogart | 2007-11-06 22:20:34 -0800 (Tue, 06 Nov 2007) | 3 lines

Added first step rough diagrams to describe the new ideas behind realization-interpretation.

---

r148 | bbogart | 2007-11-06 22:36:26 -0800 (Tue, 06 Nov 2007) | 2 lines

added notes for thesis ideas, reference etc..

---

r149 | bbogart | 2007-11-06 22:52:22 -0800 (Tue, 06 Nov 2007) | 6 lines

Added LyX file for the most recent paper which the thesis will be based on. The file was not altered to use the local references in the svn directory, that will have to be fixed. Also added the source files and PDFs for the figures used in the document, whose references will also need to be fixed.

```
------------------------------------------------------------------------
r150 | bbogart | 2007-11-07 10:04:31 -0800 (Wed, 07 Nov 2007) | 2 lines

Added ideas about diagrams to generic "notes" file.

------------------------------------------------------------------------
r151 | bbogart | 2007-11-07 10:13:19 -0800 (Wed, 07 Nov 2007) | 3 lines

Notes about timing being too fast, lack of seizure activity in
perpetual waking  state, possible additions to system.

------------------------------------------------------------------------
r152 | bbogart | 2007-11-07 10:14:25 -0800 (Wed, 07 Nov 2007) | 2 lines

See diff for more info

------------------------------------------------------------------------
r153 | bbogart | 2007-11-07 10:53:16 -0800 (Wed, 07 Nov 2007) | 2 lines

Changed "Realization" and "Interpretation" Nodes into named links.

------------------------------------------------------------------------
r154 | bbogart | 2007-11-07 11:01:53 -0800 (Wed, 07 Nov 2007) | 2 lines

Add notes about outline currently on wiki.

------------------------------------------------------------------------
r155 | bbogart | 2007-11-07 14:15:20 -0800 (Wed, 07 Nov 2007) | 2 lines

Added a few more commands for locating events (in date terms) in plots.

------------------------------------------------------------------------
r156 | bbogart | 2007-11-08 12:52:21 -0800 (Thu, 08 Nov 2007) | 4 lines

Added log-file for last run, stopped to change screens to VGA
(facilitating working with projectors) and make some changed to the
software.

------------------------------------------------------------------------
r157 | bbogart | 2007-11-08 13:29:40 -0800 (Thu, 08 Nov 2007) | 12 lines

Changed the interval from 10s to 12s to give a longer pause in between.
Increased pix_resize to 400x300 pixels (that is 480,000 sensors per
node)
Turned sleep back on.
oprofile seems dead:
opreport error:
```

```
/var/lib/oprofile/samples/current/{root}/usr/bin/find/{dep}/{anon}/12099.0x805b000.0
    x807c000/CPU_CLK_UNHALTED.100000.0.all.all.all:
Invalid argument
Should add a CPU event for CPU load compared to watchdogs. I don't
remember which PD object to use.
```

```
r158 | bbogart | 2007-11-08 13:52:26 -0800 (Thu, 08 Nov 2007) | 2 lines

Added notes about projector experiment from today.
```

```
r159 | bbogart | 2007-11-08 16:40:38 -0800 (Thu, 08 Nov 2007) | 2 lines

Last minute inspiration
```

```
r160 | bbogart | 2007-11-13 12:28:56 -0800 (Tue, 13 Nov 2007) | 3 lines

Updated the schedule based on Maia's comments and make it reflect the
current thesis outline.
```

```
r161 | bbogart | 2007-11-14 09:49:54 -0800 (Wed, 14 Nov 2007) | 5 lines

Added soos.log for the last run of installation, which was running fine
without the sleep state, has now crashed. Will do an analysis on the
resulting data. The .data file is less useful due to the crash as the pd
watchdog messages do not come up.
```

```
r162 | bbogart | 2007-11-14 09:53:28 -0800 (Wed, 14 Nov 2007) | 4 lines

Added parsed data-file to repository for analysis, remember both this
file and the raw log are in the repository to see if the cause of
the crash can be found.
```

```
r163 | bbogart | 2007-11-14 09:56:13 -0800 (Wed, 14 Nov 2007) | 3 lines

Changed the format of the .data parser so that it looks for the soos.log
in the parent.
```

```
r164 | bbogart | 2007-11-14 10:10:04 -0800 (Wed, 14 Nov 2007) | 4 lines
```

```
This file contains all the time-stamps that precede a pd watchdog
message. This should give an indication of the context where the system
went out of control.

------------------------------------------------------------------------
r165 | bbogart | 2007-11-14 15:56:39 -0800 (Wed, 14 Nov 2007) | 3 lines

Added comments about Surrey art gallery and the crash (freeze) of the
system.

------------------------------------------------------------------------
r166 | bbogart | 2007-11-16 14:55:53 -0800 (Fri, 16 Nov 2007) | 2 lines

Added an idea for having the zoom range controlled by the learning rate.

------------------------------------------------------------------------
r167 | bbogart | 2007-11-16 15:11:12 -0800 (Fri, 16 Nov 2007) | 2 lines

Diagram that shows the R->I loop nested with specific details.

------------------------------------------------------------------------
r168 | bbogart | 2007-11-16 15:45:37 -0800 (Fri, 16 Nov 2007) | 3 lines

Added a short note to add to tomboy at home about one of the sources
(Minon Kwan) I had to return. Annotated bib makes the most sense
in tomboy.

------------------------------------------------------------------------
r169 | bbogart | 2007-11-17 12:30:26 -0800 (Sat, 17 Nov 2007) | 8 lines

Added a directory in SVN to store tomboy notes. Unfortunately it seems
the format is different than the way they are locally stored in the home
directory, so that means that I will not be able to read these notes in
the older version of tomboy in the machine at school.

I wonder if I can upgrade to a newer version of tomboy (from the new
testing?) that has the sync included.

------------------------------------------------------------------------
r170 | bbogart | 2007-11-17 12:33:07 -0800 (Sat, 17 Nov 2007) | 2 lines

Fixed spelling error. (Why did that not come up in OpenOffice 2.0 on
sr-00149?)

------------------------------------------------------------------------
r171 | bbogart | 2007-11-19 11:15:00 -0800 (Mon, 19 Nov 2007) | 3 lines
```

```
Added some comments about measuring the mean/min/max memory of the
system.

------------------------------------------------------------------------

r172 | bbogart | 2007-11-20 11:40:26 -0800 (Tue, 20 Nov 2007) | 2 lines

Added a new tomboy note for the annotation of "One Place After Another"

------------------------------------------------------------------------

r173 | bbogart | 2007-11-20 12:17:52 -0800 (Tue, 20 Nov 2007) | 3 lines

Added a note about storing all the images for one node, rather than
replacing older images with newer ones.

------------------------------------------------------------------------

r174 | bbogart | 2007-11-21 13:32:48 -0800 (Wed, 21 Nov 2007) | 5 lines

MAM crashed again, this time it was not X but PD itself. This crash is
logged by 157B as its the same revision (no code changes) as the last
time I ran it (157). Now I really need something to measure ram usage
and CPU usage from inside PD, I guess with Shell. Some mix of top and cut
probably. A set of PD externals that deal with such issues would be ideal.

------------------------------------------------------------------------

r175 | bbogart | 2007-11-21 14:13:15 -0800 (Wed, 21 Nov 2007) | 6 lines

Added a Zotero based bib, and moved the older one (which could have many
errors) to old-bib. The new bib is now called "bibliography.bib".
Added (at this time incomplete) notes on Nick Kaye's book to add to your
annotated bib. Why did "One place after another" stay in svn after I
removed it? Maybe I forgot to commit at home.

------------------------------------------------------------------------

r176 | bbogart | 2007-11-22 11:26:48 -0800 (Thu, 22 Nov 2007) | 2 lines

Notes for Nick Kaye's book about site-specific practises. Will put in
tomboy and remove this file.

------------------------------------------------------------------------

r177 | bbogart | 2007-11-22 12:37:23 -0800 (Thu, 22 Nov 2007) | 1 line

------------------------------------------------------------------------

r178 | bbogart | 2007-11-22 13:02:50 -0800 (Thu, 22 Nov 2007) | 2 lines

Disconnected the suspend object to see if that was causing the crash
after a few days.
```

```
------------------------------------------------------------------------
r179 | bbogart | 2007-11-23 10:45:50 -0800 (Fri, 23 Nov 2007) | 4 lines

Added Kaye's book to tomboy notes, removed original placeholder file.
Auto-generated bibtex keys for data exported from Zotero at school.
Committed the deletion of the placeholder for one-place-after another.

------------------------------------------------------------------------
r180 | bbogart | 2007-11-23 13:12:47 -0800 (Fri, 23 Nov 2007) | 2 lines

Added Schoen to bibliography.

------------------------------------------------------------------------
r181 | bbogart | 2007-11-26 11:41:01 -0800 (Mon, 26 Nov 2007) | 6 lines

Added logs for the next run of MAM. So I've disconnected the "suspend"
stuff and its still crashing. Why? Double check the acpi stuff in the
kernel, I wonder if the machine just freezes after being up for some time.
sr-00149 has been working fine for 17 days. Will take a look at this new
log and see if I can figure out what is going on. For the SAG installation
I'll just be running the system without sleep mode, which seemed to
work for a while previously.

------------------------------------------------------------------------
r182 | bbogart | 2007-11-26 12:13:02 -0800 (Mon, 26 Nov 2007) | 2 lines

Added notes about current crashing issues (due to 400x300 pixel
down-sampling?)

------------------------------------------------------------------------
r183 | bbogart | 2007-11-26 12:14:41 -0800 (Mon, 26 Nov 2007) | 2 lines

Added note about revision references.

------------------------------------------------------------------------
r184 | bbogart | 2007-11-26 12:17:14 -0800 (Mon, 26 Nov 2007) | 2 lines

Added log-files for rev157 in the R data file. Was not helpful is
finding the cause of the failure.

------------------------------------------------------------------------
r185 | bbogart | 2007-11-26 12:32:17 -0800 (Mon, 26 Nov 2007) | 3 lines

Changed sub-sampling to 100x75 pixels, to see if that helps with
stability, also added a 4s delay to when the image is passed to the SOM and
the memory activations happen.
```

---

r186 | bbogart | 2007-11-26 13:09:32 -0800 (Mon, 26 Nov 2007) | 3 lines

Added abstractions that print the amount of CPU and RAM usage (as %)
to see about crashing causes. Changed the arrangement of the parent patch
slightly to make it a little more readable.

---

r187 | bbogart | 2007-11-26 13:18:47 -0800 (Mon, 26 Nov 2007) | 3 lines

Removed the 4s delay from [s BMU] it caused artifacts and did not help
with the problem it was aiming to solve.

---

r188 | bbogart | 2007-11-26 13:38:44 -0800 (Mon, 26 Nov 2007) | 8 lines

Changed the SOM size to match the (new) 100x75 pixel size.
I'm seeing a strange flash each 12s, no idea where it is coming from,
looks like it could be the gemhead for the video input, but that is at
100 render priority. In fact there is nothing to flash, that chain has
no geometry. Changed the schedule for no sleeping. Other recent changes
are the CPU/logging stuff (see if "log 0") fixes it. What else has
changed?

---

r189 | bbogart | 2007-11-26 13:43:38 -0800 (Mon, 26 Nov 2007) | 5 lines

Flashing still happening with "log 0" now trying to go back to 40x30
pixels. Next step could be to remove "shell" from PD, to see if that
helps stability. Oh and it is PD, not X crashing now, so we can rule out
the OS.

---

r190 | bbogart | 2007-11-26 14:09:22 -0800 (Mon, 26 Nov 2007) | 4 lines

Removing "shell" seems to have worked!!! no flashing, seems stable.
(that was without using "log 1" though) Moved resolution back up to
400x300 and lets see what happens.

---

r191 | bbogart | 2007-11-26 15:01:55 -0800 (Mon, 26 Nov 2007) | 8 lines

Added xset s off so we have no screen-saver at all (since shell is now
gone)
went back to the 100x75 pixel sub-sampling just in case, I had another
kernel crash, this time from "start"ing again after "stop" I think that
is probably due to being close to an edge of disaster, I hope this works

for a least a few days. I did not go for the 80x40 because it takes too long to fill the memory screen for such a short installation.

---

r192 | bbogart | 2007-11-26 15:30:33 -0800 (Mon, 26 Nov 2007) | 3 lines

minor changes, just moved Gabora outside of MAM. Don't have the concentration to work on this currently.

---

r193 | bbogart | 2007-11-26 16:27:05 -0800 (Mon, 26 Nov 2007) | 2 lines

Added notes for the reflective practitioner

---

r194 | bbogart | 2007-11-29 13:13:24 -0800 (Thu, 29 Nov 2007) | 2 lines

Final changes for installation at SAG for eMixer.

---

r195 | bbogart | 2007-11-29 13:20:47 -0800 (Thu, 29 Nov 2007) | 2 lines

Added some updates about MAM, and also a note about input space, output space and PTZ space.

---

r196 | bbogart | 2007-11-30 11:28:21 -0800 (Fri, 30 Nov 2007) | 5 lines

Redesigned the move-camera abstraction to make it cleaner and add the ability to ignore certain regions in the ptz space. All is implemented, but ptz-ignore has not been tested. Remember to mark these PD patches as binary soon to keep problems from happening!

---

r197 | bbogart | 2007-11-30 16:48:39 -0800 (Fri, 30 Nov 2007) | 5 lines

Tested method of ignoring a region. This should work, the question is how big the area is that we are ignoring. I think its really big in this case. Added a second outlet to ignore that can be used to gang multiples and still generate a signal to generate another random position.

---

r198 | bbogart | 2007-12-03 15:17:01 -0800 (Mon, 03 Dec 2007) | 2 lines

Added a note about the damn activations that wrap around grid borders.

```
r199 | bbogart | 2007-12-05 15:21:17 -0800 (Wed, 05 Dec 2007) | 5 lines

Added revision for final run of MAM in 3400.
When I arrived the network was full of the same image and the camera
stuck looking in the corner. I hope this is not a camera failure, but
some power glitch or something.

------------------------------------------------------------------------
r200 | bbogart | 2007-12-06 19:31:26 -0800 (Thu, 06 Dec 2007) | 2 lines

Added note about a method to get rid of the free-association buffer.

------------------------------------------------------------------------
r201 | bbogart | 2007-12-07 15:51:26 -0800 (Fri, 07 Dec 2007) | 6 lines

Added missing Dec 5th note.
Changed panning and zooming to trandom (now seeded on loadbang), tilting
still [random].
Changed parameters for 3950 installation.
Removed "ignore" logic.
Added events for pan/tilt/zoom.

------------------------------------------------------------------------
r202 | bbogart | 2007-12-08 13:38:02 -0800 (Sat, 08 Dec 2007) | 2 lines

Added notes about randomness.

------------------------------------------------------------------------
r203 | bbogart | 2007-12-09 19:28:27 -0800 (Sun, 09 Dec 2007) | 6 lines

Added notes for the books I need to return to the library:
The Reflective Practitioner
Surroundings Surrounded
The Embodied Mind
Art Practice as Research

------------------------------------------------------------------------
r204 | bbogart | 2007-12-09 19:51:07 -0800 (Sun, 09 Dec 2007) | 4 lines

Added a few of the books which I had to return to the library to the
SVN bibliography. Removed the reflective-practitioner since the notes
are included in tomboy.

------------------------------------------------------------------------
r205 | bbogart | 2007-12-10 14:41:01 -0800 (Mon, 10 Dec 2007) | 2 lines

Added the log file for the short installation run in 3956 for Bing Thom.
```

---

r206 | bbogart | 2007-12-11 08:35:38 -0800 (Tue, 11 Dec 2007) | 4 lines

Added patch for long-term testing of ignore code. After running
this over 3 million iterations I saw no failures like we did in 3400
on the last day. Will test when MAM moves into its new location.

---

r207 | bbogart | 2007-12-11 09:05:51 -0800 (Tue, 11 Dec 2007) | 3 lines

Added log-files for the shorter and longer term tests running on Insitu
at home after the failure in 3400. Have not had a chance to look
at these in R yet.

---

r208 | bbogart | 2007-12-18 16:23:15 -0800 (Tue, 18 Dec 2007) | 3 lines

Removed ptz log-files as its clear how the ptz-ignore was failing, see
notes log as of December 18th.

---

r209 | bbogart | 2007-12-18 21:20:44 -0800 (Tue, 18 Dec 2007) | 5 lines

Looked at the pan/tilt data, realized that the ignore is not working
properly (OR rather than AND) could not find any periodicity of the
randomness, nor any difference between the time seeded and non time
seeded values.

---

r210 | bbogart | 2007-12-18 21:21:37 -0800 (Tue, 18 Dec 2007) | 2 lines

The notes are not so "initial" anymore!

---

r211 | bbogart | 2007-12-18 22:32:13 -0800 (Tue, 18 Dec 2007) | 5 lines

Figured out how to make a phase-plane plot in R. [trandom] and [random]
look identical (which makes sense since time-seeding should make no
difference). The plot is totally different than rnorm() in R, the phase
space is a hexagon shape.

---

r212 | bbogart | 2007-12-20 17:19:52 -0800 (Thu, 20 Dec 2007) | 5 lines

Fixed the ptz-ignore code by combining both pan and tilt tests into a
single expression. Plots show that regions are ignored and the

re-randomize method when a value is ignored seems to work fine. Also tested gang to ignore multiple regions.

---

r213 | bbogart | 2007-12-21 15:38:46 -0800 (Fri, 21 Dec 2007) | 6 lines

Changed autosave so that one image is archived for the top and bottom of each learning cycle. This way the level of organization can be determined. Also connected memory to the test-images abstraction so I don't have to remember to change that value (as long as there are enough images created, I think there are just 400 currently.)

---

r214 | bbogart | 2007-12-27 13:22:39 -0800 (Thu, 27 Dec 2007) | 8 lines

Removed code for panonly mode (used when the tilt drive was broken).
put sleep back in, but left shell out so the camera gets a rest overnight.
added a file that shows the correlation between number of nodes and baseline CPU usage, where does the baseline CPU come from? Changed loadbangs in nodes to r loadbangs to see if that helps with baseline CPU (seems to make no effect). 20x20 nodes currently for one last installation in 3400.

---

r215 | bbogart | 2007-12-27 14:54:41 -0800 (Thu, 27 Dec 2007) | 8 lines

Changed back to pan/tilt ranges from rev199.
Put ptz-ignore back in.
a 20x20 memory grid killed this machine, certainly slower than Insitu used 12x12 for this next run.
Put the pan/tilt/zoom ranges back into the move-camera abstraction so that its easier to put them into the system (without loading the parent patch).

---

r216 | bbogart | 2007-12-27 16:07:16 -0800 (Thu, 27 Dec 2007) | 4 lines

Sorted the data by number of Nodes and added column headers.
Found this tutorial on R regression, but I've been unable to make it work on my data: http://mercury.bio.uaf.edu/mercury/R/R.html

---

r217 | bbogart | 2008-01-02 08:56:35 -0800 (Wed, 02 Jan 2008) | 2 lines

Added notes pertaining to final installation in 3400.

```
r218 | bbogart | 2008-01-02 10:47:41 -0800 (Wed, 02 Jan 2008) | 2 lines

Added log-file for final run in 3400 for fall 2007.

------------------------------------------------------------------------
r219 | bbogart | 2008-01-02 13:58:25 -0800 (Wed, 02 Jan 2008) | 8 lines

Added notes for the last day of installation in 3400.
Figured out a better way to store SOOS data to minimize the number
of variables laying around. The solution is to use a dataframe to store
the time and values for a particular event. Made an R function for
the phase-plane plot, though it almost always shows a similar structure,
which makes me wonder if I'm doing the algorithm incorrectly.
Changed move-camera so that the log shows what was sent to the camera,
not what was initially sent before ptz-ignore!

------------------------------------------------------------------------
r220 | bbogart | 2008-01-10 10:03:11 -0800 (Thu, 10 Jan 2008) | 3 lines

Added updates of figures used in chapter for "Computational Arts and
Creative Informatics" and documents for Ethics.

------------------------------------------------------------------------
r221 | bbogart | 2008-01-11 14:18:12 -0800 (Fri, 11 Jan 2008) | 14 lines

Figured out how to measure the amount of time SOOS will remember things.
Here is the summary for rev217:

   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
      0     264     786    1885    1909   39550

in seconds. So that max is about 11 hours. This is including the
sleeping state! So things are remembered longer overnight. I should be
able to threshold out these really long memories. 30min is the mean,
which should not be effected much by the long memories, which only
happened once per day. Still a sense of the memory not including
sleeping time would be nice.


------------------------------------------------------------------------
r222 | bbogart | 2008-01-11 23:40:53 -0800 (Fri, 11 Jan 2008) | 4 lines

"which" works better than "grep" to return the index of a certain value.
Added a "learn" dataframe for measuring the time of the cycle (about
1 hour).

------------------------------------------------------------------------
```

```
r223 | bbogart | 2008-01-14 15:21:07 -0800 (Mon, 14 Jan 2008) | 5 lines

Changed the name of the R data file so that it automatically gets loaded
when R starts in the "data-analysis" directory. Also saving the
"current" workspace at the end of the session should save the changes to
the appropriate file.

------------------------------------------------------------------------
r224 | bbogart | 2008-01-14 15:33:36 -0800 (Mon, 14 Jan 2008) | 3 lines

Whoops, forgot to add the .data file that corresponds to the raw log from
PD.

------------------------------------------------------------------------
r225 | bbogart | 2008-01-14 22:21:17 -0800 (Mon, 14 Jan 2008) | 4 lines

The conflict between the dream-association BMU and the real BMU from the
SOM means the the memory length of the system is hard to get at. Will
need to use a different name for the dream BMUs than the real ones.

------------------------------------------------------------------------
r226 | bbogart | 2008-01-14 22:41:03 -0800 (Mon, 14 Jan 2008) | 2 lines

Some more work on extracting proper BMUs, not working yet.

------------------------------------------------------------------------
r227 | bbogart | 2008-01-15 13:06:44 -0800 (Tue, 15 Jan 2008) | 21 lines

Made a branch for the new work on Free-Association where
the free association representation's logic is in the same system as the
SOM nodes. This was the approach to keep the activation and free
association representation in sync. The reason why it is not in trunk is
that the CPU usage is very high for some reason, too many PD objects?
I'll need to look at the baseline CPU usage problem, as I'm not even
rendering and the branch shows 50% CPU. The old version showed about 50%
while rendering!

Here are the notes during the process:
double check name "dream-bmu"
make sure that the event for dream-bmu works.
make an event that shows the string of nodes which get activated.

53 new som
55 no free assoc
50 no gemhead- [pd gemhead] instead

So where is the extra load coming from?
```

```
The dream stuff has not been double checked.

------------------------------------------------------------------------
r228 | bbogart | 2008-01-17 20:52:11 -0800 (Thu, 17 Jan 2008) | 3 lines

Attempting to resolve issues with having committed a mistake!
I hope this is converting rev226 into 228?

------------------------------------------------------------------------
r229 | bbogart | 2008-01-17 21:05:33 -0800 (Thu, 17 Jan 2008) | 2 lines

Almost there!

------------------------------------------------------------------------
r230 | bbogart | 2008-01-17 21:12:57 -0800 (Thu, 17 Jan 2008) | 3 lines

This had better work this time! I had to manually copy the files out of
the working copy and then copy them back into after an update.

------------------------------------------------------------------------
r231 | bbogart | 2008-01-18 10:06:45 -0800 (Fri, 18 Jan 2008) | 2 lines

Added files for testing the baseline CPU problem.

------------------------------------------------------------------------
r232 | bbogart | 2008-01-21 16:12:55 -0800 (Mon, 21 Jan 2008) | 7 lines

Added a file to test pix_write.
For some reason pix_write is not working on sr-00150 anymore! I have
done more updates on this machine, including a newer kernel and a new
Nvidia driver (same version, new compilation). I've already tried the
Gem binary from sr-00149 and I've also tried recompiling in case the
new libs (tif,jpg) are different. I've also tried -nogui, no matter
what I can't make Gem save anything but a 100% black image. Now I'm
not sure if it works on sr-00149, so I'm now testing on that machine.

------------------------------------------------------------------------
r233 | bbogart | 2008-01-21 16:15:41 -0800 (Mon, 21 Jan 2008) | 2 lines

Added a directory to store test images (testing for SOM sorting, not
placeholders as in test-images.

------------------------------------------------------------------------
r234 | bbogart | 2008-01-21 22:33:22 -0800 (Mon, 21 Jan 2008) | 2 lines
```

```
Work on ethics process.

------------------------------------------------------------------------
r235 | bbogart | 2008-01-21 22:50:21 -0800 (Mon, 21 Jan 2008) | 2 lines

Added reference for one of the meta-creation readings.

------------------------------------------------------------------------
r236 | bbogart | 2008-01-21 23:16:04 -0800 (Mon, 21 Jan 2008) | 2 lines

Added file to store quotes for Thesis.

------------------------------------------------------------------------
r237 | bbogart | 2008-01-23 11:13:31 -0800 (Wed, 23 Jan 2008) | 3 lines

Added directory for Papers referenced in thesis (if available) also for
quotes.

------------------------------------------------------------------------
r238 | bbogart | 2008-01-23 11:19:16 -0800 (Wed, 23 Jan 2008) | 2 lines

Added "The Garden of Chances" to bib.

------------------------------------------------------------------------
r239 | bbogart | 2008-01-23 11:23:25 -0800 (Wed, 23 Jan 2008) | 2 lines

Added reference to paper for Quote.

------------------------------------------------------------------------
r240 | bbogart | 2008-01-23 11:28:35 -0800 (Wed, 23 Jan 2008) | 2 lines

Updated notes.

------------------------------------------------------------------------
r241 | bbogart | 2008-01-23 13:33:06 -0800 (Wed, 23 Jan 2008) | 3 lines

Added test files of the city for testing of SOM in SOMLab (learning
functions and ability to categorize).

------------------------------------------------------------------------
r242 | bbogart | 2008-01-24 12:40:39 -0800 (Thu, 24 Jan 2008) | 3 lines

Changed "Artist" to "Artist Agent" and changed "Representation" to
"External Properties".

------------------------------------------------------------------------
r243 | bbogart | 2008-01-24 14:26:33 -0800 (Thu, 24 Jan 2008) | 2 lines
```

```
Changed "Process" to "Embodied Process"

------------------------------------------------------------------------
r244 | bbogart | 2008-01-24 15:58:50 -0800 (Thu, 24 Jan 2008) | 2 lines

Updated these files for MAM, and a little cleanup.

------------------------------------------------------------------------
r245 | bbogart | 2008-01-24 15:59:16 -0800 (Thu, 24 Jan 2008) | 2 lines

Renamed files to reflect new name.

------------------------------------------------------------------------
r246 | bbogart | 2008-01-24 16:35:33 -0800 (Thu, 24 Jan 2008) | 3 lines

Changed "Memory Nodes" to "Memory Network" and "Free Association" to
"Free-Association Network"

------------------------------------------------------------------------
r247 | bbogart | 2008-01-28 10:48:39 -0800 (Mon, 28 Jan 2008) | 3 lines

Added references used in "Computational Arts" chapter to thesis bib.
Made some changes to the architecture diagram for the same publication.

------------------------------------------------------------------------
r248 | bbogart | 2008-01-28 10:52:08 -0800 (Mon, 28 Jan 2008) | 3 lines

Added another set of images for testing, these ones based on images of
the mountains from the 14th floor.

------------------------------------------------------------------------
r249 | bbogart | 2008-01-28 11:10:08 -0800 (Mon, 28 Jan 2008) | 3 lines

Renamed the branch with the montage code in each node, this is because I
did not realize what it was initially! This name is much more clear.

------------------------------------------------------------------------
r250 | bbogart | 2008-01-28 11:20:49 -0800 (Mon, 28 Jan 2008) | 4 lines

Added a patch to generate images for SOM testing.
Added 400 of the resulting images (from 0 to 180 degrees of rotation).
Will now work on the SOM stuff.

------------------------------------------------------------------------
r251 | bbogart | 2008-01-28 16:33:52 -0800 (Mon, 28 Jan 2008) | 11 lines
```

```
I've made lots of changes here. First fixed the problem with pix_write,
it was due to the path saving the files to being too long. I've removed
[shell] for all items that were using it. shell has been replaced with a
python external that implements a popen call to run commands. It is not
of general use yet, but should work for the needs of MAM. pdcpu and
pdmem are then reattached and should work (logging of those values is
now enabled.) changed run so that the path of pypopen.py (trunk) is
known. test-py-system is the test patch for the pypopen stuff. Suspend
is now also working then, will do a stress-test overnight to see how
much more stable the python method is than the shell method.

------------------------------------------------------------------------

r252 | bbogart | 2008-01-28 19:01:59 -0800 (Mon, 28 Jan 2008) | 5 lines

Made some progress with the SOMLab.
Seems the SOM is blind to exactly horizontal and vertical, they are seen
as the same (which makes sense as a repeating pattern. These test
images are not very real-world. Trying a big SOM with 10 images.

------------------------------------------------------------------------

r253 | bbogart | 2008-01-29 11:16:27 -0800 (Tue, 29 Jan 2008) | 2 lines

Added a directory so we can have multiple sets of test image data.

------------------------------------------------------------------------

r254 | bbogart | 2008-01-31 11:07:24 -0800 (Thu, 31 Jan 2008) | 4 lines

Updates for running MAM in 14710.
Removed the code to autosave the som (now a command) due to very high
CPU usage during the writing.

------------------------------------------------------------------------

r255 | bbogart | 2008-01-31 15:55:20 -0800 (Thu, 31 Jan 2008) | 7 lines

Lots of work on the SOMLab and data-analysis.
I now know how to view the code-books of nodes as images in R!
The good news is that the patch does work! of course there are no clusters
in my test-image-data! How to make clusters? slight random variation in the
spacing of lines? Will get very interesting when testing with
captured image-data.

------------------------------------------------------------------------

r256 | bbogart | 2008-02-01 21:50:22 -0800 (Fri, 01 Feb 2008) | 2 lines

Changes to diagram for Colloq presentation in Late Feb 2008.
```

---

r257 | bbogart | 2008-02-04 17:04:20 -0800 (Mon, 04 Feb 2008) | 8 lines

Added output cases for various SOM configurations. None have yet shown the structure apparent in the MAM installation. Why? Is it the complexity of lots of data coming in showing more structure? What is so different about this test case? Another step will be a patch that gives a distribution of input data, say three clusters of a certain size. Apply these as the image indices of stripes, and then I should be able to see clustering. Parallel development should be happening with colour values. (from the old som-explorer project).

---

r258 | bbogart | 2008-02-04 18:15:58 -0800 (Mon, 04 Feb 2008) | 2 lines

renamed the cosine case, slight changes to patches.

---

r259 | bbogart | 2008-02-07 14:52:58 -0800 (Thu, 07 Feb 2008) | 3 lines

Update of MAM paper for Computational Arts Chapter.
Made this commit before I cut out large sections.

---

r260 | bbogart | 2008-02-07 16:05:42 -0800 (Thu, 07 Feb 2008) | 2 lines

Changes for the Computational Arts Publication.

---

r261 | bbogart | 2008-02-08 17:09:44 -0800 (Fri, 08 Feb 2008) | 2 lines

More changes for the Computational Arts Chapter.

---

r262 | bbogart | 2008-02-08 21:54:05 -0800 (Fri, 08 Feb 2008) | 6 lines

Small changes to data-analysis, need to figure out how to show code-book vectors in colour. Wait, if the SOM is not sensitive to structural information, then the code-book vectors should not show structural information should they? Would they not be in the wrong order if the code-book elements are ignored in terms of position and only recognized in terms of value?

---

r263 | bbogart | 2008-02-08 21:56:01 -0800 (Fri, 08 Feb 2008) | 6 lines

Seems the SOM is not sensitive to structure, the order of colours does

not matter, just the values! Quite a disappointment. Could the colours in the installation really allow the seemly structural sorting of images? Or is something else going on? If the SOM is not sensitive to structure, why do the code-book vectors show image structure? Position of the values must be important.

---

r264 | bbogart | 2008-02-08 23:15:24 -0800 (Fri, 08 Feb 2008) | 4 lines

Took a quick look at some code-books for case 8 and 9, strange thing is there seems to be some clustering structure there. Need to figure out how to use ggplot2 and make a matrix of code-books to compare with the output feature map.

---

r265 | bbogart | 2008-02-10 11:42:49 -0800 (Sun, 10 Feb 2008) | 5 lines

Added a version of numSquares in colour to see if that made any difference to the SOM sorting. It did not. The colours for each box stayed the same over all the iterations of increasing numbers of squares. What is the best way to test the colour stuff?

---

r266 | bbogart | 2008-02-10 21:00:55 -0800 (Sun, 10 Feb 2008) | 8 lines

Even using coloured circles the SOM does not appear any different than random! Will try solid filled coloured images next, to see if that makes any difference. I have a feeling something strange is happening and the SOM is not being properly presented. I should make a PD patch that uses the final training data from each set to reconstruct the feature-map, perhaps it will not resemble the saved feature-map and be a hint as to some rendering bug?

---

r267 | bbogart | 2008-02-10 21:03:23 -0800 (Sun, 10 Feb 2008) | 3 lines

renamed file so that a new non-circle colour generator can use the original name.

---

r268 | bbogart | 2008-02-10 22:22:23 -0800 (Sun, 10 Feb 2008) | 6 lines

So its clear that using colour for images is not working. It was working sending raw colour values to the som, so it could be something to do with the image stuff, or some other bug. Need to compare the learning results to the feature-map image saved at the end of training. Something is fishy.

---

r269 | bbogart | 2008-02-11 11:48:47 -0800 (Mon, 11 Feb 2008) | 10 lines

Turns out "1" is not the max radius (based on the code) actually this is not yet confirmed as "neighbourhood" may be scaled before it gets to the learning function in the code. Nevertheless I've committed a hard multiplier of 8.5 (the max radius for a 12x12 network) to the neighbourhood function, and keeping it in time sync with the learning function. I wonder if this will solve the strange problems with the algorithm not behaving. It does not explain the difference between the behaviour of the installation and the test-case, nor how well the code-books appear to resemble the input patterns.

---

r270 | bbogart | 2008-02-12 11:14:53 -0800 (Tue, 12 Feb 2008) | 3 lines

More editing for chapter, hopefully the change of page size for the R>I will make it render properly in latex! Added more references.

---

r271 | bbogart | 2008-02-12 21:05:06 -0800 (Tue, 12 Feb 2008) | 2 lines

A few extra changes with the math to generate a PDF.

---

r272 | bbogart | 2008-02-14 12:01:39 -0800 (Thu, 14 Feb 2008) | 4 lines

Changes to bib (ann_som) so that the pdflatex can work (underscores are trouble) slight changes "chapter" to "text". Will make a branch of the text in order to remove the meaning stuff to see if that makes sense.

---

r273 | bbogart | 2008-02-14 16:05:39 -0800 (Thu, 14 Feb 2008) | 3 lines

Added another Gabora paper (not read yet) and changed Unicode to latex codes for the umlauts.

---

r274 | bbogart | 2008-02-14 16:11:26 -0800 (Thu, 14 Feb 2008) | 3 lines

Added a branch of MAM chapter without the meaning junk, which is still in "trunk".

---

r275 | bbogart | 2008-02-17 10:34:59 -0800 (Sun, 17 Feb 2008) | 8 lines

```
Even after 1e07 iterations 30,000 sensors simply do not optimize down to
a 144 node network. Things still appear no better than random. Will run
another long term test on sr-00149 with cosine to see if that is any
better, then will see what ratio works, back to 4:3 to start? It is
clear (for thesis?) that ann_som needs to be rewritten for better some
features, like a Gaussian neighbourhood function, and random
initialization of weights.

------------------------------------------------------------------------

r276 | bbogart | 2008-02-17 10:39:03 -0800 (Sun, 17 Feb 2008) | 2 lines

Used MAM learning in some-feeder for long-term testing.

------------------------------------------------------------------------

r277 | bbogart | 2008-02-17 20:42:53 -0800 (Sun, 17 Feb 2008) | 2 lines

Added an image for white for SOMLab testing.

------------------------------------------------------------------------

r278 | bbogart | 2008-02-19 09:09:51 -0800 (Tue, 19 Feb 2008) | 3 lines

Changed the order of the authors of Gem, as they did not show up
properly in the output.

------------------------------------------------------------------------

r279 | bbogart | 2008-02-19 21:28:10 -0800 (Tue, 19 Feb 2008) | 3 lines

I've finally added the PDF to go with the LyX file for the computational
arts chapter!

------------------------------------------------------------------------

r280 | bbogart | 2008-02-20 11:48:50 -0800 (Wed, 20 Feb 2008) | 7 lines

Added .R files for cases 13 and 14. History includes some work with
ggplot, which looks like the ticket (especially faceting!) but I could
not figure out how to use RGB components, so I emailed the author. I
have not generated the test data to see how the SOM behaves for each
iteration, as I have yet, no way to visualize the results. How big is
the neighbourhood really?

------------------------------------------------------------------------

r281 | bbogart | 2008-02-20 11:49:36 -0800 (Wed, 20 Feb 2008) | 2 lines

Added some notes on plotting tiles with ggplot.

------------------------------------------------------------------------

r282 | bbogart | 2008-02-21 00:07:04 -0800 (Thu, 21 Feb 2008) | 3 lines
```

```
Added more ggplot notes, including how to render 30,000 sensors for each
unit. Now I just need to figure out how to facet all the units.

------------------------------------------------------------------------

r283 | bbogart | 2008-02-21 09:27:30 -0800 (Thu, 21 Feb 2008) | 4 lines

Added two new cases where the weights are saved for each iteration to
examine the difference between the cosine and linear learning methods.
(and the small neighbourhood)

------------------------------------------------------------------------

r284 | bbogart | 2008-02-22 12:49:01 -0800 (Fri, 22 Feb 2008) | 6 lines

Lots of work on looking at how the SOM progresses. The lack of Gaussian
neighbourhood and random initialization is probably really hurting it. We
go through about half the iterations with all nodes with the exact same
code-book. Put a R method of measuring that progress (number of
different code-books in a node) into R-Reference. (not committed yet)

------------------------------------------------------------------------

r285 | bbogart | 2008-02-22 12:49:35 -0800 (Fri, 22 Feb 2008) | 2 lines

Added lots of notes on ggplotting, loading som steps and such.

------------------------------------------------------------------------

r286 | bbogart | 2008-02-24 19:58:33 -0800 (Sun, 24 Feb 2008) | 4 lines

Slight changes in working on SOM training. Need to rework the som
external for the next steps. Which for MAM could be complex, since
we only store images on the BMU.

------------------------------------------------------------------------

r287 | bbogart | 2008-02-25 15:22:27 -0800 (Mon, 25 Feb 2008) | 4 lines

3-4 days until the deadline. Fixed lots of typos and inconsistencies. I
hope I got most of them. Very few awkward sentences, could be some comma
splices left around.

------------------------------------------------------------------------

r288 | bbogart | 2008-02-26 11:59:22 -0800 (Tue, 26 Feb 2008) | 3 lines

Got as far as "Memory System" on this draft, will print for Philippe to
look at on break today.

------------------------------------------------------------------------

r289 | bbogart | 2008-02-27 11:06:53 -0800 (Wed, 27 Feb 2008) | 3 lines
```

```
Got down to "machine intended to be creative" changes to math syntax and
little fixes here and there. Ran spell checker again.

------------------------------------------------------------------------

r290 | bbogart | 2008-02-29 09:39:47 -0800 (Fri, 29 Feb 2008) | 4 lines

Added a new simplified diagram of system arch for computational arts
chapter, this one has less detail but more consistent with the text and
gets the idea of the major components.

------------------------------------------------------------------------

r291 | bbogart | 2008-02-29 14:27:36 -0800 (Fri, 29 Feb 2008) | 2 lines

Final work on Computational Arts chapter complete, and submitted.

------------------------------------------------------------------------

r292 | bbogart | 2008-03-12 09:18:48 -0700 (Wed, 12 Mar 2008) | 4 lines

Silly typo means that the installation currently running in installation
in 14-710 will save a som for each unassigned command! "write-som" will
then show the available commands.

------------------------------------------------------------------------

r293 | bbogart | 2008-03-12 09:24:42 -0700 (Wed, 12 Mar 2008) | 3 lines

Added som files from the physical installation to examine the code-book
vectors (to compare to the simple test-case).

------------------------------------------------------------------------

r294 | bbogart | 2008-03-12 11:18:16 -0700 (Wed, 12 Mar 2008) | 4 lines

Added CSthesis latex style files. Now configured for LyX on aporia.
The file is not very LyX compatible, lots of custom ERT. I suppose the
outline will be a start and see how things go.

------------------------------------------------------------------------

r295 | bbogart | 2008-03-12 11:39:12 -0700 (Wed, 12 Mar 2008) | 4 lines

Added an outline for thesis. Uses the sfuthesis style and is more or
less the outline of the computational arts chapter. Have not yet changed
it to take into account the wiki outline.

------------------------------------------------------------------------

r296 | bbogart | 2008-03-12 23:00:35 -0700 (Wed, 12 Mar 2008) | 4 lines

Changes to code for reading SOM images in R.
```

```
Added directory for some_weights, since the data-analysis directory
seems to use up too much RAM!

------------------------------------------------------------------------

r297 | bbogart | 2008-03-14 16:27:37 -0700 (Fri, 14 Mar 2008) | 6 lines

Added files to SVN for the sfuthesis latex style. Made a master document
to store sections, based on the current outline, and filled in content
from Computational Arts chapter were available. Next steps will be to
read through SVN log for future work, and process fodder, and then to
start to fill in more details that did not fit into book chapter.

------------------------------------------------------------------------

r298 | bbogart | 2008-03-20 11:45:23 -0700 (Thu, 20 Mar 2008) | 11 lines

Made a new file for the high-level notes (Notes.txt in SOOS trunk)
called Appendix 1. The Notes were lightly edited and footnotes to
make it make sense to someone other than me were added. Also added a raw
listing of the raw SVN log, without edits. Each of these appendices is
about 10,000 words each, without adding anything, using the SFU thesis
format, the whole thesis is already 76 pages, which is seemingly
stupidly long. Perhaps some of the details from the raw SVN log, could
be included in the footnotes for the other section? Made a slight change
to the bibliography which was putting James Tittle II's name as "II"
only. Hopefully the curly braces help.

------------------------------------------------------------------------

r299 | bbogart | 2008-03-20 15:17:33 -0700 (Thu, 20 Mar 2008) | 3 lines

Some changes to intro, master and section2 towards thesis. Left off in
section 2.1, which needs much more detail.

------------------------------------------------------------------------

r300 | bbogart | 2008-03-20 15:18:33 -0700 (Thu, 20 Mar 2008) | 3 lines

Apparently I made some other changes. Note to self, quit LyX to save all
changed files and THEN commit.

------------------------------------------------------------------------

r301 | bbogart | 2008-03-21 00:44:36 -0700 (Fri, 21 Mar 2008) | 8 lines

Lots of work and changes for plotting feature-maps. Why are the maps
from the installation looking backward/upside down, when the smaller
test patches (shapes for example) are right side up? The memory images
scarcely resemble the code-books themselves. The installation som in
particular is oddly simple considering the amount of data thrown at it,
simple meaning there seem to be few clusters at all. perhaps too the
```

```
neighbourhood being too small?

------------------------------------------------------------------------
r302 | bbogart | 2008-03-26 11:06:57 -0700 (Wed, 26 Mar 2008) | 1 line

More work towards thesis, worked on sections 2 and 3, will rename the
sections without numbers and rearrange them so that the MAM stuff
flows better, with theories of creativity after the intro. Or is it
better to introduce those ideas after MAM in terms of a creative
analysis? Either way renaming the sections will make it easier to
rearrange them.
------------------------------------------------------------------------
r303 | bbogart | 2008-04-01 09:37:11 -0700 (Tue, 01 Apr 2008) | 3 lines

Added a little note about a new idea as how to drive the camera from the
free-association.

------------------------------------------------------------------------
r304 | bbogart | 2008-04-03 11:32:32 -0700 (Thu, 03 Apr 2008) | 3 lines

Made a gradient-free version of the nested R>I diagram, which will look
better in B+W and fits better with the other diagrams.

------------------------------------------------------------------------
r305 | bbogart | 2008-04-03 11:51:46 -0700 (Thu, 03 Apr 2008) | 2 lines

Added a diagram for the viewer's interpretation of meaning of the work.

------------------------------------------------------------------------
r306 | bbogart | 2008-04-04 15:29:17 -0700 (Fri, 04 Apr 2008) | 7 lines

Changed "artist / agent" to "artist" in AP R>I diagram.
Removed xml2latex, as it was not useful for converting the SVN log to
latex.
Got the csthesis format to finally work, choose the APA citation format,
Made an attempt at a new abstract that Thecla was not very fond of. Will
work on that for the next revision.

------------------------------------------------------------------------
r307 | bbogart | 2008-04-04 15:35:26 -0700 (Fri, 04 Apr 2008) | 2 lines

Now there is a PDF in SVN of the whole thesis.

------------------------------------------------------------------------
r308 | bbogart | 2008-04-04 16:44:23 -0700 (Fri, 04 Apr 2008) | 3 lines

Full rewrite of the Abstract for Thecla, made a real effort to talk
```

about the research, over and above the artifact.

---

r309 | bbogart | 2008-04-08 12:10:11 -0700 (Tue, 08 Apr 2008) | 2 lines

Initial changes to the Rokeby section as per his feedback.

---

r310 | bbogart | 2008-04-09 23:54:42 -0700 (Wed, 09 Apr 2008) | 3 lines

Seems the resulting som plot and the som code-books themselves do not match! why? Look again carefully at how images are saved to slots.

---

r311 | bbogart | 2008-04-10 11:23:29 -0700 (Thu, 10 Apr 2008) | 3 lines

Made some changes based on Rokeby's feedback.
Have not heard if this draft is fine to send to the publisher yet.

---

r312 | bbogart | 2008-04-10 11:25:19 -0700 (Thu, 10 Apr 2008) | 7 lines

More work on the chapter, I left off in growing form from context, and added some notes in the creativity theory part. Thecla gave me lots of comments on the current draft and It looks like I'll need to restructure. I'll take a look at her suggested structure and the one Philippe sent out. Make a new outline, get comments and then start filling the content into there.

---

r313 | bbogart | 2008-04-11 18:08:30 -0700 (Fri, 11 Apr 2008) | 18 lines

I made a patch to visualize the learning_results file. These files show the same structure as the visualization itself, which makes sense, but is very different than the code-books. With the 100x75 pixel colour images even after only 100 iterations the SOM is trained, in terms of its code-books, but the visualization does not reflect this. My latest theory is that the variance of the time that images are saved in the SOM means that the structure visualized is actually a jumble of almost all the states the SOM has been through, up until training stops. The new learning result file includes a counter to show how old some of the data is, the variance is very large. The next step will be to make a ggplot where the age of the unit is also visualized, that way I can see that at least the newest units (associated later in training) match the visualization.

Even if my hypothesis is correct, its hard to say what the solution

would be, a visualization based on the code-books combined with the associated values?

---

r314 | bbogart | 2008-04-11 18:09:04 -0700 (Fri, 11 Apr 2008) | 2 lines

Added a function for loading the new learning_result file into R.

---

r315 | bbogart | 2008-04-13 22:03:51 -0700 (Sun, 13 Apr 2008) | 12 lines

Even when plotting the data with the age counter I was unable to say the more recent learning results are more similar to the final visualization. It was difficult to evaluate the age, as alpha and the size of a circle glyph were both not very precise. I tried to make a dataframe from the learning result for ggplotting, but got stuck on merging the grid 12x12 and the 36 nodes that were associated with input patterns. See R-Reference for details, maybe I just need a flesh eye on it. The idea is that the loop should check if the x,y matches the unique index value then put the data-points in the grid at that point, otherwise put NAs in those points. Why can't I figure out how to do it? Should be something like x*12+y.

---

r316 | bbogart | 2008-04-15 11:41:46 -0700 (Tue, 15 Apr 2008) | 3 lines

Made revisions to the section about Rokeby's work based on discussions with him over email.

---

r317 | bbogart | 2008-04-15 18:46:12 -0700 (Tue, 15 Apr 2008) | 4 lines

Added some changes to the R code for loading code-books from ann_som so that values outside of the 0-1 range don't cause the whole loading of the SOM to fail.

---

r318 | bbogart | 2008-04-16 23:57:57 -0700 (Wed, 16 Apr 2008) | 12 lines

So the theory that the lack of structure in the MAM visualization was due to the BMUs reflecting the SOM are multiple levels of training seems very close. The colours chosen for the BMUs are reflected in the weight structure. What is unclear is that I don't see the correlation between the position of the BMU and the some weights at that iteration. Could it be that most weights are the same, so very small (invisible) differences between the weights are making the BMU appear random?

Indeed this behaviour does not seem to happen with the low-res colour. So what is the difference? Confirm that the behaviour is different with another set of SOM steps.

---

r319 | bbogart | 2008-04-17 09:58:37 -0700 (Thu, 17 Apr 2008) | 5 lines

Restructured the outline based on a fusion of Thecla's suggestions and the Pasquier's suggestions. Included revised parts of Rokeby's section to match the computational arts chapter revisions. Added own references to bibliography, PDCON07 and Computational Arts Chapter.

---

r320 | bbogart | 2008-04-18 15:50:22 -0700 (Fri, 18 Apr 2008) | 3 lines

Starting restructuring based on Thecla's feedback.
First crack at a intro done.

---

r321 | bbogart | 2008-04-23 23:30:04 -0700 (Wed, 23 Apr 2008) | 3 lines

Added lots of code for loading and plotting colour (rather than image) SOMs.

---

r322 | bbogart | 2008-04-23 23:38:01 -0700 (Wed, 23 Apr 2008) | 10 lines

1. "instar" appears to be a Gaussian neighbourhood, so no need to implement that. Or perhaps its a linear radius fall-off. Still it optimizes much faster.

2. Its very clear, based on the image to colour som comparison that the BMUs reported by ann_som actually do match the neighbourhoods! So something is broken in the code that associates input patterns to storage nodes. Should be somewhat easy to find as the BMUs are probably just used in the wrong place.

---

r323 | bbogart | 2008-04-28 16:37:22 -0700 (Mon, 28 Apr 2008) | 3 lines

Finished the first draft on the second "artistic practise" chapter.
Added lots of references.

---

r324 | bbogart | 2008-04-30 15:50:24 -0700 (Wed, 30 Apr 2008) | 3 lines

Changed the som.image reading and plotting code to show nodes in the

```
right place!

------------------------------------------------------------------------
r325 | bbogart | 2008-04-30 23:01:06 -0700 (Wed, 30 Apr 2008) | 3 lines

Added some more files for the meeting with Philippe tomorrow. I'm not
sure what is causing what appears to be going on.

------------------------------------------------------------------------
r326 | bbogart | 2008-05-02 20:44:34 -0700 (Fri, 02 May 2008) | 7 lines

So I made a version the same as the 30,000 pixel version, and guess
what, it behaves exactly the same as the 30,000 sensor version, not the
way the colour version works. Try adding that 4th sensor with a "1"
value and see how that works. Also the plotting code for 1x1 pixel SOMs
does not work, check R-reference and I can't get the axis flipping
correct. Appears my strange problems are at the patch level.

------------------------------------------------------------------------
r327 | bbogart | 2008-05-02 20:54:19 -0700 (Fri, 02 May 2008) | 2 lines

This is the results I was referring to in the last log entry.

------------------------------------------------------------------------
r328 | bbogart | 2008-05-02 21:24:10 -0700 (Fri, 02 May 2008) | 4 lines

Did another simplified test using RGB and RGBA 3/4 sensors. They do
behave differently, but they do not show the strange iteration bug.
Perhaps this is a problem with pix_dump?

------------------------------------------------------------------------
r329 | bbogart | 2008-05-02 23:00:21 -0700 (Fri, 02 May 2008) | 6 lines

This looks like the deal-breaker! There needs to be a delay between
saving the image to the buffer and dumping it into the SOM. I choose
250ms arbitrarily, based on the simplified test. Will run a big full
test with colour images with many iterations and see the results in the
morning!

------------------------------------------------------------------------
r330 | bbogart | 2008-05-03 22:14:46 -0700 (Sat, 03 May 2008) | 4 lines

Changed memory system to match the new work done with SOMLAB. including
a delay before pix_dump and INSTAR rule, rather than Kohonen rule.
R-Reference shows some extra notes.

------------------------------------------------------------------------
```

r331 | bbogart | 2008-05-06 19:49:36 -0700 (Tue, 06 May 2008) | 2 lines

This file was not included in the case3 commit.

---

r332 | bbogart | 2008-05-06 19:51:20 -0700 (Tue, 06 May 2008) | 2 lines

These files were also missing from the server copy.

---

r333 | bbogart | 2008-05-06 19:53:09 -0700 (Tue, 06 May 2008) | 3 lines

I'm getting tired of manually adding these files. That's what I get for creating so many and not committing often enough.

---

r334 | bbogart | 2008-05-06 19:53:57 -0700 (Tue, 06 May 2008) | 2 lines

Another missing file.

---

r335 | bbogart | 2008-05-06 19:54:31 -0700 (Tue, 06 May 2008) | 2 lines

Another missing file.

---

r336 | bbogart | 2008-05-06 19:55:11 -0700 (Tue, 06 May 2008) | 2 lines

This entire case was missing from the server copy!

---

r337 | bbogart | 2008-05-06 20:11:20 -0700 (Tue, 06 May 2008) | 4 lines

Added some new cases, and since I'm running out of space removing the old cases from the working copy. To get them just update to the revision before this commit!

---

r338 | bbogart | 2008-05-06 21:52:03 -0700 (Tue, 06 May 2008) | 3 lines

All that mess about some files in the working copy not being committed caused there to be some leftovers.

---

r339 | bbogart | 2008-05-09 14:52:47 -0700 (Fri, 09 May 2008) | 4 lines

```
Lots of additional work on the Creativity Theory chapter. I've made a
draft, spell checked and ran through twice and will send it to Gabora
for (dis)agreement.
```

```
r340 | bbogart | 2008-05-16 10:19:07 -0700 (Fri, 16 May 2008) | 4 lines

Made a really rough draft of next chapter (creative machines) next one
will take more than a week (artistic enquiry / art-research) will need
to get on reading. Getting there slowly.
```

```
r341 | bbogart | 2008-05-23 17:43:09 -0700 (Fri, 23 May 2008) | 4 lines

Failed efforts to methodology chapter. After meeting with Thecla maybe
I'll remove the whole thing, and integrate the process of the work into
the "the artistic artifact" chapter. Would be better for flow that way.
```

```
r342 | bbogart | 2008-05-25 21:08:47 -0700 (Sun, 25 May 2008) | 5 lines

In order to remove the old case files (using up too much disk space) I
need to commit this revision first. I'm just working on a few comparison
runs to show the difference between MAM and normal learning, as well as
a more optimized MAM learning.
```

```
r343 | bbogart | 2008-05-25 21:10:41 -0700 (Sun, 25 May 2008) | 2 lines

Removed a few old cases to save some space, enough?
```

```
r344 | bbogart | 2008-05-29 16:30:31 -0700 (Thu, 29 May 2008) | 5 lines

Lots more work on methodology chapter. Integrated work from old
"artistic enquiry" chapter into "the artistic artifact". Will rename old
chapter to "deprecated" and rename artistic artifact to
"growing-form-from context.
```

```
r345 | bbogart | 2008-05-29 16:42:41 -0700 (Thu, 29 May 2008) | 4 lines

Final changes to rename and rearrange the chapters into the new form. I
feel a lot better about the thesis now, just have to hammer through the
MAM chapter and future work and conclusion and refine refine refine!
```

```
r346 | bbogart | 2008-05-29 16:45:58 -0700 (Thu, 29 May 2008) | 3 lines

Forgot to change this script when rearranging the order of chapters and
removing one.

------------------------------------------------------------------------
r347 | bbogart | 2008-06-04 19:17:39 -0700 (Wed, 04 Jun 2008) | 2 lines

First draft of whole thesis!

------------------------------------------------------------------------
r348 | bbogart | 2008-06-10 17:54:59 -0700 (Tue, 10 Jun 2008) | 5 lines

Changes to thesis up to meeting with Thecla.
Will move the "future work" from chapter to section in MAM.
Form from Context will be the last chapter and sum up the ideas.
That chapter needs to be a lot more concise.

------------------------------------------------------------------------
r349 | bbogart | 2008-06-12 14:29:15 -0700 (Thu, 12 Jun 2008) | 4 lines

More work towards defence.
Spell checked and changed en dashes to em dashes.
not fully checked, but close.

------------------------------------------------------------------------
r350 | bbogart | 2008-06-14 12:21:59 -0700 (Sat, 14 Jun 2008) | 4 lines

This is damn close to being done.
Will spend some time editing the appendices and make minor revisions to
growing form from context and conclusion.

------------------------------------------------------------------------
r351 | bbogart | 2008-06-15 11:34:02 -0700 (Sun, 15 Jun 2008) | 2 lines

Added a note about ideas for "meditative machine".

------------------------------------------------------------------------
r352 | bbogart | 2008-06-15 16:42:17 -0700 (Sun, 15 Jun 2008) | 3 lines

A few more revisions before printing the hard-copy on Monday.
Will regenerate the SVN log and edit that.

------------------------------------------------------------------------
r353 | bbogart | 2008-06-15 16:47:45 -0700 (Sun, 15 Jun 2008) | 2 lines

Added gimp files for free-association figures.
```

Appendix D

SOM Test Data