The 12th White House Papers Graduate Research in Cognitive and Computing Sciences at Sussex

Edited by
Darren Pearce

CSRP 512

November 1999

ISSN 1350–3162

UNIVERSITY OF
SUSSEX
AT BRIGHTON

Cognitive Science
Research Papers

Contents

Preface

The Isle Of Thorns is a village near Haywards Heath and, each year around June, COGS research students spend three days socialising and playing sports. Presentations of research ideas have also been known to take place surprisingly frequently...

This year marked the 12th Isle of Thorns workshop and to complement the presentations, the White House Papers typically collects together recent publications of COGS postgraduate research students. It aims to highlight their research interests and demonstrate what new students may well be doing in the very near future. However, this year sees a change in this approach aiming to make the White House Papers more generally accessible. These proceedings consist entirely of short papers mostly 1–2 pages in length except for the disobedient few... If readers are interested in the material then they can obtain more information and full publications (in some cases) from the relevant author. Also new to this year's proceedings are the inclusion of photographs of the authors as well as biographical information.

We would like to thank Shimon Edelman for the Isle of Thorns funding and Phil Husbands for the funds to distribute this publication. Many thanks are also due to John Halloran who organised the Isle of Thorns Workshop for the second year running. Last and by no means least, it goes without saying that this year's Isle Of Thorns and White House Papers could not have been the success they have been had it not been for all the contributors. Putting up with my incessant hassling and strict submission criteria could not have been easy.

Darren Pearce
COGS Post-Graduate CSAI Student Representative
Compiler and Editor

By way of advertisement...

Biographical Information

Zahra Al-Rawahi is currently a Ph.D. student in the School of Cognitive and Computing Sciences at University of Sussex. She is working in applying Gardner's theory of Multiple Intelligences in connection with the Cognitive Apprenticeship approach in designing computer-assisted learning tools for medical students. She obtained a BSc. degree in Physics from Sultan Qaboos University (Sultanate of Oman) and then went on to a Masters degree in Education (Computer-based learning) from New Mexico State University (US) in 1995. She also worked for two years in medical physics and in laboratory teaching.

Saleh Al-Shidhani graduated in 1992 with a BSc. degree in Physics from the Sultan Qaboos University (SQU). He finished a Master's degree in Astronomy in 1996 at Sussex University. He is currently working towards his Ph.D in CSAI. His research interests spread across several interdisciplinary fields including HCT, ILE/ELE, AI in education, cosmology,...etc

Saleh Al-Shidhani

Adnan M. Al-bar obtained a Bsc degree in Computer Science from King Abdulaziz University in Jeddah, Saudi Arabia in 1990. He then went on to complete an Msc degree in Computer Science in the area of software engineering at The George Washington University, Washington DC, USA in 1997. His research interests include distributed systems, OOP and network programming.

Adnan Al-bar

James Allen graduated from Loughborough University in 1998 with a 2.1 in Human Psychology. He then worked in a shoe shop for 3 months. He subsequently completed the Autumn term at The University of Sussex studying for a Psychology M.A., but transferred to a D. Phil in Psychology in January 1999. He is also a DJ on Brighton Palace Pier.

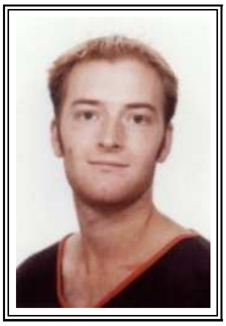

James Allen

Kate Cavanagh studied Psychology at the University of Manchester (1992-1995). She then moved to the University of Sussex - and has been here ever since.

Kate Cavanagh

Stephen Clark graduated from Gonville and Caius College, Cambridge University, in 1994, having started out as a mathematician but ending up as a philosopher. He then worked for two years as a trainee actuary. Actuaries work in the fields of pensions and insurance, using mortality and accident statistics to try and predict future liabilities. After that he studied for an MSc in Cognitive Science at the University of Manchester which led to his current status as a DPhil student in Computer Science and Artificial Intelligence. He specializes in the field of Statistical Natural Language Processing.

Stephen Clark

Glossing over solid working-class background, Robert Clowes took a first degree in English and Philosophy at Keele and unexpectedly Florence University. He spent some time teaching English in Romania before taking a Masters in Applied Informatics at Middlesex and Barcelona Universities and spent about eight months building half-intelligent hyper-instruments before Brian Eno cornered the market. He also taught and wrote half of a Middlesex University course in IT for social scientists. He became an Internet Imagineer for a couple of years but failed unluckily in making his first Million in an AI children's Internet project. He came to Sussex University to do a DPhil in a difficult-to-define subject on the borderline of AI and Cognitive Linguistics.

Robert Clowes

In 1992 Claudia Gama obtained a degree in Computer Science from the Federal University of Bahia (Brazil). She then started an MSc at the Federal University of Rio de Janeiro (Brazil) in Collaborative Systems in Medical Education. From 1993 to 1998 she worked as a lecturer in the Department of Computer Science at the Federal University of Bahia. In 1998 she started a PhD in COGS.

Claudia Gama

Theresa Gannon studied psychology at the University of Birmingham and graduated in 1998. She then continued her studies at the University of Sussex where she is currently researching the role of cognitions in child sex offending.

Theresa Gannon

Miguel Angel Garcia-Ruiz obtained a first degree in Computer Engineering and a Masters degree in Computer Science at the University of Colima (UdeC) in Mexico. Miguel worked in the Image Processing and Multimedia departments at UdeC. Currently, he is in the first year of a PhD in Computer Science and Artificial Intelligence at COGS.

Miguel Garcia

Rory Graves obtained a BSc. (Hons) in Computer Science and Artificial Intelligence at the University of Sussex. He then stayed on to do a DPhil.

Rory Graves

John Halloran was awarded a degree in Photography from Bournemouth and Poole College of Art in 1982. He spent the next thirteen years as a photographer and graphic artist and then as a teacher of English as a foreign language. His first (academic) degree was BSc in Psychology from the Open University (1992-1995). An MSc in Knowledge-based Systems was gained in 1996 from the University of Sussex. He is about to enter the third year of his BT-sponsored DPhil in Human-Computer Interaction at the University of Sussex, due to complete in the year 2000.

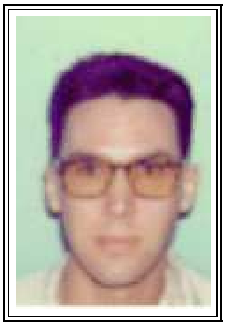

John Halloran

Tudor Jenkins studied for a BA in Artificial Intelligence at Sussex University from 1992 to 1995. He began his DPhil in January 1998.

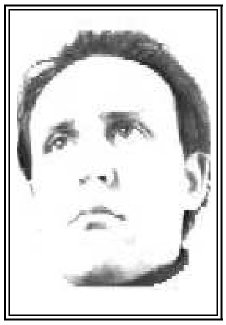

Tudor Jenkins

Kristy Lascelles went on to study Psychology BSc at the University of York straight after obtaining A-Levels in Chemistry, Biology, Psychology and General Studies. She graduated with a 2:1 in July 1999 just after she began her Research Fellowship and DPhil in Evaluative Conditioning at the University of Sussex.

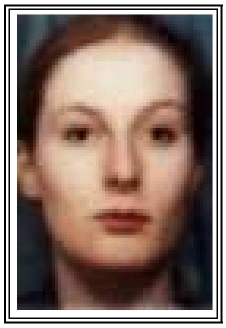

Kristy Lascelles

KK Lau studied Mechanical Engineering at the Hong Kong University of Science and Technology. He then spent 2 years in the manufacturing industry after which he returned to study Psychology at University of Durham in 1997. In 1999, he joined the Psychology group of COGS in the University of Sussex.

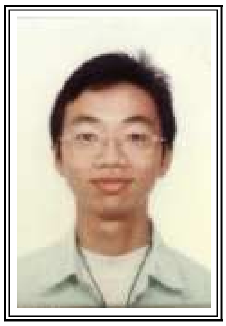

KK Lau

Ann Light is part of the Human Centred Technology (HCT) group in COGS, in the final year of her doctorate on the perception and use of interactive components on websites and how they mediate user/producer relationships. Before starting this, she took the KBS MSc at Sussex and many years ago she completed an English degree here too. She has worked as a teacher and a journalist.

Ann Light

Nuno Otero studied Psychology at Instituto Superior de Psicologia Aplicada, Lisbon. He spent one year working at a training department of an insurance company and afterwards he taught Organisational Psychology at Instituto Superior the Economia e Gestao for one year. His last big work-related decison brought him to COGS to enroll in a DPhil programme.

Nuno Otero

Tim Owen worked as a Mechanical Engineer for 6 years before coming to the University of Sussex in 1995 to do a B.Sc. in Computer Science. He then continued directly onto the DPhil. programme, supervised by Dr. Des Watson.

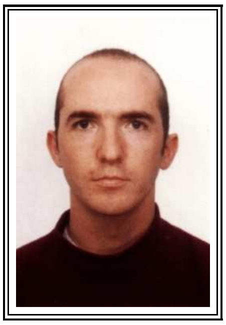

Tim Owen

Georgia Panagiotaki studied Psychology at the university of Crete, Greece. She then completed a MSc in Health Psychology at City University in London.

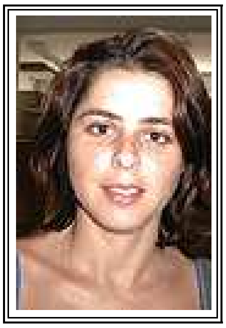

Georgia Panagiotaki

Darren Pearce studied Computer Science at Churchill College, Cambridge from 1994–1997. During the following year he undertook to stay on and sacrifice all social-life and sporting activity to do an MPhil in Computer Speech and Language Processing. In October 1999, he joined COGS in the Natural Language Processing group.

Darren Pearce

Andrew Philippides studied pure mathematics at King's College, Cambridge. After a year off he came to Sussex and did an MSc in Knowledge-Based Systems. He then had another year off before starting his DPhil and is about to enter his third year.

Andrew Philippides

Ahti Pietarinen is a DPhil student in CSAI, under the supervision of Dr Bill Keller. He received an MPhil in Computer Science from University of Turku, Finland in 1997, and an MPhil in Philosophy in the same year. In 1999–2000 he is undertaking fieldwork at University of Helsinki.

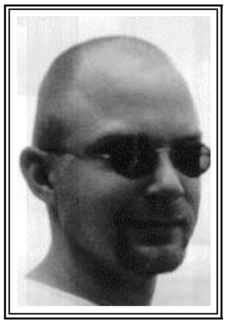

Ahti Pietarinen

Jorge A. Ramirez Uresti studied Computer Systems Engineering and an M.Sc. in Computer Science at ITESM-CEM, Mexico. He worked part-time at the same institution as a research assistant for 2 years while studying for the Master's degree. Afterwards, he worked as a lecturer in Computer Science for 1 year at the same university. In October 1995 he started a Ph.D. in CS & AI at COGS.

Jorge Ramirez-Uresti

Fabrice Retkowsky spent 2 years studying Maths and Physics in Versailles (France), before getting a degree in Telecommunications in Brettony. He then moved to Brighton to do an MSc in Human-Centered Computer Systems in COGS before starting a PhD on "The Cognitive Aspects of Software Reuse".

Fabrice Retkowsky

I did a first degree in Software Engineering at UDLA, Puebla, Mexico. Then it all started to go wrong, I did an MSc. in Knowledge Based Systems at Edinburgh University, and to make things worse I ended up doing a DPhil at Sussex University.
Despite my hopeless late vocational choices, there was a time when I had clear goals. Back in Mexico I worked as a taxi driver and door-to-door salesman among other sensible jobs.

Pablo Romero

Hanson Schmidt-Cornelius studied Computer Science and Artificial Intelligence at the University of Sussex from 1993 to 1996. Then he completed a Masters in Knowledge Based Systems (Intelligent Robotics) at the University of Edinburgh from 1996 to 1997. In 1997 he returned to the University of Sussex and started a DPhil, researching Active Computer Vision Controllers.

Hanson Schmidt-Cornelius

Anil K Seth studies Natural Sciences at King's College, Cambridge, specialising in Experimental Psychology. He worked as a research assistant in Cambridge and Sydney before taking the Knowledge Based Systems MSc at Sussex University. He is currently in the third year of a D.Phil. supervised by Phil Husbands and Hilary Buxton, funded by the EPSRC.

Anil Seth

Nicholas joined COGS as an undergraduate in 1994 studying computer science and artificial intelligence. After graduating in 1997 he took a DPhil position in COGS researching the use of AI techniques in the development of adaptive network protocols.

Nick Sharples

Tom Smith studied Natural Sciences at Cambridge University (1991–1994) specialising in Theoretical Physics. He worked for two years in reasearch groups based at Oxford University, Royal Free Hospital School of Medicine, Institute of Education and OFSTED before taking the MSc in Knowledge Based Systems at Sussex University (1996–1997). He is currently in the first year of his DPhil, after one year working with Dr Phil Husbands on biologically-inspired approaches to neural networks.

Tom Smith

Helen Startup completed a degree in a Applied Psychology from the University of Sussex in the summer of 1997. She is currently undertaking a DPhil with Professor Graham Davey, exploring the psychological causes of catastrophic worrying.

Helen Startup

Andrew Stevens completed his undergraduate degree in Computer Science at the University of Sussex in 1995. After two years of working as a Software Engineer he returned to Sussex to start a DPhil., specialising in compiler optimisations for Java. He is about to enter his third year of studies and hopes to finish his research next year.

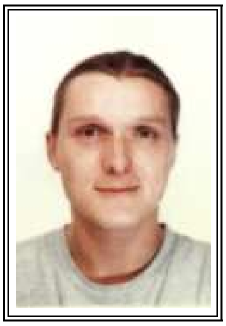

Andrew Stevens

Sian Williams graduated from Nene College, Northampton in 1996 with a II:i in Behavioural Science. She then worked for a year to finance a year travelling through South-East Asia, Australasia and the Pacific Islands. She is currently going into her second year of a DPhil in COGS for which she receives a demonstrator bursary.

Sian Williams

Virtual Hospital Round: A Cognitive Tool for Clinical Teaching

Zahra Al-Rawahi

zahraar@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

This paper investigates the use of an intelligent tutoring system to overcome problems with clinical teaching. Although, there are a number of medical diagnosis expert systems which have been designed to help medical students and medical practitioners in deciding about diagnosis, little of these systems studied students difficulties with clinical diagnosis. This paper will report on some of these problems such as: the anchoring problem, forcing the diagnosis and endless enquiry.

Whilst a number of systems have been developed based on a cognitive apprenticeship approach, and a significant body of research has been carried about learning with multiple representations, nothing is known about the effectiveness of multiple intelligences (MI) on students' learning. To address this, a prototype was built and some experiments will be performed to study the effectiveness of MI in comparison with multiple representations. Gardner's theory of Multiple Intelligences and cognitive apprenticeship approach have influenced the design of the Virtual Hospital Round (VHR). The VHR is a practice based learning tool in Cardiovascular diseases.

The VHR prototype allowed for a unique blend of intelligences (learning styles) in each student, and assessed their development. It enabled students to explore a topic according to their learning styles preferences but simultaneously encouraged them to conduct their interaction on a more systematic basis. Also as both approaches used in VHR are based on the belief that learning generally takes place in the context of social interaction, we are going to study students' ability to transfer their learning through sharing their experience by being a reviewer or a critic of each other's learning.

Investigating the Use of the Laurillard MoTLP as a Pedagogical Framework for the Cosmic Simulator

Saleh Al-Shidhani

salehsh@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

The Cosmos-Simulator is a proposed interactive system as apart of a simulation based interactive learning environment (SbILE). The SbILE would provide an interactive means of simulating cosmological models of the universe, exploring and studying the outcome in depth and quantifying the results. The system is aimed to promote conceptual understanding of the cosmological principle and confront several misconceptions associated with the expansion of the universe as a domain.

The Expansion and Evolution of the Universe are characterized (Al-Shidhani, 1998) as follow:

1. an open-ended process that depends on the initial physical conditions and the forces involve.

2. a complex and dynamic process that changes with time and results from continuing interactions.

3. highly speculative and a controversial topic in Astronomy without direct observations, and described by several theoretical models.

Consequently, many students have difficulty in understanding the implications of the cosmological principles of homogeneity and isotropy because they need higher order spatial reasoning and imagination. This reasoning and imagination need to be transformed from the highly speculative to a solid, coherent, and more realistic processes.

The Laurillard Model of the Teaching-Learning Process (MoTLP) is concerned with promoting conceptual understanding and hence was chosen as a pedagogical framework of the system. The model (shown in the diagram below, after Laurillard, 1993) represents the various activities between a student and a teacher (the system in our case) and the stages of conceptual learning via experiencing constructed worlds. This general model provides a framework within which to view the Cosmos Simulator; for example the system presents the learner with a concept and through the interactive interface the learner can express his/her view and understanding of that concept. Consequently, the system should pick a relevant a simulation outcome as a ”constructed world” for the learner to explore and validate the presented statement, or conceptual implication. Laurillard, 1995 stated that ”Interaction between the learner and the world is a vital part of the learning process because it is this that situates academic knowledge in our experience of the world”.

The other advantage of this model is that it considers the exploration of the simulated world as a process parallel to constructive discussion between the teacher and the learner. Moreover, it outlines the information processing tasks carried out by the learner and the valuable trace that the teacher needs to tune his/her feedback and discussion. The main disadvantage of the model is that it is descriptive and presents no mechanisms by which one could make the current media be reflective, discursive and adaptive. Some of the suggested model’s activities are very challenging to implement; e.g. the dynamic generation of teachers reactions, specially the one depending on understanding the learner’s conceptions and how they relate to the ”target” conception.

Here we propose two frameworks, the first addresses learners’ naive misconcepts by providing online explanatory information with argumentative text and visual illustrations of the concept and its implication. The second framework addresses the more profound difficulties in envisaging the above concepts. The knowledge rectification scheme constitutes of the following procedure: a) review the key-points of the concepts (i.e. ensure no lack of knowledge), and discuss the implications and cause/effect relations. b) suggested exploring some simulation outputs, c) ask the learner to compare antithetical simulation outputs and to check the validity of related hypothesis.

In conclusion, Laurillard Model of the Teaching and the Learning Process has a lot of potential as a framework for the cosmic simulator but at the same time the model implementation is very demanding on codes to achieve a reasonable level of interactions, discussion, response analysis, and adaption. As a result the implementation of such system for addressing the conceptual understanding of a complex domain requires a lot of time and efforts.

References

Al-Shidhani, S. (1998). *The cosmos simulator.* (Presented at The Isle of Thorns (IOT98) Postgraduate Workshop, COGS, University of Sussex)

Laurillard, D. (1993). *Rethinking university teaching: A framework for the effective use of educational technology.* London, UK: Routledge.

Laurillard, D. (1995). Multimedia and the changing experience of the learner. In *Programme on learner use of media.* Institute Of Educational Technology, Milton Keynes. UK: Open University. (Paper No. 52)

Distributed Directory Services for Active Network

Adnan Al-bar

aalbar@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

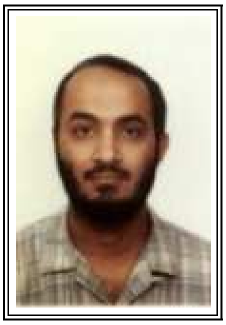

Here is a very general idea of what is my research interest is all about.

The first time I was exposed to network operating system was 7 years ago when I and a colleague tried to setup a network for the Telecom college in which we were both teaching. We installed Novell 3.11 as our network operating system. We had a lot of trouble trying to figure out what is going on. Since then, I got interested in networking in general and network programming in specific.

Now, here I am in the School of Cognitive and Computing sciences, Cogs, trying to study and research on an advanced topic in network programming. I am interested in distributed directory services and active network. How can we use active network to improve directory services over our networks. In another word, *how active network can improve distributed directory services?*.

Another thing I wanted to do is to see how that can be applied to a running project in the network lab group which is *SafetyNet*. To be able to do that, I need to do the following for the rest of the year in. While doing that I think I will be able to get the topic off the ground and then submit a *proposal* about it around Jan. 00.

1 Reading

1.1 Literature survey and review of

1. Distributed system theory
2. Distributed Directory Services
3. Active Network
4. SafetyNet project

2 Programming Skills

1. Doing more with Java
2. Network programming
3. Jini technology

3 Courses

3.1 Autumn term 99

1. Advanced Programming
2. Distributed System
3. Research methods

3.2 Spring term 00

1. Technical communication skills

Unconscious processing in anxiety and the repressive coping style

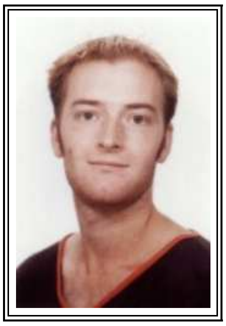

James Allen

jamesa@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Background Theory

My research follows directly on from Derakshan's (Derakshan & Eysenck, 1997, 1999), in that there was evidence from her work that the strategies used by repressors to avoid threat are below the level of conscious awareness. Thus I am trying to identify the cognitive mechanisms involved in the unconscious processing of threat related material. A cognitive approach examining such unconscious biases for threat related information is crucial to understanding the cognitive mechanisms of anxiety, and may give insight into the cause of the RCS.

2 What I have achieved so far

2.1 Experiment 1: Unconscious detection of emotional content words Allen and Derakshan (1999)

I have just finished running an experiment, which presented various emotional content words below the level of conscious awareness (20 ms) and then asked participants to try to recognise the word by typing it. The words were masked with random letter fragments for 10ms before and after the 20ms presentation time to make the task even more difficult. The sample used were 53 participants identified as extreme scorers on self report measures of trait anxiety and defensiveness, i.e. the four-fold classification.

The words were divided into neutral, positive, sad, anxiety and angry words. Threat words are the combination of sad, anxiety and anger words. Words were presented in an order randomised for each participant. If the word was not typed in correctly then 20ms was added to the presentation time. This way, thresholds for various words could be established, for example a person correctly identifying the word 'love' after three trials would have a threshold of 60ms.

The analyses looked at the number of tries required to correctly identify each word, which is really the same as looking at the threshold just described. It was found that the anxiety groups did not differ on the number of tries required to identify positive or neutral words, although positive words were identified significantly faster than all other word types. This last finding is in line with previous literature.

It was hypothesised that the four anxiety groups would differ on the number of tries required to correctly identify sad, anxiety and anger words, specifically that LA would take longer than HA. We were also interested in how repressors would respond: would they act like high or low anxious or would they act differently to all other groups? Analyses examining whether the groups differed on sad, angry and anxiety words separately was not significant. The pattern of responses was, however, similar for each group of threatening words (sad, anxiety, anger) which led us to combine them to form a 'threat' group of words.

There was a visual trend indicating that LA and REP groups were taking longer than HA and DHA to identify threatening words. To ascertain whether the differences were significant, a repeated measures ANOVA was carried out using a polynomial contrast for word type (positive, negative) and a simple contrast for anxiety group to ascertain how large this difference was. The main effect for anxiety group was non significant, although a contrast comparing the HA and DHA with LA and REP on threatening words was approaching significance, indicating that LA and REP were detecting threatening words slower than the other two groups. The implication of this result is that REPs act like the LA at an unconscious level of detecting threatening words, showing an avoidance pattern. The high degree of similarity between LA and REP, and HA and DHA is of particular significance.

3 What I aim to achieve (probably in the next year)

3.1 Experiment 2: Unconscious processing of anxiety related words using an Affective Priming Paradigm (APP). Allen and Davey (1999)

I have just finished writing a specification for an APP experiment. The design of the experiment borrows heavily from the work of Hermans, De Houwer & Eelen (1994), who employed an affective priming paradigm. The authors replicated the procedure employed in the seminal paper of Fazio, Sanbonmatsu, Powell & Kardes (Fazio, Sanbonmatsu, Powell, & Kardes, 1986) who demonstrated that the time needed to evaluate target words as positive or negative decreased if they were preceded by a similarly valenced prime word, but increased when preceded by a prime of opposite valence. The authors related the findings to Bargh, Chaiken, Govender & Pratto's (Bargh, Chaiken, Govender, & Pratto, 1992) allegation that the automatic activation effect is a pervasive and unconditional phenomenon.

The current experiment will use anxiety related words (e.g. personal and social threat, positive, negative and neutral words). It is unknown whether REPs will maintain their defensive strategies and act similarly to the LA or whether REPs will behave similarly to HA. To clarify a possible discrepancy between unconscious processing of anxiety and physiological reactions, a heart rate monitor will be used concurrently with presentation of the stimuli.

References

Allen, J., & Davey, G. (1999). *Unconscious processing of anxiety related words using an affective priming paradigm (APP).* (Forthcoming)

Allen, J., & Derakshan, N. (1999). *Unconscious detection of emotional content words.* (Forthcoming)

Bargh, J. A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the attitude activation effect. *Journal of Personality and Social Psychology*, *62*, 893–912.

Derakshan, N., & Eysenck, M. W. (1997). Interpretive biases for one's own behaviour and physiology in high-trait anxious individuals and repressors. *Journal of personality and Social Psychology*, *73*, 816–825.

Derakshan, N., & Eysenck, M. W. (1999). *Manipulation of focus of attention and its effects on anxiety in high trait-anxious individuals and repressors.* (Submitted)

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, *50*, 229–238.

The Development of Fears and Phobias

Kate Cavanagh

katecav@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

Specific phobias are extremely common in the general population. Epidemiological surveys indicate that more than 11% people suffer from an intense and disabling phobia at some time in there lives. There is considerable evidence from both laboratory and clinical studies that specific phobias (as well as other anxiety disorders) are underpinned by expectations and beliefs about aversive or traumatic consequence of interacting with feared stimuli. However, there is very little research into either the kinds of expectations and beliefs phobics have about stimuli other than their phobic object, or how the exaggerated beliefs phobics evidence might arise. My research in COGS tackles these two issues.

Examples of my current research include the following study:

2 Background

Recent research has indicated that some animal phobias (spiders, bugs, snakes, rats) might be associated with disgust.

3 Aim

Does experience of disgust increase peoples expectations of aversive outcomes when encountering animal stimuli?

4 Method

Participants completed a simple measures of fear and outcome expectations before and after undergoing an autobiographical-recall mood induction procedure which induced disgusted or neutral mood.

5 Results

Following the disgust mood induction procedure participants reported significantly greater fear and significantly greater aversive outcome expectations of spiders and bugs. No differences were found following the neutral mood induction procedure.

6 Conclusions

These findings indicate that experienced disgust might play an important role in inflating expectations of aversive and traumatic outcome following encounters with some animals. This adds to our understanding of the development of common and pernicious animal fears.

A Probabilistic Approach to Sentence Disambiguation

Stephen Clark

stephecl@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

My research is in the area of Natural Language Processing (NLP). Rather than try and describe NLP in one line myself, I will give this description from a well-known text book in the area (Natural Language Understanding, James Allen):

> The goal of [NLP] research is to create computational models of language in enough detail that you could write computer programs to perform various tasks involving natural language.

Examples of tasks we might want a computer to perform are translating sentences from one language to another, retrieving information from a large body of on-line text (such as the WWW), checking spelling and grammar, and interpreting commands written in natural language.

In order to do something useful with a sentence of natural language, a system must have some idea of what the sentence means; it must have some kind of meaning representation of the sentence. The central problem in arriving at a meaning representation for a sentence of natural language is the problem of ambiguity. Ambiguity pervades the whole of natural language, at all levels of analysis. To give some examples, there is ambiguity at the word level, as some words have more than one meaning; there is ambiguity at the part of speech level, as some words belong to more than one syntactic category; and there is ambiguity at a sentential level, as complete sentences can often be interpreted in more than one way. It is ambiguity at a sentential level that I am concerned with. The aim of my research is to give a system some knowledge which it can use to disambiguate sentences.

An example of an ambiguous sentence is the following:

Fruit flies like bananas

Presumably we are able to rule out the unlikely interpretation for this sentence because we know that fruit doesn't normally fly, and things that do fly don't fly like bananas. It is exactly this kind of knowledge - what kinds of things tend to appear in particular argument positions of particular verbs - that I try and get a system to use in order to disambiguate sentences. I get the system to acquire this knowledge automatically, by scanning through a corpus, a large body of online text, and counting the number of times that certain things appear as the arguments of verbs. These counts can then be used to estimate probabilities (such as the probability that a kind of fruit appear as the subject of *fly*), and these probabilities can be used to give a score to each potential interpretation of a sentence. The hope is that the interpretation with the highest score is the correct interpretation, or at least that the interpretation with the highest score is close to the correct interpretation.

A Comparative Approach to Language Trained Apes and Self-Organising Language Games. Steps to Building Perceptual Symbol Systems - or Tudor's Dream (Jenkins, 1999)

Robert Clowes

robertc@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

My research starts from the premise that symbol systems derive their particular nature from being simultaneously communicative and cognitive (Vauclair, 1997). A hypothesis I attempt to evaluate through the prisms of cognitive ethology, situated robotics and language trained apes. The hope is to provide some empirical evidence for this hypothesis by modelling communicative perceptual symbol systems with simulated and robotic agents using inspiration from ape language training research. Preliminary work exploring this hypothesis is discussed.

Traditional AI is largely concerned with the properties of physical symbol systems and how researchers have tried to use them to build cognitive systems with such properties as inferential power, productivity and the ability to formulate truth conditional propositions; in other words humanlike systems. Modeling these systems has run into a number of problems however. A primary problem with these symbol systems has been, be that through transducing knowledge from sensory apparatus into symbols, grounding it back again in motor activity, or reasoning with any degree of alacrity about anything other than toy scenarios.

Parallel to the history of AI the semiotic tradition which finds it origins in the work of Peirce and Saussure, tries to explain symbols systems not primarily as cognitive systems but as systems of representational communication. The question I am trying to tackle is whether the theoretical account of these communicative symbol systems can provide new insight into the grounding of cognitive symbol systems. One way of looking at these symbol systems is to look at the borderline cases, and apes are a good example.

Much recent literature in ape language training has focused on what types of communicatory systems it is possible to teach apes. First of all are these systems symbolic, and then, do these newly taught communication abilities enable other cognitive abilities, enhancing perhaps problem solving strategies or the voluntary direction of attention. The various possibilities about what might be going on here are complex and contentious, but a major hypothesis is that learning a symbolic communication system may allow the representation of knowledge in a new and more abstract representational structure (Premack, 1983). Perhaps these proto-words we find that have been taught to apes are in effect mind tools (Dennett, 1996); ways of building symbolic processing capacity.

Recent work at the VUB AI lab and the Sony Research Laboratories in Paris, under the direction of Luc Steels have attempted to model some of the ways these types of communication games can be born, in communities of simulated and robotic agents (Steels, 1996a, 1996b; Steels & Vogt, 1997). The main drive behind these experiments is to show that an ontology and a communicatory system can mutually bootstrap each other into existence; the process by which this happens having been called discrimination, and language games. The dynamics of these games allow the agents to bootstrap both a referential communication system and share a grounded representational system which encodes facets of the world which they inhabit.

The robots and the apes are of very different behavioural and cognitive orders; and we must be careful about drawing false parallels, yet there seems nonetheless to be some value in adopting a comparative approach. Particularly in using the ape research to indicate how we might begin to inculcate some sort of symbolic order in the work on agent simulations. I intend to exploit these parallels to push the agent modelling work forward. Particular research focuses are how syntactic structures may build from symbolic reference (Savage-Rumbaugh, 1986) and how these may be used to enhance comparison operations that unaugmented apes seemingly can't achieve (Premack, 1986).

In the following year I will be concentrating on these parallels in the simulation work I will be pursuing. My continuing research aims to create a roadmap toward the integration of communicative and perceptual symbol systems (Barsalou, 1999) through building communicating problem solving agents. The aim ultimately is to instantiate in a group of interacting agents a communicatory symbol system that can be pushed toward a fully cognitive one.

References

Barsalou, L. W. (1999). *Perceptual symbol system.* (Behavioral and Brain Sciences (Forthcoming))

Dennett, D. (1996). *Kinds of minds.* Phoenix Books.

Jenkins, T. (1999). Evolving communication mechanisms in a complex environment. In D. Pearce (Ed.), *The 12th white house papers.* University of Sussex, Brighton.

Premack, D. (1983). The codes of man and beasts. *Behavioral and Brain Sciences*, *6*, 125–167.

Premack, D. (1986). *Gavagai! or the future history of the animal language controversy.* MIT Press.

Savage-Rumbaugh, S. E. (1986). *Ape language. from conditioned response to symbol.* New York: Columbia University Press.

Steels, L. (1996a). Perceptually grounded meaning creation. In *Icmas.* Kyoto.

Steels, L. (1996b). Self-organising vocabularies. In *Alife v.* Nara Japan: The MIT Press, Cambridge MA.

Steels, L., & Vogt, P. (1997). *Grounding adaptive language games in robotic agents.* Cambridge MA: The MIT press.

Vauclair, J. (1997). *Primate cognition: From representation to language.* —.

The Reflection Assistant: Improving Reflection and Metacognitive Skills in Interactive Learning Systems

Claudia Gama

claudiag@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

Abstract This research proposes the design, development and evaluation of a Reflection Assistant to be used with computer-based learning environments for complex domains (such as medicine). The main goal is to develop students' metacognitive and self-monitoring skills through interactions with the Reflection Assistant during problem-solving activities. The purpose is to investigate better means to improve metacognitive knowledge, including awareness on the learning process and capacity of transferring skills to new learning situations.

1 Introduction

The argument that drives this research is that metacognition plays a major role in cognitive development (Flavell, 1979; Derry & Hawkes, 1993). Enabling students to develop conscious, explicit model of their metacognitive skills by means of reflective activities, should facilitate the improvement of both cognitive and metacognitive expertise.

In complex domains the student needs to organize a series of different facts in a coherent knowledge structure, to analyse situations, generate hypotheses and decide paths of solutions. Reflective activities encourage students to analyse their performance, contrast their actions to those of others, abstract the actions they used in similar situations, and compare their actions to those of novices and experts.

The general question that leads this research is: *How metacognitive knowledgee is best learned?* More specifically, the research aims to investigate the following points:

1. In which moment is it more appropriate to provide metacognitive scaffolding and reflection on the problem-solving process?
2. Which balance between reflection on learning task (very focused on the task) and reflection on general metacognitive skills (broader view) is more effective for a particular situation?
3. In which extent the metacognitive skills are applicable across tasks or are specific to domains or content areas?

2 The Reflection Assistant

The Reflection Assistant is designed as a general tool that can be configured accordingly to the specific domain that it will be applied. The general architecture is shown below (Figure 1).

To provide a more flexible use the main components contain several parameters that can be modified depending on the circumstances under which the assistant is going to be used. Besides these components, a communication mechanism (or protocol) that makes possible the integration with learning systems is defined.

3 Conclusion - Research Methods

As a first step, the Reflection Assistant will be developed to be used with an existing medical case-based system (named PATSy). PATSy (Patient Assessment Training SYstem) is a case-based learning environment that keeps a database of medical cases in the domain of aphasiology (see (deductive reasoning, patsy system, medical education, 1999) for more details). The insights obtained from this study will be used to identify key components of the Reflection Assistant that can be applied in other complex learning environments.

References

Cox, R. (1999). Modelling the diagnostic reasoning skills of expert clinicians and students

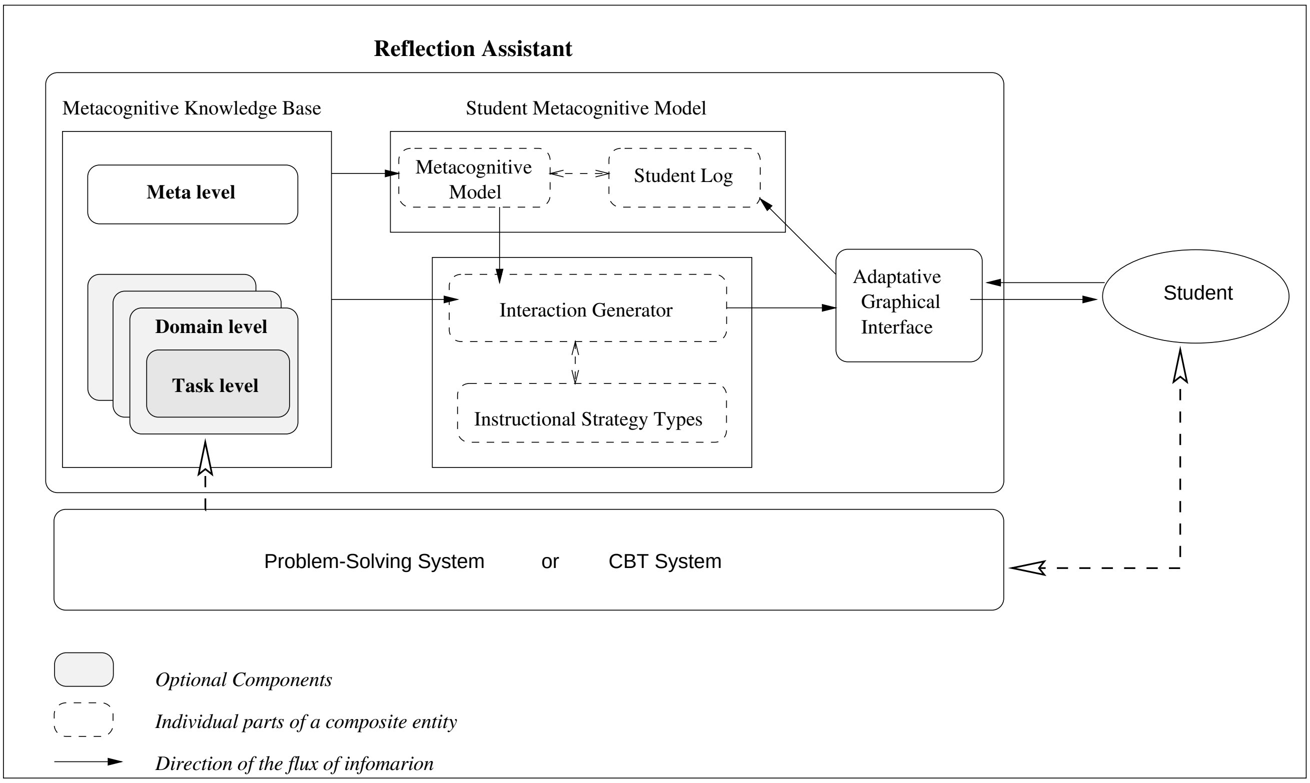

Figure 1: General Architecture of the Reflection Assistant

in the domain of aphasiology. In P. Brna, M. Baker, & K. Stenning (Eds.), *C-lemmas roles communicative interaction in learning to model in mathematics and science.*

Derry, S. J., & Hawkes, L. W. (1993). Local cognitive modeling of problem-solving behavior: An application of fuzzy theory. In S. P. Lajoie & S. J. Derry (Eds.), *Computers as cognitive tools* (pp. 107–140). Lawrence Erlbaum Associates.

Flavell, J. H. (1979). Metacognition and cognitive monitoring. *American Psychologist*, *34*(10), 906–911.

The Role of Cognitive Distortions in the Perpetration of Sex Offences

Theresa Gannon

theresag@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

A major component inherent in theories of sex offending is that of cognitive distortions'. This was a term coined by Abel, Becker, and Cunningham-Rathner (1984) which they use to refer to the justification of behaviour typically displayed by child sex offenders. Examples of cognitive distortions' proposed by Abel et al. (1984) include the following:

- The lack of physical resistance of a child indicates a willingness to engage in sexual relations with an adult.
- The sexual interaction with an adult is educational for the child.
- Children do not tell others about the secret activity because they enjoy it.
- No harm to the child is brought about through merely fondling the child.
- Children's questions about sex are really an indication that the child would like to engage in sexual relations with the adult.
- Child-adult sexual relations will be acceptable to society in the future.

Many current treatment programs for child sex offenders actually attempt to alter this distorted thinking. Indeed, Morrison, Erooga, and Beckett (1994) argues that cognitive-behavioural treatment is the most popular method of treating child sex offenders in the UK. The primary aim of this treatment is to eradicate the offender's distortions through educating offenders on the connections between distortions and behaviour, challenging the distortions and engaging in role reversal exercises.

Although cognitive distortions' are continually targeted in therapy, very little is known about the offender's cognitions. A key problem, still unresolved is separating out what sex offenders actually believe as opposed to what they say. One of the main problems with the theory of cognitive distortions' is that existing evidence in favour of such distorted cognitions is based upon clinical observations of such justifications during assessment (e.g. Abel et al. (1984)). To date, there are no standardisable procedures for assessing convicted sex offender's beliefs. Indeed, it is possible that cognitive distortions' may actually be the end product of sexual offending as opposed to the cause and maintenance of it.

Memory is an important aspect of self, which provides valuable insight into underlying cognitive structures. Bartlett (1932) was the first to illustrate that the recall process did not support the previous assumption of a rigid and lifeless trace (Oldfield, 1954). In a series of experiments requiring subjects to recall a passage created from a differing cultural perspective, it was concluded that literal recall was only produced in exceptional circumstances.

Further evidence suggesting that misremembering is a natural and informative process is proposed by Neisser (1981). Neisser conducted a detailed case study on the memory of John Dean who was counsel to President Nixon during the Watergate scandal. A comparison of Dean's testimony with recorded conversations revealed that Dean's memory was a mixture of systematic distortion and basic accuracy. Neisser's study reveals that even when Deans's statements were inaccurate there was a strong sense that he was being truthful. A comparison of Dean's testimony with the transcript for the particular meeting of September 15 1972 illustrated that Dean's testimony of the meeting was not true' on comparison with the actual conversation. For example, Dean recollects that the president asked him to sit down and complimented him on doing a good job, however, the transcript reveals no such conversations. Neisser (1981) comments that cross examination of the testimony indicated that Dean was doing the best he could to be honest so where did these memory distortions originate?

Neisser (1981) argues that there are two main avenues of explanation. Firstly, many of the distortions can be explained through the use of common scripts'. These are everyday scripts which are known by the majority of people and may be used to aid recall. In this instance, Dean may have been using this script when he remembered that the president asked him to sit down (which was not mentioned on the transcript, although it is possible that a simple gesture was employed). Secondly, many of the distortions displayed by Dean may have been a reflection of Dean's own character and needs. In this case, Dean's testimony represents his own preference for what he thought should have happened (for example, he may have thought that the president should have praised him for his work). Important implications arise from Dean's testimony; Dean's internal ambitions distorted his memory for the events that took place. Apparent attempts to tell the truth are met with distorted recollections of the inflated importance of Dean in every affair. Despite this exaggeration of events and roles, Dean retained the common thread of the situation across events. He did not distort this aspect of the testimony but merely embedded this theme within slightly different events. Dean's memory did not retain the events in a rigid fashion, but moulded them around his own character whilst producing the basic message of the Watergate scenario.

2 Study One

2.1 Methodology

Study one will use this information on memory and schemas to assess the underlying beliefs of sex offenders prior to treatment. Vignettes have been constructed which depict various sexual offences against children. These will be presented to the offender and they will be asked to comment on the thoughts, and reactions of the perpetrator, before, during and after the offence. However, this procedure cannot tell us whether the assessment reveals the true nature of the offender's beliefs. Therefore offenders will be issued with a surprise memory task asking them to recall the original vignettes.

Analysis of the results will proceed at two levels. Firstly, a comparison will be made at stage one in which similar and different crimes will be compared to reveal possible differences in the offender's cognitions. Secondly, analysis will involve examining possible drift in memory over the two week interval. It is proposed that offenders will recall ambiguously presented information to fit in with their existing schemas. Therefore, if these schemas are distorted, this should be reflected in the information they recall in the surprise memory task. A comparison of these stages may provide a way of tapping into differences between the offender's over edited utterances and tacit underlying beliefs.

The proposed methodology will be informative in a number of ways. It will produce a standardisable procedure for assessing sex offender's beliefs which will be more reliable than anything else currently in use. Consequently, this study will provide valuable evidence for existing models which presume the existence of cognitive distortions' in the absence of empirical evidence.

Following the successful development of this methodology, it will be used to evaluate the offender's beliefs post treatment. This stage will involve recruiting the same offender's following satisfactory completion of cognitive-behavioural therapy. In order to indicate whether the treatment has affected the offender's beliefs, it is necessary to repeat the procedure outlined earlier to assess the infomation recalled after a period of two weeks. If the offender's show evidence of recalling information to fit in with the cognitive distortion' schema (whilst displaying a lack of these distorted attitudes in the initial assessment of the vignettes), it will illustrate that the treatment program is only affecting the offender's statements and not the underlying beliefs.

The methodology will then be used to comparatively evaluate different types of treatment regimes currently in practise. For example, a comparison could be made of the improvement in belief system of offenders who have completed cognitive-behavioural therapy with offenders who have completed other treatment regimes. On identification of the most successful type of treatment, it would be essential to use the proposed methodology to identify which parts of the treatment program are most sucessful at altering the beliefs of sex offenders. These parts could then be extended for future programs to ensure the successful treatment of sex offenders. The potential benefits of such advances are far reaching for offenders, potential victims and, ultimately, society as a whole.

References

Abel, G., Becker, J., & Cunningham-Rathner, J. (1984). Complications, consent and cognitions in sex between children and adults. *International Journal of Law and Psychiatry*, *7*, 89–103.

Bartlett, F. C. (1932). *Remembering: a study in experimental and social psychology.* Cambridge: University Press.

Morrison, T., Erooga, M., & Beckett, R. C. (1994). Cognitive behavioural treatment of sex offenders. In *Sexual offending against children: Assessment and treatment of male abusers*. London: Routledge.

Neisser, U. (1981). John dean's memory: a case study. *Cognition*(9), 1–22.

Oldfield, R. C. (1954). Memory mechanisms and the theory of schemata. *British Journal of Psychology*, *45*, 14–23.

An Interactive 3D Learning Environment for Protein Chemistry

Miguel Garcia

miguelga@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

Abstract Students of introductory chemistry courses learn intricate three-dimensional organic molecules, especially amino acids and proteins, through abstract representations. Traditional techniques for teaching molecular chemistry still have some limitations. This research project, which is currently in progress, poses the development of a 3D tutoring system, which will involve virtual environments (VE) technology and an intelligent agent for teaching basic information on amino acids and proteins.

The main purpose of this research project is to assess the presence and interaction between the student and a *virtual learning companion* which will be a simulated pedagogical agent within the virtual environment. Both student and companion will personify avatars (virtual personifications). The companion will give the student support and exchange of ideas in constructing and inspecting virtual molecules, given by the tutoring system as an exercise to be solved. Also, the companion will act as a student with equal or more knowledge on the domain than the real one.

1 Introduction

One of the major concerns in teaching introductory chemistry courses is the elucidation of abstract information, such as three-dimensional structures, formulae, components and functions of amino acids and proteins. Thus, conventional courses have been taught using either wooden or plastic 3D molecule models, textbooks, photographs, posters, and transparencies, amongst other didactic materials, with acceptable results (Jones, 1996).

Also, Television programmes, taped chemistry video courses, interactive CD-ROMs and the Internet have helped teachers in conveying molecular chemistry concepts with fairly good results. However, most of these techniques are one-way, and the information is limited only to the contents of the video clip.

Nowadays, novel techniques like virtual environments can be used to allow the students to manipulate and visualise a virtual molecule as active participants, at their own pace.

1.1 Definition of Virtual Environments (VE)

As Kalawsky (1993) explained, virtual environments (VE), also known as virtual reality, is a computer system that generates a three-dimensional graphical ambient known as virtual world, where the user experiences an effect called immersion (the sense of presence within the VE world), and he/she navigates through the virtual world and interacts with the graphical objects that reside within it, using special input/output devices.

2 Statement of the Problem

According to Swift and Zielinski (1997), chemistry textbooks, molecule sketches drawn on blackboards, and three-dimensional molecule scale models have important constraints in chemistry teaching, such as lack of interactivity, no 3D spatial representation in some techniques, and inaccuracy of represented molecular structures.

A word of caution, sometimes it is better to use planar (2D) representations of information, such as diagrams, for explaining chemical formulae and certain physical properties of chemical components (Scaife & Rogers, 1996).

The main purpose of this research project is to assess the presence and interaction of a *virtual learning companion*, which will be a pedagogical agent that will reside within the VE world. It will give support and guidance to the student in constructing and inspecting a virtual molecule. Thus, a research study will be conducted to assess the effectiveness of a 3D learning environment for protein chemistry, involving a simulated agent that will act as a student companion with equal or more knowledge on the topic than the real student.

3 The Implementation of an Intelligent Tutoring System and a Learning Companion

An intelligent tutoring system (ITS) can be defined as a computer system based on artificial intelligence techniques that tutor a student on a specific knowledge domain (Preece et al., 1994; Yazdani, 1987).

As Chan and Baskin (1988) commented, a Learning Companion System (LCS) is the equivalent to an intelligent tutoring system (ITS), whereas the LCS involves a pedagogical agent, or learning companion. The agent can learn from the real student, and also the student can explain concepts to it (hence, the student learn *by teaching*). Both the student and the learning companion are engaged in solving problems proposed by the tutoring system, exchanging ideas and strategies on how to solve them. The learning companion can be *weak* or *strong*, depending on the extent of knowledge that the agent has on the domain (Ramirez-Uresti, 1998). Also, the theory of Constructivism will be addressed as a guiding philosophy (Bruner, 1960).

4 Methods

In order to collect information on participant's actions and motivation in the research experiment, four methods can be used: questionnaires, oral tests, video recording, and the sensing of the participant's movements within the virtual world.

5 Conclusion

This paper has presented an outline of a doctoral research project, which is relevant to the state-of-the-art Artificial Intelligence and Human-Computer Interaction research. The proposed project explained in this paper needs further refining. Specialists and teachers on molecular chemistry will be consulted with suggestions concerning to information on the domain that the tutoring system will address.

References

Bruner, J. S. (1960). *The process of education.* Cambridge, MA: Hardvard University Press.

Chan, T.-W., & Baskin, A. B. (1988). "studying with the prince" the computer as a learning companion. In *Proceedings of intelligent tutoring systems, its-'88.* Montreal, Canada.

Jones, L. L. (1996, June). *The role of molecular structure and modeling in general chemistry.* CHEMCONF'96, an On-Line Computer Conference. (Available at `http://www.inform.umd.edu:8080/EdRes/Topic/Chemistry/ChemConference/ChemConf96/Jones/Paper3.html`)

Kalawsky, R. (1993). *The science of virtual reality and virtual environments.* Wokingham: Addison-Wesley.

Preece, J., et al.. (1994). *Human-computer interaction.* Reading, MA: Addison-Wesley Publishing.

Ramirez-Uresti, J. A. (1998). Teaching a learning companion. In G. Ayala (Ed.), *Proceedings of the international workshop "current trends and applications of artificial intelligence in education"* (p. 83-89). Monterrey, Mexico.

Scaife, M., & Rogers, Y. (1996). External cognition: how do graphical representations work? *International Journal of Human-Computer Studies*, *45*.

Swift, M. L., & Zielinski, T. J. (1997). *Using computers in chemistry and chemical education.* Washington, DC: American Chemical Society.

Yazdani, M. (1987). Intelligent tutoring systems: An overview. In R. E. Lawler & M. Yazdani (Eds.), *Artificial intelligence and education* (Vol. 1). Exeter, England: Ablex Publishing Corporation. (Second International Conference on Artificial Intelligence and Education)

An Introduction to Networks

Rory Graves

roryg@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

Abstract Networks like computers are thought of as very complex. And yet, in reality computers and networks just appear complex by doing lots of very simple things very fast. This paper gives a brief introduction to networks and how they work using a road network as a metaphor.

1 Introduction

To explain the workings of networks such as the Internet, I am going to start with a very simple metaphor. As the paper progresses I will add more features to the explanation expanding into protocols and active network concepts. As the paper progresses I will introduce more and more of the *technical terms* until I sneakily end up talking about networks in networking terms. This paper should hopefully give a novice user a fairly good understanding of what is going on *under the hood* when they use email and browse the web or other network based activities.

2 Basic Model

Imagine the below picture is a road network, there are no telephones, no radio, the only way for people at different ends of a road to communicate is to put the message in a car and send it off in the direction of the recipient. The drivers of these cars are very stupid; they do not know how to get to where they are going. Fortunately at each junction stands a traffic cop, as each car arrives at a junction the cop asks the driver where they are going, and directs them to the correct exit. So if Jim (At point A) wants to send a message to Bob (At point B) this is what would happen:

- Jim (at A) puts his message in a car and sends it to junction A.
- At A the car is directed to the A-B road.
- At B the car is directed to the B-C road.
- At C the car is directed to the C-D road.
- At D the car is directed off to the recipient Bob.

So how does this tie up with real networks? Rather closer than you might expect.

A message (or packet) takes a certain amount of time to pass down a road (known as latency). The traffic cops at each junction takes a certain amount of time to direct each piece of traffic (the switching latency). Roads may be longer or shorter, but each road can only handle so much traffic in each unit time (called bandwidth). If there where enough road for everybody this would be the end of the discussion. Unfortunately, just like real roads there is often too much traffic and not enough road leading to jams. And just like real roads, just adding more tarmac often does not work, the more road you supply the more people will use it. The next section explores how we can deal with these traffic jams (or congestion control).

3 Dealing with Jams

Lets take a simple example, two roads coming into a junction with more traffic coming in than the one road leaving can manage. At a real junction we would have tailbacks. In networking these are banned, the entrance to each junction must be clear. So how can we deal with this? Fortunately we can arbitrarily destroy cars if we have too many. In a real road system this would probably cause some complaints! When there are too many cars for a road they are stacked up on the slip road (if the jam is just temporary then no cars will be lost). If however the slip road fills too then any new cars that can not fit will be obliterated. If the data in the car is that important the send must keep a copy and make sure it gets through as described in the next section.

4 Coping with loss

I wish to send a large message, it will not all fit in one car so I must break it across a number of vehicles. I also want to know for certain that all of the message arrives and is in the right order. How do I go about this? First lets assure that all of the smaller messages arrive in the correct order, so the large message makes sense. This can be achieved by adding a sequence number to each message. This means that even if the messages arrive out of order at the receiver he can rebuild the message.

So lets start with a really simple *stop-go* protocol. Sender A sends his first message numbered 1. When receiver B gets this message he sends back a reply saying *got message 1*. When the reply arrives back at the sender he sends message 2 and so on. If after sending a message the sender receives no reply he will wait a short while (messages take time) and then resends the message. This protocol works well, but has one minor drawback, it takes a long time. If each message took 3 hours to travel between the sender and the receiver the sender could only send one message every 6 hours.

Can we improve upon this? Yes, by realising that more than one part of the message can be sent at a time. It may still take 6 hours for a message to be sent and a reply to come back, but potentially hundreds of cars could be moving along the road at the same time. TCP (Transmission Control Protocol), one of the mainstay protocols on the Internet uses this approach. All of its messages are numbered as in the previous example but unlike the previous example, a number can be sent simultaneously. The protocol starts by sending a given number of messages one after the other. For each reply received the next message is sent, if messages are missing they will be resent after a timeout. TCP is actually even cleverer than this, it tries to find the biggest window it can and continually adjusts it to compensate for loss. Much work has been done on tweaking the numbers to make this work the best it can under differing conditions.

5 Conclusions

So, does this stuff actually do anything useful? If you use the web, email, newsgroups or nearly any other Internet tool the answer is a resounding yes. TCP is the mainstay protocol that makes all of the Internet applications that you use actually work. This paper gives a brief overview of what is really going on underneath the hood in real networks. It all turns out to be just a bunch of old packets...

Artefact Reconfiguration for Extended Groupwork

John Halloran

johnhall@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

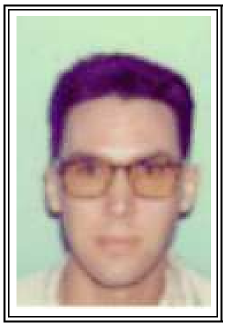

Note: The issues in this short paper are more fully discussed, and referenced, at: http://www.cogs.susx.ac.uk/users/johnhall/UG.html

1 Overview

My research is motivated by the need to understand how web-based computer artefacts are used to support extended groupwork - groupwork which is one-off, relatively unstructured, and extemporary - and how the artefact can be reconfigured to better support such work. Artefact reconfiguration means redesign of the computer application, customization, or training. I am working on a version of Activity Theory (AT) built around the concept of the activity space to analyse artefact use in the domain of student groupwork. This framework can be used to indicate what sorts of reconfiguration should take place.

2 The Activity System

Student groupwork is a domain in which we are unable to predict or specify exactly how artefacts will be used. AT is recognised as a potentially useful conceptual framework for analysis since it provides a holistic approach to development of activities in the context of which artefact use can be understood.

In AT, examples of groupwork can be treated as activity systems. An activity system features a collective subject, a common object, history, close coupling, and evolved artefacts. For example, at a primary care clinic, the different staff (including doctors, nurses, receptionists) work toward a common object, the provision of healthcare. Because the staff are jointly oriented toward this aim, they work together: a collective subject or community. Specific subjects interact with the object by means of mediating instruments (see Figure 1). These can be physical, symbolic, external

and internal; including for example X-Ray, medical records, treatment-related concepts and methods, depending on the subobject of the common object they are addressing; for example, making a diagnosis, or accessing a patient case. These are evolved artefacts. Close coupling refers to the reciprocal co-evolution of subject, object and mediation, and to rational, developed links between them. Close coupling is reflected in a division of labour - the division of tasks of power and status; and a set of rules - regulations, norms and conventions that constrain actions and interactions within the activity. These characteristics arise when the activity has a grounding in history, and they position, stabilise, and make explicable the use of artefacts; and help to realize the object as an outcome.

Contradictions in an activity system refer to developmental anomalies between nodes of the system; for example, when new objects in the shape of new types of health problem present for which current mediation is inadequate. Breakdowns are what occur when an artefact resists a subobject of the common object - for example, a receptionist is unable to search for a patient case because the system does not explain its search criteria. Thus, we are able to take an artefact perspective, which means looking at specific usability issues of the artefact and suggesting reconfiguration in relation to the relevant object. The breakdown in search criteria is something that affects a subobject of the common object, and we can see that by fixing it we progress the whole activity system.

The Activity Space These conceptual tools break down when applied to student groupwork, because the characteristics of activity systems are missing. In order to analyse student groupwork, an alternative conceptual framework is developed based on the concept of the activity space. Activity spaces, unlike activity systems, do not have an established historical ground and are pre-systematic. Activity spaces feature contradictions within and between their nodes (subject, object, mediation) shown as parallel lines in Figure 2. These propagate in complex ways around the space which

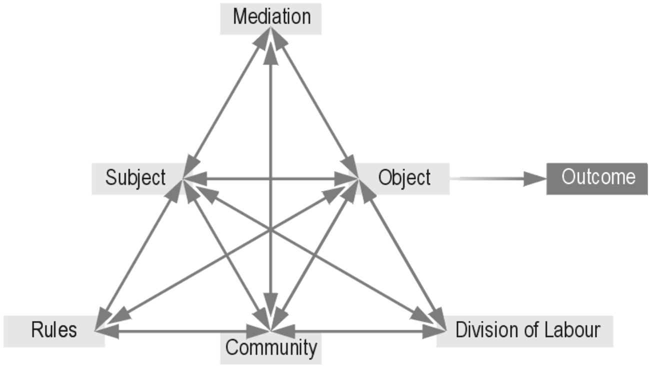

Figure 1: Activity System

means we can only understand artefact use and think about reconfiguration given a characterisation of the whole activity space and the contradictions within it. Applying the concept of the activity space also requires us to take an intentional rather than cultural-historical perspective dedicated to characterising contradictions not as developmental anomalies but as mismatches between what the different subject groups within the activity space are attempting to achieve.

3 Example case study

Research was carried out on groups of students undertaking compulsory undergraduate groupwork as part of a software design and evaluation course and using Lotus Notes, a messaging system with shared space for document storage, to support their work. Tutors pointed out that use of Lotus Notes was accessible to them, and material there would be used to assess group management. The course was mapped as an activity space. As a first-time activity, it does not have a grounding in history. A basic contradiction was found: students wished to maintain privacy while minimally observing the guidelines on Lotus Notes use; tutors hoped for full, elective use of this system, which would be accessible by them. Thus, there is no common object; and the subject is not collective because the different subject groups are oriented to different objects. These contradictions propagated to mediation. There was a conflict between the use of Lotus Notes and the teaching network. In order to satisfy course guidelines, some materials had to be inspectable on Lotus Notes. But to preserve privacy and save work (since Lotus Notes had to be learned), the teaching network was chosen. The teaching network did not allow inspectability where this was needed either by the tutor group, or by other students in the same group, necessitating the creation of special shared spaces on the network. Thus extra work was created both in terms of creating new functionality on the teaching network and in needing to make inspectable postings to Lotus Notes. Thus, neither system supported student objects; or tutor objects, and neither is an evolved artefact, because it is not fitted to a common object or a collective subject. This also implies that there is not close coupling. This context makes it difficult to discern the set of rules, community, or division of labour which characterise an activity system and are associated with close coupling. We can begin to see, then, how distributed contradictions in an activity space can explain artefact use; and we can refer to the activity as an incongruent activity space.

Using the concept of the activity space implies considering the entire space when attempting to reconfigure the artefact. This means attempting to remove or lessen the contradictions between subjects, objects, and mediation. Thus we need to improve Lotus Notes functionality in terms of, for example, a better front end for categorization of messages and support of threading. However, we also need to consider network congruence. Changes to Lotus Notes as an artefact should respect the need for students to use the teaching network in tandem with it; finding ways to notify of postings, and to offer file transfer between the two systems. Such changes, however, imply a need to reduce object contradictions. Tutors might need to drop the requirement for inspectability, and concentrate on promoting groupwork through useful software (the creation of special shared spaces on the teaching network is an example of students' need to do this). Inspectability could be made voluntary: students could be asked to produce evidence of group management and shared work optionally rather than compulsorily including material on Lotus Notes.

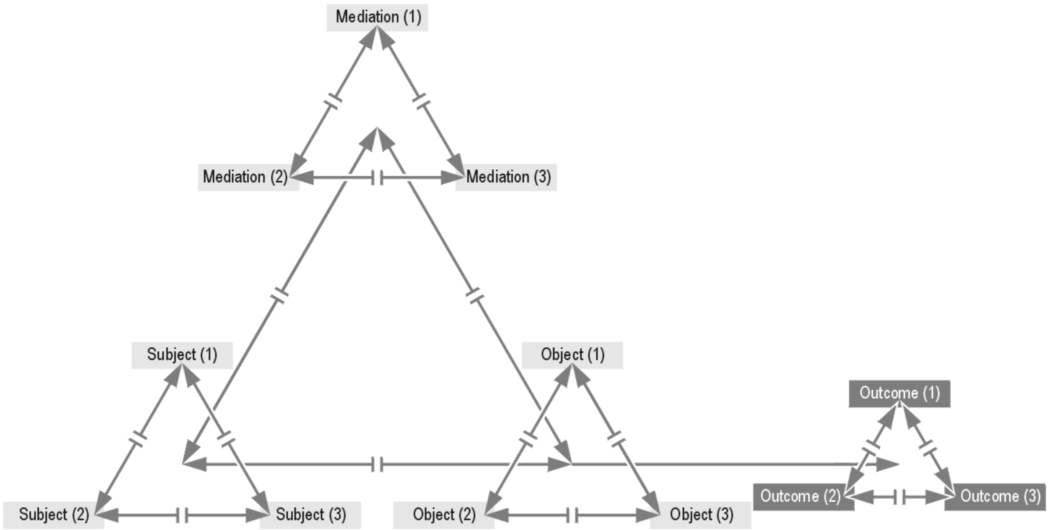

Figure 2: Activity Space

4 Conclusion

Extended groupwork has properties which mean we have to consider it as a pre-systematic space which must be characterised in terms of its contradictions in order for artefact reconfiguration to be undertaken. The concept of an activity space can be used to make this characterisation and indicate what contextual issues are involved and should be considered. The approach means that an artefact perspective is likely to be insufficient: we need to think about reducing contradictions at all of the poles of the activity space as well as their propagation. In this way we might be able to introduce greater systematicity into the space and to evolve an artefact within it. The research, as well as being oriented towards artefact solutions, also has a theoretical motivation: finding a way to operationalise activity theory and attempting validate an activity-theoretic conceptual framework through artefact reconfiguration spanning successive presentations of the same groupwork.

Evolving Communication Mechanisms in a Complex Environment

Tudor Jenkins

tudorj@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

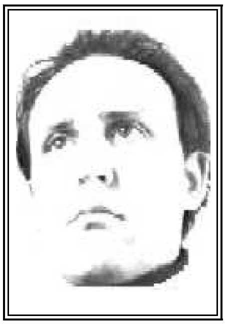

This paper describes an ongoing research program to assess the viability of using situated adaptive behaviour (SAB) techniques in the development of inter-agent communication. Particularly whether genetic algorithms can be employed as a suitable search method to design complex cooperative communication in a non-trivial environment.

New communication skills are developed by building on existing capabilities in an incremental manner by exposing an adaptive population of agents to an environment undergoing a gradual augmentation of complexity. Generating the initial functionality requires some path through an enormous search space that arises from the need for 'matching' transmitter and receiver strategies. A solution to overcome this problem is discussed in the next paragraph and experimental details and results are presented later. With simple signalling behaviour established focus will turn to the exploitation of a principle akin to darwinian evolution but applied to the realm of signals based on Dawkins (1976) concept of memes. Variations in signals through noise and transmitting errors along with selection of specific signals by the population as useful acts of communication will provide the necessary conditions for adaptation and development of signals for ever greater complexity of communication. Suitable control architectures and signaling systems will be assessed for viability.

Initial work has concentrated on finding methods for agents to develop the use of appropriate sensors and actuators required for signaling interactions. The major obstacle common to many learning problems is one of credit assignment. An act of cooperative communication requires five separate stages to be in place for any reward to be accrued by either of the two participants (and that is assuming some group selection can be applied to reward the transmitter). These are to be aware of signaling opportunity (transmitter); produce appropriate actuations/ signal (transmitter); signal transmission; detect and discriminate signal (receiver); actuate appropriate response (receiver). Unfortunately within the context of signaling, the absence of a single stage will result in not even a partial benefit from the presence of the other stages. Clearly, some other motivation for these stages is required. Appealing to the idea developed by ethologists (Krebs & Dawkins, 1978), that communication can arise as a means of manipulating other animals, I have been attempting to exploit this idea by providing an environment where one faculty (reception behavior) of the communication system is developed independently and the secondary side (transmission) then develops so as to exploit the first faculty for means other than what it was first developed for. Results are reported along with experimental details.

Existing work in this field has already been undertaken on problems of limited complexity. Agents have been given a prescribed set of symbols to use and negotiate some coherence between transmitting and receiving populations (Werner & Dyer, 1991; MacLennan, 1991; Oliphant, 1996; Cangelosi & Parisi, 1996; Fyfe & Livingstone, 1997; Saunders & Pollack, 1996).

1 Experiments & results

Experiments are performed on populations of virtual robots with dynamics loosely modeled around the characteristics of kehpra robots. A signaling system analogous with sound was chosen rather than one more closely tied in with the rest of the agents behavior such as movement and vision in order to ease the analysis of any signaling that occurs. Dissambiguating signaling from other behaviors has been notoriously difficult (Noble & Cliff, 1996) so the choice of a medium that can have few other uses eliminates the problem. Conversely, the use of this otherwise 'unneeded' channel causes many of the problems of integrating it into the behavioral repetoir. Arbitrary numbers of actuators and sensors that can vary over a continuous range of 'sounds' have been employed to limit any future restrictions on the range of vocabulary available to agents. Populations of these robots are placed together in a maze environment where they were free to navigate around, collecting resources and returning them to a base area which causes an augmentation of their fitness level. This fitness level is then used on a periodic basis by a genetic operator.

The maze environment started in its simplest form as a T-junction with resources placed periodically at the end of one or another of the corridors and the third corridor acting as a base. Resources only being placed at the end of one of the corridors, selected randomly, when all have been collected from the currently filled one. The reason for this arrangement being to provide an advantage to the group of agents that share information about which corridor currently contains the resources thus minimizing the time each agent has to spend exploring in comparison with other groups that fail to do this. Adding further corridors and junctions off from these provides a simple way to incrementally add to the complexity of the environment and the information agents might need to convey.

With a population of 60 agents and a genetic operator employing both crossover and mutation, populations evolved within 100k generations to move about their maze, navigate towards resources, return with them to base and also develop efficient search strategies. With this capability established, the development of sound reception and behavior linking this with what might later be deemed useful behavior with respect to communication was needed. In order to achieve this, different signals were emitted whose frequency depended on the corridor which currently held the resources, thus allowing the agents to take advantage of the correlation between the various signals and the position of resources. After several trials, no added benefits to the agents fitness were observed from this seemingly useful information available to them. When a record was made of all signals within the environment during these trials, a possible reason for the failure became apparent. Each of the agents was capable of producing signals just from random activations of their signal actuators and as a result of this the environment was swamped with signals, so whether or not a particular signal was emitted by the environment to indicate the position of resources, the signal would be present almost continuously from other agents. Two solutions to overcome this problem were proposed. One has been tried, trials involving the other are underway.

Firstly, a cost in terms of fitness was attached to signaling. After several trials it was possible to eliminate interfering signals from the agents but only with the imposition of quite severe costs for signaling which had the effect of converging the populations early on in the trials. This convergence causing problems for development of the rest of the agents abilities. On one trial with a considerably longer run (400k generations) it was possible to evolve agents that did respond to the signals from the environment. A problem with this approach stems from the very discrete nature of the control architecture. By eliminating at an early stage the agents ability to emit signals, a problem is being stored up for later experiments where costs of signaling would be relaxed in order for communication to come about. With the production rules causing signaling removed from the population, it will be a hard task to reintroduce the ability. This result suggests the need for a new control architecture that allows for more gradual change, thus allowing signaling to be just switched off with the opportunity for this being reversed with relative ease.

The second solution being developed relies on the principle that trials involving many less agents (from larger population) would generate less interference, thus eliminating the need for such high costs being imposed on signaling. Taking a small set of agents and assessing their performance with the additional benefit of rewarding these as a group so that interfering signals from one agent rather than just lowering each of the agent's fitness but keeping them on an equal footing, lowers the group's performance with respect to other groups and thus adds to the pressure to stop producing interference. The approach of using small groups to eliminate interference will probably be a feature of all further work. With the restricted size of the environment available to the agents it is not possible for some spatial distribution of the agents to erradicate enough interference of others. Results are currently unavailable from this experiment but similar consideration is being paid to the limitations of the current control system.

With the reception skills established, the next phase is to reduce the information available to the agents from the environment thus creating an opportunity for them to generate similar signals under appropriate circumstances to confer a similar advantage to their group members. Of particular interest will be the problem of how one can allow useful signals needed to generate the receiving behavior to diminish and maintain the receiving behavior whilst allowing sufficient time for signaling to arise to exploit this behaviour.

2 General remarks

The development of faculties required for communication needs to be integrated with many other areas of agent control mechanism. To find suitable conditions for achieving this, one has to look beyond the act of communication to create functionality that might then have a role in this behavior. The idea of developing agents with complex linguistic capabilities is not the goal of this work. As a mobile robot designer working within the SAB paradigm might start simply with any motion, followed by more efficient movement then consider directional capability etc., so too must this work address more fundamental aspects of communication. Linguistic capabilities are a distant horizon and without the development of primitive communication skills first including the acquisition of grounded symbols, the building blocks will not be in place for such grand schemes.

References

Cangelosi, A., & Parisi, D. (1996). *The emergence of a 'language' in an evolving population of neural networks* (Tech. Rep. No. NSAL-96004). National Research Council, Rome: Neural Systems and Artificial Life Group.

Dawkins, R. (1976). *The selfish gene.* Oxford: Oxford University Press.

Fyfe, C., & Livingstone, D. (1997). Developing a community language. In *Proceedings of the fourth european conference of artificial life.*

Krebs, J., & Dawkins, R. (1978). Animal signals: information or manipulation? In J. Krebs & N. B. Davies (Eds.), *Behavioural ecology: An evolutionary approach.* Oxford: Blackwells.

MacLennan, B. (1991). Synthetic ethology: an approach to the study of communication. In C. G. Langtona, C. Taylor, J. D. Farmer, & S.Rasmussen (Eds.), *Artificial life ii: The second workshop on the synthesis and simulation of living systems* (Vol. X, pp. 631–658). Redwood City, CA: Addison-Wesley.

Noble, J., & Cliff, D. (1996). On simulating the evolution of communication. In P. Maes, M. Mataric, J. A. Meyer, J. Pollack, & S.Wilson (Eds.), *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behaviour* (pp. 608–618). MIT Press.

Oliphant, M. (1996). The dilemma of saussurean communication. *BioSystems*, *37*(1–2), 31–38.

Saunders, G. M., & Pollack, J. B. (1996). The evolution of communication schemes of continuous channels. In P. Maes, M. Mataric, J. A. Meyer, J. Pollack, & S.Wilson (Eds.), *From animals to animats 4: Proceedings of the fourth international conference on simulation of adaptive behavior* (pp. 580–590). MIT Press.

Werner, G., & Dyer, M. G. (1991). Evolution of communication in artificial organisms. In C. G. Langtona, C. Taylor, J. D. Farmer, & S.Rasmussen (Eds.), *Artificial life ii: The second workshop on the synthesis and simulation of living systems* (Vol. X, pp. 659–687). Redwood City, CA: Addison-Wesley.

Investigating the Associative Properties of Evaluative Conditioning

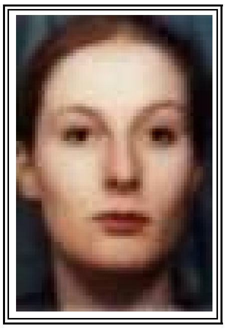

Kristy Lascelles

kristy@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

In evaluative conditioning (EC), the contiguous pairing of an affectively neutral stimulus (NS) with an affectively positive (AS+) or negative (AS-) stimulus leads to a change in the affective evaluation of the NS in the direction of the AS with which it was paired. So pairing a NS with an AS+ will result in the NS becoming more positively evaluated (liked), whereas pairing a NS with an AS- will result in the NS becoming more negatively evaluated (disliked). EC has been argued by Baeyens (Baeyens, Eelen, Bergh, & Crombez, 1989; Baeyens, Eelen, Crombez, & Bergh, 1992) to be a qualitatively distinct form of learning compared to Pavlovian conditioning in that:

1. an awareness of the NS-AS association is not necessary for EC to occur;
2. EC appears to be resistant to extinction;
3. EC is based on a contiguous NS-AS relationship and not on an NS-AS contingency.

However, Field and Davey (1997, 1999) have criticised much of the work conducted on EC, mainly due to the failure of most EC experiments to use adequate between-subject controls. Therefore, it cannot be determined whether any results obtained in the EC paradigm are due to an associative process, or not. An alternative nonassociative account, the exemplar-comparison model [(ECM) see Field and Davey (1999)], has been suggested to account for EC by Davey (1994), in which he proposes EC type effects occur via conceptual learning and transfer. That is, the NS is placed into either category ('liked' or 'disliked') according to feature overlap. So, if the NS consists of a number of features, the majority of which are also seen in AS+, then it will be categorised as 'liked'. Likewise, if the majority of feature seen in the NS are also seen in AS-, then the NS will be categorised as 'disliked'.

2 Aims And Objectives

1. The ECM predicts that EC type effects should be influenced by stimulus similarity. The more similar a NS and AS are, the more feature overlap should be seen, and the higher the likelihood that the NS is evaluated in the same direction as the AS has been evaluated. Another prediction made by the ECM is that EC type effects will not only occur with affectively evaluated stimuli (e.g. like-dislike), but will also be seen when other conceptual dichotomies are employed. This research will further test these and other predictions made by the ECM.

2. The characteristics of EC which distinguish it from other forms of conditioning (e.g. conditioning without any awareness of NS-AS associations and the apparent resistance to extinction) will be further examined, but using appropriate between-subjects controls such as the block/sub-block control in conjunction with a no-treatment control, as suggested by Field and Davey (1999).

3. A counterbalanced design procedure, which is exempt from artifactual EC type effects (resulting from stimuli selection and pairing biases) will be used to investigate the associative nature of EC.

4. This research will examine whether any selective learning effects can be obtained in EC.

References

Baeyens, F., Eelen, P., Bergh, O. V. den, & Crombez, G. (1989). The influence of cs-ucs perceptual similarity/dissimilarity on human evaluative learning and signal learning. *Learning & Motivation*, *20*(322-333).

Baeyens, F., Eelen, P., Crombez, G., & Bergh, O. V. den. (1992). Human evaluative conditioning: acquisition trials, presentation schedule, evaluative style and contingency awareness. *Behaviour Research & Therapy*, *30*, 133–142.

Davey, G. C. L. (1994). Defining the important theoretical questions to ask about evaluative conditioning: A reply to martin and levey. *Behaviour Research & Therapy*, *32*, 307–310.

Field, A. P., & Davey, G. C. L. (1997). Conceptual conditioning: Evidence for an artifactual account of evaluative learning. *Learning & Motivation*, *28*, 446-464.

Field, A. P., & Davey, G. C. L. (1999). Reevaluating evaluative conditioning: A nonassociative explanation of conditioning effects in the visual evaluative conditioning paradigm. *Journal of Experimental Psychology: Animal Behaviour Processes*, *25*, 211-224.

Diagrammatic Representation

KK Lau

kamkeung@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

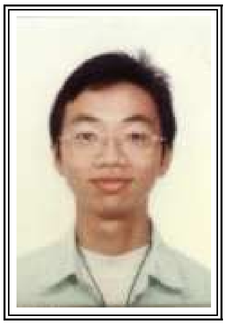

My research study focuses on the uses of diagrammatic representation in information. While the uses of language have been investigated in linguistics for a long time, the uses of diagrams has not been systematically explored. My present research will focus on the cognitive behaviour of reading diagrams. I am interested in the relevant mechanisms and phenomena in diagrammatic cognition.

At this moment, I have basically finished my literature review. The next step is to locate my exact area of study and to propose appropriate hypotheses. In addition, I am trying to conceive some ideas of possible experimental designs to test proposed characteristics of cognitive patterns of diagrammatic perception. Although exact ideas about research are not quite clear, the direction of work is almost confirmed.

One important possible implication of the research will rest on the improvement of diagram use in information transmission. Information transmission involves both traditional graphical representation using paper and advance computer information applications including virtual reality.

Beyond the Interface

Ann Light

annl@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

The functions offered by websites have increased. With this increase has come a range of entry devices allowing users to specify their needs so sites can respond appropriately. These devices are clearly interactive - in that the behaviour of users affects the response they receive - but this is not the point and click' interactivity of following links. So is there a change in the quality of interaction going on? And, if so, which constructs might be appropriate for describing and exploring users' behaviour in these new contexts?

This paper puts forward evidence that users go about their business on websites with two levels of awareness:

- that of the interface,
- and, when users become involved in entering text, that of the social context beyond the interface as well.

In entering text as part of interacting with a site, users start to perceive their behaviour in terms of person-to-person rather than person-to-machine relationships, even though they may know that they are addressing software.

2 Theoretical background

The paper takes its theoretical grounding from Goffman's work on social cues and also Levinson's interpretation of his 'participant roles'. It also draws on arguments developed in Reeves and Nass' book The Media Equation', which finds that individuals' interactions with computers, television and new media are fundamentally social and natural, just like interactions in real life' (p5).

There is considerable work in the field of computer-mediated communication, some of which applies to discussion sites on the Web, and we argue that the study reported here has pushed such research into an area traditionally analysed in terms of interface design and usability, with metaphors drawn from space. We want to explore design language for the interface using terms and metaphors drawn from communication.

3 Methodology

In approaching this study, we wanted to know:

- What thoughts went through the mind of the users as they approached and started the task of entering text into websites?
- How did these thoughts compare with their thoughts when using other parts of the site?
- Was there justification for using a communication metaphor to describe interactions with the site?

Thoughts are, of course, impossible to collect in unmediated form. However, we decided to collect users' accounts of the task of using a website, regarding the thinking we wanted to investigate as a series of mental activities to be reported on. This enabled us to use discourse analysis to explore language use and look for patterns across accounts and variations within them (Potter & Wetherell, 1987) but raised the problems discussed below.

The basis of the accounts It became apparent in tests that if we controlled the conditions of the experience we could affect the content of the account we received, particularly in the area of motivation. Since motivation in conducting an activity is a controlling factor in the perception of the activity, this immediately reduced the study's scope and credibility.

Collecting the accounts We dismissed using concurrent verbal protocols as the detail we required interfered with conducting the Web task (Ericsson & Simon, 1993). We found the weakness of prompting by reviewing the task on video and concluded that retrospective questioning was the best means of getting the accounts. Vermersch's explicitation interviewing technique

To this end, we adopted the interviewing technique developed by Vermersch, and used for several years to evoke cognitive processes retrospectively, particularly in the French educational system. (For a description of the technique, see (Vermersch, 1994, 1999; Depraz, Varela, & Vermersch, 1999)).

Design of study 20 Web users were interviewed who had entered text into a site of their choosing. Search sites, such as Excite, were excluded. This produced nearly 12 hours of audio interviews for analysis.

Participants This were chosen as familiar users of the Web, and demographically typical of 'experienced' users according to the latest GVU survey (GVU, 1998), with a larger ratio of women. All participants were European English speakers.

Participant roles used in the analysis Applied from Levinson we used 'target' 'ratified' and 'unratified' recipients to describe categories of addressees. We also used the term 'producer' in a different sense to Levinson, to mean the originator of the website.

Transcription notation Punctuation has been used to give a sense of the utterances and { } included where words have been omitted in quotes from transcription.

4 Findings

Transcriptions were divided into segments and these compared inter and intra account. The main body of this paper quotes extensively from these transcriptions as evidence of the assertions below.

Overview There were a number of different conditions involved with users visiting sites, entering text and thinking about who received the information sent.

Motivation Filling in text on websites was only a deliberate activity for two users, an irrelevance for those forced to register and an expected part of a greater transaction for the remainder. Most people gave most of their attention to completing their primary activity on the site.

Awareness of the interface Interviewee's showed peripheral awareness of the interface in their accounts.

Awareness of others Interviewees showed peripheral awareness of others in giving their accounts of entering or thinking about entering text. This was manifested particularly in the switch from using 'it' to 'they' in referring to elements of the site.

Perceptions of 'they' 'They' referred primarily to the producer of the site, but could also refer to someone or something receiving information from the user.

Social behaviour Not only those deliberately addressing a human audience behaved as if engaged in an interpersonal interaction.

Other audience Unratified audience appeared in the accounts as overhearers in the system and as others given access to user-submitted information, unbeknownst to the user.

5 Discussion of findings

Interviewees move swiftly and easily between two states of reference, showing an awareness of the interface and also shadowy figures beyond it. This second kind of awareness develops when text is entered. This makes a strong case for adopting a communication metaphor for interactive site components.

Mode of Entry Text entering requires selecting from a greater choice of words than a menu offers. Once users are forced to think about how to describe something - including themselves - they start to think about how to present it. Interviewees make sense of this by drawing upon self-presentation behaviours that are familiar to them from person-to-person interactions. It follows that users' perceptions of the recipient of their information - and their requirements - will affect what the users do and say next, just as happens in exchanges between people. Where users do not have a clear impression of the recipient, one might expect a less confident response to the site.

Target recipient The implication from the study above is that users, left to their own devices, will expect to enter a dialogue once entering text. The kind of dialogue is then a matter for development by the designers so that the voice of the site speaks clearly and makes transparent with whom the user is engaged.

Information Content Imparting personal information brought with it concerns that were not only about security. Differences in users' responses can be attributed , in part, to trust. Where users were interacting with strangers' there were more questions raised in the accounts about who wanted to know the information and why.

Purpose The relationship between the user's purpose and the information to be entered was important. If entering information imposed activity for no gain to the user it was unwelcome.

Branding The perceived audience has an identity constructed from a combination of:

- what users bring with them in terms of expectations of the brand,
- users' experience with the site, especially with parts where the users' contribution is solicited,
- users' purpose in being there.

'Homepages' - identified with an individual - were treated with less suspicion and more indulgence in their information requests than any other kind of site, whereas companies, in general, were seen as unreliable or looking to exploit the user.

6 Discussion of method

This study produced material of an appropriate content and granularity to meet the ambitious research agenda without elaborate staging or equipment. While setting out with a fixed procedure and a structured interviewing technique produces a more manageable quantity and kind of data, in this context it would have risked simplifying the relationship between language, thought and behaviour in ways that eliminate much of what was most interesting.

References

Depraz, N., Varela, F., & Vermersch, P. (1999). *On becoming aware: the pragmatic of experiencing.* (To be published by MIT, Cambridge, CA)

Ericsson, K. A., & Simon, H. (1993). *Protocol analysis: Verbal reports as data* (2nd ed.). Cambridge, MA: MIT. (First edition 1984)

GVU. (1998). *10th survey.* `http://www.gvu.gatech.edu/gvu/user_surveys/survey-1998-10`. (As at 10/6/99)

Potter, J., & Wetherell, M. (1987). *Discourse and social psychology: Beyond attitudes and behaviour.* London: Sage.

Vermersch, P. (1994). *L'entretien d'explicitation.* Paris: ESF. (In French)

Vermersch, P. (1999). Introspection as practice. *Journal of Consciousness Studies*, *6*(2–3), 17–42.

3D Virtual Environments and Learning

Nuno Otero

nunop@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

The research that I intend to pursue concerns to the assessment of 3D virtual environments usefulness for learning. In particular, I am focusing on the different forms of visualising and manipulating graphical representations that this technology allows and its relations with conceptual learning. There has been much hype in education about these technologies but the validation of the benefits is still not established, and typically, the research only takes the form of comparative studies. Little structured cognitive analysis has been carried out to explain the real advantages of applying VE's to learning. Moreover, it seems that there is a big gap between the investigations about interactivity properties/design principles of virtual environments and its use to support learning.

This research will try to contribute towards an understanding not only of how to design virtual environments for learning, but also when to design virtual environments for learning and how to use 3D interactive representations in conjunction with other types of representations for an effective learning process.

To carry out this kind of research, however, is not easy. Even the apparently simple task of choosing the domain seems complex. In fact, it is difficult to clearly establish how a certain concept or problem must be represented in a 3D format and whether this same concept would be better represented in another format. The chosen domain to my research is geometry and I will try to understand the benefits of using 3D representations and the possibility to manipulate them for learning the stereographic projection, a concept that involves three dimensions. The theoretical point of departure is the external cognition framework (Scaife and Rogers, 1996; Rogers and Scaife, 1998) but the literature review will also cover the following themes:

- the use of virtual environments for conceptual learning (where 3D manipulable graphical representations clearly inscribe);
- the role of external representations in cognition and in the learning process in particular;
- the role of diagrams in learning geometry.

The experimental part consists of a cognitive analysis of the interactivity of different systems that will use different types of graphical representations (basically, 2D or 3D graphical representations). The notion of computational offloading will be central to the analysis. In a first step, the analysis involves understanding how variations on the graphical representations and the explanatory text affect the cognitive load and how this will affect the learning process. The second step will be related to the introduction of dynamic links between the different types of representations, trying to establish how this modifies the apprehension of the concept.

A Runtime Object Model for Polymorphic Languages with Stack-Allocation

Tim Owen

timothyo@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

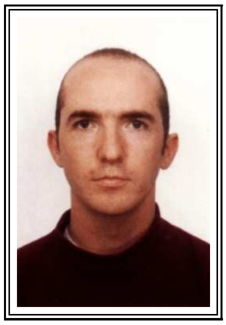

1 Research Overview

The aim of my research is to investigate ways of implementing object-oriented programming languages that contain *both* parametric polymorphism and stack-allocated objects. The purpose of this work is to enable more efficient implementation of advanced language features that give the programmer more expressive power. Further details can be found on my web page.[1]

Parametric polymorphism allows programmers to abstract over 'some type' of data, by using a type variable to represent the unknown actual type. This is analogous to the way an ordinary function can abstract over some value, by specifying a named parameter which represents the actual values supplied by the caller. Combining parametric polymorphism with the traditional (subtype) polymorphism found in object-oriented languages gives the programmer considerable expressive power, and is particularly useful for writing libraries of generic, reusable code.

Work on language design and design patterns has indicated that a clear distinction can be made between value-based and reference-based data. Value-based data is typically small and immutable, making it ideal for storage on the runtime stack rather than incurring the overhead of heap-allocation. In addition, compiler optimisations can make good use of knowledge about value-based data. Some languages provide support for such lightweight objects, but vary in the way they interact with polymorphism.

Combining parametric polymorphism and value-based data presents some challenges to efficient implementation of the language. In particular, the runtime object model (which describes the organisation of memory for object storage) requires careful design. Existing work in this area is mainly based on functional programming languages, which do not have subtyping. My contribution is to apply and extend research results into object-oriented languages.

My interest is in using runtime type information in the object model, to support a language including the features just described. In a wider context, runtime types enable a range of language technologies such as dynamic method dispatch, garbage collection, dynamic linking, reflection and object persistence. The details of some of these are considered 'further work', but are to be born in mind during my research.

[1] http://www.cogs.susx.ac.uk/users/timothyo/research

Children's understanding of astronomical ideas

Georgia Panagiotaki

georgiap@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

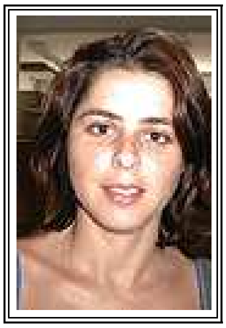

Research in developmental psychology, cognitive science and educational psychology has shown that children construct an intuitive understanding of the world that is based mainly on their everyday experience and observations. The nature of this naïve understanding differs significantly from the currently accepted scientific explanations of various natural phenomena. Research in the last decade has focused on a number of different areas such as biology (Carey, 1988; Carey & Smith, 1993), physics, mechanics, psychology, observational astronomy (Nussbaum & Novak, 1976; Sharp, 1996; Sneider & Pulos, 1983) etc.

Several studies have been conducted for the investigation of school children's knowledge of astronomical ideas such as the shape of the earth, the sun and the moon, the gravity, explanations of day and night cycle and the four seasons. Findings suggest that young children construct a number of consistent mental models to make sense of their physical environment. Mental models of the earth shape that have been identified (Vosniadou & Brewer, 1992; Baxter, 1989) are:

1. The initial model, based mainly on children's everyday experience (i.e. the earth is flat with people living only on top of it),

2. The synthetic model, that combines elements of intuitive understanding and scientific explanations (i.e. the earth looks like a hollow sphere and people live inside its flat surface)

3. The scientific model, that follows the currently accepted scientific explanation of the earth shape. It has been argued that children enter school holding initial (intuitive) models that are gradually replaced by synthetic and finally by scientific ones. This is achieved through the accumulation of new information and the replacement of various misconceptions that children hold. Instruction and culturally transmitted information play a very significant role in this process.

The methodology of previous work has been based on structured open interviews, in combination with children's drawings for the support of their verbal responses. My first study is a replication of the previous research with emphasis on the earth shape, the location of the sky and the day and night cycle. It also focuses on two methodological issues. The main tool was a forced choice questionnaire offering two alternatives a scientific and an intuitive one in contrast to an open questionnaire that may reinforce repeated questioning and often confusion of the respondents. It was also used to compare verbal responses given in addition to 2-D drawings with verbal responses given in addition to the manipulation of 3-D physical models of the earth that were presented during the interview. It has been argued (Butterworth, Siegal, Newcombe, Dorfman, 1998) (Butterworth, Siegal, Newcombe, & Dorfmann, 1998) that children have difficulties in representing the earth in two dimensions and this can lead to more flat earth responses -initial models. Seventy two 6 to 12 year-old school children were tested. All participants answered the same questionnaire. Half of them produced drawings while the other half manipulated concrete models that represented various earth shapes.

The results of the present study did not fully support Vosniadou et al's (1992, 1994) (Vosniadou & Brewer, 1992, 1994) findings in relation to the process of conceptual change through the initial, synthetic and finally scientific models of the earth shape. The majority of participants (44%) held consistent scientific/spherical models, while only a very small percentage (4%) held consistent intuitive/flat models. Mixed models (characterised by no consistency of answers) were held by 28% of the children. The rest of the participants held mainly scientific models with one accepted deviation.

No significant differences were found in the comparison between drawings and the manipulation of physical earth models.

References

Baxter, J. (1989). Children's understanding of familiar astronomical events. *International Journal of Science Education*, *11*, 502-513.

Butterworth, G., Siegal, M., Newcombe, P., & Dorfmann, M. (1998). *Young children's knowledge of the shape of the earth in australia and england.* (In preparation)

Carey, S. (1988). Conceptual differences between children and adults. *Mind and Language*, *3*, 167–181.

Carey, S., & Smith, C. (1993). On understanding the nature of scientific knowledge. *Educational Psychologist*, *28*, 235–251.

Nussbaum, J., & Novak, J. (1976). An assessment of children's concepts of the earth utilizing structured interviews. *Science Education*, *60*, 537–550.

Sharp, J. G. (1996). Children's astronomical beliefs: A preliminary study of year 6 children in south-west england. *International Journal of Science Education*, *18*, 685–712.

Sneider, C., & Pulos, S. (1983). Children's cosmographies: Understanding the earth's shape and gravity. *Science Education*, *67*, 205–221.

Vosniadou, S., & Brewer, W. (1992). Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology*, *24*, 535–586.

Vosniadou, S., & Brewer, W. F. (1994). Mental models of the day-night cycle. *Cognitive Science*, *18*, 123–183.

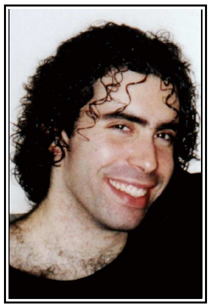

Automatic Spelling Correction

Darren Pearce

darrenp@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

Systems that attempt to automatically process text usually expect the input to be free from errors. In reality, this is far from the case due to possibilities such as grammatical mistakes, spelling mistakes or typing errors. The material presented in this brief paper details one way in which simple 'machines' can be used to capture and potentially correct common typing errors.

2 Error Operations

Typing errors can be divided into four *error operations*:

- **Deletion** where a character is missed out:

  `example` → `exaple`

- **Insertion** where an extra character is inserted:

  `example` → `extample`

- **Substitution** where a character is *replaced* by another

  `example` → `examble`

- **Transposition** where two adjacent characters swap positions:

  `example` → `exmaple`

The chances of substitution and insertion errors are particularly influenced by the adjacency relations between keys on a keyboard. For example, on a qwerty keyboard, the letter `G` (shown in bold) has eight adjacent letters (shown in grey):[1]

On other keyboard designs such as an alphabetic or 'split' qwerty, the adjacency relations would be different:

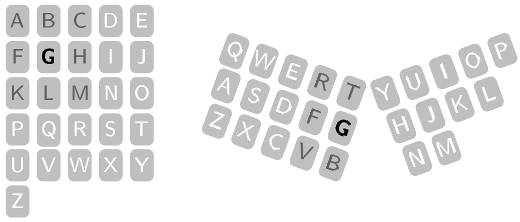

3 Spelling Machines

In order to be able to recognise spelling errors automatically, one approach is to define 'spelling machines' each of which recognises a single word.

These can be drawn using black dots and arrows between them. For example, for the word *hand*, the machine could be:

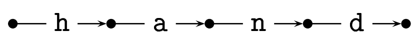

A machine is traversed by starting at the left-most dot and following the arrows remembering the letters on the arrows in sequence. At the moment, this machine recognises only the sequence of characters `hand` as representing the word *hand*. By adding more arrows to this example machine, it is possible to represent the error operations described above.

[1]For the purposes of explanation in this paper, the keyboard is considered to consist solely of the letters A to Z although the treatment given here can be easily extended to cover all keys.

Deletion If arrows are not required to have letters on them it is possible to capture deletions:

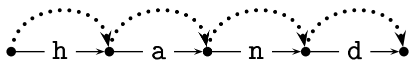

When a dotted, letter-less arrow is traversed, no letter is remembered but a new dot is still reached. So the above machine would recognise `had` and `an` as both (possibly) representing the word *hand*.

Insertion If dots can be connected to themselves, the insertion operation can be represented as well:

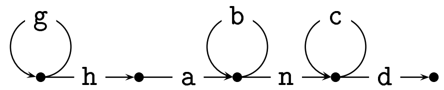

This machine would therefore recognise `ghand` and `habnd` as possible misspellings of *hand*.[2]

Substitution This error operation can be represented by adding additional arrows between each pair of points:

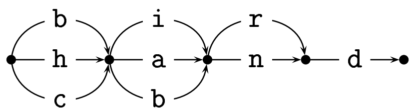

So, `bird` and `cand` would both be admitted by this machine as the word *hand* with errors.

Transposition Representing this keyboard error using these machines is slightly more complicated since extra dots are needed. This is so that the machine can remember which letters are being swapped.

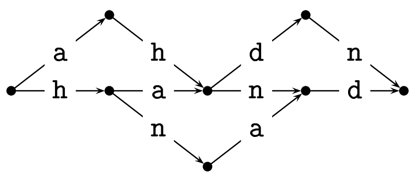

Thus `ahnd` and `hadn` are recognised as *hand* with transposition errors.

4 Generalising

Each machine would typically consist of all the necessary arrows (and dots) to capture the four error operations. Several of these machines could then be concatenated to handle the recognition of spelling errors on more than one word.

Consider the text `onside`. It could be matched by the machine for *on* followed by the machine *side* with a deletion error for the space (◇):

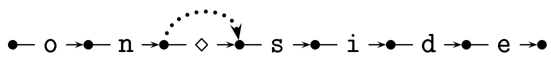

A competing match would be for the word *inside* with the initial `i` replaced by an `o`; a substitution error:

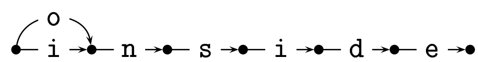

5 Conclusions

A system that implemented the approach described would score the match for each word and the best would be selected as the correction. In an interactive environment, a list of possible corrections would be presented to the user best-first and the user could select as appropriate. Automatic spelling correction without any user intervention is more problematic since if there are several possible corrections all with similar scores, there is risk that the text may be corrected to something that is still wrong or possibly worse!

[2] Each dot would in fact have connections to itself for every letter on the keyboard. For the sake of clarity, these are not all shown here.

Diffusible Neuromodulators in Real and Artificial Neural Networks

Andrew Philippides

andrewop@cogs.susx.ac.uk

Centre for Computational Neuroscience and Robotics

Abstract Recent results have shown the importance of the freely diffusing gas nitric oxide (NO) in modulation of synaptic activity. I present a review of my current research into the role of diffusing neuromodulators in both real and artificial neural networks. Results of modelling NO diffusion from a realistic structure are given and two important features are identified. The remainder of the paper describes the incorporation of a diffusible neuromodulator into an artificial neural network (ANN).

1 Introduction

The discovery that the gas Nitric Oxide (NO) is a neuronal signalling molecule has radically altered our thinking about how information is transmitted in the brain (Hölscher, 1997). Traditionally neurotransmission is thought to be spatially and temporally restricted and from the pre-synaptic to the post-synaptic neuron. However the release of NO does not require specialized point-to-point synaptic contacts and unlike traditional neurotransmitters, NO can diffuse through cell membranes. NO may therefore act without the need for conventional synaptic connectivity and its action is not necessarily locally confined to the immediate post-synaptic neuron (Hartell, 1996). Once synthesized, NO diffuses in three dimensions away from the site of synthesis regardless of intervening cellular or membrane structures. A model of NO spread could, therefore, potentially provide a theoretical framework for evaluating the signalling capacity of NO in the brain.

2 Modelling the diffusion of Nitric Oxide in Real Neural Networks

Although models of NO diffusion in the brain have been published (Wood & Garthwaite, 1994; Lancaster, 1994) almost all of this work has concentrated on modelling the instantaneous activation of NO synthesis and the spread of NO from point sources. In this approach, one assumes that the source structure is unimportant and that diffusion of NO from it is as if from a point-source situated at its centre. However, such approximations are far too gross leading to flawed conclusions (Philippides, Husbands, & O'Shea, 1998; Vaughn, Kuo, & Liao, 1998) and are unable to model hollow sources - many of which are found in the brain (O'Shea, Colbert, Williams, & Dunn, 1998) - correctly. Moreover, such an approach is unnecessary as methods developed in the field of thermodynamics (Carslaw & Jaeger, 1959) are readily applicable to modelling diffusion from realistic structures.

These techniques were used to model a hollow spherical NO source which approximates a spherical neuron whose cytoplasm synthesises NO but whose nucleus does not, as are found in some insect brains (Park, Straub, & O'Shea, 1998). Three main results were obtained from this work: firstly, the crucial role of source structure in the diffusional process was highlighted; secondly, it was seen that a 'reservoir' of NO can build up in the centre of sources which leads to a high central concentration which persists long after the causal signal; thirdly, this reservoir was shown to cause a delay in the rise in NO concentration at points away from the source. A full description of these and other results and of the techniques used to generate them see Philippides et al. (1998), Philippides, Husbands, and O'Shea (1999).

3 From Real to Artificial Neural Networks

In the same way that artificial neural networks (ANNs) have been used as an abstract model of the traditional connectionist picture of neural signalling, the concept of a diffusible neuromodulator has been incorporated into an ANN giving rise to a new style of networks called 'GasNets'. Here, in addition to the underlying ANN in which positive and negative 'signals' flow between units, an abstract process loosely analogous to the diffusion of gaseous modulators is at play. Some units can emit 'gases' which diffuse and are capable of modulating the behaviour of other units by changing parameters governing the transfer functions of the nodes. This form of modulation allows a kind of plasticity in the network in which the intrinsic properties of units are changing as the network operates. The networks function in a 2D plane; the geometric layout is crucial to the way in which the 'gases' diffuse and affect the properties of network nodes. Each unit has a genetically set position on this plane, and 'diffusion' from a unit spreads out from this site position to affect other units. See Husbands, Smith, Jakobi, and O'Shea (1998), Husbands (1998) for fuller details.

GasNets have been successfully used as robot controllers for a number of tasks (Husbands et al., 1998; Husbands, 1998). However, the model of diffusion used is very simplistic with gas building up and decaying linearly on the same time-scale as the causal signal. Having seen that this is not the case in real systems, I am currently incorporating a more realistic and potentially more powerful concept of diffusion into the GasNet framework. This involves unlinking the time-course of the build-up and decay of NO from the causal signal, incorporating a delay in the rise in concentration for points away from the emitting node and allowing greater areas of the network to be influenced by gas the longer the source node emits gas for. Provisional results for these networks applied to the problem of generating walking controllers for hexapod robots are very promising.

4 Conclusion

A description of results from modelling diffusion of NO in real systems has been given and it has been shown how these results can be incorporated into a network utilising gas driven modulation of the node properties. It is hoped that this enterprise, at the interface of neuroscience and engineering, can not only aid our understanding of how brains work but can also help spawn a new generation of adaptive machines.

Thanks to the BBSRC for grant IR3521-1 and BT for sponsorship of the BBSRC CASE studentship.

References

Carslaw, H. S., & Jaeger, J. C. (1959). *Conduction of heat in solids.* Oxford University Press.

Hartell, N. A. (1996). Strong activation of parallel fibers produces localized calcium transients and a form of ltd that spreads to distant synapses. *Neuron*, *16*, 601–610.

Hölscher, C. (1997). Nitric Oxide, the enigmatic neuronal messenger: its role in synaptic plasticity. *TINS*, *20*, 298–303.

Husbands, P. (1998). Evolving robot behaviours with diffusing gas networks. In P. Husbands & J. A. Meyer (Eds.), *Evolutionary robotics* (Vol. 1468, pp. 71–86). Springer-Verlag.

Husbands, P., Smith, T., Jakobi, N., & O'Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, *10*(4), 185–210.

Lancaster, J. (1994). Simulation of the diffusion and reaction of endogenously produced nitric oxide. In *Proceedings of the national academy of science usa* (Vol. 91, pp. 8137–8141).

O'Shea, M., Colbert, R., Williams, L., & Dunn, S. (1998). Nitric oxide compartments in the mushroom bodies of the locust brain. *NeuroReport*, *9*, 333–336.

Park, J.-H., Straub, V., & O'Shea, M. (1998). Anterograde signalling by nitric oxide: characterization and *in vitro* reconstitution of an identified nitregic synapse. *Journal of Neuroscience*, *18*(14), 5463–5476.

Philippides, A., Husbands, P., & O'Shea, M. (1998). Neural signalling: It's a gas! In L. Niklasson, M. Boden, & T. Ziemke (Eds.), *Proceedings of the 8th international conference on artificial neural networks.* Springer-Verlag.

Philippides, A., Husbands, P., & O'Shea, M. (1999). *Four-dimensional signalling by nitric oxide: a computational analysis.* (Submitted)

Vaughn, M., Kuo, L., & Liao, J. (1998). Estimation of nitric oxide production and reaction rates in tissue by use of a mathematical model. *American Journal of Physiology*, *274*, H2163–H2176.

Wood, J., & Garthwaite, J. (1994). Models of the diffusional spread of nitric oxide: Implications for neural oxide signalling and its pharmacological properties. *Neuropharamacology*, *33*, 1234-1244.

Games Logic Plays

Ahti Pietarinen

`ahtivp@cogs.susx.ac.uk`

School of Cognitive & Computing Sciences

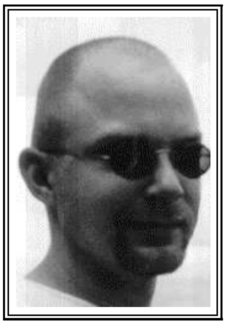

Games are an inherent part of our intellectual, social and cognitive life. Their pull can be perceived everywhere, and we indeed are experienced in dealing with (ir)rational decisions, exploiting strategies, as well as acting on the basis of our and others' preferences.

Logic, on the other hand, has a lot to gain from game theory. Usual semantic methods of logic only describe the conditions under when the statements of logic are true and when false, without giving any explanation why the formulas obtain their truth-values. Traditional semantics also is static and linear. It does not directly uncover the underlying evaluation processes, nor is it adaptive enough for changing contexts or dynamic regulations on how information flows within the statements.

This picture is rapidly changing. In my work, I investigate what game theory, especially imperfect information games, can do for logic. Essentially, a game G is defined for each model $\mathcal{M}$ and a formula φ of a language L between two players, the verifier V (trying to show that the sentence is true in $\mathcal{M}$) and the falsifier F (trying to show that the sentence in false in $\mathcal{M}$). One can then show that the sentence φ is true (false), iff V (F) has a winning strategy in $G(\varphi)$.

By varying the characteristics of the game one can get new and interesting languages whose formulas are not reducible to traditional linear statements. One can also extend games to modal logics by setting possible worlds amongst the players' semantic resource, and derive novel languages accordingly, with expressive powers coming close to second-order logic.

Besides being a powerful tool in logic, games have methodological advantages over traditional semantics. We can now safely say that all human cognitive activities can be modelled, if not explained, by playing logical games. This is because of the abstract role of the players as playing the game of seeking and finding: existence becomes a cognitive ability of uniting experiences in consciousness, given the stimuli ('there exists x' = 'one can find x'); connectives become rational decisions (a disjunction $\vee$ = 'follow only one path,' a conjunction $\wedge$ = 'work at both fronts'). All this and much more can now be captured by purely logical means.

The rest of the my work involves such concepts as the nature of games of imperfect information and games of imperfect recall (memory) in logic, information flow in games, informational independence in logic and in natural language, information flow and mutual dependencies in quantifiers and modal notions, concurrent and dynamic games in logic, semantic games for epistemic logic (logics of knowledge and belief), and various linguistic issues such as semantics of negative polarity items, branching quantifiers, and intentional identities.

LECOBA: A Learning Companion for Binary Boolean Algebra

Jorge Ramirez-Uresti

jorgeru@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

My research is in the area of Learning Companion Systems (LCS) (Chan & Baskin, 1988). An LCS is a variation of an Intelligent Tutoring System (ITS) where besides the tutor and the student a third agent is added: a Learning Companion (LC). The role of the LC is to be a peer for the human student and help her as another student would do. For example, the companion could be a role model, both students could collaborate and compete as equals, the companion could be an student of the human student, the companion could be a source of advice, etc. LCSs are relatively new systems so there are many questions to be answered. In particular, the expertise and behaviour of the companion must be carefully chosen so it can help a human student in her learning activities.

Research has shown that students learn more and better when they have the opportunity to teach other students (Berliner, 1989; Goodlad & Hirst, 1989). A student who teaches another student will have to revise, clarify, organize, and reflect on her own knowledge in order to be able to teach it, i.e. the student will need to master the knowledge. Based on this, I want the explore the hypothesis that a LC with less knowledge than the human student will help the student to learn by encouraging her to teach the LC.

I have developed a LCS in the domain of Binary Boolean Algebra called LECOBA (Ramírez Uresti, 1999). The tutor teaches the students the rules and laws of boolean algebra and how to use them to simplify boolean expressions. Two types of LC are implemented: one with low expertise (weak) and one with high expertise (strong). The issue of how to motivate the student to put effort into teaching a weak LC is tackled in LECOBA. There are two modes of interaction between the companion and the student: motivated and free. In 'motivated interaction' scores are used to encourage the student to interact with the LC. In 'free interaction' there is no such pressure; the student is only told that it is beneficial for her to interact with the companion. In either case, either the student or the companion solve problems and seek/propose justifications for each move from the other agent. So the student will be able to give the LC suggestions, ask it for justifications and, most importantly, to teach it (Ramírez Uresti, 1998).

Experiments using a 2x2 design (weak and strong companions vs. motivated and free interactions) have been conducted and are being analysed. It is expected that motivated interaction between a weak LC and a human student will be the most beneficial of the interactions by encouraging the student to teach the LC and therefore learn more efficiently.

Acknowledgements

The author wishes to thank *Consejo Nacional de Ciencia y Tecnologia (CONACYT)* of México for its financial support.

References

Berliner, D. (1989). Being the teacher helps students learn. *Instructor*, *98*(9), 12-13.

Chan, T.-W., & Baskin, A. B. (1988). "Studying with the prince" the computer as a learning companion. In *Its-88: Intelligent tutoring systems* (p. 194-200). Montreal, Canada.

Goodlad, S., & Hirst, B. (1989). *Peer tutoring: A guide to learning by teaching.* London: Kogan Page.

Ramírez Uresti, J. A. (1998). Teaching a learning companion. In G. Ayala (Ed.), *Proceedings of the international workshop "current trends and applications of artificial intelligence in education". the fourth world congress on expert systems* (p. 83-89). Monterrey, México: ITESM.

Ramírez Uresti, J. A. (1999). The LECOBA Learning Companion System: Expertise, Motivation, and Teaching. In L. Baggott & J. Nichol (Eds.), *Intelligent computer and communications technology: Teaching & learning for the 21st century. peg99, ninth international peg conference* (p. 193-201).

Everything about a je-ne-sais-quoi of "pas grand chose" (Abridged)

Fabrice Retkowsky

fabricer@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

Abstract bla bla bla PhD bla bla experiments bla bla bla writing up bla bla bla subjects bla bla quite good fun in the end.

1 Introduction

I started my PhD in January 1997. Instead of talking in too many details about my current experiments, my IoT talk was about, well, what I did in the last 2 1/2 years (a.f.a. work is concerned, that is).

2 First year

I spent most of it trying to overcome a certain feeling of under-motivation: I had started this PhD because I loved Brighton and wanted to stay here for a little bit longer. I still managed to collect loads of ideas, too many ideas, and couldn't choose what my research would focus on (e.g. which experiments to conduct). After one year I decided to build a program that would help me conducting experiments. I called it Igor (don't bother).

3 Second year

Spent developing the program, and starting my first experiment: evaluating it. Most frustrating aspect: whatever I do, whichever small aspect of my domain I limit myself to, I will only be able to 'glimpse' at it (that's the vain effort to explain the talk's title). Another problem: finding subjects!

4 Third year

Finished the evaluation, and currently doing my second experiment. I should have time to conduct another one this autumn. Writing-up... Everybody told me if would take *six* months, and it does! Best thing to do is to start very early (e.g. beginning of third year), and do it in parallel to research work, little by little.

5 Conclusion

In 4 months, 4 days, 12 hours and 20 minutes. Hopefully.

NB: the subject of my PhD is "The Cognitive Aspects of Software Reuse".

Focal structures in Prolog

Pablo Romero

juanr@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

My original interest had to do with developing a learning environment for Prolog programming. I soon found that, in trying to build such environment, a good knowledge of the mental model programmers build of the programs the are working with is vital in order to help them in their learning process. However, little is known about this mental representation for the case of Prolog. In this way, the objectives of my project changed soon early on and I became more interested in investigating this relatively unexplored area. The global aim of my research could be expressed as an investigation into the characteristics of the mental model Prolog programmers build when doing program comprehension. In particular, which could be a good model to explain the structural side of these mental models.

2 Program comprehension and focal structures

Several studies have proposed that a concept known as 'Programming Plans' (Pennington, 1987; Gilmore & Green, 1988; Davies, 1990) can be used to explain aspects of the programmers' structural knowledge. These studies suggest that there is a strong link between the mental model that a programmer builds and an organisation of the code according to a Plan-like structure. Plans are proposed as the external analogue of the programmers' internal structural knowledge that is used to organise the program as a hierarchy of meaningful units. These units are considered as frames that comprise stereotypical programming procedures and whose slots can be filled with variables related to the specific problem being solved. In this way, Plans can be seen as Data Structures that represent generic concepts stored in memory. Studies of Programming Plans have considered mainly procedural languages. Some studies have tried, without much success, to find evidence of a relationship between Plans and Prolog programmers' mental models (Bellamy & Gilmore, 1990; Ormerod & Ball, 1993). There are alternative structural models for Prolog. Bowles and Brna (1993) and Gegg-Harrison (1991) propose that Prolog programmers' structural knowledge is related to 'Prolog Techniques' and 'Prolog Schemas' respectively.

Some studies have suggested that the elements that comprise these units of meaningful information have different degrees of relevance or saliency for programmers (Rist, 1989; Davies, 1993, 1994; Rist, 1995). The elements that directly encode the goal of a particular Plan are said to be focal for experienced programmers. For example, if the programming task is to accumulate data and compute an average, the key or focal element of it will be the place where the division between the running total and the number of items takes place (Rist, 1995). Although Plans have been related to data-flow as well as to functional information, focal elements of Plans seem to be highly related to function only.

According to Pennington (1987), program comprehension depends on the detection of several types of information. Therefore, it might be the case that there is not only one but several models that can explain the Prolog programmers' mental model. The choice of which aspects or types of information are relevant for Prolog programmers might depend on the nature of this language as information structure, but also on the programmers' level of skill and the objective of the comprehension activity among other factors.

Although this is not strictly a comparison between different programming paradigms, I have related my work to findings in the procedural paradigm. I have been comparing the issue of the nature of beacons in the code and their relation to focal structures with the findings of Wiedenbeck (1986) and Davies (1994) and the relation of this structural knowledge to performance with the work of Vessey (1989) and Gilmore (1991). Trying to explain the differences, if any, to their results is important in order to have a broader picture of the programming activity as a cognitive skill.

3 Focal structures in Prolog program comprehension and debugging

My overall research plan comprises three experiments and the reanalysis of experimental data derived from a study by Gilmore and Green. I have performed and reported two experiments and the reanalysis of the experimental data.

The reanalysis of Gilmore and Green's experimental data was useful to establish a specific methodology for the three experiments of my project. This analysis tried to find out whether the debugging and comprehension performance of programmers was related to focal structures. That is, whether errors in focal lines of programs were spotted more easily than errors elsewhere. This was found to be the case, and therefore an investigation into the structural part of the mental model of Prolog programmers using focal structures as one of its main features seemed promising.

The first experiment of my research project tried to characterise the nature of the mental models that programmers build when comprehending Prolog code by finding out which of several structural models are most relevant for the case of this programming language. To measure the relevance of a specific structural model, the experiment considered a recall task similar to the one in Wiedenbeck (1986). Subjects were asked to understand and memorise a small Prolog program, and then recall what they could of it. This code was analysed in terms of the different models of structural knowledge of Prolog and their associated key segments. The relative success of recollection of the different key segments was compared against the relative success of recollection of the rest of the program to establish the relevance of these structural knowledge models for Prolog programmers. The results suggest that Prolog Schemas, a structural model related to data structure relationships, is of central importance to Prolog programmers. This result contrasts with those obtained for procedural languages, where Programming Plans, a concept related to functional information, seems to be the dominant model.

The results of this experiment needed to be confirmed and its importance needs to be related to a more common programming task such as debugging or program modification. The experimental programs of the first experiment used cryptic variable names, and Prolog is very sensitive to naming style. It can be argued that the results of the first experiment have to do with the effect of naming style more than with a property of the programming language used. So the second experiment tried to find out whether the results of the first one can be confirmed even when different naming styles are considered.

The second experiment therefore looked at the influence of naming style on the choice of the relevant information type for Prolog programmers. This experiment tried to replicate the results of the first one, this time by applying a recognition instead of a recall task taking into account only two kinds of information types, Schemas and Plans. For this experiment, two groups of Prolog programmers, novices and experienced, tried to understand and memorise several Prolog programs. After this study time, they tried to discriminate program segments as belonging to the studied programs or not. There were significant differences for these two groups when comparing their recognition accuracy. Experienced programmers were able to recognise segments categorised as focal better than novices, regardless of kind of structure (Plans or Schemas) or of naming style. The results of this experiment indicate that at least for the case of Prolog, experienced programmers restructure the organisation of their Plan and Schema knowledge to allow that certain elements of these structures become more relevant than others.

The reason for plan information not being apparently important in the first experiment seems to be due to the difference in the two experimental tasks used in this study, recall and recognition. It has been argued that recall facilitates high rates of accuracy for Schemas but not for Plans. Therefore, it is possible that some other structural models and information types besides Plans and Schemas might be important for Prolog programmers.

4 Conclusions and further work

Experienced programmers seem to construct several views of the program, according to the different types of information implicit in the program text. It is likely that their structural knowledge for each of these aspects of the code shows an organisation in which specific elements are more relevant than others. This would imply that the notion of focal structures involves not only function but other types of programming information.

More experimentation is needed to verify these conclusions. The last experiment of my project will look at these issues from the point of view of Prolog program debugging. It is of central importance to verify the findings from the previous experiments with tasks that have practical application such as debugging. If this experiment shows that the notion of focal structures is valid for practical tasks, then its application to programming teaching can be explored and hopefully exploited.

References

Bellamy, R. K. E., & Gilmore, D. J. (1990). Programming plans: Internal and external structures. In K. Gilhooly, M. T. G. Keane, R. H. Logie, & G. Erdos (Eds.), *Lines of thinking: Reflections on the psychology of thought, vol 1.* London, U.K.: Wiley.

Bowles, A., & Brna, P. (1993). Programming plans and programming techniques. In P. Brna, S. Ohlsson, & H. Pain (Eds.), *World conference on artificial intelligence in education* (pp. 378–385). Edinburgh, UK: Association for the advancement of computing in education.

Davies, S. P. (1990). The nature and development of programming plans. *International Journal of Man-Machine Studies*, *32*, 461–481.

Davies, S. P. (1993). Models and theories of programming strategy. *International Journal of Man-Machine Studies*, *39*, 237–267.

Davies, S. P. (1994). Knowledge restructuring and the acquisition of programming expertise. *International Journal of Human Computer Studies*, *40*, 703–726.

Gegg-Harrison, T. S. (1991). Learning Prolog in a schema-based environment. *Instructional Science*, *20*, 173–192.

Gilmore, D. J. (1991). Models of debugging. *Acta psychologica*, *78*(1).

Gilmore, D. J., & Green, T. R. G. (1988). Programming plans and programming expertise. *Quarterly Journal of Experimental Psychology*, *40A*, 423–442.

Ormerod, T. C., & Ball, L. J. (1993). Does design strategy or programming knowledge determine shift of focus in expert Prolog programming? In *Empirical studies of programmers, fifth workshop* (pp. 162–186). Norwood,NJ: Ablex.

Pennington, N. (1987). Stimulus structures and mental representations in expert comprehension of computer programs. *Cognitive Psychology*, *19*, 295–341.

Rist, R. S. (1989). Schema creation in programming. *Cognitive Science*, *13*, 389–414.

Rist, R. S. (1995). Program structure and design. *Cognitive Science*, *19*, 507–562.

Vessey, I. (1989). Toward a theory of computer program bugs: an empirical test. *International Journal of Man-Machine Studies*, *30*(1), 23–46.

Wiedenbeck, S. (1986). Processes in computer program comprehension. In E. Soloway & S. Iyengar (Eds.), *Empirical studies of programmers, first workshop* (pp. 48–57). Norwood,NJ: Ablex.

THE EYE

Hanson Schmidt-Cornelius

hanson@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

The Human eye is positioned by six muscles. Two **oblique muscles** which roll the eye around the line of sight, and four **rectus muscles** which control the horizontal and vertical eye movement. The following summary introduces a mechanical eye which simulates the operation of the **rectus muscles**. An outline of the experimental setup is also provided, giving a small impression of the additionally required hardware.

2 THE EYE

There exist many active camera systems that can perform movements, similar to these found within the human extra ocular motor system. However, their control and operation often does not have much in common with natural vision systems.

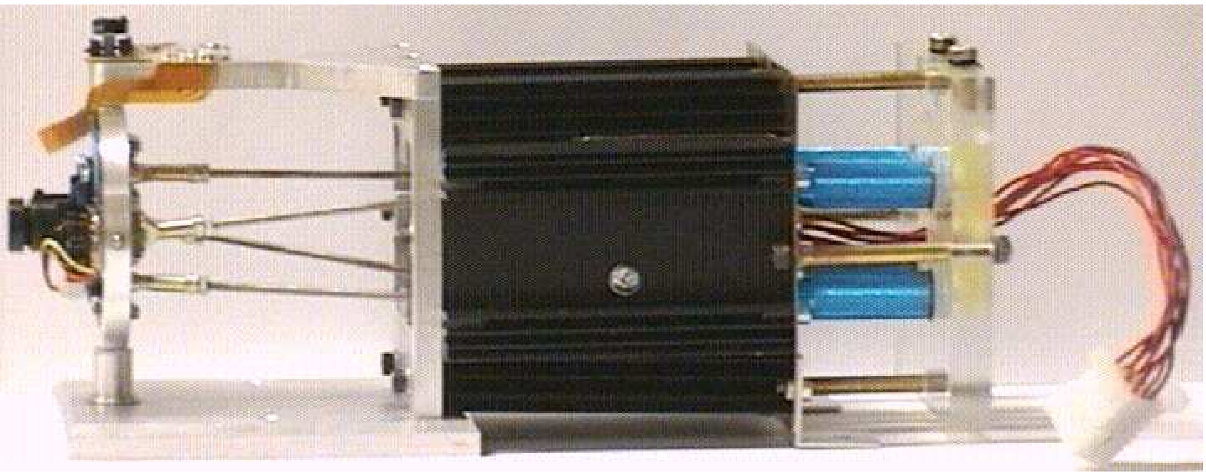

Figure 1: THE EYE

THE EYE is modelled specifically on the human eye and can subsequently perform most of the movements the human eye can perform. In figure 1 the architecture of THE EYE is clearly visible. It consists of three units. To the left, one can see the platform on which a small camera is mounted. This platform forms the first unit. From there, four rods extend which are connected to the second unit. Hidden behind black heatsinks are four powerful actuators which apply linear movements to the rods. The third unit performs mechanical damping, which can be adjusted as required. These three units actually form a small part of a large experimental setup.

3 The Setup

To activate and process information to and from THE EYE, much more hardware is required. The experimental setup in figure 2 shows briefly what is required, without elaborating how every component works and interacts with each other. The computer forms the central part, as it controls the movement of THE EYE and processes the images it returns. It also controls the location of the laser on the screen. Monitor 2 is connected to a remote controlled active camera, which allows the supervision of experiments, without the risk of eye dam-

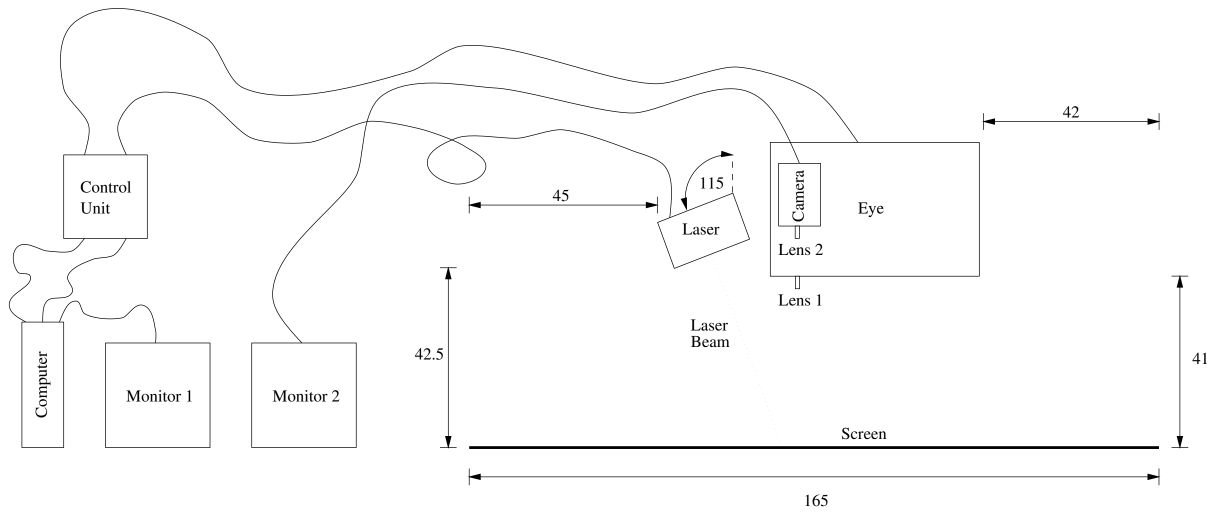

Figure 2: Experimental Setup

age due to laser radiation.

4 The Operation

The laser is used to provide visual stimulation, through illuminating parts of the screen. Images on the screen are then returned to the computer, through THE EYE. Simple image processing operations can determine where in the image the laser is located. Using this information as input, appropriate eye movements can be applied. Such as constantly following the laser dot or performing fast saccades to the location of illumination.

Evolving Mechanisms of Choice Behaviour

Anil Seth

anils@cogs.susx.ac.uk

Centre for Computational Neuroscience and Robotics

Abstract A brief description of my research is presented, the focus of which concerns minimal conditions on internal mechanisms underlying choice behaviour in animals and robots. The central themes are (a) that understanding mechanism requires an appreciation that behaviour is a joint product of agent, environment, and observer, and (b) that the complexity of internal mechanisms can be assessed in terms of the complexity of the environment in which they operate. These themes are explored by using genetic algorithms (GAs) to evolve internal mechanisms for "choosing" agents in robotic, animat, and game-theoretic contexts.

In the context of artificial life (AL) choice has been largely synonymous with the problem of *action selection,* defined by Maes (1994) as "the problem of what to do, at any given time, in order to further progress toward multiple, time-varying goals." Approaches to this problem have generally assumed the existence of internal correlates for behaviours (relatively independent of both environment and observer), and some explicit arbitration between them by some action selection "mechanism". This perspective derives from the early ethological work of Lorenz and Tinbergen, and criticism has frequently been levelled at a tendency to ignore more recent biological and ethological developments.

Further background can be found in Seth (1998a) in which I use a GA to evolve a minimal mechanism for action selection in a virtual robot (an animat). This work demonstrates that action selection can occur without the internalisation of distinct behaviours, and without any explicit internal arbitration. In Seth (1999a) I explore the *matching law,* a well-established observation (but unfamiliar in AL) that animals (and humans) tend to allocate the frequency of their responses to different stimuli (in other words, to choose) in proportion to the reinforcement obtained from each stimulus type. Importantly, the pattern of choice predicted by the matching law may often *not* be optimal, (Herrnstein, 1997). I again evolve animats with absurdly simple internal mechanisms to demonstrate that matching behaviour does not necessitate the kind of internal complexity often assumed in the biological literature. Much of the burden of explanation can be shifted onto the environment and the observer, just as in the case of action selection.

This work also provides an approach to understanding *why* organisms often match, given that matching is often sub-optimal. It turns out that matching may indeed be optimal in environments populated by *groups* of agents; irrational individual behaviour can perhaps then be considered as a result of a social evolutionary heritage. Current work is building on this foundation by looking closely at the concept of the *ideal free distribution,* which is a way of predicting the *optimal* distribution of groups of agents in environments of non-homogenous resource density. An important part of this effort is the development of a formal relationship between the use of GAs to evolve animat mechanisms and the framework of optimal foraging theory in behavioural ecology.

A practically distinct but theoretically related thread of my research has to do with *complexity;* specifically, how the evolution of internal complexity is influenced by the structure of the environment. The philosophical framework for this position derives from Ashby (1956) and Godfrey-Smith (1996), and my contribution has been to flesh out this framework with practical examples from evolutionary game-theory and evolutionary robotics. In Seth (1997) and Seth (1999b) I demonstrate that the presence of noise in an evolutionary Iterated Prisoner's Dilemma context can encourage the evolution of more complex strategies. In Seth (1998b) and Seth (1998c) I complement this work by evolving controllers for a Khepera mobile robot in environments with different structural properties, analysing the complexity of the evolved behaviours and underlying neural dynamics.

In summary, I am attempting to (a) do some real biology and psychology, using AL to investigate the links between mechanism, individual behaviour and group behaviour, and (b) fit this into a theoretical perxpective that respects the distinction between behaviour and mechanism, and relates internal complexity to environmental structure. I do *not* claim that my simple evolved creatures work in the same way that animals work, but I *do* claim that they illustrate some of the prejudices and assumptions that cloud our judgement of intelligence can arise from material of any sort.

References

Ashby, W. (1956). *An introduction to cybernetics.* Chapman Hall.

Godfrey-Smith, P. (1996). *Complexity and the function of mind in nature.* Cambridge University Press.

Herrnstein, R. (1997). *The matching law.* Harvard University Press.

Maes, P. (1994). Modelling adaptive autonomous agents. *Artificial Life*, *1*(135–162).

Seth, A. (1997). Interaction, uncertainty, and the evolution of complexity. In P. Husbands & I. Harvey (Eds.), *Proceedings of the 4th european conference on artificial life (ecal 97)* (pp. 521–530). MIT Press.

Seth, A. (1998a). Evolving action selection and selective attention without actions, attention, or selection. In R. Pfeifer, B. Blumberg, J. Meyer, & S. Wilson (Eds.), *Proceedings of the 5th international conference of the society for adaptive behaviour (sab 98)* (pp. 139–147). MIT Press.

Seth, A. (1998b). Noise and the evolution of complexity; a study in evolutionary robotics. In P. Husbands & J. Meyer (Eds.), *Proceedings of the 1st european workshop on evolutionary robotics (evorob 98).* Springer-Verlag.

Seth, A. (1998c). The evolution of complexity and the value of variability. In C. Adami, R. Belew, H. Kitano, & C. Taylor (Eds.), *Proceedings of the 6th international conference on artificial life (alife vi)* (pp. 209–221). MIT Press.

Seth, A. (1999a). Evolving behavioural choice: an investigation of herrnstein's matching law. In D. Floreano, J.-D. Nicoud, & F. Mondada (Eds.), *Proceedings of the fifth european conference on artificial life (ecal99).* (In press)

Seth, A. (1999b). *A cybernetic perspective on the role of noise in the iterated prisoner's dilemma* (Tech. Rep. No. 506). University of Sussex.

Evolutionary Approaches to Adaptive Protocol Design

Nick Sharples

nichs@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Introduction

This work explores the possibility of applying evolutionary search strategies to the synthesis of network communications protocols. Specifically, the adaptive algorithms employed by transport layer protocols for packet switched networks. The primary objective is the development of a reusable methodology, or conceptual framework, which will ease the development of this *adaptive* element. Obviously, a protocol designer must anticipate a wide variety of scenarios and incorporate strategies for coping with them as they arise. However, it is extremely difficult to develop a protocol which behaves optimally and is non-detrimental to the network as a whole. Since, the emergent dynamic from the interaction of the communicating processes becomes difficult and maybe impossible to predict. This paper will begin with a short description of network congestion; a well know phenomenon which if undetected can drastically reduce network bandwidth. The algorithms used by the sender process to avoid and control congestion are definitively adaptive and the synthesis of such protocols using this technique is the ultimate goal of this work. The paper will conclude with a brief discussion of the first stage in development of this idea and some initial results.

2 Objective

At a purely software level the Internet is composed of many small interacting processes, each of which has a goal that is clearly defined at the level of that process, i.e. that of communicating data to a receiving peer process. At a higher level the collective interaction of those processes is a direct result of the rules that govern a single communication. But, unless we are careful at the lower level the emergent dynamic at that higher level may be unpredictable, leading to network congestion and finally collapse.

Congestion occurs when a router begins to receive more incoming packets than it can service on its outward path (a bottleneck). As more packets arrive, the router's buffer begins to fill to the point where all the space is allocated. Any new packets arriving after this point are discarded and we are now in a state of congestion. Without congestion control, a non-adaptive sender would attempt retransmition of the lost packets which will only add to the state of congestion. If this continues, a state know as congestion collapse occurs in which the network experiences a sudden factor-of-thousand drop in bandwidth. So, failure to infer loss as congestion only compounds the problem, leading to this exponential decrease in the end-to-end route throughput, know as congestion collapse.

To avoid congestion we require a protocol which has the ability to adapt its behaviour to the changing networks conditions. On detecting loss it must slow the injection of packets into the network, but it must also make optimal use of the precious bandwidth available to it. It is this later point, which conflicts so strongly with congestion, that makes the development of adaptive protocols such a difficult task for the designer.

3 Hypothesis

From the perspective of a protocol, the network can be viewed as its virtual environment. The protocol being a communicating entity of that environment, whose actions are dictated by its underlying control architecture. An end point protocol communicating with its peer can infer congestion in the end-to-end route, through lost packets. That is, it has the ability to sense changes in its environment. Since a protocol, in its simplest form, is a control architecture. Can we systematically search the space of possible architectures to discover optimal solutions?

4 Methodology

The methodology proposed here, centres around the ability to search the design space of possible protocols. The key element then, is the method used to explore that space. A popular search strategy, widely employed for the development of control architectures, is genetic algorithms (GAs). GAs are a non-domain specific adaptive search strategy, and are used as a highly effective optimisation tool. The concept, inspired by biological evolution, has become increasingly popular recently for the optimisation of dynamical problems. The process begins with an initial population, of randomly generated individuals. Each individual is a suitably encoded solution to the problem at hand. Typically, the solution, or genome, is encoded as a binary string to allow easy manipulation by genetic operators, and converted to a usable representation, the phenome, prior to evaluation. For linear optimisation each individual is evaluated against a fitness function, which provides a measure of its ability to solve the problem. For non-linear problems, such as those faced by adaptive systems we must evaluate the functionality it exhibits in the working environment. However, for many applications, but notably not the application we propose, this is impractical due to the amount of time an evaluation may take. Instead, an individual's functionality is assessed in simulation. Individuals are now selected for breeding, selection being proportional to fitness. That is, fit individuals have a better chance of selection over poor scoring individuals. This is important to allow a thorough exploration of the search space. During breeding fit solutions are combined, and possibly mutated, to produce an offspring. Each offspring is placed into a new population and breeding continues until the new population is full, at which point the old population is replaced with the new. Successive generations produce fitter and fitter solutions will converge on an optimal solution.

At the heart of the methodology proposed is the concept of placing an operational representation of a protocol in a simulation of its working environment. By allowing the simulation to execute for a number of time steps, as it would in the real world, its performance can be evaluated either during or after the simulation, depending on task. It is important to recognise that the simulation need only model those aspects of the system which are relevant to the particular behaviour we endeavour to produce. Under certain conditions this may mean that some aspects of the real working environment are abstracted, provided the abstraction can be justified. We must also be careful to ensure that each aspect parameter we model is equal for each execution of the simulation.

Using the simulation as a method of evaluation we now require a representation of the working protocol which can be manipulated through the search space of possible protocols. The obvious choice for this representation is a variation on the simple finite state machine, aptly named the communicating finite state machine. Since the beginnings of protocol engineering the communicating state machine has provided a simple model to describe the execution of communicating processes. It also lends itself easily to encoding schemes which can be genetically manipulated in the search space.

5 Stage one

The goal of the first stage of development was to produce a system which could generate protocols capable of communicating data over an unreliable medium. To that end two experiments were performed, the first ensured consecutive transmition attemps always succeded and was used to test the optimisation side of the software produced. The results of this experiment can be seen in figure 1. The second allowed message loss over consecutive transmitions and is presented in figure 2.

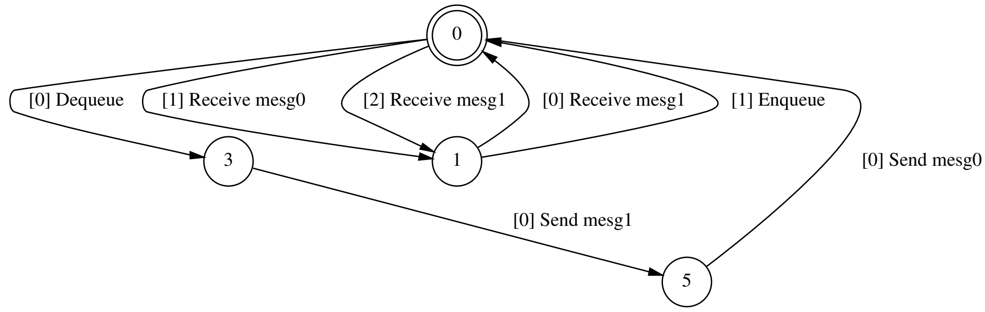

Figure 1: Evolved finite state machine, messages are only lost once

As you can see in the following figures each transition begins with a numerical value in square brackets which indicates the order in which each transition is evaluated. In effect each transition forms a boolean expression. If the expression holds the machine makes the transition to the next state and remains in its current state if no transition can fire. Each transition also has an associated action:

- Send : Place a message on the communications channel, taking data from memory if required. False if no data available to send, in other circumstaces true.
- Receive : Receive a message from the communications channel and place the data into memory. True only if message available on channel.
- Enqueue : Place the item in memory, into receive buffer. True only if item in memeory.
- Dequeue : Place the first element of the transmit buffer into memory. True only if item available in transmit buffer.
- Null : Always true.

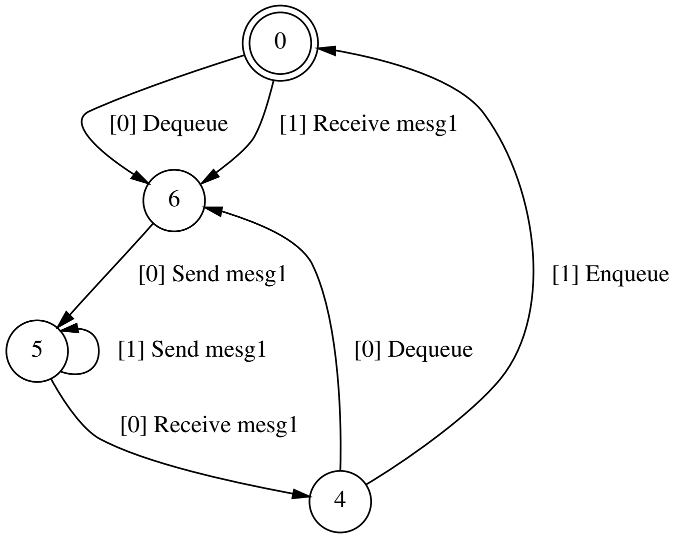

Figure 2: Evolved finite state machine, messages are lost multiple times

Biologically Inspired Techniques for Robotics

Tom Smith

toms@cogs.susx.ac.uk

Centre for Computational Neuroscience and Robotics

This PhD aims to explore, in the context of generating complex behaviour, the use and incorporation of biologically inspired techniques and processes to augment the standard models of Artificial Neural Networks. Two areas in particular will be emphasised. First, the development of networks based on spiking neurons (Maass & Bishop, 1999), formally shown both to be equivalent to integrate-and-fire neurons (Gerstner, 1999), and pound-for-pound more powerful than the continuous input-output neurons used in more conventional neural networks (Maass, 1997). Second, 'lifetime' learning, or adaptation, will be explored using a number of different scenarios; associative Hebbian-style synaptic adaptation will be looked at alongside less well explored models of neuronal property modulation (LeMasson, Marder, & Abbott, 1993; Levy, Colbert, & Desmond, 1990; Husbands, Smith, Jakobi, & O'Shea, 1998). Various forms of input adaptation will also be investigated. Other areas of interest outlined will include investigation of networks of heterogenous neurons, morphogenesis and compact encoding of networks, and network architecture.

The driving force behind the project is the development of robotic controllers able to operate over extended periods of time in a noisy dynamic environment. These are not easy to design. The problems lying behind the hand-coding of controllers able to cope with such noisy data are immense. Even at a single time point, the system is likely to need to interpret data not describable by low-order statistical parameters, with non-Gaussian distributions and nonlinear relationships between data elements. A bigger difficulty is predicting how this input will change over time (the underlying statistical distribution of the data may also be changing): systems will thus need to adapt over their lifetime of operation. The class of Artificial Neural Networks (ANNs), loosely based on simple models of biological nervous systems, possess those properties required to deal with time-changing noisy and ill-defined data (Kohonen, 1997; Bishop, 1995; Hertz, Krogh, & Palmer, 1991; Rumelhart, McClelland, & PDP Research Group, 1986), and are thus suitable for control tasks and nonlinear estimation in which more traditional probabilistic methods fail.

Recent work has mainly concentrated on 'one-shot' mapping of an input to output without any retained 'memory' of activity over time: the ANN property of being able to approximate any function (Kolmogorov, 1957; Funahashi, 1989), and proofs that training algorithms will converge to solutions (Rumelhart et al., 1986) show that any single set of vectors can be mapped onto another arbitrary set. These proofs combined with a growing body of empirical knowledge on the use of such methods have meant that ANNs are now a powerful tool in pattern recognition, function approximation and other problems with well-defined goals. However, it must be emphasised that the same cannot be said for the generation of complex behaviour over time. Major difficulties arise from having a much vaguer knowledge of the required input-output mapping at any one time, with the added possibility that such mappings may also be dependent on recent activity. The usual training methods are not easily applied to networks used as robot control systems.

However, the loss of good algorithmic training methods should not be viewed with despair. The class of *generate-and-test* random search processes can provide solutions where *gradient descent* methods fail, and are well equipped to avoid local sub-optimal solutions (Aarts & Lenstra, 1997; Kirkpatrick, Gelatt, & Vecchi, 1983). Such methods do not come under the umbrella of algorithmic training, in that no method is employed to predict how to reduce the output error (the gradient descent mentioned above). A solution (network) is encoded, and evaluated on some performance criterion. The encoding undergoes some kind of random change to produce a second network which is typically compared with the first, and the best kept. Without the need to predict how to reduce the output error, we are not restricted to well-understood networks amenable to formal analysis, and such methods as multi-start hill climbing (Aarts & Lenstra, 1997), simulated annealing (Aarts & Korst, 1989; Kirkpatrick et al., 1983), and evolutionary computation (Holland, 1992; Koza, 1992) can in principle be used to produce networks of very different operation to more traditional techniques.

The research in this PhD is targeted at identifying useful neural network mechanisms for the generation of complex behaviours, over and above the traditional processes of activity flowing along 'wires' between nodes. In particular, we are free to explore alternative models based on other biological processes: recent work in this area has invesigated models of diffusing neurotransmitter overlaid on a more traditional network (Husbands et al., 1998; Husbands, Smith, & O' Shea, 1999). It is unclear what processes may be useful for the generation of such behaviour, with little published research, but plausible mechanisms in real biological systems include more complex integrate-and-fire neurons, modulation of network properties via secondary processes, and lifetime adaptation. This PhD aims to investigate such processes.

References

Aarts, E., & Korst, J. (1989). *Simulated annealing and boltzmann machines.* John Wiley and Sons.

Aarts, E., & Lenstra, J. (Eds.). (1997). *Local search in combinatorial optimization.* John Wiley and Sons.

Bishop, C. (1995). *Neural networks for pattern recognition.* Oxford University Press.

Funahashi, K. (1989). On the approximate realization of continuous mappings by neural networks. *Neural Networks*, *2*(3), 183–192.

Gerstner, W. (1999). Spiking neurons. In W. Maass & C. Bishop (Eds.), *Pulsed neural networks* (pp. 3–53). MIT Press / Bradford Books.

Hertz, J., Krogh, A., & Palmer, R. (1991). *Introduction to the theory of neural computation.* Addison-Wesley.

Holland, J. (1992). *Adaptation in natural and artificial systems* (2nd ed.). MIT Press/Bradford Books.

Husbands, P., Smith, T., Jakobi, N., & O'Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, *10*(4), 185–210.

Husbands, P., Smith, T., & O' Shea, M. (1999). Robot nervous systems based on evolved gas modulated networks. *Robotics and Autonomous Systems*. (In Press)

Kirkpatrick, S., Gelatt, C., & Vecchi, M. (1983). Optimization by simulated annealing. *Science*, *220*, 671–680.

Kohonen, T. (1997). *Self-organizing maps* (2nd ed.). Springer-Verlag.

Kolmogorov, A. (1957). On the representations of continuous functions of many variables by superpositions of continuous functions of one variable and addition. *Dokl. Akad. Nauk SSSR*, *114*, 953–956.

Koza, J. (1992). *Genetic programming.* MIT Press/Bradford Books.

LeMasson, G., Marder, E., & Abbott, L. (1993). Activity-dependent regulation of conductances in model neurons. *Science*, *259*, 1915–1917.

Levy, W., Colbert, C., & Desmond, N. (1990). Elemental adaptive processes of neurons and synapses: A Statistical/Computational perspective. In M. Gluck & D. Rumelhart (Eds.), *Neuroscience and connectionist theory* (pp. 187–236). Lawrence Erlbaum.

Maass, W. (1997). Noisy spiking neurons with temporal coding have more computational power than sigmoidal neurons. In M. Mozer, M. Jordan, & T. Petsche (Eds.), *Advances in neural information processing systems 9* (pp. 211–217). Cambridge: MIT Press.

Maass, W., & Bishop, C. (Eds.). (1999). *Pulsed neural networks.* MIT Press / Bradford Books.

Rumelhart, D., McClelland, J., & PDP Research Group the (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. volume 1: Foundations.* MIT Press.

Stop Rules for Catastrophic Worrying

Helen Startup

helenst@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

Abstract Catastrophic worrying involves the process by which worriers perceive progressively worse and worse outcomes to a specific worry topic, and this is usually the result of them posing automatic questions of the 'what if....?' kind. This paper describes a pair of studies which show that the perseverative iterative style that worriers posses is linked to an interaction between a) the implicit stop rules inherent in a task and b) negative mood.

1 Study 1

Study 1 employed a mood induction procedure to investigate the effects of induced negative, positive and neutral mood states on the tendency of participants to both catastrophise and 'reverse' catastrophise a hypothetical scenario (being the Statue of Liberty). The catastrophising interview involves the experimenter asking the participant 'what is it that worries you about being the Statue of Liberty?'. If the participant replies, for example, "people will be walking around my head", the experimenter then asks "what is it that worries you about people walking around you head?". The participants' answer is thrown back in the form of a question until the participant fails to come up with any more responses. For the 'reverse catastrophising' procedure a similar format is adopted, however instead of asking what worries the participant about a posed scenario the experimenter asks them " what is it that is good about being the Statue of Liberty" and so on, until no more responses are offered.

Results indicated that compared to the positive and neutral mood induction procedures, the negative mood induction resulted in significantly higher post-induction ratings of 'negative affect' (composite of sadness and anxiety ratings), and lower ratings of happiness. These group differences were reflected by participants in the negative mood induction generating significantly more steps than the positive and neutral mood groups in both the catastrophising and 'reverse catastrophising' procedures.

The finding that negative mood induction increases the number of steps emitted in a 'reverse catastrophsing' procedure runs contrary to the predictions of a mood congruency explanation (Sinclair & Mark, 1992). Why should participants in a negative mood persist at iterating what is good about something? The findings, however, are consistent with a mood-as-input explanation (Martin, Ward, Achee, & R.S.Wyer, 1993). If the catastrophising and 'reverse catastrophsing' procedures are viewed as problem-solving tasks in which participants are attempting to analyse a problem and this generates an 'as many as can' stop rule typical of systematic processing (Martin & Davis, 1998), then negative mood will influence judgements about whether the individual believes they have satisfactorily completed the task or not. Specifically, the negative mood may provide the individual with information that the task may not yet have been satisfactorily completed, resulting in perseveration regardless of the valency of the iterative task.

2 Study 2

Study 2 attempted to further tease out the role of negative mood on the iterative potential of worriers and nonworriers. Specifically, this study attempted to asses a) the implicit stop rules that worriers utilise during catastrophising and b) the valency of their mood during catastrophising. For example, for worriers to persist in generating catastrophising steps for longer than non worriers, an implication is that a) worriers are using an 'as many as can' stop rule and b) are in a more negative mood than nonworriers (the less intuitively plausible hypothesis is that worriers are using a 'feel like continuing' stop rule and are in a more positive mood than nonworriers).

Theses hypothesis were tested by asking groups of high and low worriers to perform either a positively valenced item generation task (generate reasons why being the manager of a hotel is a good thing) or a negatively valenced item generation task (generate reasons why being a manger is a bad thing). And to persist either according to an 'as many as can' stop rule or a 'feel like continuing' stop rule. In addition levels of trait and state mood were recorded for each participant.

The results of study 2 offered further support for a mood-as-input explanation of catastrophic worrying. Using a standardised item generation task, under the 'as many as can' stop rule, worriers generated significantly more reasons than non worriers, and this finding was unaffected by the valency of the task. Conversely, under a 'feel like' continuing stop rule, worriers generated fewer reasons than nonworriers. Analyses of state and trait mood measures taken also suggested that worriers have elevated levels of both state and trait negative affect and lower levels of positive affect compared to low worriers.

The results of the present studies provide evidence for a mood-as-input interpretation of catastrophic worrying (Martin & Davis, 1998). This approach moves beyond the suggestion that processing is governed solely by mood congruency mechanisms and suggests that mood may provide information about continuing or terminating a problem-solving task which is interpreted in the context of the stop rules for the task. It is argues that the negatively valenced affect worriers bring to the problem-oriented task interacts with the implicit stop rules they use to generate a perseveration that results in exacerbated catastrophising.

References

Martin, L., & Davis, B. (1998). Beyond hedonism and associationism: A configural view of the role of affect in evaluation, processing, and self-regulation. *Motivation and Emotion*, *22*, 33–51.

Martin, L., Ward, D., Achee, J., & R.S.Wyer. (1993). Mood as input: People have to interpret the motivation implications of their moods. *Journal of Personality and Social Psychology*, *63*(317–326).

Sinclair, R. C., & Mark, M. M. (1992). The influence of mood state on judgement and action: Effects of persuasion, categorisation, social justice, person perception and judgmental accuracy. In L. Martin & Tesser (Eds.), *The construction of social judgements* (pp. 165–193). Hillsdale, NJ: Lawrence Erlbaum.

HIPO - Hardware Independent Parallel Optimisation

Andrew Stevens

andrewst@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

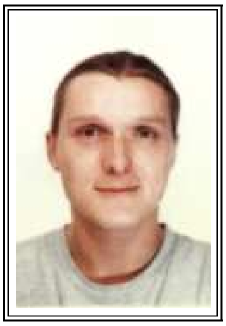

Abstract Modern computer architectures are tending towards parallel processing for increased performance and so modern compilers must attempt to make use of these resources. My research looks at compiler optimisation techniques applied to Java programs that can automatically extract parallelism from sequential code. This paper will outline the need for this research and give an overview of my work to date.

1 Introduction

Due to the physical limits in the advances of processor technology, future architectures will make use of multiple CPUs for increased performance to meet the demands of modern software (Zima & Chapman, 1990; Kamin, 1995; Vries & Lee., 1993; Keleher & Tseng, 1996). The operating systems for these machines will handle the efficient deployment of processes amongst the processors, but the job of exploiting the benefit of parallel execution in each process rests with the programmer and/or compiler. Since most applications are written either without the knowledge of the target architecture or are designed to be platform independent, there is a clear advantage to give the job of extracting parallelism to the compiler. This would also mean that current software could be recompiled for faster execution without any time consuming software modifications.

2 Compilers in General

A typical compiler is a program that translates source code from language X to executable code for architecture Y. It is required to retain the semantic meaning of the source code (defined by the language specification) and produce code to maximise the resources of the target architecture. The two main stages involved in compilation of a program are: syntax/semantic analysis of the source code, and code generation. The syntax analysis verifies the grammatical structure of the program whereas the semantic analysis verifies that the meaning does not break any of the language rules. Once these stages are complete the compiler will usually have created an internal representation of the program from which machine code can be generated. The language analysis stage of the compiler is commonly known as the front-end. The code generation is known as the back-end.

It is a trivial exercise to develop a compiler that generates poor quality code but it becomes a very difficult task to produce high quality machine code. There has been much research into different code optimisation techniques, some of which have been incorporated into commercial compilers. A problem in the field is that there are so many different and potentially very good techniques being implemented and tested independently. These bespoke systems can be difficult and time consuming to merge, especially if the optimisers are language or machine dependent. It would be far more beneficial to pool the research of the entire compiler community into one machine independent compiler system. This would allow new research methods to benefit from any current methods (Wilson et al., 1994). There is one such compiler system that does this, it has been developed by a research group at Stanford University and defines a common compiler intermediate representation called Stanford University Intermediate Format (SUIF) (Amarasinghe, Anderson, Lam, & Tseng, 1995). The SUIF system allows the intermediate code to pass through a number of machine independent optimisations before being passed to the back-end code generators. As the SUIF format is machine independent, all the optimisations applied to this format are also machine independent. This means that the system can be used to cross compile from any of the supported front-end languages to any of the supported back-end code generators. The SUIF compiler toolkit provides libraries and optimisations that can be used and applied to this representation. I have chosen to use these tools since they will enable my research to concentrate on machine independent parallel optimisations without having to implement time consuming code generators. The code generators are necessary in order to produce executable code and test the effectiveness of the optimisations.

3 Choice of Language

Java was chosen as a basis for my optimisations since the language design contains in-built support for parallel units (threads) and synchronisation. Also, since Java is usually executed on a virtual machine, it can suffer from slow execution speed. Since this virtual machine is architecturally neutral it is safe to assume that the language representation will be similarly independent. Such independence suits my research into parallel optimisations that are not tied to a specific machine.

The first stage of my work involved deciding on the Java representation to analyse. The choices being the source language or the Java Virtual Machine (Lindholm & Yellin, 1997) instructions called bytecodes. As all Java programs make calls to standard class libraries and these libraries are only available in bytecode format, I chose bytecodes. Starting at the bytecode level allows any referenced classes to be analysed and incorporated into the optimisations without any extra work. The main disadvantage lies with the high-level language information that is lost in the translation from source to bytecode. Such information may have been helpful in the optimisation analysis. The syntax and semantics of these bytecodes are well defined by the Java Virtual Machine specification, which is essential when attempting to automatically restructure code without changing its meaning.

4 Optimisation Analysis

In order to automatically extract parallel threads, a Java program must first be analysed. This analysis can point to areas of the program that would benefit from parallelisation and verify that such alterations would be possible. Since this analysis takes place without running the program, there is a great deal of information that is unavailable. Most programs change their behaviour when dynamic runtime variables change. In an attempt to extract some of this runtime or dynamic analysis information, I am designing a Java bytecode instrumenter. This program inserts profiling instructions into a Java program that will store dynamic information when the program is executed. The profiling information can then be made available to assist in the compile-time or static analysis. Information can also be retained once the bytecodes are translated into the SUIF format. An additional SUIF optimisation stage can then be guided by this dynamic and static analysis data.

The results of the optimisations can be gauged by conducting benchmarking tests using a variety of SUIF optimisations along with generating code for different architectures. The effect of the optimisations will be weighed against the cost of those optimisations and improvements that they achieve. Costs include the size of the code generated, any increase in runtime memory usage and any increase in execution time due to extra thread synchronisation. These costs can also be calculated for the compile-time stage to extract information such as the compilation time and memory usage. All such factors contribute to the usability of a compilation system.

5 Conclusion

I hope to prove that my optimising compiler can automatically improve the resource usage of a Java program in a multiple processor environment and so increase the speed of execution. I also hope to show how such a system would be of potential benefit to both software developers and hardware manufacturers alike.

References

Amarasinghe, S., Anderson, J., Lam, M., & Tseng, C.-W. (1995). An overview of the suif compiler for scalable parallel machines. In *Siam conference on parallel processing for scientific computing.* San Francisco, CA.

Kamin, S. (1995). Report of a workshop on future directions in programming languages and compilers. *acmsigplan*, *30*(7), 9–28.

Keleher, P., & Tseng, C. (1996). Improving the compiler/software dsm interface: Preliminary results.

Lindholm, T., & Yellin, F. (1997). *The Java virtual machine specification.* Addison-Wesley.

Vries, D. D., & Lee., C. G. (1993). *A vectorising suif compiler* (Tech. Rep.). University of Toronto: Department of Electrical and Computer Engineering.

Wilson, R., Wilson, C., Amarasinghe, S., Anderson, J., Tjiang, S., Liao, S.-W., Tseng, C.-W., Hall, M. W., Lam, M., & Hennessy, J. L. (1994). *Suif: A parallelizing and optimizing research compiler* (Tech. Rep. No. CSL-TR-94-620). Computer Systems Laboratory, Stanford University.

Zima, H., & Chapman, B. (1990). *Supercompilers for parallel and vector computers*. ACM Press/Addison-Wesley.

Assessing organisational ability and planning effectiveness

Sian Williams

sianw@cogs.susx.ac.uk

School of Cognitive & Computing Sciences

1 Aims

1. To construct a scale to measure organisational ability.
2. To investigate the difference in planning styles of organised and disorganised individuals and assess the effectiveness of these styles in leading to successful task completion.
3. To investigate how organisational ability moderates the motivational and emotional effects of failing to complete one's tasks as planned.
4. To see how different environments can help/hinder people of different organisational ability levels in completing their tasks as planned.

There is an abundance of popular literature that commends the benefits of being organised and managing one's time effectively. This stems from a belief that poor allocation of time leads to an increase in stress and a decrease in performance. However, the techniques associated with achieving these goals have received very little empirical research. One study by Macan (1995) assessed three popular time management techniques, (setting goals and prioritising; mechanics; and preference for organisation), and found them to be ineffective. The results showed that while time management training was not effective in increasing the adoption of time management behaviours, individuals who set goals and prioritised, and had a preference for organisation reported greater job satisfaction and reduced stress tensions. Macan suggests that this link is moderated through 'a perceived control over time', and that the ineffectiveness of the time management training was due to the variations in behaviour being prevalent before the training.

Research into the ability to organise and meet one's goals and plans is an important area of study with both applied and theoretical importance. Failure to accurately estimate completion times can have economic, social and personal costs. There is a great history of over-optimistic, and even unrealistic predictions in construction projects. For instance, in 1957 the Sydney Opera House was originally estimated to be completed early in 1963 for $7 million. A scaled down version finally opened in 1973 at a cost of $102 million. This is by no means an isolated example and current planners seem to make the same errors in their predictions. This over-optimism and neglect of past experiences is also prevalent at a personal level with more day to day activities. One example well known to academics is the person who takes a stuffed briefcase home at the weekend to complete, fully aware that on similar past occasions the work remained uncompleted by Monday.

Intuition supports the idea that some people are more effective planners than others. But there is very little literature specifically testing individual differences in the organisational ability and planning ability. The first aim therefore is to construct a scale that measures organisational ability. From my pilot study I have identified four possible factors of this dimension:

1. preference for organisation;
2. mental organisation;
3. physical organisation; and
4. planning ability.

Studies of individual differences in cognitive styles have revealed that people do indeed differ in their preference for structure and organisation (Thompson, Naccarato, & Parker, 1989; Neuberg & Newsom, 1993). Past research has focused on the extent to which people limit the amount of information they process to be consistent with their beliefs and goals. These studies address how people differ in their preference for cognitive structure. They neglect the distinction between mental organisation and physical organisation, and fail to investigate individuals' needs for organisation in their workspace and environment. Therefore a second aim of the proposed research is to investigate the relationship between preference for organisation and the other facets of organisational ability. To see whether the person who has a preference for organisation achieves it, and if not what effect this has on the individual.

Planning is arguably one of the most valuable domains of cognitive ability. It enables us to cope with complex and changing situations; it guides decision-making through anticipation of possible futures; and it can reduce complexity by introducing representational structures. It is especially important in large scale organisations where many tasks have to be co-ordinated. Hayes-Roth and Hayes-Roth (1979) define planning as "the predetermination of a course of action aimed at achieving some goa". Sine planning requires the ability to anticipate, and to a certain extent, predict the future, many human judgment errors can be encountered.

There is a growing body of literature that supports the idea that people's plans are influenced by subjective factors such as motivations, wishes, beliefs, and affective states (Buehler & Ross, 1998; Hoch, 1984; Johnson & Sherman, 1990). People tend to generate a single or very small number of possible outcomes in making their predictions and assume its validity. Attention is focused on that pattern of events which interferes with their ability to generate alternative scenarios. Further, when people imagine the occurrence of a future event they become convinced tat it will happen. Work on the 'Planning Fallacy' has revealed that people show an optimistic bias in predicting their behaviour and that in generating predictions will focus on their plans and goals rather than relevant past experiences (Buehler, Griffin, & Ross, 1994).

The work done to date on plans and predictions has ignored the possibility of individual differences, assuming that everyone is open to the same biases and errors. There hasn't been research linking cognitive styles and preferences, planning, and actual success in meeting those plans. The purpose of this program is to provide such a linkage.

References

Buehler, R., Griffin, D., & Ross, M. (1994). Exploring the planning fallacy: Why people underestimate their task completion times. *Journal of Personality and Social Psychology*, *62*(4), 538–551.

Buehler, R., & Ross, M. (1998). Identity through time: constructing personal pasts and futures. In *Blackwell handbook of social psychology.* Blackwell. (Chapter in progress)

Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognotive model of planning. *Cognitive Science*, *3*, 275–310.

Hoch, S. (1984). Availability and interference in predictive judgement. *Journal of Experimental Psychology: learning, memory, and cognition*, *10*, 649–662.

Johnson, M., & Sherman, S. (1990). Constructing and reconstructing the past and the future in the present. In E. Higgins & R. Sorrentino (Eds.), *Handbook of motivation and social cognition: foundations of social behaviour* (Vol. 2, pp. 482–526). New York: Guildford.

Macan, T. H. (1995). Time management: test of a process model. *Journal of Applied Psychology*, *79*(3), 381–391.

Neuberg, S., & Newsom, J. (1993). Personal need for structure: individual differences in the desire for simple structure. *Journal of Personality and Social Psychology*, *65*(1), 113–131.

Thompson, M., Naccarato, M., & Parker, K. (1989). Assessing cognitive need: the development of the personal need for structure and personal fear of invalidity scales. In *Annual meeting of the canadian psychological association.* Halifax, Nova Scotia. (Cited in Neuberg and Newsom (1993))