The 16th White House Papers
Graduate Research in Cognitive and Computing Sciences at Sussex

Edited by
Kingsley Sage

CSRP 552

June 2003

ISSN 1350–3162

UNIVERSITY OF
SUSSEX
AT BRIGHTON

Cognitive Science Research Papers

Preface

Each year, towards the end of June, postgraduate researchers from COGS at the University of Sussex get together to discuss their research and (probably more importantly) socialise.

The three-day Isle Of Thorns Annual Graduate Workshop, named after a former conference centre of the University of Sussex, is taking place this year at Herstmonceux Castle for the third time. Each delegate will present a brief overview of their research, gaining valuable research experience as well as the opportunity to take advantage of the interdisciplinary nature of COGS.

Named after one of the buildings at the original conference centre, the White House Papers was formally the follow-up publication for the workshop. Following the precendent set by last year's workshop, the White House Papers are being published in advance of the conference and are intended to provide a wide snapshot of postgraduate research in COGS.

Intended to be easy reading, the 16th White House Papers consists of typically two-page submissions accompanied by author photographs and short biographies. Due to the breadth of research being carried out in COGS, much of the emphasis at the workshop and in the papers is to present to an audience whose members come from a large variety of backgrounds. This means that many of the papers serve as a general or non-specialist introduction to the author's work. If readers are interested in the material then they can obtain more information and full publications (in some cases) from the relevant author.

Again following the precedent set last year, all of the papers have been reviewed by members of a Program Committee and authors given the opportunity to amend their papers in response to their reviewer's comments. I would like to take this opportunity to thank the Program Committee for their invaluable contributions and help which I am sure will continue throughout the workshop. Also, on behalf of the postgraduate research body, I would like to thank Phil Husbands for the funding of the IOT conference and the distribution of this publication.

Lastly, it should not be forgotten, however, that the Isle Of Thorns conference and the White House Papers could not be such successes without the contributors themselves.

Kingsley Sage
Wednesday 4th June 2002
COGS Postgraduate CSAI Student Representative
Compiler and Editor

Program Committee

Mike Beaton
Hanneke De Jaegher
Jason Harrison
Joanne Lawson
Tony Morse
Sebastian Rasinger
Kingsley Sage
John Sung

Contents

Biographical Information

Elly Adams is a former Open University student. In 2000 she began the MRes course and changed to a part-time DPhil in 2001. In 2003 she became a Research Fellow. Her main areas of interest are memory and eyewitness testimony.

Elly Adams

Marzieh Asgari-Targhi is a DPhil student at COGS in her second year. Her first degree was in Philosophy at University of London. She then did a Masters in History and Philosophy of Science at the London School of Economics and Political Sciences (University of London). Her research interests lie mainly in analysis of causal reasoning in Artificial Intelligence.

Marzieh
Asgari-Targhi

Caroline Auric graduated with a BA in Psychology in 2002, having completed her first degree at COGS. She is currently in the first year of her DPhil studying implicit cognitions in the eating disorders.

Caroline Auric

Mike Beaton has a BA in Physics from Oxford University and an MSc in Artificial Intelligence from Edinburgh University. He is now studying for a DPhil in the philosophy of cognitive science. He is interested in exploring the requirements for a scientific explanation of consciousness, and hopes to relate this problem to a discussion of the current health or otherwise of functionalism, of the neuroscience of emotion and of the nature of scientific explanation itself.

Mike Beaton

Originally from Brighton, Victoria Bourne came to COGS to do her undergraduate degree in Psychology 1997-2000. Being a glutton for punishment, she decided to stay on to do her DPhil with Graham Hole investigating hemispheric specialisations and cooperation in face processing. She is now in her third year.

Victoria Bourne

On his third degree after studying, Physics and Mathematics. Seeking to answer the challenging question of what really is he doing, he is working on perfecting the art of making it look as if he working really hard.

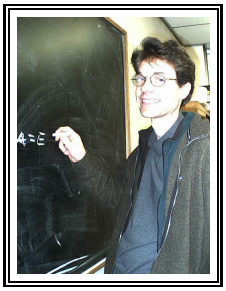

Andy Bowery

Diane Brewster took a BD(Hons) in 1977, followed by a PGCE in 1979 and a MA(Theol) in 1982. She then taught and lectured for almost 20yrs, mainly inModern and Feminist Theology and Interfaith dialogue. From 1994 - 2000 Diane worked for a City of London Insurance broker (systems support and online research). She then returned to acedmia and took an OU BSc(Hons) in 2000 and a MSc in HCCS at COGS in 2002. She started her DPhil at COGS in October 2002.

Diane Brewster

After many years employment in a variety of ‘character building’ jobs, Martin Coleman finally enrolled at City University (London) in the summer of 1998. Successfully completing a first degree in psychology Martin was then lucky enough to secure a post at the University of Sussex to read for his DPhil. His twin ambitions are to both push back the boundaries of human knowledge whilst simultaneously avoiding the horrors of a ‘proper’ days work.

Martin Coleman

Alice Eldridge gained a BSc in psychology from Leeds university in 1999, after working in Brighton for a few years, she was sucked back to academia by the MSc in Evolutionary and Adaptive Systems. She is now in the first year of a DPhil. in CSAI. Research focuses on exploring applications of complex adaptive systems in audio, for generative art, software engineering and scientific visualisation.

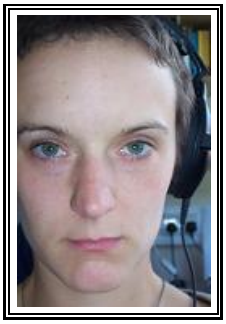

Alice Eldridge

David Ellis holds a BSc(Hons) in Computer Science and Artificial Intelligence from the University of Sussex. He is now studying towards a DPhil in the network lab at Sussex University under the supervision of Ian Wakeman.

David Ellis

Eunice Fajobi obtained her first and second degrees in English/Education at the Obafemi Awolowo University, Nigeria. At the completion of the latter, which was on ESL, she took up appointment with the University as a lecturer in the Department of English. Prior to this, she had taught at various levels of education in Nigeria, ranging from the nursery school through to the Teachers' Higher College of Education. She is committed to the improvement of the teaching and learning of Phonetics and Phonology. She is also interested in aspects of first and second language acquisition.

Eunice Fajobi

Dr Chrisantha Fernando did a Medical Degree at Wadham College, Oxford, worked as a house officer, and is deferring an S.H.O rotation in Psychiatry until he explains how genes make brains.

Chrisantha Fernando

Rowanne Fleck graduated with her first degree in Artificial Intelligence and Psychology at the University of Edinburgh in summer 2002. After a short spell working as administrator on a project looking at developing a distance learning course using mobile technology, she came to Sussex to find out more about human computer interaction.

Rowanne Fleck

Benoit Gaillard did a telecommunication engineering degree in an engineering school, in Brittany (France). For the final year of this degree, last year, he did the master of science in intelligent systems, in COGS, within the ERASMUS scheme. He's now starting a Phd in Computer Science and Artificial Intelligence in COGS. His research interest are the neurons.

Miguel Garvie studied Systems Analysis for one year in Buenos Aires at the Centre de Altos Estudios en Ciencias Exactas and after getting a BA in Computer Science from Cambridge he's now doing research in Evolutionary Electronics under Adrian Thompson in the Centre for Computational Neuroscience and Robotics.

Miguel Garvie

Amanda Harris did her BA in Applied Psychology in the School of Cultural and Community Studies at Sussex (2001). After a year of working for Brighton and Hove's Social Service department she began an EPSRC funded DPhil studentship. The studentship is attached to the Riddles Project in COGS is in its first year.

Amanda Harris

After following his joint BA in History and Classics with a BA(Hons) in Classics at the University of Natal, South Africa, Jason Harrison moved to the School of Oriental and African Studies in London, ostensibly to pursue an MA in Linguistics, but gaining an MA in Eastern Christianity
instead. He is currently in the first year of his DPhil in Linguistics where he is attempting to combine Theological and Linguistic interests into his research.

Jason Harrison

Maggie Kerridge graduated from Sussex with a degree in English Language in 1999 and is still here. She is now researching the relationship between changing social concepts and lexical change under the supervision of Professor Richard Coates.

Maggie Kerridge

Ira Konstantinou did her first degree in Psychology at the American College of Greece in Athens. She was then accepted to do a DPhil with Professor John Gardiner at City University in London. When he tranferred to Sussex University she followed. She is now at her third year, completing her thesis on face recognition and memory awareness.

Anu Koskela was born and grew up in Finland. She moved to Britain in 1997 and gained a BA in Linguistics from the University of Sussex in 2001. She then worked as a translation project coordinator in the computer games industry before starting her DPhil in Linguistics in October 2002.

Anu Koskela

Kristy Lascelles is currently a 4th year (thesis in preparation) psychology D.Phil. student researching the human learning phenomenon: Evaluative Conditioning with Graham Davey. She has also been involved with some memory research with John Gardiner. Her previous career involved a BSc(Hons) in Psychology from the University of York.

Kristy Lascelles

Joanne Lawson started her part-time DPhil on the effects of information on the development of children's fear beliefs when she began work as a research assistant in COGS in March 2002. Before that, from 1998 to 2001, she studied Experimental Psychology at Oxford.

Joanne Lawson

Chi-Ho Li got his first degree in Philosophy from the Chinese University of Hongkong and then a Masters degree in Artificial Intelligence from the Edinburgh University. Afterwards he worked in a company developing Chinese-to-English translation software. There he had to manually write the grammar rules for the MT system, and thus became interested in grammar induction.

Chi-Ho Li

Paul Loader did his first degree in Philosophy at the University of Kent at Canterbury. After a gap of many years he then went on to do an MSc in Information Technology at Queen Mary, University of London (QMUL), graduating in September 2001. In January 2003 he began his DPhil studies in Philosophy of Cognitive Science.

Paul Loader

Ian Macinnes is a member of the Centre for Computational Neuroscience and Robotics at Sussex University where he is taking a doctorate in artificial intelligence. His interests include evolutionary robotics, embodied cognition, neural networks, artificial life, theories of non-Mendelian inheritance such as maternal effects, hardware evolution, and autopoiesis together with dynamical theories in cognitive science.

Ian Macinnes

Barbara Maidment trained as a journalist at RSP (now University of Westminster) and worked as a printers' rep and in publishing. She combined motherhood with working part time as a headhunter's researcher and then came to Sussex (EAM) in 1995 to read English Language, graduating in 1998. She is also currently historian and guide at Groombridge Place Gardens, near Tunbridge Wells. She has appeared on several local radio stations and ITV Digital, and look what happened to that.

Barbara Maidment

Marek McGann is a DPhil student working with Prof. Steve Torrance on the Philosophy of Cognitive Science. In particular, He is looking at the use of goal-directed concepts such as intentions, purposes and deliberate actions in current theories in Cognitive Science (or the lack of such concepts, as is often the case).

Marek McGann

Mark McLauchlan is a DPhil student at Sussex University. He studied German, Economic History and Computer Science at Victoria University in New Zealand before studying a Masters degree in Artificial Intelligence at Edinburgh University. His research interests lie in applying machine learning techniques to problems in Natural Language Processing.

Mark McLauchlan

Erika Annabel Martinez-Miron obtained a BSc degree in Computer Science at Universidad Autonoma de Puebla (Mexico) in 1996. Then she moved to Mexico City to study at the Universidad Nacional Autonoma de Mexico where she obtained a MSc degree in Computer Science in 2001. Now she is her second year DPhil at COGS.

Erika Annabel Martínez Mirón

Lectured English Syntax (College of Education l996), lectured English Syntax (University of Botswana l997). On study leave since 2000, holder of BA (1991) Post Grad Dip Ed. (1993) from UB, MA Appl Linguistics, Sheffield 1996, MA theoretical Linguistics Essex 2002, first year linguistics PhD student (Sussex), other qual. Certificate & Diploma (Library studies).

Setumile Morapedi

Anthony Morse gained a degree in Cognitive Science at the University of Hertfordshire graduating in 1999. He then completed an MSc in Evolutionary and Adaptive Systems at Sussex University where he has now in his finalyear of a DPhil investigating the development of structure in neural mechanisms.

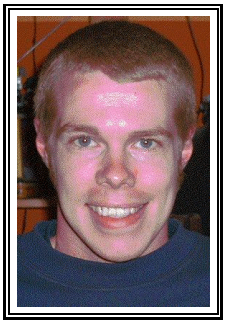

Anthony F. Morse

Darren Pearce started his DPhil in 1998 and has been trying to finish for several years now. He is now juggling work on the thesis with full-time research in the IDEAs Lab in COGS.

Darren Pearce

Sebastian Rasinger is a 2nd year DPhil student in Linguistics. Under the supervision of Dr Max Wheeler and Dr Melanie Green, he is working on the acquisition of English syntax by the Bengali community in East London. Being a semi-trained social scientist, he has a secret passion for statistics, and is also reading for a postgraduate certificate in social research methods at Sussex University.

Sebastian Rasinger

Genaro Rebolledo-Mendez is 2nd year DPhil student at Sussex University. His first degree was in Computer Sciences in Veracruz (Mexico), he then did a MSc in Human Centred Computer Systems at Sussex University. His interests include educational technology, motivation and affective computing.

Genaro Rebolledo-Mendez

Jon Rimmer teaches and researches in the area of Human-Computer Interaction (HCI). He has also had several years of professional work in the area - gathering data about people in order to inform the design of better interactive products, environments and experiences.

Jon Rimmer

Enrico Rossoni graduated from University of Rome (Italy) in 2000, with a degree in Physics. Since then, he has been involved in neuronal modelling and vascular fluid dynamics. In 2002, he started a DPhil in Computational Neuroscience at COGS. His research currently focuses on developing computational models of neuroendocrine cells.

Enrico Rossoni

John Rowston attained an Honours degree in Psychology from the Open University in 1994. In 2002 he completed his Masters in Evolutionary and Adaptive Systems at Sussex University. This paper serves to outline the beginnings of his DPhil research, which has culminated from an interest both in sound and in dolphins stretching back many years.

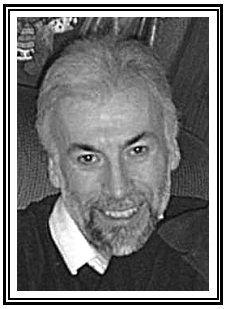

John Rowston

Kingsley Sage is a research fellow working on the EU ActIPret project with Hilary Buxton and is in his second part-time year as a DPhil student. His thesis area is in the learning of structure and temporal dynamics for computer vision systems to build on-line vision systems that are task controlled.

Kingsley Sage

Brenda Smith graduated from the University of Sussex in 1999 with a BA in Applied Psychology. She started her DPhil in 2001, under the Supervision of Professor John Gardiner investigating memory processes in individuals with Asperger's Syndrome.

Brenda Smith

Sampsa Sojakka is a DPhil student at the Centre for Computational Neuroscience and Robotics (CCNR) at Sussex. He is studying controller architectures for autonomous agents under the supervision of Inman Harvey and Ezequiel Di Paolo.

Sampsa Sojakka

John Sung has BS and MS from Carnegie Mellon University in Electrical and Computer Engineering. His primary interest were in fault tolerant computing and microcomputer architecture. After working for the Alpha Verification team for 3 years, he moved onto a software data base company as a software engineer. After that, he decided to obtain an MS in Computer Science working on Aspect Oriented Programming Languages. Currently, he is working on understanding the process of scientific theory development in Genetics.

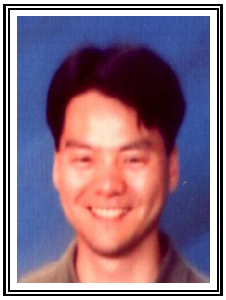

John Sung

Aisha Thorn is in the second year of her DPhil, researching methods for analysing mammal vocal repertoires. She is supervised by David Young in COGS and Karen McComb in EP. She studied CSAI in COGS as an undergraduate after a three-year break from education. Having sampled 'the real world' she now hopes to remain a student indefinitely.

Aisha Thorn

Robert Vickerstaff graduated from Cambridge in 1999 with a BA(Hons) in Natural Sciences, specialising in Genetics. He further completed a C programming course before coming to COGS for the Evolutionary and Adaptive Systems MSc in 2000/1. After a year of voluntary work Robert began his DPhil in the CCNR in October 2002.

Robert Vickerstaff

Xinglong Wang is a DPhil student in Natural Language Processing group at the University of Sussex. He obtained his first degree at Harbin Institute of Technology, China and a MSc degree in software engineering at the University of York. His research interests include cross-language information processing, word sense disambiguation, collocation extraction, etc.

Xinglong Wang

Dawn graduated from the Mount Alison University, Canada, in 1995 with a BSc in Psychology and Biology. In 1999, Dawn completed her MSc in Developmental Psychology at the University of Sussex. After working in Canada as a support worker/counsellor for a year, Dawn returned to Sussex to begin her DPhil in the area of childrens social cognitive development. Dawn is now working as a research fellow on the Children's Social Behaviour Project with Robin Banerjee.

Dawn Watling

Julie Weeds graduated from Trinity Hall, Cambridge in 1998 with a 1st in Computer Science. Whilst training in London as a patent attorney, she realised that she didn't like the real world very much and so fled back to Cambridge a year later to do a MPhil in Computer Speech and Language Processing. In October 2000, she arrived at COGS to begin her DPhil with the Natural Language Processing Group under the supervision of David Weir. And to the non-believers out there, she would now like to say that there is light at the end of the tunnel....

Julie Weeds

In 2001, Hanneke De Jaegher obtained her degree of "Licentiaat in de Wijsbegeerte" (cf. MA in Philosophy) from the Department of Philosophy at the Free University of Brussels (VUB) in Belgium, where she was also connected to the AI-Lab. Her licentiate's thesis was about research in cognitive science; the thesis compared and integrated Maturana's theory of life and cognition, approaches to investigating cognition in the field of AI and aspects of findings on autistic thinking. Autism sheds an interesting light on central issues in cognition, such as perception, affect and social cognition, development, emotion, the meaning of and relationship between low- and high-level cognition and the role of rationality. Hanne started a DPhil in COGS in October 2001, in which she continues to research the nature of autism and thinking.

Hanneke
De Jaegher

Improving the Reliability of Automobile Memory

Elly Adams

ellya@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Memory for automobiles is of critical importance for many types of forensic investigation. For example, in child abductions, traffic accidents, and crimes using 'get away' cars, eyewitnesses are often asked about their memories of the automobiles involved. While much research has examined memory for people, and in particular faces, there has been little work on memory for automobiles. This lack of research is surprising given the frequency of automobile crimes and parents' fears of child abduction (Davies, Kurvink, Mitchell, & Robertson, 1996). Recognizing the importance of automobile identification, police in the UK have devised 'Motorfit' (Grantham, 1989), a method for eyewitnesses to construct a picture of an automobile from memory in a manner similar to those used in many computer assisted face identification packages. While showing potential, Motorfit proved difficult to implement. The main aim of this research is to help to bridge the gap between what is known in the scientific literature about automobile memory and the needs of the police investigators and courts. The research will address five principle questions:

1. Is it possible to predict who will be accurate at memory for automobiles?

2. Are there some automobiles and situations which led to particularly reliable memory reports?

3. What do people recall about an automobile?

4. Are there methods to improve the reliability of reports, particularly for those witnesses with unreliable memory?

5. How can these methods be best implemented in computer software?

One concern about much eyewitness research is that the stimuli and situations are too artificial to make realistic generalizations to crime scenes. Artificial stimuli and situations are often used to increase experimental control. In order both to explore causal hypotheses about memory process and generalize to real situations it is necessary to use a range of stimuli (Wright, 2002). This research uses three distinct methods. The first is presenting multiple automobiles to participants, recording various aspects of these automobiles and of the participants, and exploring memory. This includes showing multiple computer stimulations of automobile events. The second uses 'walk-about' methods which involve participants walking around a car park and later being asked about several of the automobiles. Finally, individual crime reconstructions will be made and stored digitally. These will then be shown to participants. The research will not answer all questions regarding automobile memory but will provide a foundation for academics and police investigators.

References

Davies, G. M., Kurvink, A., Mitchell, R., & Robertson, N. (1996). Memory for cars and their drivers: A test of the interest hypothesis. *Basic and applied memory research: Practical applications, 2*.

Grantham, S. (1989). In *Motorfit.* Home Office, Police Research Group, UK.

Wright, D. B. (2002). In *First steps in statistics.* Sage publications.

Causal Reasoning in Artificial Intelligence

Marzieh Asgari-Targhi

marzata@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Primarily AI people are interested in the study of the computations that make it possible to perceive, reason and act. Reasoning, in its simplest definition, is the set of processes that enable us to go beyond the information given. The processes could be logical (inductive or deductive), emotional, intiutionistic, etc. A strong part of our reasoning is causal. It is dominant in our everyday commonsense thinking and speech. In fact the sole meaning of many verbs and verb phrases is the description of an effect of some unspecified cause. Also very common in scientific thinking from geology to psychology. Since causal reasoning seems to be an important factor in human behaviour, it stands to reason that researchers in AI should strive to understand it and make use of it in their theories.

2 Research aims

The goal of my research is to examine causal reasoning in Artificial Intelligence. Are AI scientists interested in causality? If so, why? More importantly how is causal reasoning used in AI theories and projects? I will argue that advances in AI are highly relevant to the philosophical problem of causation and philosophical investigation of causation will be enriched by considering AI scientists's approaches. Their scientific goals, ways of asking questions, kind of answers they accept, conceptual and technical tools are for the most part different from those of the philosopher. This means that the very least they might be able to do is to ask new questions about the old problem if not actually come up with solution. My own contribution, hopefully, will be to find out what role causation plays in AI, and what its significance is for the philosophical understanding of causation.

3 Existing approaches

It is the case that most theories of causation in AI, so far, have been informal; the approach has been that we understand intuitively what causation is, i.e in any simple analysis an event is responsible for the occurrence or coming to being of another event, and therefore rather than waste time defining it one should embody it in a program or possibly in a slightly more abstract computational framework. The recent attempts by (Pearl, 2000), (Glymour, Spirtes, & Scheines, 2000), (Shafer, 1996) and (Bell, 1999) in which efforts are made from different perspectives at rigorous definitions of the concept, have paved the way for many researches to try to understand causation with scientific tools. Their approaches can be broadly divided into the following two groups.

3.1 Pragmatic approach

As I mentioned above this is the view dominant among scientists (until recently) that conceptual definition of causation is the philosopher's job. Instead they use intuitive understanding of causation in their projects, i.e instead of defining causation they embody it in a program or a more abstract computational framework. A good example would be people who work in vision learning, they just assume some form of dependency (e.g, correlation between statistical data) and use belief Bayesian networks learning and reasoning to identify cause and effect.

3.2 Interdisciplinary approach

- Holistic approach: causation has many different facets; uncertainty element, logical element (necessary and sufficient elements) and qualitative element which is unique to any individual causal situation. The focus of this approach is not on any single aspect of causation but on all of them (Pearl, 2000). encapsulates all the above aspects.

- Bayesian approach: applying Bayes nets on causal inference and causal reasoning, thereby developing the causal interpretation of Bayes nets (Glymour et al., 2000).
- Probabilistic approach: Causation can be identified with its uncertainty element, therefore provide a probabilistic theory of causation (Shafer, 1996).
- Logical approach: The formalisation of philosophical theories. In particular the logical formalisation of Mackie's theory; his INUS account of causation by (Bell, 1999).

4 Philosophical significance

All the above school of thoughts in AI attempt to answer two important philosophical questions: (a) How do we know one event is the cause of another event? (b) What is the nature of the causal mechanism involved between the two events? Philosophical theories, in general, try to reduce the problem of causation to one of these two questions. But an AI causal model deals with both types of questions. For example, in robotics researchers, in one particular application are trying to build a highly capable and autonomous robot. The robot has to be able to navigate with minimal human intervention i.e, have a sufficient knowledge of its environment (that, precisely, is dealing with type (a) questions). Robots also must be able to detect anomalies and deal with them effectively and use some sort of causal mechanism (that is dealing with type (b) questions). Moreover they must be able to manage their limited resources, including power and computation and use them in an efficient manner. Finally, they must integrate all these capability into working reliable systems, that is making a model that answers both types of the questions. I believe if they succeed making that kind of model, then the built model will capture an important pattern of our causal reasoning.

References

Bell, J. (1999). Primary and secondary events. *Working notes of the 4th workshop on Non-monotonic Reasoning, Action and Change IJCAI,* 65-72.

Glymour, C., Spirtes, P., & Scheines, R. (2000). Causation, prediction and search. In *Logical foundation of proabability.* The MIT Press, Cambridge, MA.

Pearl, J. (2000). In *Causality, models, reasoning, and inference.* Cambridge University Press.

Shafer, G. (1996). In *The art of causal conjecture.* MIT Press, Cambridge, MA.

Anorexia and Bulimia: disorders of control?

Caroline Auric

carola@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Background

Clinicians are increasingly recognising that, although the symptoms of eating disorders revolve essentially around food, weight and body shape concerns, the fundamental (and often causal) problems usually lie in other areas of the sufferers' life and/or personality (Dalgleish et al., 2001). Indeed it has been suggested that eating disordered individuals perceive shape and weight as an 'index of control' (Fairburn, Shafran, & Cooper, 1999); control they feel they do not possess over other areas of their lives. In line with the emerging focus on control as an underlying motivation in the eating disorders, (Vitousek & Hollon, 1990) carried out an investigation into the schematic content and processing in the cognitions of eating disordered subjects. They recognise that weight and shape concerns are central to the psychopathology of AN (anorexia nervosa) and BN (bulimia nervosa), but propose that individuals with eating disorders use these quantifiable outcomes in such a way as to organise and control their behaviour, perceptions, thoughts and affect. As such, individuals with eating pathology develop organised cognitive structures (schemata), which can affect the way in which they process information in certain domains (e.g. stimuli relating to food and weight).

The recurrent appearance of the implication of control in the Eds (eating disorders) has led to further investigation into this area in recent years. Researchers have generally used one of two methodologies in an attempt to acquire further insight into the use of food and weight by ED sufferers to establish a sense of control over some aspect of their lives. The first, and most frequently used method of collecting such data is that of self-report measures such as the Locus of Control of Behaviour Questionnaire, and the self-control inventory (Lugli-Rivero & Vivas, 1997). The use of self-report measures however has generated such mixed findings (Dalgleish et al., 2001); (Sandbeck, 2001) that questions have been asked about the validity of these measures, along with the usual concerns about the transparency of self-report measures (demand characteristics, self presentation). The need for less transparent techniques is particularly relevant to ED research because sufferers often strive to preserve their maladaptive beliefs and behaviours (Vitousek & Hollon, 1990). Some ED sufferers may deny they have a problem, whilst others refrain from expressing, or may not be consciously aware of the motivations and thoughts that underlie their maladaptive behaviour. For this reason, in an attempt to expose the hypothesised implication of cognitive biases in the eating disorders I will employ an experimental technique called the Implicit Association Test (IAT) (Greenwald, McGhee, & Schwartz, 1998) in my research.

The IAT is essentially a categorization task that assesses the differential association of two target concepts with an attribute dimension. This is achieved by measuring the response times and numbers of errors made on categorization tasks where the target-concept and attribute dimension combination are more or less 'compatible'. For example, given the target concepts summer and winter, and the attribute dimensions hot and cold, individuals are likely to perform at a faster rate when summer and hot are combined than when winter and hot are combined because the former should be more strongly associated (unless you live in Australia!). Thus, the more associated the concepts, the easier it is to respond, resulting in faster reaction times and less errors made.

2 The experiments

The proposed experiments will compare three groups of females (AN, BN, and normal controls) on their performance on the IAT tasks. These groups will be matched for age and gender and selected on the basis of clinical diagnosis (for the two clinical samples) or a low score on the Eating Attitudes Test (for the normal controls; 0-1, where 20+ is considered an appropriate cut-off for identifying

people with eating pathology (Garner, 1997)).

The target concepts for the first experiment will be food and friendship, and the target attributes controlled and uncontrolled. These concepts were chosen on the basis that they are both areas with which issues of control may have become entangled for ED sufferers. Research reviewed here suggests that the interpersonal domain (e.g. friendship) may be an area over which individuals with EDs feel they have little control (Rezek & Leary, 1991), whereas food is thought to be the only domain over which ED sufferers feel they do have (or strive for) control. The pairings of food with control, and friendship with uncontrolled represent the compatible trials, and of food with uncontrolled, and friendship with controlled the incompatible trials. The eating disorder groups are expected to respond faster and make fewer errors than the control group on the compatible trials.

References

Dalgleish, T., Tchanturia, K., Serpell, L., S.Hems, deSilva, P., & Treasure, J. (2001). Perceived control over events in the world in patients with eating disorders: a preliminary study. *Personality and Individual Differences, 1*, 453-460.

Fairburn, C. G., Shafran, R., & Cooper, Z. (1999). A cognitive behavioural theory of anorexia nervosa. *Behaviour Research and Therapy, 37*, 1-13.

Garner, D. M. (1997). Psychoeducational principles in treatment. *Handbook of Treatment for Eating Disorders*.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74(6)*, 1464-1480.

Lugli-Rivero, Z., & Vivas, E. (1997). Trastornos de alimentacion y control personal de la conducta. *Salud Publica Mex, 43*, 9-16.

Rezek, P. J., & Leary, M. R. (1991). Perceived control, drive for thinness, and food consumption: Anorexic tendencies as displaced reactance. *Journal of Personality, 59*, 129-142.

Sandbeck, T. J. (2001). In *The deadly diet.* Oakland, New Harbinger Publications, Inc.

Vitousek, K. B., & Hollon, S. D. (1990). The investigation of schematic content and processing in eating disorders. *Cognitive Therapy and Research, 14(2)*, 191-214.

Can We Still Use The Argument From Multiple Realizability?

Mike Beaton

mjsb20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Within the philosophy of mind, there are many different opinions as to how the mental relates to the physical. In this paper I will discuss some aspects of functionalism, one of the most widely held but also most widely criticised of these views.

The basic idea behind functionalism is that the defining feature of a mental state is the functional role which it plays in the organism which has it. Thus, for instance, a mental state is a pain state if it is such as to cause the organism which has it to respond aversively to the stimulus which caused it. This is without doubt over-simplified, but a functionalist would claim that there is *some* description of this sort which captures everything there is to capture about pain.

A common opposing view is that functionalism can never capture what pain actually feels like. Typically, it is held that the precise physical matter of which an organism is made determines what its mental states feel like. Thus, human pain feels the way it does because we are made of neurons. A (hypothetical!) martian's pain could never feel the same as a human's pain, however functionally similar he was to us, because he is made of marti-ons and not neurons.

2 Multiple Realizability

Probably the most basic argument underlying the functionalist view is the argument from multiple realizability. It goes something like this. Imagine that you had cut your finger off. Now it seems perfectly possible that in the future we could produce an artificial replacement finger which could respond to motor signals and generate sensory signals in a way very close to your own lost finger. If the new finger was well enough made (if it responded to movement commands and sensory stimulation in the same way as your real finger) then there seems to be good reason to believe that it would feel the same as your real finger used to feel, even if it was not made of flesh and neurons.

Now it seems that the level at which you would need to analyze 'in the same way' in the above is a functional level. What role do the signals from your brain play in telling the finger to move? What role do the signals from your finger play in telling your brain about touch, damage, heat, cold etc.? There doesn't seem to be any reason why we couldn't work out these functional roles, and so why we couldn't build such an artificial finger.

If you're prepared to concede that we could build a silicon finger (say) that felt roughly the same, but not *exactly* the same, then you are on a slippery slope. For surely the level of analysis which would tell us why the finger only felt roughly the same would still be functional? If the finger felt different then in order to explain why, we would want to find out exactly the (functional) situations in which its signals and responses were different.

The above argument is easier to accept for something peripheral like a finger than it is for, say, your brain. But if the argument is correct, then why stop at fingers? Why believe that there is any part of your nervous system, or indeed of your whole body, such that you would be able to feel a difference if it was replaced by something functioning exactly the same? Daniel Dennett makes this argument repeatedly in (Dennett, 1991).

3 The Churchlands

Paul and Patricia Churchland have, over more than twenty years, consistently made a clear case against functionalism.

Patricia Churchland apparently believes that the above argument from multiple realizability is completely discredited (e.g. (Churchland, 1986), (Farber, Peterman, & Churchland, 2001)). However, some highly respectable philosophers such as Dennett continue to use it, and it is far from clear that Churchland herself has discredited it.

For example (Farber et al., 2001) defines functionalism as the "philosophical hypothesis that mental states are defined by their role in a func-

tional (usually, computational) economy of other such states, independently of their physical instantiations" and goes on to conclude that the "functionalist ... hypothesis is simply incompatible with the data".

The text in brackets in the first quote is important, because almost all of the Churchlands' direct arguments against functionalism are actually arguments against the hypothesis that the mechanism of human cognition can be analyzed as a compositional symbol system. I accept arguments against the symbol system hypothesis, but I do not believe that this tells us anything about functionalism. A commitment to the idea that all there is to say about pain is its functional role, is completely independent of a commitment to the claim that things playing such a role need to combine and interact as computational symbols do. At times (e.g. (Churchland, 1986), p.358) the Churchlands seem to accept this point. Nevertheless, they then continue to treat functionalism as identical to the symbol system theory of mind, and thus as discredited.

4 To Reduce Or Not To Reduce

A common position, held especially amongst neuroscientists and amongst philosophers basing their ideas on this work, is that the mental should be fully reducible to the physical (Crick & Koch, 1990); (Churchland & Churchland, 1990); (Churchland, 1996); (Farber et al., 2001). Examples, from other fields, of the kind of physical reduction these people are talking about include the explanation of the macroscopic properties of gases in terms of the microscopic behaviour of molecules, or the explanation of biological inheritance in terms of genes (which are themselves explained in terms of biochemistry, which is itself explained in terms of physics).

It is often assumed that functionalism is completely opposed to this reductionist position. The Churchlands, as usual, provide a very clear expression of this anti-functionalist point of view (Churchland, 1986), ch.7; (Churchland & Churchland, 1990); (Churchland, 1996). I have no space here to go into details, but I wish to contend elsewhere that the type of reduction that the Churchlands outline is in fact a reduction manqué, without the explanatory force of a reduction which embraces that the idea that all physical description is intrinsically functional. (Other authors who have previously argued similarly include (Richardson, 1979) and (Sober, 1992).)

5 Conclusion

I have outlined one of the key arguments for functionalism, and touched on just a few of the key points made against it in one of the clearest and most sustained attacks; that made over some twenty years by Paul and Patricia Churchland.

The only point I have been able to express in any detail here, in a short paper which also outlines the functionalist position, is that the Churchlands persist in treating functionalism *per se* as discredited when a detailed examination of their arguments shows that they themselves accept that they have only discredited the symbol system theory of mind.

I have also touched on the point, which I hope to be able to expand on in much more detail elsewhere, that I believe that despite the Churchlands considerable contribution to the debate on the nature of scientific reduction, they have in fact failed to capture the true nature of reduction, and thus failed to realise that a functionalist description of the mental is indeed fully compatible with scientific reductionism.

References

Churchland, P. (1986). *Neurophilosophy.* MIT Press.

Churchland, P. (1996). The rediscovery of light. In *Journal of philosophy.* 211-228.

Churchland, P., & Churchland, P. (1990). Intertheoretic reduction: A neuroscientist's field guide. *Seminars in the Neurosciences*, 249-256.

Crick, F., & Koch, C. (1990). Toward a neurobiological theory of consciousness. *Seminars in the Neurosciences, 2*, 263-275.

Dennett, D. (1991). *Consciousness explained.* Little, Brown & Co.

Farber, I., Peterman, W., & Churchland, P. (2001). The view from here: The nonsymbolic structure of spatial representation. *Foundations of Cognitive Science, ch. 4.*

Richardson, R. (1979). Functionalism and reductionism. In *Philosophy of science, 46* (p. 533-558).

Sober, E. (1992). Learning from functionalism: Prospects for strong artificial life? *The Philosophy of Artificial Life*, 361-378.

Hemispheric repetition priming of familiar faces

Victoria Bourne

victoria@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

If a visual stimulus has previously been seen, subsequent recognition is faster. This phenomenon is known as repetition priming. Repetition priming has been used in the three experiments reported here to investigate hemispheric specialisations in face recognition. It is widely known and accepted that the right hemisphere (RH) is the dominant hemisphere for processing and recognising faces. Although the left hemisphere (LH) is seen as dominant for other functions, such as language function, evidence from normal and clinical populations has suggested that the LH is also involved in face processing and recognition. My DPhil has investigated the distinct contribution of each hemisphere to face recognition. It has examined the possibility that the RH is specialised in processing global or configural facial information, whereas the LH is specialised in processing featural facial information. These specialisations have been investigated using various manipulations of facial information. For example, blurring a face, which reduces featural information whilst leaving configural information intact, disrupts LH processing but not RH processing. Inversely, presenting the disjointed features of a face, which disrupts the facial configuration whilst leaving the featural information intact, disrupts RH processing, but not LH. The RH dominance for face recognition and the independent contribution of each hemisphere to face recognition was investigated using repetition priming. In these experiments faces were initially presented unilaterally in each visual field. Participants had to make a familiarity decision. The same faces were then presented again bilaterally. Recognition of these faces was used to examine the specified hemispheric specialisations. In the first experiment faces initially presented to the RH significantly facilitated subsequent recognition. In contrast those initially presented to the LH caused no priming effect and were recognised no quicker than unprimed faces. This finding supports the RH specialisation for face recognition. In the second experiment unmanipulated and blurred faces were initially presented to each hemisphere. The same faces were then presented again, bilaterally and all unmanipulated. The third experiment was identical, but presented disjointed facial features rather than blurred faces. Experiment two supported the suggestion that the LH is specialised in processing featural information as blurred faces presented to the LH did not cause priming, presumably as the LHs main source of information was degraded. Blurred and unmanipulated faces presented to the RH caused similar priming effects. Results of experiment three were comparable with the manipulation reducing the priming effect of disjointed features presented to the RH, but not to the LH. The findings of these three experiments replicate the findings of the earlier experiments in my thesis, but using a different experimental paradigm of repetition priming. They suggest that the RH is dominant for face recognition and specialised in processing configural facial information, whereas the LH has a secondary but important role in face recognition processing featural facial information.

The Significance of Correlations in Neuronal Spike Trains

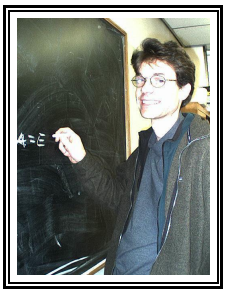

Andy Bowery

abowery@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

Correlation between afferent spike trains of neurons in the brain have been observed in the activity of neurons in-vivo and in-vitro. This is due to the highly convergent/divergent nature of the connectivity of the cortex, common input and recurrent connectivity. This structure also leads to overlap in the receptive fields. Synchrony and concerted oscillations between individual spike trains can be viewed as two subgroups of correlation, in both the activity is temporally correlated. When two neurons are correlated in activity, then the firing of the neuron is not independent, if one neuron fires then there is an increased likelihood then that the other will fire. What is the significance of correlation? Does it play a role or is it just an effect of the high degree of connectivity? After all correlation in the activity could reflect nothing more than the connectivity and receptive field overlap. There a number of results (Salinas & Sejnowski, 2001), (Salinas & Sejnowski, 2002) both theoretically and experimentally that suggest a role for correlations, together with a number of speculations. It has been suggested that correlations may have a role regulating Spike- Timing Dependent Plasticity (STDP). Correlations between spikes could be used for information processing. It has been proposed that correlations gate the flow of neural information. Correlations may act as switch turning on and off transmission. Synchronous spikes to a neuron have been found to evoke a stronger response from the neuron than independent spike trains. This has been demonstrated theoretically, correlated input to a neuron has been shown to impact the response of the postsynaptic neuron, dependent upon the balance of excitatory to inhibitory inputs to the neuron. Correlated fluctuations to a balanced neuron increase the likelihood of firing of the postsynaptic neuron. Correlations also affect the variability in the output train, producing an increase in the variability. Correlations have been suggested to account for the observed variability in the output spike trains. Spike trains in awake animals have been found to be highly variable, however spike generating mechanisms are highly reliable. The CV(ISI) of the typical cortical neuron is close to 1, however this number should be much lower for an integrator that adds up small contributions in order to fire (this is however in the absence of inhibition). With inhibition the number is increased, though still lower than recorded data. Several lines of evidence have pointed to correlations as a source of the missing variability in the observed output spike train (Feng & Brown, 2000). Perhaps correlations can be viewed as a further coding dimension for internal representation, additional to variations in firing rate. As such, correlations be controlled independent of firing rate (deCharms, 1995). Correlations have been found to signal the presence of a change in stimulus without sustained changes in firing rate. If correlations can change without a change in firing rate this could reflect a change in internal state, with perhaps different correlation patterns reflect different internal states. Such independent modulations have been suggested to be used for object representation. Correlation has been found to covary with expectation, attention and response latency, all processes that affect the transit of information, but not the meaning. Correlation has also been found to be stronger when, for example, the sensory discrimination task is more difficult. The degree of interplay between firing rate and correlation is largely unknown, will changes in correlation make a difference if it is coupled with a large change in input firing? The cortical microcircuit can be analysed in two dimensions, intensity (in terms of mean firing rates) and coherence across neurons (in terms of synchrony and cross-correlation). There is strong evidence that correlations are important dynamical features of a microcircuit. Can such correlation in neural circuits be controlled by other circuits to perform useful operations? The correlation structure of a neural population may change dynamically, and determine responses of the downstream targets. There are two main questions that remain unsolved: whether

correlations have a specific function role (such as encoding stimulus features, gating of information or participate in all functions as firing rates do), and whether correlations can be controlled independently of firing rates? A group of neurons may change another group by changing the firing rates or the local correlations, but are these changes independent of each other, a framework is needed.

References

deCharms, R. C. (1995). Information coding in the cortex by independent or coordinated populations. In *Proceedings of the national academy of science.* USA.

Feng, J., & Brown, D. (2000). Impact of correlated inputs on the output of the integrate-and-fire model. *12*, 671-692.

Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience, 2*, 539-550.

Salinas, E., & Sejnowski, T. J. (2002). Integrate-and-fire neurons driven by correlated stochastic input. *14*, 2111-2155.

eHABITATS: Context and Affectivity in Computer Mediated Communication

Diane Brewster

dianeb@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

In December 2001 the UK Prime Minister and the Secretary of State for Education announced a 50 million package of funding for ”curriculum online”. The concept of e-learning has become as central for schools and higher education as it had become for commerce a few years earlier.

The promise of relatively cheap, fast delivery of information and teaching resources is being held up as an ideal to be worked towards. On the other hand the dream does not seem to be being easily attained. A recent article in the newsletter from the Southampton Education Authority states that “Electronic learning is to education as electronic commerce is to economic activity: a disappointment in its current state of development ...” and continues “...of the $2.7 billion invested in e-learning in 2000, an inordinate sum is gone. Most of the e-learning companies founded in the last three years have failed. In particular, the attempt to use the Internet to reform American education from kindergarten through 12th grade has been ruinously expensive and fruitless.” (`http://www.southampton.gov.uk/education/education_services/classroom_of_the_future/elearning.htm`).

A recent ethnographic study in a UK high school (Brewster, 2002) has further confirmed the results of wider studies throughout the world that the actual take-up of technology in e-learning has failed to capture the imagination of teachers. In the words of Larry Cuban (Cuban, 2001), computers in classrooms are largely “oversold and underused”. The response of governments to this lack of enthusiasm has been to set more targets and increase the budget for training teachers in the use of technology.

Computer Mediated Communication, in the form of Computer Supported Collaborative Learning (CSCL), is a way of engendering peer collaboration and teacher support through online conferencing systems. It might be expected that this use of new technology to facilitate ongoing discussion and student support outside of the classroom, would have been taken up with some enthusiasm by at least a substantial minority of teachers. This does not appear to be the case, out of a staff of 44 at the school in question, only four had any personal experience of online groups - and none of them had experienced its use within an educational context. This finding supports that of the wider surveys cited above

In contrast to this the university sector has begun to embrace the use of technology for online delivery of course materials and student support, with systems such as WebCt `www.webct.com` gaining ground, with its sellers making great claims for its ability to “Transform the educational experience” but here again the use of CSCL is not the norm.

One exception to this is the UK’s Open University which, as a distance-learning organisation, has pioneered the use of CSCL software. First this was using a system called CoSY and more recently Centrinity’s *FirstClass* conferencing system. CoSy was a text-only system accessed via an offline reader called Wigwam, while limited to text it did offer excellent threading. *FirstClass* on the other hand has a poor threading system but provides opportunities for e-mail, file exchange and sharing and real time chat, as well as the online asynchronous conferencing. The (password protected) network is now also available through a browser interface, which is much slower than the *FirstClass* software but offers most of the same functions.

Exploring the various system providers leads to the conclusion that they are offering very much the same model with similar functionality. Some facility for file exchange and live chat (and / or instant messaging), with a “conferencing facility” for asynchronous conversation. Being web based has the advantage of not needing bespoke software, but the major disadvantage of being slow, particularly on a dial up connection. Any user selection / action requires the browser to refresh or change the current page - or open a new one, all of which takes

time and adds to user frustration. The user's attempt to communicate with others / collaborate is frustrated by the tool being used rather than facilitated by it - and yet there are many examples of users persevering with frustrating software and interfaces in order to "get back into" an existing virtual community, it seems that once the group is functioning well the sociability benefits will outweigh usability difficulties.

The interesting question, however, is what does this process of group formation need in order for it to happen successfully?

Highly motivated individuals, such as distance learning students needing help, or fans of a TV show looking for information, are probably willing to overcome usability problems and put up with less than perfect interfaces because of their high motivation and engagement with their subject. But what about those people, for example school students, for whom CSCL is being suggested as a tool to enable or indeed create that motivation and engagement?

Do participants in CSCL groups perceive themselves to be in some kind of virtual "space"? If so what kind of space is it? What kinds of visual clues should we be giving to users about the function of the space to be used and what constitutes appropriate behaviour within it? When we walk into any real room we have some idea of its function because of its design and contents, a coffee bar would not be mistaken for a classroom - yet we often use the same virtual layout for spaces used for very diverse functions. With the advent of newer, more mobile, devices such as PDAs and Tablet PCs we have the added issue of "work / study" being done in areas which have traditionally been seen in "leisure" terms - such as in front of the TV or in the open air. We need to find an underpinning framework / pedagogy that will support the development of CSCL across these diverse technologies.

References

Brewster, D. (2002). Towards a whole school ICT policy: An ethnographic study of the 'teacher as user'. In *MSc (HCCS) dissertation.* University of Sussex, School of Cognitive and Computer Science, UK.

Cuban, L. (2001). In *Oversold and underused: Computers in the classroom.* Harvard University Press, Cambridge, MA.

Sunk Cost and the Rating of Facial Attractiveness

Martin Coleman

mdc20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Commitment escalation, or investing more resources into an originally poor choice, is a behaviour familiar to most people. For example, the rule of thumb, "when in trouble, double" is common among UK stock-traders. The phenomenon has also been demonstrated empirically. (Staw, 1976) found that individuals committed greater amounts of money to a failing investment when they, themselves, were personally responsible for the original poor choice. Such work points towards previous losses (sunk costs), not becoming psychologically sunk but, continuing to motivate future decisions. The sunk cost effect ("The greater tendency to continue with an endeavour once a prior investment of time, money or effort has been made," (Arkes & Blumer, 1985), p.124) has also been shown in both behavioural and financial situations. For example, Arkes and Blumer found that customers paying full price for theatre season tickets attended significantly more plays, over a 6-month period, than those subscribing at a discounted rate. Similarly, they found that participants were more likely to invest $1 million dollars in an economically unsound project to build a "radar blank" plane if they had already invested $9 million in the same project. That is, finishing the project so that previous investment was not "wasted". Indeed, the authors argue that all such behaviours are based on "the desire not to appear wasteful".

It can be argued that a self-justification motive prompts both the sunk cost and escalation effects. In the case of escalation individuals are said to commit more funds to an original poor choice in the hope of "turning the situation around". By doing so they hope to prove the ultimate rationality of the original decision. In the case of sunk costs individuals are said to self justify by avoiding the appearance of wastefulness.

In all the works cited thus far the ways for individuals to self-justify have been severely restricted. In almost all experiments to date participants have been presented with purely economic decisions. However, if self-justification does lie behind the sunk cost and escalation effects then, people should be willing to self-justify in other ways. The subjective rating of facial attractiveness is highly consistent between individuals e.g. (Hansell, Sparacino, & Rondi, 1982); (Dongieux & Sassouni, 1980). These ratings, however, are not absolute. Individual judgments can be influenced by a variety of factors including context, radiation, labelling, mood, age and familiarity. This paper examined whether attractiveness ratings are also influenced by sunk costs. The impact of sunk cost on social and perceptual judgments was recorded. 74 female participants aged 18-30 years (mean 22.5, S.D. 3.4) took part in a computer- simulated blind date. The computer simulation necessitated participants to invest either a negligible, moderate or high amount of time, money and effort in getting to see their date. It was hypothesized that participants would increase both ratings of attractiveness and estimated likelihoods of date continuance as sunk costs rose. Both hypotheses were supported (rating of attractiveness ($p=0.01$) and likelihood of date continuance ($p=0.02$)). This was the first empirical study to find both a social perceptual influence of sunk cost.

References

Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processe, 35*, 124-140.

Dongieux, J., & Sassouni, V. (1980). The contribution of mandibular positioned variation to facial esthetics. *Angel Orthodontist, 50*, 334-339.

Hansell, S., Sparacino, J., & Rondi, D. (1982). Physical attractiveness and blood pressure: Sex and age differences. *Personality and Social Psychology Bulletin, 8*, 113-121.

Staw, B. M. (1976). Knee deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance, 16,* 27-44.

Listen to your System: The potential for Auditory Display in Adaptive Systems Research

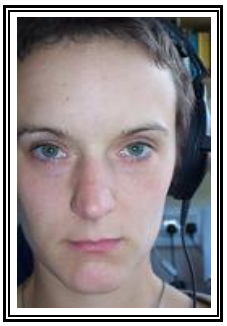

Alice Eldridge

alicee@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Ten years ago, Scaletti and Craig, (Scaletti & Craig, 1991) argued that properly designed auditory displays could potentially increase the amount of data that a human can simultaneously process beyond that achievable with traditional visual displays. The potential for data sonifications is now being explored by a multidisciplinary community of researchers (psycho-acousticians, composers, computer scientists, digital synthesis specialists) in a variety of contexts (assistive technologies, pedagogical aids, HCI, scientific visualisations etc). A substantial amount of research is needed in areas such as auditory attention, multi-modal perception and data-sound mapping techniques before a theory driven approach to sonification can be developed. However, the existing body of research in auditory perception not only supports the potential of sonification as a data analysis technique, but reveals particular characteristics of audio that promote consideration of situations in which its use may be particularly advantageous.

2 Motivation

One obvious advantage of auditory presentation is that unlike visual stimuli, there is no requirement for specific user orientation or attentional focus. This makes auditory display ideal for monitoring abnormalities in background processes, and suggests the possibility of parallel listening - enabling the monitoring and processing of multiple data sets simultaneously. The literature addressing complex, dynamic auditory patterns in speech and music (Bregman, 1990); (MacAdams & Bigand, 1993) reveals basic features of auditory perception such as sensitivity to temporal characteristics and facilitation of detection of small changes in frequency of continuous signals. This suggests that audio representations may be particularly useful in comprehending fast changing or transient data, discriminating between periodic and aperiodic events, and monitoring complex temporal data, or data embedded in other more static signals which may be too noisy to detect in a visual representation. The fact that salient musical patterns can be recognized and recalled even when subject to radical transformations enables the perception of relational patterns that are more difficult to perceive in raw or graphically presented data.

3 Application in Adaptive Systems Research

Many of these characteristics suggest that sonification may be particularly useful in comprehending high-dimensional complex dynamic systems, such as those deployed in artificial life and adaptive systems research. Studies are currently underway to examine the comparative efficacy of auditory and visual displays in users' ability to classify certain qualitative cellular automata (CA) states. The final states of these discrete dynamic deterministic systems fall into one of three classes: ordered (point or limit cycles), random, or complex (Wuensche, 1997). It is felt that classification of these dynamic patterns may be amenable to auditory analysis.

The outputs of the one-dimensional binary CAs are represented in audio by taking the states (on or off) of individual cells to drive the production of tones (play or rest): the dynamic graphical patterns are transformed into rhythmic patterns. In addition, statistics describing the frequency distributions of the specific production rules are used to define the pitch. Four successive states are played simultaneously, thus the current state is always accompanied by the 3 previous states, providing a context similar to that of 2 dimensional visual displays.

Initial work confirms that ordered states can be recognised from the rhythmic representation alone. Although accuracy rates of pure audio displays are not yet available, early work suggests that for pattern recognition tasks, a considerably smaller

amount of information need be presented for accurate classification.

Although promising, the current implementation is not considered to surpass a visual display in efficacy in this context. Future work aims to develop more sophisticated mappings that capitalise on special features of auditory perception, for example methods of transforming the data are being explored such that aperiodicities appear as interaural discrepancies.

References

Bregman, A. S. (1990). In *Auditory scene analysis.* Cambridge, MA.

MacAdams, S., & Bigand, E. (1993). In *Thinking in sound: The cognitive psychology of human audition.* Oxford: Clarendon Press.

Scaletti, C., & Craig, A. B. (1991). Using sound to extract meaning from complex data. In *Proceedings of processing, display and interaction ii.*

Wuensche, A. (1997). Attractor basins of discrete dynamic deterministic networks CSRP 461. In *PhD thesis.* School of Cognitive and Computing Sciences, University of Sussex, UK.

Distributed Denial Of Service; A Modern Problem

David Ellis

davide@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Distributed Denial Of Service (DDOS) is a fairly recent phenomenon although the ideas behind it are nothing new. Denial Of Service means: "An explicit attempt by an attacker to prevent legitimate users from using a service". DDOS is a DOS where the attackers are distributed world wide and all target a particular service. DDOS attacks have become a real problem in the last 4 years with attacks against Yahoo, EBay, GRC, UK and US government websites and CNN; each attack lasting as long as a week!

The are two main types of Denial Of Service in general these are:

- Flooding the Network Bandwidth
- Starvation of a Servers Resources

Server resources which include: memory; cpu time; disk space; running programs; among others are the services which attackers attempt to exhaust. Network bandwidth is simply another resource which can be the target of an attack. This paper will focus on DDOS attacks against network bandwidth.

DDOS attacks may be focused at a number of different levels: At a home PC user, a large dot.com company or maybe even an entire ISP. Because of these varying levels of granularity there have been a lot of different architectures and mechanisms designed to try to combat such attacks.

The next section will describe the DDOS tools which are in use. Section three will then discuss some recent research in the area of stopping DDOS and the final section will conclude the paper.

2 How to acheive DDOS

Most DDOS attacks work by "flooding" the bandwidth of the target. This means they send more data to the target than the target can handle, thus rendering the target unable to receive requests from legitimate clients. Figure 2 shows the general idea of DDOS attack.

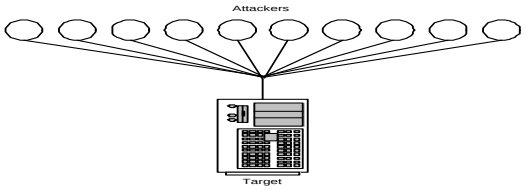

Figure 1:

Many tools exist for automating such attacks including Shaft (Dittrich, Long, & Dittrich, 2002), Trinoo (Dittrich, 1999b) and Stacheldraft (Dittrich, 1999a). Each tool is similar and provides the general framework above, the only real differences between the tools are the varying usage of encryption and the use of different types of packet flood.

A typical DDOS attack follows the cycle shown if Figure 2. This set of events can be made to happen automatically using viruses and worms such as code red (team, 2003), it can also happen semi-automatically; whereby an attacker has a script which performs the tasks automatically, finally this can be acheived manually although this is rare.

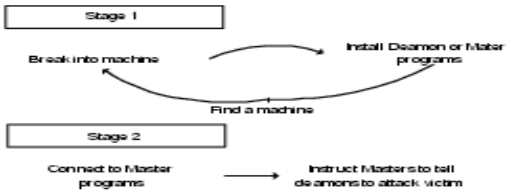

Figure 2:

There are a number of stages at which we could stop people launching DDOS attacks. The first is the exploitation stage, the attacker must first acquire a number of "zombie" machines; If all machines on the Internet were secure so people couldn't hack them, zombies couldn't be installed

and nobody could launch DDOS attacks. The second stage is the identification of control traffic generated when the attacker is talking to the zombies. The final stage is when the attack is actually happening.

3 Current Research

There has been a lot of research in the area of congestion control in networks such as RED and CBQ. None of these however apply directly to the problem of DDOS but are nonetheless useful. Some research into the use of advanced routers is being conducted in (Ioannidis & Bellovin, 2002) and (Gil & Poletto, 2001); these use algorithms in an attempt to identify the attackers and pass this information to more upstream routers. One major problem in a DDOS attack is that the source address of the attacker or the zombies can be forged; this is known as spoofing. Many people have been working on anti-spoofing technologies such as probabalistic packet marking, icmp traceback messages and other router level protocols. More recently people have begun thinking about using peer to peer overlay networks such as RON (Andersen, 2001), SOS (Keromytis, Misra, & Rubenstein, 2002) and MayDay(Andersen, 2003). This kind of network is a layer above the normal TCP/IP layer; This provides us with a number of advantages over standard routing: more accurate metrics mean faster routes; if one node in the overlay network is lost through DOS or otherwise the others can still provide the route to the server.

4 Concluding Remarks

The paper has examined the recent threat of Distributed Denial Of Service attacks prompted by the rise in incidents over the past few years. We have examined the tools the people are using, and how the attackers are able to use these tools. Finally we present some ongoing research areas in this area.

References

Andersen, D. (2001). *Resilient overlay networks.*

Andersen, D. (2003). *Mayday: Distributed filtering for internet services.* http://nms.lcs.mit.edu/papers/mayday-usits2003/paper.html.

Dittrich, D. (1999a). *The "stacheldraht" distributed denial of service tool* (Tech. Rep.). University of Washington.

Dittrich, D. (1999b). *The dos project's "trinoo" distributed denial of service tool* (Tech. Rep.). University of Washington.

Dittrich, D., Long, N., & Dittrich, D. (2002). Analyzing distributed denial of service tools: The shaft case. In *Usenix.*

Gil, T. M., & Poletto, M. (2001). MULTOPS: A Data-Structure for bandwidth attack detection. In *Usenix* (pp. 23–38). citeseer.nj.nec.com/gil01multops.html.

Ioannidis, J., & Bellovin, S. M. (2002). Implementing pushback: Router-based defense against DDoS attacks. In *Proceedings of network and distributed system security symposium, catamaran resort hotel san diego, c alifornia 6-8 february 2002.* 1775 Wiehle Ave., Suite 102, Reston, VA 20190: The Internet Society.

Keromytis, A., Misra, V., & Rubenstein, D. (2002). *Sos: Secure overlay services.*

team caida. (2003). *Caida analysis of code-red* (Tech. Rep.). Caida.

Title

Eunice Fajobi

eunicef@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

To be inserted ...

Genes Don't Make Cakes

Chrisantha Fernando

ctf20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Evolution acts incrementally and compositionally to construct systems consisting of hierarchical autonomous primitives (HAPs). The existence of HAPs has been eluded to in Simon's "nearly-decomposable systems", in work on the evolution of modularity and in recent work subsumed under the term "dynamical hierarchies". Degeneracy, canalization, and biological complexity exist as a consequence of the existence of HAPs. HAPs can be utilized to help solve the competing conventions problem so facilitating effective recombination. HAPs underlie the massive ontogenic re-organizability seen in nervous systems, in genomes themselves, and indeed in any biological system and so confers ontogenic robustness which facilitates increased evolvability. Evolutionary approaches to neural network design have not utilized HAPs, and my research seeks to demonstrate their existence.

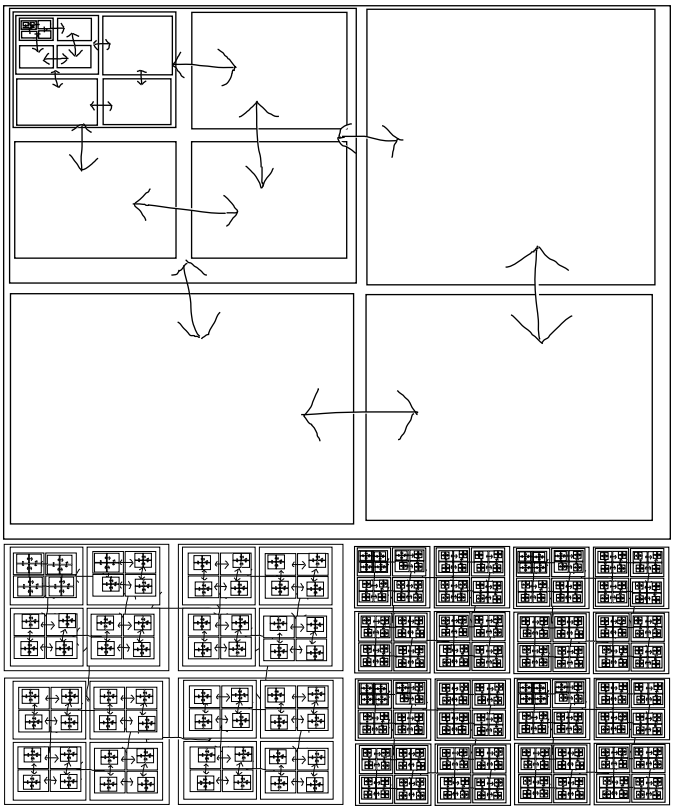

Each box represents an Autonomous Primative. Arrows represent interactions between primatives within a level (i.e. between boxes of the same size).

Figure 1: A Diagram Showing Hierarchical Autonomous Primitives

2 Research

Metaphors are essential and misleading. Genetics has been highly influenced by culinary and architectural metaphors; genes are the "recipes" or "blueprints" for things. This immediately implies an end point, i.e. a cake or a building, which having been constructed presumably we either eat, or live in. This is nice, but it bares little relation to the truth. Rather, we claim that what is produced is a hierarchy-of-societies of living primitives. The hierarchy of primitives is shown in figure 1 and consists of a recursive Russian doll arrangement of boxes.

However, unlike the Russian doll, each box (primitive) contains four smaller boxes, and each of these smaller boxes consists of four smaller boxes and so on ad infinitum. The bi- directional arrows represent interactions between primitives. There is incomplete connectivity, constrained by two principles; only adjacent boxes of the same size interact directly. One representation this diagram can sustain is as follows. Let the smallest box size be the genotype. Constrained interactions (represented here as adjacency constraints) at the level of the genes results in the formation of amino acids. Let the amino acids be the boxes one size larger than the gene boxes. Constrained interactions between amino acids produce proteins. Let the box one size larger than the amino acid boxes represent proteins. Constrained interactions between different proteins result in increasingly polymerized and complex assemblies, e.g. microtubules, actin-myosin chains, and sub-cellular compartments. Let us assume several levels of organization whereby increasing orders of object

are produced until eventually we reach a cell, and so let the box one size larger than proteins represent cells. Constrained interactions between cells of similar type produce tissues and so let the box one size larger than the cell represent tissues, similarly tissues form organs, and organs form organisms, organisms form groups of similar organisms, and multiple groups of different organisms form ecosystems. At each level there exist particular characteristic patterns of interactivity. This diagram gives a flavour of hierarchy but it does not convay the proposed autonomy of the primitives. By autonomy we mean that each primitive can be treated as an agent, envisaged as attempting to satisfy some cost function, i.e. undertaking some goal directed behaviour, on the basis of a value function, i.e. a set of "rules" governing behaviour [1].

An operational definition of autonomy is given in Figure 2. Quite simply, it makes sense to define something as autonomous if we can gain insight into its behaviour by treating it like an agent that makes "decisions" for itself. We start by showing how water molecules provide a moot, but critical, example of autonomous agents.

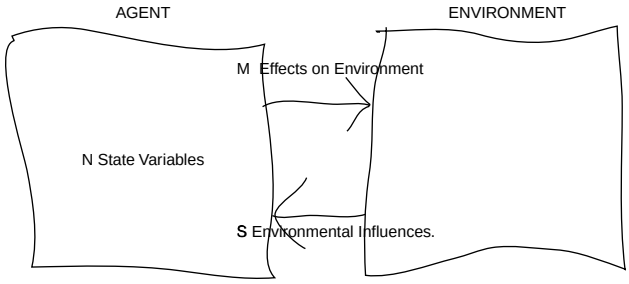

An autonomous agent is defined as one which seeks some goal. The greater the number of perturbations it is able to overcome to achieve that goal the more autonomous it is. Consider the above coupled dynamical system. If we treat the agent as an autonomous dynamical system (in the technical sense), then the greater the number of parameters we are able to alter in that agent whilst not affecting the achievement of the agents' goal, the more autonomy we say that the agent possess. There are several mechanisms whereby an agent could achieve such autonomy. One method is to increase the number of state variables within itself. The greater the amount of internal re-organisation, plasticity, or self-organising ontogenetic potential, the larger is the number of environmental perturbations that can be compensated for. This is effectively Ashby's "Law of Requisite Variety."

Figure 2: An Operational Definition of Autonomy

Further information can be found on `www.chrisantha.com`.

[1] The term "rules" is not used to mean a discretely represented set of I/O mappings, but to be a continuous dynamical system which defines the behaviour of that primitive when coupled to a particular environment. The hierarchical arrangement may be seen as effectively constraining the nature of this environment such that primitives tend to be most influenced by adjacent primitives of the same type, and only slightly influenced (i.e. be affected by only a small number of parameters) from primitives of different layers.

Emotion Cubes: An exploration into pleasurable experience with interactive technology

Rowanne Fleck

rmmf20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

The field of Human Computer Interaction is expanding, now trying to create technologies that are not only more usable, but that are a pleasure to use too. Emotion Cubes is a system that uses a tangible interface to explore some of the issues concerned with developing a pleasurable user experience with interactive technology. In particular it was used to explore how the amount of control users have over their experience effects how much they enjoy it.

2 A pleasurable user experience

There are a number of views on the amount of control people should have in order to secure an enjoyable experience. The theory of flow states that optimal experience occurs most frequently when people feel they have control over their actions (Csikszentmihalyi, 1997). Laurel also argues that to have a pleasurable experience with interactive technology, the user must be able to input into the system frequently, have a large range of choices open to them at these times and that these choices should have an impact on the outcome of the interaction (Laurel, 1986); (Laurel, 1991). The more control then they have over their experience, the more enjoyable it will be. However, it is also suggested that we are interested in what is novel, surprising or uncertain in our environment, that these situations raise our arousal (readiness to react) levels and that transitory jumps in arousal can be pleasant (Berlyne, 1960). (Gaver & Beaver, 2003) talk about how ambiguity in a system can enhance user experience. Similarly Inference Machine, a piece of interactive art, deliberately only allowed people to influence rather than control the system so they could derive pleasure from reflecting upon the questions it raised (Sengers & Liesendahl, 2002). This body of research suggests that the unexpected results obtained by reducing the amount of direct control people have over their experience could also lead to greater enjoyment. Emotion Cubes varied the amount of control people had to see if any of the above findings were supported.

3 Emotion Cubes

The 2 emotion cubes are approximately 10cm cubed with embedded tags which allows their orientation to be read. One has 6 different colours on its faces and the other has 6 different textures. When both the cubes are placed on the tag reader, an animation and sound is played on a screen. There were two different conditions. In one condition the combination of upper faces of the cubes controls the animation played, providing the user quite direct and high levels of control. In the other condition, a random colour value is assigned each time the colour cube is replaced and a random texture assigned each time the texture cube is replaced. The animation associated with this combination is then played on the screen. This reduces the amount of control the cubes give over the animations. Participants were assigned either to the higher or lower control condition and asked to play with the emotion cubes thinking about the kind of emotions suggested to them by the animations. They filled in a questionnaire about their experience and discussed their views about it after the session.

4 Analysis in Progress

Observation throughout the trial and a brief look at the questionnaire results suggest that people did not find the higher control condition any more or less enjoyable than the lower control condition and did not find the lower control condition more frustrating or, contrary to predictions, more surprising. They also did not rate themselves as having more control over the animations in the higher control condition than the lower control condition. Although there did not seem to be any link between how much control people rated themselves as having in the lower control condition and how much they said they enjoyed the experience, the more

control people felt they had in the higher control condition, the more they said they enjoyed the experience. These findings seem to support the view of Laurel and Csikzentmihalyi that a greater sense of control leads to more enjoyment when direct control is offered. However, they also suggest that when the aim of the experience is not to provide direct control, the relationship between control and enjoyment is not so clearly defined. People generally said they enjoyed using the tangible interface or made no comment about it at all. A more detailed look at the data including the video footage is required.

5 Conclusion

Initial observations highlight a number of problems with the described set up to investigate the complex issues involved when creating pleasurable experiences. On further examination of the literature it becomes clear that what constitutes a pleasurable or desirable experience depends largely on the kind of task being carried out, individual differences, and many other factors not considered in the Emotion Cubes design. There is also a possibility that the lower control condition is not sufficiently different from the direct condition or well constructed enough to create higher levels of novelty, surprisingness or ambiguity and that these values are hard for people to rate on a 5 point scale. Equally hard to rate on such a scale was the amount of control people considered themselves as having with the comments they made clearly indicating different criteria were used. A number of useful lessons have been learned from this experiment and a different, more directed approach may be more appropriate to investigate the issues raised further.

References

Berlyne, D. E. (1960). In *Conflict, arousal and curiosity.* McGraw-Hill.

Csikszentmihalyi, M. (1997). In *Finding flow: the pschology of engagement with everyday life.* Basic Books, New York.

Gaver, W. W., & Beaver, J. (2003). In *Ambiguity as a resource for design.* CHI (in submission).

Laurel, B. (1986). Interface as mimesis. In D. A. Norman & S. W. Draper (Eds.), *User centred system design: new perspectives on human-computer interaction.* Hillsdale, Erlbaum.

Laurel, B. (1991). In *Computers as theatre.* Reading, Addison-Wesley.

Sengers, P., & Liesendahl, R. (2002). The enigmatics of affect. In *Designing interactive systems.* ACM, London.

Exploring the neural code

Benoit Gaillard

bg22@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

The brain is commonly accepted as being the place where our cognitive processes occur, and within it the neurons are commonly considered as playing a central part in those processes. There is an incredibly big number of neurons in a human brain: 20 billions, and they communicate through synapses that connect each neuron to many others. The number of synapses is even more impressive: 240 trillions. The communication is done with sequences of action potentials or spikes: spike trains. A spike is an electrical pulse that is sent from one neuron to the other, that is very short and always have the same shape: It doesn't carry any information in itself. What carries information is its relationship to the receiving neuron an to the other spikes: How many spikes has the neuron received, in how much time or at which precise instants... those spike trains whose spatio-temporal characteristics carry information and underlie our thoughts are the neural code. Although many interpretations have been proposed, this code is still far from being understood. So, trying to make any sense out of those spike trains is the global motivation for my work, and, as i'm just starting, i'm currently studying various approaches, by applying them to practical visual tasks.

Firing Rate

Historically, this is the first approach to the neural code (Adrian(1926), (Adrian, 1926)). It is based on the idea that, the more spikes a neuron receives, the more excited it gets, and the more spikes it sends (it fires). Thus, we can evaluate the degree of excitation of a neuron by the frequency of emission of spikes: its firing rate. I have used this code for studying the variations in the discrimination performance of a neuron, in a noisy environment, according to the ratio between inhibitory and excitatory incoming synapses. It follows The work of Jianfeng Feng ((Feng & Liu, 2002)).

However, evaluating this firing rate is too long for the brain to use it for real tasks: we need to average the time between two spikes over at least a hundred spikes, and we know that it takes around 10 milliseconds for a neuron to fire,while the brain can perform high level discrimination tasks such as making the difference between an animal and a means of transportation in 150 ms. (Van Rullen and Thorpes,2000,(VanRullen & Thorpes, 2000)).

Spike timing

This time constraint leads to the conclusion that information between neurons must be carried with only a couple of spikes, embedded in their precise timing.

Absolute spike timing Knowing the exact instant of the arrival of each spike at the synapses would carry all the information we want, however we would need a sort of absolute time reference in the brain, which is not very biologically plausible.

Rank order A more realistic model is to deal with the order of arrival of spikes, without focusing on the precise instant of their arrival. This can be easily implemented in the brain, as the spikes propagate at finite speed between the neurons. This has been used for modelling learning (Sejnowski, (2001) (Rao & Sejnowski, 2001)), and for a rapid visual processing model, by Thorpes & Vanrullen.(2002, (VanRullen & Thorpes, 2002)). I have played around with this code as well, trying to evaluate how efficient it can be at performing neural competition for motion perception.

Population code

The main problem of those spike timing neuronal codes is that they lead to models that are poorly adaptable and very sensitive to noise or variations in the input. to overcome those weaknesses, we need to come back to the study of the statistics of the spike trains (The mean firing rate being only one amongst the many statistical features). To overcome this problem of time that we've seen

with the firing rate code, we can get our statistics from the times between the very first spikes of a population many neurons at the same time, because we know that there are a lot of neurons and connections in the brain. trying to relate those population statistics to the some proposed rank order codes, by doing models that perform similar tasks, is the next direction of my research.

References

Adrian, E. D. (1926). The impulses produced by sensory nerve endings. *Journal of Physiology (London), 61*, 49-72.

Feng, J., & Liu, F. (2002). A modelling Study On Discrimination Tasks. *Biosystems.*

Rao, R., & Sejnowski, T. J. (2001). Spike-timing hebbian plasticity as temporal difference learning. *Neural Computation, 13(10).*

VanRullen, R., & Thorpes, S. (2000). Is it a bird, is it a plane? Ultra-rapid visual caegorization of natural and artifactual categories. *Perception*, 539-550.

VanRullen, R., & Thorpes, S. (2002). Surfing a spike wave down the ventral stream. *Vision Research, 42(23)*, 2593-2615.

Are you fit enough to recognize failure?

Miguel Garvie

mmg20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

Self-diagnosing hardware is important especially in mission critical systems exposed to radiation. Built-in self-test (BIST) is widely used yet commonly requires more than 100% overhead as in double redundant systems or off-line testing. Evolutionary methods applied to hardware have produced circuits comparable to those designed by experts and also unconventional circuits in which hardware resources are used extremely efficiently. Moreover, many evolved systems in nature exhibit self-diagnostics such as the immune system.

All this led to the prospect that evolutionary methods could explore areas of design space which reuse hardware components so that they contribute both to the circuit's main functionality and its BIST, leading to a low overhead on-line solution. Below we describe the first ever attempt at realizing this possibility of evolving self-diagnosing hardware (Brill & Helweg, Accessed 20/05/02), (Gaunaurd, Brill, Huang, Moore, & Strifors, 1998).

A generational genetic algorithm (GA) was used with a population of 32 individuals. The amount of conventional design knowledge used to set up the fitness evaluation function and the mapping from genotypes to circuits was kept at a minimum. Evolving circuits were made up of two input logic gates and were evaluated in a simple digital logic simulator where noise was introduced in order to facilitate transfer to real hardware. Hardware faults were simulated by sticking a gate output at 0 or 1, a model well established in industry.

Small circuit tasks were chosen as good starting points to establish a proof of principle that BIST functionality could be evolved for them: a one bit full adder, a two bit multiplier and an edge triggered D-latch. The fitness function evaluated a number of circuit properties here listed in decreasing priority order: perform the desired task, off-line BIST, on-line BIST, minimize gate count. BIST behaviour was evaluated by testing if an extra output E went high when the task outputs were incorrect due to an induced fault. And so a process of *survival of the meekest* was commenced.

From a population of random individuals, after 14100 generations of evolution, there emerged an individual performing the adder task using the minimum 5 gates and having 90% fault coverage off-line BIST using only 2 extra gates overhead. This circuit performs a hybrid of on-line/off-line self-diagnosis which could be implemented in a BIST system with 31% the overhead of the conventional off-line solution. About 15000 generation later a full (100% coverage) on-line BIST solution for the adder was found using only 50% the overhead of the conventional on-line solution. Another run that imposed extreme noise conditions arrived at an on-line solution that includes a low-pass filter to iron out glitches at the output. In effect, this circuit could be clocked at twice the speed as the conventional on-line BIST solution.

A new run was seeded with a hand designed multiplier using the minimum 7 gates. Nearly 150000 generations later it suffered one modification while 4 gates were annexed performing full off-line BIST requiring 36% the overhead of the conventional equivalent. A multiplier with full on-line BIST was also evolved from a population of random genes after roughly 4 million generations. This circuit used 64% the overhead of the conventional on-line solution and its unconventional structure is shown in Fig. 3. An on-line self-diagnosing edge-triggered D-latch was also evolved after 3 million generations and had the same structure and overhead as the conventional solution.

Self-diagnosing circuits have been evolved for the first time and are competitive with conventional ones in terms of fault coverage and gate count overhead. Evolved circuits exploit conventional design principles such as voting and design diversity and also unconventional principles, such as computing checksums while cascading outputs. These principles which allow them to reuse logic for the main

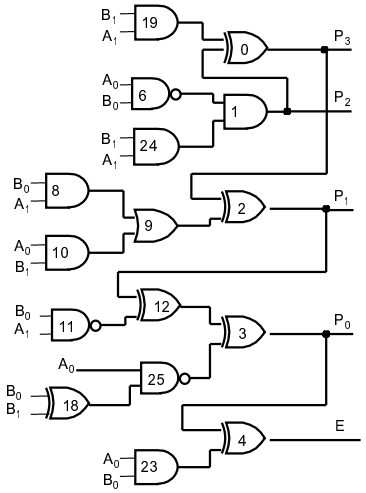

Figure 1: Evolved two bit multiplier performs full on-line self-diagnosis with only 64% the overhead of the conventional equivalent by reusing logic for the main task and self-test.

task and BIST could prove useful if adopted by designers. Some circuits were extremely modular in structure while others were inscrutable. The reason for this is unknown, but then again *evolution moves in mysterious ways*.

Previous work (Sansone & Harackiewicz, 2000) suggests larger circuits are riper to evolutionary optimization yet unfortunately computational power is the limiting factor when evolving them (you can easily contribute your unused CPU time to this project at `http://www.cogs.susx.ac.uk/users/mmg20/contrcpu.htm`). Our current efforts include the evolution of BIST for industry sized modules, of self-diagnosing analog circuits perhaps under varying operating conditions (Luckin, 1998) and of circuits capable of testing the tester under multiple faults.

References

Garvie, M., & Thompson, A. (2003a). Evolution of combinational and sequential on-line self-diagnosing hardware. *In press.*

Garvie, M., & Thompson, A. (2003b). Evolution of self-diagnosing hardware. *Proc. 5th Int. Conf. on Evolvable Systems (ICES '03): From biology to hardware, 2606 of LNCS*, 238.

Thompson, A., Harvey, I., & Husbands, P. (1996). Unconstrained evolution and hard consequences, towards evolvable hardware. *vol. 1062 of LNCS*, 136.

Vassilev, V., & Miller, J. (2000). Scalability problems of digital circuit evolution: Evolvability and efficient designs. *The Second NASA/DoD workshop on Evolvable Hardware*, 55.

Let's talk about it: Exploring the qualities of effective peer collaboration

Amanda Harris

alh23@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

My research aims to examine the skills and abilities children bring to a collaborative learning environment and in doing so identify the circumstances that lead to successful collaboration. I am working as part of the Riddles project which is concerned with improving children's reading comprehension skills through training that draws attention to the ambiguity in joking riddles.

Linguistic awareness is the ability to treat language as an object and to reflect on its structural features such as phonology and syntax. (Yuill & Oakhill, 1991) identify a deficit in this ability among poor comprehenders. A poor comprehender is a child with a reading comprehension age at least 6 months behind their reading accuracy age. The riddle is a linguistic device which requires being able to reflect on the use of ambiguity in language in order to appreciate its humour. (Yuill, 1998) found that linguistic awareness can be developed by drawing children's attention to the ambiguity in riddles and encouraging discussion on the different interpretations of the language. Significant improvements in children's reading comprehension skills were found after exposure to the riddle training. The training has been developed into a multimedia software package called 'Joke City', which presents riddles to pairs of children whose task it is to identify the ambiguity within each riddle. For example:

Question: How do you make a sausage roll?
Answer: Push it down a hill.

In order to 'get the joke' children have to disambiguate the word roll and understand it has two meanings. Joke City draws their attention to this by asking them to 'click on the part of the word that has two meanings'.

In the current project Joke City is being developed with particular emphasis on exploring the role of linguistic awareness in text comprehension, the use of multiple external representations of meaning, the use of peer collaboration and the implementation of software scaffolding. My research focuses on how best to support peer collaboration in this context in order to maximise the learning benefits it can offer.

Joke City provides an ideal collaborative context; two children working together towards a common goal. However, observations of videotaped data of children interacting with the software have revealed that interactions within this context are not always collaborative. Some pairs collaborate very effectively while others consistently fail to work together productively. I am interested in exploring the qualities that make these interactions so different and as a result aim to identify the circumstances that may contribute to a successful collaborative experience.

Defining collaboration has proven a difficult task as the term is used very differently both within and between different disciplines (Dillenbourg, 1999). The definition provided by (Teasley & Roschell, 1995) below is perhaps the one most commonly used and quoted in the literature.

> Collaboration is a coordinated, synchronous activity that is the result of a continued attempt to construct and maintain a shared conception of a common problem. (p. 70)

Underpinning their emphasis on 'coordinated, synchronous activity' is the assumption that all members of the collaborative group participate equally in the task and there is no division of labour amongst participants. However, this implies that all members of the group are both equally able and equally willing to participate in the collaborative process (Burton, Brna, & Treasure-Jones, 1997). I would argue this is not necessarily the case and a clearer notion of the different qualities participants contribute is vital to a better understanding of the nature of collaboration.

Identifying what contributes to effective collaboration has implications for both classroom practice and software design. For example, the way in which collaborative contexts are constructed will be informed by this work. In addition, if the learner

model is able to assess the quality of the interaction, scaffolding could be provided in order to support the collaborative process.

References

Burton, M., Brna, P., & Treasure-Jones, T. (1997). Splitting the collaborative atom: How to support learning about collaboration. *Artificial Intelligence in Education: Knowledge and Media in Learning Systems,* 135-142.

Dillenbourg, P. (1999). In *Collaborative learning: Cognitive and computational approaches.* Pergamon, Oxford, UK.

Teasley, S. D., & Roschell, J. (1995). The construction of shared knowledge in collaborative problem solving. *Computer Supported Collaborative Learning.*

Yuill, N. (1998). Reading and riddling: Special issue on language play. *Cahiers de Psychologie Cognitive, 17,* 313-342.

Yuill, N., & Oakhill, J. (1991). In *Children's problems with text comprehension: An experimental investigation.* Cambridge University Press. Cambridge.

Aspects of Semitic Tense

Jason Harrison

jgh20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

Much of the work on the problem of how to interpret the verbal system of the Semitic languages has been focused on Biblical Hebrew and Arabic (both Qur'anic and Modern). This problem derives from attempting to impose Indo-European-language-centric interpretations onto a non-Indo-European group, in particular grammatical terminology from Greek and Latin. This leads to the claim that languages like Hebrew, Arabic, and Syriac have no "tense" when in fact a distinction should be made between having no tense and not marking tense on the verb. This in turn has led to confusion in the prescriptive grammar books of the abovementioned languages, with some using the term "perfect" to indicate a past tense, others a completed action; and the "imperfect" as a future (or non-past encompassing both present and future), or as an incomplete action. This category of complete/ incomplete has nothing to do with tense, but is instead an aspectual distinction and is binary in nature.

In my research I will look at the discussion of the historical explanations of the Semitic verbal system (with Biblical Hebrew as the prime example) and will look at the issue of linear time as represented in general linguistic theory, which, as mentioned before, for the large part has been occidental and Indo-European language focused. Following from this I will examine the different theories of tense and aspect as applied to the Semitic verbal system (Goldfajn, 1998), (Driver, 1936), (Anderson, 1974), as well as a survey of the literature in particular on the Syriac verb. Parallel occurrences of identical verb forms in the Syriac Gospel versions will especially be treated to discern whether their translation from Greek is identical and whether they can be (should be) interpreted with an equal sense. My research will involve an examination of expressions of tense and aspect within the verbal system of Syriac focusing on a specific corpus, that of the Old Syriac, Peshitta and Harklean versions of the New Testament Gospels, and how they relate to the same expressions within the (supposed) original *Koine* Greek. I will argue that certain innovative structures had to be developed to accommodate for translation between an Indo-European language with its complex verbal system allowing for both tense and aspect and a Semitic one in which the verbal system is traditionally seen as being purely aspectual.

Within biblical Hebrew, Arabic, and Syriac the Past/ Perfect is generally regarded as the base upon which the other senses and meanings are built and is sometimes termed the suffix conjugation (which is a more temporally neutral term), as it takes suffixes to indicate person. The Future/ Imperfect is termed the prefix conjugation as it takes prefixes to differentiate it from the suffix conjugation as well as some personal suffixes. This would seem to indicate a bipolar differentiation of the verbal system. But this distinction is more complex than simply a past/ non-past distinction.

A survey (McFall, 1983), found that in the translations of the Hebrew "tense" forms in the Revised Standard Version Bible the suffixed **qṭl** form could be translated either by a past, present tense or even a future. Similarly with the prefixed **yiqṭl** form; while most had a future sense translation, the present and also the past could be used. So clearly something other than tense is marked on the verb. This something is normally taken as being aspect. The theory that the "tense" forms are actually aspectual had been put forward as early as the 1820s (Ewald, 1827) , with the past representing perfective (meaning completed) and the future, imperfective (meaning incompleted). This, rather than clearing up the confusion only served to add to it particularly concerning the role of the proclitic **w**-conjunction, or *vav* conversive, usually translated as "and", which when added to a **qṭl** form, seemed to give it a future tense, and when added to a **yiqṭl**, gave it a past tense sense. A similar occurrence appears in Syriac, which will be looked at more thoroughly during the course of my research.

As the tense systems of Biblical Hebrew and Qur'anic Arabic have been assumed to be virtually identical and to be representative of the Semitic tense system, it is held that whatever theory holds for one language will hold for the other language and by extension for the Semitic language family.

Out of this work on the two abovementioned languages (Binnick, 1991):435 says that there are two main theories concerning Semitic "tenses"

1. those which look at purely semantic meanings within which are those involving aspect, relative tense, absolute tense or other and,

2. those which involve pragmatics.

Comrie in his work *Aspect* (1976) suggests a mixture of the two.[1] Relatively little research has been done on the verbal system of Syriac, which for a number of reasons is an extremely important Semitic language, from linguistic, theological and historical perspectives. Syriac was used as a bridge language in the translation of texts from Greek to Arabic. Greek philosophical, medical and scientific texts were first translated into Syriac and then from Syriac to Arabic, either incidentally (that is, they already existed in Syriac translation) or as part of the translatory process.

Debate continues to rage over whether the first Gospel portions were composed in Koine Greek, or an Aramaic dialect (possibly Syriac) [2] , which could have profound theological implications, the examination of which does not fall within the scope of my current research or the remit of this paper.

References

Anderson, F. I. (1974). *The sentence in biblical hebrew.* The Hague: Mouton.

Binnick, R. (1991). *Time and the verb: A guide to tense and aspect.* Oxford: Oxford University Press, UK.

Black, M. (1967). *An aramaic approach to the gospels and acts.* Oxford: Oxford University Press, UK.

Comrie, B. (1976). *Aspect: an introduction to the study of verbal aspect and related problems.* Cambridge: Cambridge University Press, UK.

Driver, G. R. (1936). *Problems of the hebrew verbal system.* Edinburgh: T & T Clark.

Ewald, H. (1827). *Kritische grammatik der hebräischen sprache.* Leipzig: Hahnschen Buchhandlung.

Goldfajn, T. (1998). *Word order and time in biblical hebrew.* Oxford: Oxford University Press, UK.

McFall, L. (1983). *The enigma of the hebrew verbal system.* Sheffield: Almond.

[1] (Comrie, 1976):63

[2] Matthew Black (Black, 1967):120 suggests that the statement of Eusebius in *Historiae ecclesiae* IV.22.8 that Hegessipus quoted from the Syriac could refer to a Syrian Gospel from before c. 170 CE ∀[Hegessipus] sets down certain things form the Gospel of the Hebrews and the Syriac (Gospel) and in particular from (writing in) the Hebrew tongue.

Evaluative Lexis in the Religious Domain of Early Modern English

Maggie Kerridge

maggiek@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

My interest is in changing ideological positions and how these can be deduced from changing language use. I am investigating the role of evaluative vocabulary in Early Modern English with reference to Martin's Appraisal Systems (Martin, 2001). Unlike current government pronouncements, which are intended to change popular thinking through persuasion or spin, 16th and 17th century Royal Proclamations were commands and carried the ultimate charge of high treason for non-compliance. During the reign of Henry VIII the Church in England broke from Rome. Given Henry's consequent preoccupation with affairs of religion and my discovery from preliminary research into various areas of vocabulary that religious terminology was the most vitriolic in Early Modern English, this seemed a profitable area to investigate for evidence of ideological change. Fear of social unrest and its threat to their hold on power seems to dominate Tudor Proclamations concerning religion. I plan to investigate whether this continues throughout the period studied.

To discover whether social attitudes generally appear to have been regulated, I am looking at changing language use in Jonathan Culpeper's unpublished Corpus of English Dialogues (1560-1760), to which I have kindly been allowed access. The trial proceedings section in particular will provide a fertile source of contemporary attitudes, since they constitute a dialogue between those in authority holding the approved or politically correct line and those who oppose this. Fairclough recently explained the use of politically correct language thus:

> 'people interact…they represent to themselves and each other what they do…what they do is then shaped and reshaped by their representations…this is the basis of theories of social life as socially constructed as an effect of discourses. Such processes of intervention…can be seen as attempts to change discourses on the assumption that changing discourses…may lead to changes in other elements of social practices through processes of dialectical internalization.'
>
> (Fairclough, 2003):22

My hypothesis is that the application of politically correct language is far from new. Its persuasive purpose has been more lightly disguised than that of other evaluative words but its function is the same. We are now more aware of the concept because the diffusion of social influence has led to the sources of words that succeeded in influencing usage becoming more diverse. The use of politically correct terminology is an attempt to make the attitude of the user appear to be the norm and is emulated to show solidarity but

> 'relatively successful enactment does not guarantee relatively successful inculcation…people may acquiesce to new discourses without accepting them - they may mouth them rhetorically for strategic and instrumental purposes.'
>
> (Fairclough, 2003):25

In the 16th century these purposes would have been political expedience and the attitude was frequently derogatory, whereas modern politically correct usage tends to substitute a less distasteful word for a more offensive one. Given the tendency of derogative words to lose their negative impact over time, do any of the terms under consideration lose their negative evaluation over the period and is any such change in rhetoric reflected by a change in reality?

References

Fairclough, N. (2003). *Political Correctness in Discourse and Society, 14:1*.

Martin, J. R. (2001). Beyond exchange: Appraisal systems in english. In *Evaluation in text oxford.* Oxford University Press, UK.

Face Recognition and Memory Awareness

Ira Konstantinou

irak@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

The research conducted for my Dphil was based on the notion that recognition memory reflects two components of conscious awareness operationalized by remember and know responses. Every positive recognition decision is followed by either a remember or know response by the participant. Some researchers have proposed that remembering reflects the ability to mentally travel back in time and become aware of past experiences. Knowing is proposed to involve the awareness of knowledge about the world without reference to one's own experiences. For these researchers remembering and knowing reflect memory systems rather than processes, and they do not bear a one- to-one relationship to processes during recognition. Other investigators directly identify remember with the slower, controlled process of recollection and know with the faster, automatic process of familiarity during recognition. In order to address this difference in views, my research employs the remember-know paradigm along with a procedure designed to control the time between the test item presentation and recognition judgments. If short deadlines allow only the familiarity process to influence recognition and long deadlines allow the more controlled recollection processes to govern recognition, then this should be apparent in the remember and know responses. Previous research has shown that under this procedure know responses do not directly identify with a fast, automatic process, as participants give more know responses when given more time to respond.

Using novel and famous faces as stimuli, further evidence regarding remembering and knowing at short and long deadlines was sought by manipulating study time or the level of processing during study, resulting in the effects of these manipulations being obvious in remembering and not in knowing. Moreover, with the longer deadline, a parallel increase in knowing was found when levels-of-processing manipulation was used with the famous faces but not with the novel ones. Thus, my research provides additional evidence for the support of the view that there is no one-to-one relationship between these two responses and the aforementioned processes, and contributes to the extension of evidence to different kinds of materials, namely faces. It also suggest that the greater opportunity for conscious control over recognition provided by the longer deadline benefits knowing, as well as remembering, only when the faces to be retrieved are famous; a finding that is intriguing and interesting as it provides further evidence for differences in the encoding and retrieval nature of novel and famous faces. Another type of encoding manipulation was used in two further experiments, where faces of same and different racial origin - in relation to the participants - where studied either one or three times under the assumption that remembering as well as knowing would increase after three study presentations. It was found that participants remembered overall more faces of their own racial origin, and remembered more faces of both origins after they were exposed to the faces three times. Knowing was not affected. The selective effects on remembering lend further support to the claim of dissociation between remembering and knowing and pose a question for the familiarity hypothesis aiming to explain the phenomenon of own-race bias.

In conclusion, both series of experiments have produced results that replicate and extend previous findings in the area of recognition memory and can be informative for other specialists in the areas of education, clinical psychology, expert eyewitness testimony and equal opportunities policies.

References

Gardiner, J. M., Ramponi, C., & Richardson-Klavehn, A. (1999). Response deadline and subjective awareness in recognition memory. *Consciousness and Cognition, 8*, 484-496.

Jacoby, L. L., & P. O. Dolan, T. D. J. abd. (1998). Two effects of repetition: Support for a dual-process model of know judgements and exclusion errors. *Psychonomic Bulletin and Review, 5*, 705-709.

Wright, D. B., Boyd, C. E., & Tredoux, C. G. (2001). A field study of own-race bias in South Africa and England. *Psychology, Public Policy, and Law, 7(1),* 119-133.

Autohyponymy and Mental Representations

Anu Koskela

amk21@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

My research area and the topic of this paper concerns the phenomenon of autohyponymy - where one word can be used with a more general meaning and with a more specific one that is in a relationship of hyponymy or class inclusion to the general meaning. In this paper I will discuss this phenomenon briefly and outline a cognitive linguistic account of lexical representation following work by (Langacker, 1987) and (Tuggy, 1993).

2 Concepts and meanings in cognitive linguistics

One of the key ideas in cognitive linguistics is that, as experimental evidence shows, human conceptual categories are not definable by necessary and sufficient conditions and they do not have fixed boundaries, but rather exhibit a prototype structure in terms of more representative and more marginal members and variable boundaries (Rosch, 1978). It is also assumed that the mental lexicon does not provide fixed literal meanings for word forms. (Langacker, 1987) has argued that linguistic items act as "access nodes" into an open-ended network of encyclopaedic knowledge; thus different aspects of our knowledge may be activated on different occasions of use as motivated and constrained by the linguistic conventions of the speech community and the context cf. (Croft & Cruse, to appear); (Fauconnier, 1997).

3 Autohyponymy

As mentioned, autohyponymy is a phenomenon where one word has a more generic and a more specific use. A classic example of this is *dog*, which can mean either a canine in general or more specifically a male member of the canine species. The same phenomenon has also been called specific-generic polysemy in anthropological literature on folk categorisation of natural kinds (Berlin, Breedlove, & Raven, 1973). My particular research interest is autohyponymy of concrete object names - cases such as *coat* which can mean an overgarment in general (including jackets) or a prototypical long, warm overgarment that contrasts with jackets. As mentioned, research into human categorisation has shown that category boundaries may be construed differently by different speakers and on different occasions of use. Therefore, for any category with a prototype structure, we can potentially always find instances where in some contexts the category boundaries are construed to be more narrow or more extended. However, given observations such as that *coat* has a more prototypical and a more generic use, one of the questions we can ask is whether this is simply because the category boundaries are contextually construed differently or whether the different extensions could be said to constitute separate mental representations.

4 Senses and mental representations

From the linguistic point of view, we can approach this question by asking whether the generic and specific uses of a word constitute different senses of the word. Traditionally a distinction has been made between three modes of representation of meaning in the mental lexicon: homonymy, where a word has two or more distinct meanings that are not related (e.g. bank), polysemy, where the word has distinct and unrelated meanings which are semantically related (e.g. *run* a marathon and *run* a business) and monosemy, where a word has a single meaning and any different interpretations are derived pragmatically. So, the question is whether a word such as *coat* is monosemous or polysemous. In many traditional approaches to word meaning and sense boundaries it has been assumed that the distinction between words that have a single sense and those that have a number of distinct (although possibly related) senses is straightforward and clear-cut. However, work in the cogni-

tive linguistic framework suggests that homonymy, ploysemy and monosemy in fact form a continuum; (Tuggy, 1993); (Croft & Cruse, to appear). To illustrate why this is argued to be the case, consider one of the traditional "tests" for determining whether the different meanings of a word are distinct or unitary. The so-called linguistic test relies on the possibility of using the two meanings of the word in a single construction such as 1) below:

(1a) Peter is hot and so is Simon.

(1b) I saw my cousin yesterday and Simon saw his.

In (1a), if we assume that Peter is hot in terms of having the experience of heat while Simon is sexually attractive, the sentence sounds odd, like a pun. This is taken to indicate that the two meanings of *hot* are distinct senses. In (1b), on the other hand, even if my cousin is female and Simon's is male, the sentence is fine, which means that *cousin* is monosemous in as much that there aren't distinct senses for male and female cousins. However, this test is problematic in that in some cases it is difficult to determine whether the sentence sounds odd or not, different speakers may differ in their intuitions and sometimes the context may influence the interpretations. For example, Tuggy (1993) has argued that the acceptability of sentence such as (2) is often a matter of degree:

(2) Mary is painting and so is Tim.

(3) Paula wore a coat and so did Mike.

(2) seems peculiar if Mary is painting a still-life but Tim is painting the bedroom walls, but if Mary is painting a mural, the sentence sounds better. Furthermore, if we try to contrast autohyponymous senses as in (3), speakers' intuitions on whether the sentence is acceptable or not if Paula wore a long overcoat and Mike wore a summer jacket depends on the context and what level of detail is judged relevant in our construal of the category COAT.

5 A cognitive grammar account of mental representations

One model that can account for such gradable and context-sensitive intuitions is that proposed by (Langacker, 1987) and (Tuggy, 1993). In this model it is argued that lexical categories are characterised as hierarchical networks of nodes with higher level schemas as abstractions of the commonalities of more specific, lower level schemas. The more specific schemas represent distinctness of meaning and the more abstract schemas unity of meaning, but importantly both may coexist as part of a speaker's linguistic knowledge. The nodes in the network become more cognitively salient or entrenched through repeated activation in linguistic usage events. Different nodes vary in their level of entrenchment, which correlates with their likelihood of activation. Thus the continuum between monosemy, polysemy and homonymy is represented in this model so that in the case of monosemy, a schema representing the commonalities between different usages is well entrenched and the specific usages are not very well entrenched. In homonymy, the specific instances are well entrenched and the general schema is very abstract and not well entrenched at all. Polysemy is characterised by varying degrees of entrenchment of the more abstract or specific schemas - thus cases of polysemy fall on a continuum between monosemy and homonymy. The differences in people's intuitions regarding the unity and distinctness and relatedness of senses then relies on the activation of either more general or more specific schemas in particular contexts. Speakers may vary in the kinds of networks they construct, depending on what specific instances they encounter in their experience and what generalisations they extract. Importantly, however, there will be a level of overlap in the mental representations of different speakers provided by conventionalisation of linguistic usage as something shared by the linguistic community.

6 Conclusion

In relation to the mental representation of autohyponymy, we can state that it is possible that some particular construals of category extension may be entrenched as nodes in a network for a particular speaker - the level of entrenchment may vary partly explaining the uncertainty of judgements regarding the distinctness of generic and more specific usages. To make more specific claims about the status of the senses we would have to look at their degree of conventionalisation using methodology such as corpus studies. Further issues for research include the question of what factors might influence the degree of entrenchment and conventionalisation of the autohyponymous senses.

References

Berlin, B., Breedlove, D. E., & Raven, P. H. (1973). General principles of classification and nomenclature in folk biology. *American Anthropologis, 75*, 214-242.

Croft, W., & Cruse, D. A. (to appear). In *Cognitive linguistic.* Cambridge University Press, UK.

Fauconnier, G. (1997). In *Mappings in thought and language.* Cambridge University Press, UK.

Langacker, R. W. (1987). Foundations of cognitive grammar. In (Vol. 1). Stanford: Stanford University Press.

Rosch, E. H. (1978). Principles of categorisation. *Cognitive Psychology, 4*, 328-350.

Tuggy, D. (1993). Ambiguity, polysemy, and vagueness. *Cognitive Linguistics, 4-3*, 273-290.

Why I Don't Like Whiskey

Kristy Lascelles

kristy@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

What was the first alcoholic drink to make you vomit? Chances are you now avoid that flavour like the plague. It's also likely that if you were asked to indicate how much you liked or disliked that particular taste, you would reply with a profound dislike. But did you dislike that flavour before your head started spinning and your stomach did the back flip? Probably not! The process via which you came to dislike the flavour of the alcoholic drink is called evaluative conditioning (EC) and arose because you came to associate the flavour with something you didn't like - vomiting. EC is a process that can also explain how we come to like things (as well as dislike things) and is often used in advertising where a product is paired with something liked (naked women, the perfect lifestyle) so that an association is formed between the product and the liked stimulus, and consequently the product also becomes liked. As well as explaining our acquired general likes and dislikes and playing a large role in advertising, EC might also be of use in explaining phobias, fetishes and other psychopathological disorders, such as eating disorders.

A large area of research has been dedicated to the EC phenomenon and has implied that under laboratory conditions likes and dislikes can be elicited in a large range of objects, using a large range of stimuli (e.g. pictures, faces, flavours, odours and haptic stimuli) and that the effect is reliable and robust. This research has also provided a number of observations, unusual for human learning:

1. These effects can be elicited without conscious awareness of the pairing of the neutral stimulus with the liked or disliked stimulus.

2. These effects cannot be removed through the usual means of extinction (where a neutral stimulus that has become liked or disliked through the pairing process is then presented repeatedly in the absence of the liked or disliked stimulus, and this would normally lead to the stimulus regaining its original neutral status).

3. The pairings do not have to be based on a statistical contingency, but must have some form of temporal or spatial relationship. In short, this indicates that EC cannot be explained through conventional human learning mechanisms (such as Pavlovian conditioning).

However, a number of methodological flaws (inadequate controls, no baseline measures of liking prior to the pairing process in order that a change in liking might be ascertained etc.) in this research area has meant that the reported effects may have been the result of some experimental artefact, and not due to the pairing of a neutral stimulus with a liked or disliked stimulus, let alone some special form of human learning. The current research was aimed at ascertaining whether the pairing of a neutral stimulus with a liked or disliked stimulus could bring about the liking/disliking of that neutral stimulus once the methodological problems encountered in previous research were corrected for.

This research verified that EC effects could be obtained as a result of the pairing process, however the effect was not nearly as reliable or robust as was previously believed. In fact, the effect was found to be extremely sensitive to the conditions under which it would occur and effects similar to previous research could often not be replicated when similar conditions to the early research were employed. Once the experimental conditions were varied, EC could be obtained. The influential factors were found to include how many stimuli the participant was required to learn about, how many times the stimuli were paired and the amount of time between presenting the neutral stimulus and presenting the liked or disliked stimulus. However, everyday cases of EC might occur more easily due to the increased salience of everyday liked and dis-

liked stimuli, which could not be imitated in the laboratory for ethical reasons.

As well as establishing some of the experimental conditions under which EC could be established, another important finding was observed. A phenomenon known as selective conditioning, where some stimuli are more easily associated than other stimuli, was also observed. In one particular experiment pictures of food were used as neutral stimuli and these were paired with pictures of obese (disliked), normal (neutral) or extremely thin (disliked) women. Only the pictures of food paired with the pictures of obese women became disliked, despite the fact that the pictures of extremely thin women were actually more disliked than the pictures of the obese women. In addition, some learning theorists have argued that selective conditioning effects are due to innate biological preparedness. However, a follow up study revealed that the selective conditioning effects observed here were due to an expectancy bias where foods were expected to be paired with obese or normal women, more than they would be expected to be paired with extremely thin women. This indicates that selective associations, in humans at least, might not be accounted for through innate predispositions, but ontogenetic factors might also play an important role.

This research has indicated that EC is not quite the resilient effect that it was once thought to be and that selective conditioning can occur. The road ahead for EC involves firmly establishing the experimental conditions under which it will occur and explaining why it might be such a fragile learning mechanism. Other selective conditioning effects should be explored, which might uncover important preliminary links to understanding some psychopathological disorders. In addition, EC should be compared with other forms of human learning to determine whether in fact it does show the same patterns of learning, or whether it is indeed a different form of human learning entirely.

Age effects on the acquisition of fear

Joanne Lawson

jl30@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

In the 20th century, associative learning has been the dominant explanation for how we acquire fears and phobias - we start with a few naturally aversive stimuli such as loud noises or pain, which give us our first experiences of fear and the fear response generalises or extends to other stimuli, such as dogs or spiders, through associative learning. Following (Rachman, 1977), most researchers consider three pathways to fear: classical Pavlovian conditioning (something frightening occurs in the presence of a previously neutral object and the fearsomeness transfers to the previously neutral object); vicarious learning (learning by seeing someone else being scared of something); and information transfer (learning to fear by being told that something is dangerous or should be feared). Although much experimental evidence has been gathered and is being gathered to support the existence and operation of these pathways to fear, there are problems. Many of the objections to associative learning approaches centre around the fact that they over-predict the occurrence of fear: if someone's fear of swimming can be traced to the time they were very scared by watching Jaws, then we have to specify why not everyone watching Jaws becomes scared of swimming. This is a particular problem with retrospective surveys that just ask fearful people about likely causes and events at the onset of their phobia. Such studies rarely include a control group who had similar experiences but did not become phobic. It seems that there must be mediating factors that determine whether someone comes out of an associative learning event with a long-lasting fear. One factor that doesn't seem to have received much experimental attention is age, although it has been noted that different phobias have different characteristic ages of onset. The most thorough study of this was conducted by (Oest, 1987), who found that animal phobia had the earliest onset age (7 years), followed by blood phobia (9 years), dental phobia (12 years), social phobia (16 years), claustrophobia (20 years), and finally agoraphobia (28 years). It is possible that there are critical developmental periods for the acquisition of different specific fears - associative learning events occurring during these critical periods have a more profound effect. To investigate this, a highly controlled experiment designed to reflect associative learning of animal fears was carried out with two age groups of children, one at, and one older than, the hypothesised critical developmental period for the acquisition of animal fears. The question is whether children are more sensitive to information about fearsome animals at the age at which most animal phobias begin.

2 Participants

Fifty-nine children between the ages of 6 and 8 ($\mu = 92.9$ months, $\sigma = 6.6$) and 58 children aged 12 and 13 ($\mu = 152.7$ months, $\sigma = 3.9$).

3 Procedure

The participants were shown pictures of the three novel animals: the cuscus, the quoll and the quokka, and given a specially constructed Fear Beliefs Questionnaire (FBQ). This is a set of 21 questions - 7 different questions, repeated once for each animal - about fear-relevant beliefs, such as "do you think the animal would hurt you?" and "would you go out of your way to avoid the animal?". The participants answered these questions at this stage, before they knew anything about the animals. Next, the participants were told negative information about one animal and positive information about another animal by the experimenter. (Counterbalancing ensured that all animals were paired equally often with negative, positive and no information). Participants then completed the FBQ again, and two other measures of fear and attitudes towards the animals were taken. Firstly, to see if there are effects on behaviour, participants were asked to stroke the animals (actually soft toys in wooden touch-boxes) and the time taken to approach and stroke each animal was recorded. Secondly, a computerised test - the Implicit Association Test (IAT) - was completed by all participants.

The IAT is used to measure the association that people have between two particular concepts. If one constructs a classifying task with compound categories (e.g. press 'E' if the word is either a flower or a pleasant word; press 'I' if the word is either a weapon or an unpleasant word), participants will be quicker to perform the classifying task if the concepts that are grouped together are already associated in their mind. So, in this example, one might hypothesise stronger associations between {flowers + pleasant things} and {weapons + unpleasant things} than between {weapons + pleasant things} and {flowers + unpleasant things}. Consequently, one would expect quicker reaction times in a trials where the categories were {flowers + pleasant things} and {weapons + unpleasant things} than in trials where the categories were {weapons + pleasant things} and {flowers + unpleasant things}. Thus, by comparing reaction times in a classification task between trials with different concept pairings, one can get a measure of the associations that people have between those concepts. In this instance, it was used to assess whether participants have formed associations between particular animals and nice and nasty things. This measure gets around the problem that children could just have been telling the experimenter what they thought he or she wanted to hear.

Participants completed the FBQ and IAT again after a week, a month and three months.

4 Results

4.1 Fear Beliefs Questionnaire

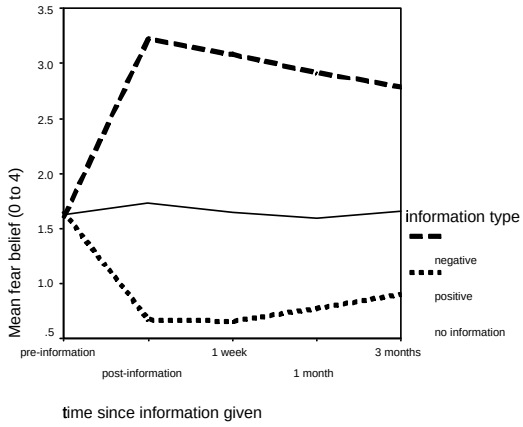

Figure 1: Mean fear belief by information type over time

Participants' scores for each animal on the Fear Beliefs Questionnaire showed a significant interaction between the effects of time and information type ($F(5.17, 485.50) = 26.17$, $p < 0.0005$). That is, the different types of information (positive, negative or none) about the animals had different effects on participants' fear beliefs. The magnitude of this effect was not significantly affected by age group (non-significant 3-way interaction: $F(5.17, 485.50)$, $p = 0.27$).

4.2 Behavioural measure - touch boxes

Time taken to approach and stroke each animal was standardised (converted to z-scores) because participants always took longer to approach whichever animal was first. Participants in both age groups took longer to approach the animal about which they had been given negative information ($F(1.85, 162.68) = 8.87$, $p < 0.0005$), and there was no interaction with age group ($F(1.85, 162.68) = 1.17$, $p = 0.31$).

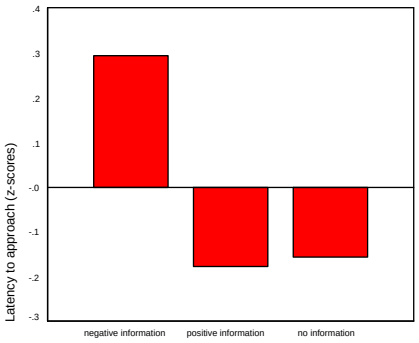

Figure 2: Mean latency to approach animal by information type

4.3 Implicit Associations Test

If participants have formed associations between a particular animal and, say, 'nice' things, then they will be quicker in a categorisation task that requires them to put pictures of that animal in the same category as 'nice' words than in the same category as 'nasty' words. The trial types are called Compatible and Incompatible, and the difference in reaction times between these trial types is taken as a measure of strength of association. There was a significant effect of trial type ($F(1, 98) = 32.53$, $p < 0.0005$), with Incompatible trials having longer reaction times than Compatible trials, as expected. There was a significant interaction of trial type with age group ($F(1, 98) = 4.40$, $p = 0.039$). Inspection of the data shows that there was a smaller difference between Compatible and Incompatible trials in the older age group.

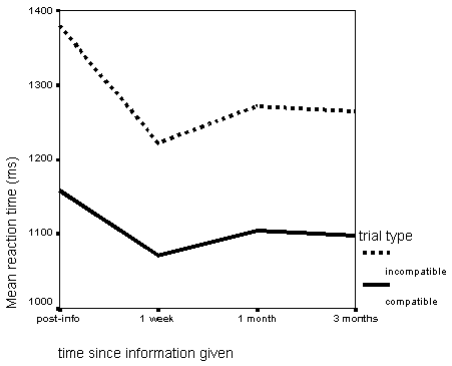

Figure 3: Mean reaction time by trial type over time

5 Discussion

The hypothesis that valenced information about animals would have a greater impact on children at the developmental period at which animal fears are most common and at which animal phobias most commonly begin was upheld for one of three measures. No difference was found between older and younger children on a direct questionnaire measure of fear-related beliefs and attitudes (the FBQ), nor on a direct behavioural measure of how long it took to approach and stroke the animals. Where a difference did appear was in the IAT - an indirect, implicit task, the results of which are not a direct result of the conscious behaviour of the participant. It may be that the difference in sensitivity to valenced information only shows up in tasks that are immune to participants' conscious strategising. This explanation could be followed up with a range of other implicit tasks, such as affective priming.

References

Oest, L. G. (1987). Age of onset in different phobias. *Journal of Abnormal Psychology*, *96*, 223-239.

Rachman, S. J. (1977). The conditioning theory of fear acquisition: A critical examination. *Behaviour Research and Therapy*, *15*, 375-387.

Automatic Grammar Induction by Distributional Clustering

Chi-Ho Li

chl20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Most natural language processing systems need syntactic analysis to some extent, and therefore require a parser and a set of grammar rules. While there has been a lot of research on various parsing algorithms, the problem of how to obtain a grammar has received less attention. This short paper introduces a framework on the automatic induction of grammar rules.

2 Supervised vs Unsupervised Learning

The various approaches to linguistic knowledge acquisition can be separated into two camps: the supervised and the unsupervised. Supervised approaches assume the existence of some golden standard to guide the learning process, while unsupervised ones do not. Such golden standard could be, for example, a large collection of text annotated with correct grammatical structure. Although supervised approaches give much better results at the current state of the art, they are undesirable for two reasons. First, building the golden standard is an extremely laborious and time-consuming task. For this reason there have not been sufficient resources like annotated text for many languages. Second, supervised approaches do not explain how human beings learn language. Obviously, no children learn their native languages by being told that such and such sequences of words/part-of-speeches are called noun phrases, etc. Therefore unsupervised approaches are more preferable in both theoretical and practical terms.

3 Parameter Learning

A widely established unsupervised grammar learning technique is the *Inside-Outside algorithm*, which is in fact an application of the *Expectation-Maximization algorithm* in statistical inference. Put it in simple terms, what the Inside-Outside algorithm does is to iteratively re-estimate one of the following figures on the basis of another:

- the probabilities of grammar rules
- how many times the rules are used in parsing a collection of text

It is proved that such iterative re-estimation will converge to a stable set of figures, though it may be a local maximum only.

There is a big problem with the Inside-Outside algorithm: what are the things whose probabilities are to be estimated? Before we estimate the probabilities of grammar rules, we should first decide which rules exist in the grammar and which rules do not. In the jargon of statistical language learning, what the Inside-Outside algorithm does is merely to learn grammar *parameters* (i.e. rule probabilities). We are still left with the problem of learning grammar *structure* (the rules themselves).

4 Structure Learning

The choice of a set of sensible grammar rules can be conceived as a problem of classification:

1. the classification of all sequences of categories into legitimate ones and illegitimate ones. (E.g. the sequence ADJECTIVE NOUN is a legitimate phrase while PREPOSITION DETERMINER is not.)

2. the classification of all legitimate category sequences into different types of phrase. (E.g. the sequence DETERMINER NOUN is a noun phrase while VERB NOUN is a verb phrase.)

In machine learning, there is an area called *clustering*, which investigates automatic classification. A clustering algorithm treats data points as *vectors*, each dimension of which corresponds to a particular *feature*. It then groups the vectors into several classes, or *clusters*, in accordance with some metric for measuring the distance between vectors. In order to apply some clustering algorithm to grammar learning, therefore, we have to represent the

candidate phrases (sequences of categories) as vectors of some features, and find a way to measure the distance between them.

In the tradition of structural linguistics, a phrase is characterized by its *distribution*, or *syntactic context*. For example, if the phrase DETERMINER NOUN is situated between a verb and a preposition, then the pair [VERB - PREPOSITION] is a context of the phrase. If there are N types of categories, then there are N^2 types of context. Each context type can be treated as a feature, and its value can be the frequency or probability that a phrase is situated in that context. Therefore a phrase can be represented as an N^2-dimensional vector of syntactic context.

As to the measures of distance, many consumers of clustering algorithms adopt the Euclidean distance. Yet the results of our pilot experiments show that this metric performs poorly. However, if the feature values of vector are taken as probabilities, then each vector is itself a probabilistic distribution over the context types. There are several metrics proposed by information theorists to measure the distance between probabilistic distributions. The pilot experiments reveal that most of these metrics give very good classifications of part-of-speeches.

5 Conclusion and Future Work

This paper introduces a feasible approach to learning grammar structure, which can be combined with the Inside-Outside algorithm to form a comprehensive strategy in grammar acquisition. Our pilot experiments show that the approach is promising. In future a complete grammar learning system based on distributional clustering will be developed.

Alienation and the Philosophy of Cognitive Science

Paul Loader

pl27@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

What bearing, if any, might the concept of 'alienation' have on issues in the Philosophy of Cognitive Science? To answer this we first need to decide what is meant by 'alienation'.

The term 'alienation' is used in a variety of ways in everyday discourse and in disciplines such as Sociology, Law and Psychology but the concept of 'alienation' I wish to examine, whilst not unrelated to many of these, is from a perhaps less familiar source - the writings of the young Karl Marx.

2 Marx's concept of 'alienation'

Marx's concept of alienation, itself a development of earlier work by Hegel and Feuerbach, resists easy summary but often a common starting point is the account found in the 1844 Manuscripts (McLellan, 1987). Here it is asserted that in Capitalist societies 'alienated labour' manifests itself in four ways:

1. alienation from the product of one's labour,
2. alienation from ones own productive activity,
3. alienation from man's "species being" (man is alienated "from his own body, nature exterior to him and his intellectual being, his human essence"), and
4. alienation from other human beings.

Ignoring for now the problem that the third of these 'dimensions' of alienation does not seem to be 'of a type' with the other three, we might get at least some idea of what Marx is saying by noting that each of these dimensions can be viewed as having a double aspect insofar as it has a practical basis (something is actually taken from the individual) and a 'psychological' consequent (that which is taken away comes to be viewed as 'separate' or 'strange'). To illustrate: in the first case, the product of the individual's labour is practically removed from him (it belongs to his employer) and as a result it confronts him "as an alien object that has power over him".

3 So what?

What has any of this got to do with Cognitive Science - Cognitive Scientists have no particular interest in the detrimental effects of factory work after all? Well the idea is that, despite having its roots in the arena of production, 'alienation' is not just confined to the workplace but can manifest itself more generally in our everyday and theoretical conceptions of the world. (Here there is some overlap with the concept of 'ideology'). If this theory is true then it seems to me at least plausible that some such 'alienated' conceptions can be found in the Philosophy of Cognitive Science. Three possible candidates are given below.

- *Strong AI* - The notion that a technological artefact can literally exhibit its own intelligence might be construed as an example of "alienation from the product of labour". On the one hand the hope for genuine strong AI speaks of a misplaced projection of one's own powers onto inanimate matter in the same way that artefacts (e.g. masks) were imbued with magical powers in earlier societies. On the other hand we might say that it speaks of a lack of recognition that the very real intelligence embodied in technological artefacts is human at source because the product of human labour power.
- *The 'Other Minds' Scenario* - This is the scenario which posits the need for inferential or other cognitive mediation before one individual can recognise that other individuals are owners of minds. It is to be found in traditional Philosophy of Mind (Ayer, 1956), less traditional Philosophy of Mind (Dennett, 1981) and is also to be found in certain interpretations of 'Theory of Mind' in Psychology

e.g. (Baron-Cohen, Tager-Flusberg, & Cohen, 2000). Such a scenario could be construed as a variant of the "alienation from other human beings" theme, and thus critics of this scenario such as (Wittgenstein, 1953) and those with '2nd Person' phenomenological perspectives (Thompson, 2001) can be viewed as part of a counter - 'de-alienating' - tendency.

- *Cyborgs* - According to Marx, while "nature is the inorganic body of a man" this state of affairs becomes inverted in capitalist production where the labourer becomes a "mere living appendage" of the machine. This version of 'alienation from body' (part of alienation from species being) implies the individual's consequent reduction to a technological artefact. In Philosophy of Cognitive Science variations on this theme are to be found in speculation about the physical integration of human consciousness with technological artefacts (Clark, 2003) and also more generally in the conception of mental states and processes as silicon realizable.

How convincing are these examples? Is it just a question of pattern matching, in an ad hoc fashion, various disparate elements of theory from the Philosophy of Cognitive Science with ill defined constructs from an outmoded ideology (Marxism), or is there something explanatorily useful to be gained by suggesting that 'alienation' is rife in the Philosophy of Cognitive Science? (Any expressions of opinion gratefully received.)

References

Ayer, A. J. (1956). *The problem of knowledge.* Pelican Books. (See especially Chapter 5).

Baron-Cohen, S., Tager-Flusberg, H., & Cohen, D. J. (2000). *Understanding other minds.* Oxford University Press, UK.

Clark, A. (2003). *Natural born cyborgs.* Oxford University Press.

Dennett, D. (1981). *True believers: The intentional strategy and why it works.* Reprinted in Mind Design II (Ed. Haugeland, 1997) MIT Press.

McLellan, D. (1987). *Karl marx selected writings.* Oxford University Press. (See especially Section 10).

Thompson, E. (2001). *Between ourselves.* Imprint Academic.

Wittgenstein, L. (1953). *Philosophical investigations.* Basil Blackwell.

Visually guided agents with evolved morphologies

Ian Macinnes

ianma@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Evolutionary robotics is a branch of Artificial Life in which the morphology and control structures of robots are altered over time to adapt to certain environments or perform certain tasks. Broadly, these evolved robots are used in the exploration of two main fields - the study of biological life and evolution, and the study of robot engineering. My research involves the use of robotics to study characteristics of living organisms. By evolving robots to duplicate behaviours and developmental processes found in nature, we can test hypothesizes about their origin and function. Most work in evolutionary robotics is performed on robots simulated in a computer, and concentrates on the evolution of their controllers. However, I concentrate on the adaption and change of their physical shapes and morphologies.

2 Evolution of Robot Brains and Bodies

In 1992, Randall D. Deer(Beer, 1996) highlighted the use of dynamical recurrent neural networks for the control of mobile robots. A neuron network is a control structure that can make choices based on the type of noisy, cloudy information that we experience as living things in the world on a day-to-day basis rather than the symbolic kind that computers normally require. Dynamcial means the controller uses time as an integral part of the way it works, and recurrent means that the outputs of the controller can feed back into its input. Beer used this controller to create a simulated walking robot in the shape of a cockroach, among other things.

Today this type of controller is heavily used within the evolutionary robotics field as it is so suited to evolutionary development by means of natural selection. A population of simulated robots is created within a computer, all with randomly configured brains. They are set to perform a task, or show a behaviour, and are graded on how well they do. The ones that do the best, even though this may not be very good at all, are allowed to reproduce and replace the robots that do not do so well. These better performing robots are said to be fitter. Each time a new child robot is created and put into the population, a small random change occurs to its controller, so it may behave slightly differently from its parent. In a lot of cases, the new robot will be graded worse than its parent, but in a few cases, it may perform better. These better ones in their turn have a greater chance of reproducing, and hence passing on their patterns of behaviour. It's also possible to evolve the physical shape of a robot.

In the early nineties, Karl Sims(Sims, 1994a, 1994b) produces simulated robots that consisted of differently shaped blocks. He evolved them to fight for possession of a box in a gladiator type contest. The robots evolve many strategies to beat their competitor and capture the box - for example, they developed strong limbs to push their competitor away and grab the box. Inspired by this, I produced similar robots to produce different locomotive behaviours. First, a population of single blocks is created. They are initially unable to move at all, giving them all a fitness of zero. Parents are randomly picked and reproduced. Whilst reproducing, mutations appear in the instructions used to specify the child robot's shape and the configuration of their brains. This makes them look slightly differently. Eventually, they develop strange shuffling motions, or limbs that allow them to walk. Some evolved single limbs to pull themselves along. I repeated the experiments with swimming robots. These evolved paddles and steadying arms.

Next I added sensors. The angle of all the blocks in relation to each other were fed into the brain, so the agent knows the relative position of all it's parts. Light detecting sensors that stimulate neurons when certain colours enter a field of view were added, giving the agents a simple form of vision. A typical task I evolved for a population to do was to move, searching for blocks of a particular colour,

approach them, and move them as far as they could by pushing or rolling it. I found that the physical morphology of the agents altered completely to enable it to perform the task better. The front of the agent changed so that when the agent and a block were in contact, the physical shape of the agent was instrument in keeping contact, rather than the vision sensors. This is an example of embodied cognition.

3 Future World

I plan to use these robots to demonstrate and research different aspects of biological organisms, such as sensor development and social behaviour. I also hope to develop a method by which simulated robots whose morphologies and controllers have been evolved with in computers can be physically built and have the same behaviour as their simulated counterparts.

References

Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behaviour. *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour*, 421-429.

Sims, K. (1994a). Evolving virtual creatures. *Siggraph Proceedings*, 15-22.

Sims, K. (1994b). volving 3d morphology and behaviour by competition. *Artificial Life, IV*, 28-39.

To manglen of the syntax up: what a nineteenth century Sussex diary might reveal about working class literacy

Barbara Maidment

barbaram@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

My research is concerned with the diffusion and practice of literacy amongst the rural working class in the nineteenth century Sussex Weald. The quantitative part of my work attempts to capture a profile of 19th century literacy levels in the Weald as outlined in last year's White House Paper (?). The qualitative part of the study focuses on studying Wealden working class literacy practices in order to help answer questions concerning people's beliefs and expectations about the function and value of literacy, both for individuals themselves and for their children.

2 Exploring a community's literacy

One way to explore a community's literacy practices is to look at what people write. A number of 19th century working class diaries, journals and letters have been published, including some from Sussex, but my aim has been to try to locate original unedited documents. In order to do so, appeals were made to descendants of villagers who make up my research group, which have produced some unpublished and hitherto unknown writings. Additionally there are copies of letters and some original diaries and journals deposited in the East Sussex Record Office. This material reveals not only the interests and concerns of the various writers, but also provides first-hand, although sometimes circumstantial, evidence about education, literacy levels, dialect, pronunciation and, importantly, people's attitudes to literacy.

This paper will focus on one newly-surfaced item; a 'Dairy Book' (sic) (?) kept between July and November 1829 by George, the 11-year old son of James White, blacksmith and farmer at Chiddingly, and his wife Mary.

George attended the village school run by Richard Lower, father of the Sussex antiquarian Mark Anthony Lower. At the time he wrote his diary George had left school and was working as a labourer on his father's farm. George died aged 17 in 1835.

The diary provides abundant information on a number of topics of interest to social historians: work, agricultural management, family life and relationships, religious observance, household routines and family visits. My interest is in the extent to which Sussex dialect and/or George's short education account for his syntax and his spelling, and what can be deduced from the diary about family literacy practices.

Stylistically the diary is repetitive. Entries are short; most contain just a few lines. These invariably start with a description of the weather followed by a report of the day's work, activities, visitors, journeys and outings. Each entry is concluded with as many numerals from one to ten as fit on the line. Nowhere in the diary does George reflect on experiences or provide any evaluative comment on activities or events. Arguably, self-analysis would be unusual in children's writing today, and so much more so in the writing of children in a far less self-conscious age.

Otherwise, the diary is notable for its equal degrees of relative sophistication and rustic or childish simplicity. George may have acquired his handwriting style during his time at school, which probably lasted four or five years, longer than was common amongst his contemporaries. His letters and numbers are formed in an italic script that was the educational hallmark, although random capitalisation of proper names, nouns and adjectives characterises his writing. The diary contains almost no punctuation.

George's writing shows a number of non-standard spellings which may indicate a localised pronunciation, e.g. 'Boreship' for 'Boship', 'Herringly' for 'Hellingly' (17 times),'Furrell' for 'Firle' and 'passel' for 'parcel'. Occasionally his non-standard spellings appear to conform with 19th century variable spelling conventions, e.g. 'stopt' for 'stopped' and 'bornd' for 'born'. Otherwise, some non-standard spellings are regular throughout the diary, e.g. 'poney' and 'Fier' which

never appear as ‘pony’ or ‘fire’. Other spellings vary, e.g. ‘Chapel’ and ‘Chaple’, ‘beens’ and ‘beans’. Among local dialect words are *Wimen* (winnowing); *Thecking* (thatching a haystack); and *Egeten* (hoeing) (?) The grammatical style of the entries is single clause, past tense declaratives (although ‘came’ is always written ‘come’). Sophistication of construction is achieved by joining two such clauses with the conjunction ‘and’, inscribed as an italic +. George’s habitual preposition + gerundival construction, as in *I went to gathering currants* is sometimes, but not consistently, written with -en rather than -ing ending, as in *I went to picken up potatoes & to Gatheren apples.* But this cannot be accounted for as merely childish idiolect since the identical construction also appears in the mature writing of others (?) and (?). Elsewhere George writes *I went to stripen hop poles & Stephen Funnell went to stacken of them*, the latter phrase being, arguably, a preposition + pronoun construction; similar versions have been identified in other southern dialects, although with the -ing rather than -en inflection, viz:“*penning of ‘em*” (Wiltshire and Dorset) and “*sharpening of en*” (Dorset) (?). A discontinuous phrasal verb construction appears in several places, for example *I went to picken up pottatoes & James & Stephen went to Digen of them up*. This construction has also been identified elsewhere; the Survey of English Dialects reports “*bagging of her up*” from Berkshire and from Cornwall “*galing of them up*” (?).

The non-standard syntax of George’s diary appears therefore to owe more to the local dialect than to lack of or a poor standard of education. His spelling is mostly inconsistent and demonstrates both local pronunciation and imperfectly learned systems.

But it is in terms of a literacy practice the diary is intriguing. No other 19th century working class child’s diary is known, which begs questions about its function. As a blacksmith and farmer, George’s father was probably as well off as many of his contemporaries, but at a time of agricultural depression and poor harvests he couldn’t afford to keep his eldest son at school beyond the age of 10, needing his labour to run the farm. Clearly the family valued schooling; on 14th August George records ‘*I went to scool in Benjamin Stead*’. The place had been paid for, but when George’s younger brother could not attend the expense was not to be wasted. Perhaps his parents encouraged George to keep the diary in order to preserve his literacy skills, which might explain the numerals following each entry, since they serve no obvious purpose except that of the practice of writing. There is no way of knowing whether the diary George wrote was a polished version, previously drafted and corrected, or whether he wrote it himself without being overseen. George wasn’t the only member of his family to keep a diary; his sister Naomi, born two months before George died and destined to become the village schoolmistress, kept hers in the 1850s. In this family, then, there appears to have been a tradition of respect for literacy and learning, which in the 19th century characterised upper working class artisans like James White and his family.

3 Acknowledgement

Coates, R (version of 2000). *The Dialect of Sussex*. (in preparation).

References

Hall, H. (1957). *A dictionary of the sussex dialect and collection of provincialisms by W. D. Parish.* Bexhill Gardner.

Maidment, B. (2002). Nineteenth century working class literacy in the weald. In (Vol. 15). University of Sussex, UK.

Russell, G. (1849). *Diary.* Unpublished, privately owned.

Standing, L. (1999). My old family diary. *Sussex Family Historian, 12-6.*

Upton, C., Parry, D., & Widdowson, J. (1994). *Survey of english dialects.* London Routledge.

White, G. (1829). *A dairy book (sic).* facsimiled and transcribed by Brian Hubble, unpublished, privately owned.

Understanding Deliberate Minds

Marek McGann

mm56@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Proponents of the embodied approach to Cognitive Science conceive of knowledge and understanding in terms of our mastery of skills for interacting with the environment. Our comprehension of the world we live in is expressed through our ability to act adaptively within that world. This, one of the fastest growing conceptions of cognition, implicitly depends upon the idea of purpose.

Skills, or mastery, seem only ever to be an ability to manipulate something reliably to one's own ends. Adaptiveness. relatedly, seems to be the altering of behaviour (often deliberately) so that one's own interests are satisfied. Interests, goals, intentions and purposes are a constellation of concepts that have been given precious little consideration within Psychology or Cognitive Science.

This lacuna in the state of the art has been touched on by numerous authors over the past few years. It has been thrown into relief by authors from the embodiment and dynamical systems camps, most notably by the philosopher Andy Clark, and psychologist Merlin Donald. In a series of articles, Clark (2002,2001), has identified a need for a means of describing deliberate, intentional action. Drawing on the neuro-psychological work of Milner (1995), Clark notes an apparent dichotomy in the manner in which the brain processes our interactions with the world. One stream of visual processing, the dorsal stream, may be adequately explained by a rich description of the detailed workings of the human organism in visuomotor action. The ventral stream of conscious visual perception appears to operate on more deliberative principles than immediate behaviour. These are principles with which contemporary Cognitive Science is ill-equipped to deal. Merlin Donald (Donald, 2001) also identifies intention and purpose as central to mental function, in the context of an incisive criticism of Psychology and Cognitive Science for ignoring the importance of these concepts. Donald argues that the skillful, intricate activities which we see routinely in people (such as holding one's side in a prolonged discussion or argument) would be inconceivable, let alone explicable, without making a person's intentions or goals central to our understanding of cognition.

2 A Outline of Goal-Directed Concepts

Based on Donald's description of levels of awareness, I offer a rough taxonomy of such concepts,. This differentiates basic biological interests (closely associated with homeostatic functions), from domain specific reasons (la Hurley, forthcoming), and more abstract purposes. Purposes are envisaged as longer term, allowing the organism to react not simply to the immediate situation, but to string reasons together in particular temporal order, to respond to situations which are drawn out over prolonged periods of time. The kinds of situations which demand purposive rather than rational behaviour are situations in which there is no single direct and immediate (that is, within the span of a short-term working memory) response to what is going on

2.1 Interests

Intentions are inevitably directed at ends, and thus begin with a creature being interested. By this, I do not mean to imply curiosity, enjoyment or any kind of reference to hobbies. I use the term more in the sense of when we say "financial interests". The Oxford English Dictionary offers: "the relation of being objectively concerned in something".

Every living thing has a range of interests, usually because they fairly directly concern their welfare or survival. These interests may be basic survival needs, reproductive motivations or other simple facts about the creature which will give certain facts about the immediate environment value. For something with interests, these are very basic forms of implication in the world. An organism's interests will be determined by the basic physical requirements of its make-up. They will generally coincide with the creature's needs, but I do not believe that the concept of need exhausts interests. An animal, for instance, may have no need to mate (for its own survival), but its biological make-up

may make such activity one of its interests, one of those things which affect the organism's behaviour without any real mediation between impact on the creature and reaction to that impact.

More complicated animals, perhaps those capable of associative learning of particular kinds, would be capable of reacting not only according to their immediate interests, but rather on the basis of a more abstract goal-directed construct: reasons.

2.2 Reasons

Susan Hurley (forthcoming) has argued that rich conceptual abilities (and the inferential promiscuity they imply) are not necessary for an animal to have reasons, and for those reasons to be the animal's own. She describes two features of actions which are sufficient for those actions to be done for reasons - holism (the differentiation between means and ends) and normativity (the possibility of action in error). An animal capable of acting in such a way acts not according to the simple changes in the states of its interests, but for reasons.

However, because of the fact that reasons are coupled to interests by a limited associative system, rationality and reasoning will be context-bound. That is, we will see what Hurley (forthcoming) refers to as "islands of rationality". The animal will learn the implications of new patterns of stimulation, but will do so within particular situations or domains. Though it will be possible for behaviour within a domain to be flexibly respondant to subtle changes in implications, cognable properties will only be perceived where there is some structure or scaffold in the environment to make reliable the relationship between the event and the implied. Some reinforcer must be present as an indicator of the relationship.

Layering even more complex and rich forms of implication into the world is possible once an animal becomes capable of teaching itself, rather than have to wait for a reliability in the relationship between a new event and a known concern. Once the creature can take command of its own reasons, it becomes capable of reacting purposefully.

2.3 Purposes

Purposes are not closely associated with particular actions. Rather, in much the way that reasons are more abstract structures of interests placed in context, so purposes are structures of reasons. Reasons are contextualized to particular domains or situations, identifying contexts given particular types of properties which are extant in the world. If an agent is capable of generating their own such context-determining properties, they are effectively capable of transforming the context of the present action, manipulating the implications of the situation and organizing reasons and actions according to longer-term goals.

Such freedom means that the purposeful agent is capable of extended practical reasoning. By fixing certain representations concerning the future, present sensations are given a new layer of implication. That future becomes a new context which provides further value and implication to the perceptual world of the agent. These future-directed representations play a rle described for intentions by Bratman (1999) in his planning theory of practical reasoning. By fixing certain future actions or end states - by commtting to them - we set the parameters of practical reasoning. Knowing these requirements, we can reason towards their satisfaction. Even given cases with two equally attractive options (so-called Buridan cases) we become capable of making a decision and forming intentions not on the basis of the merits of one option over the other, but on the merits of having to make *some* decision as required by the commitments earlier made. Intentions, for Bratman, are thus states of mind which bracket and constrain reasoning from the moment they are formed such that actions consistent with them become more likely, and the agent will act rationally (that is, according to their goals).

3 Conclusion

What is offered here is a first step toward a framework for intentions contextualized within contemporary Cognitive Science. It is expected that a change in viewpoint, revaluing the crucial rle played by goals, intentions, purposes and reasons in human cognition, will lead to some interestingly new understanding of many phenomena within Cognitive Science.

References

Bratman, M. (1999). *Intentions, plans and practical reasoning.* Stanford: CSLI Publications.

Clark, A. (2001). Visual consciousness and motor action: Are the bonds too tight? *Philosophical Review, 110*(4), 495-519.

Clark, A. (2002). Is seeing all it seems? *Journal of Consciousness Studies, 9*(5-6), 181-202.

Donald, M. (2001). *A mind so rare.* London: Norton.

Hurley, S. (forthcoming). Animal action and the space of reasons. *Mind and Language*.

Milner, D., & Goodale, M. (1995). *The visual brain in action.* Oxford University Press: Oxford.

Bootstrapping a better parser

Mark McLauchlan

mrm21@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

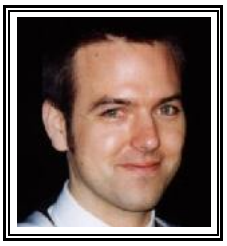

1 Introduction

One hundred years from now, people will undoubtedly be conversing freely with their computers. The obvious question is, how will those computers do syntactic analysis of each sentence? For example, I might tell my computer:

1. I'll buy the car with green propellers

It would then correctly identify *I* as the subject, *buy* as the verb and *with green propellers* as an adjunct of the object *car*. It's extremely important to get this right, since this analysis leads naturally to a semantic interpretation where I want the car with the optional green propeller blades. Unfortunately there is a second analysis where the prepositional phrase *with green propellers* modifies the verb, implying that I want to purchase a flying car using green propellers as currency. Syntax is a first step towards understanding language, and the ultimate goal of my research is to create computers that are capable of such understanding.

As this example illustrates, the key stumbling block is *ambiguity*. Armed with any reasonable grammar of English, every sentence will have hundreds or thousands of possible analyses. Parsing programs usually work by assigning mathematical probabilities to every possible analysis, based on evidence taken from a large corpus of training material. This paper describes the problems with these corpora and how we can overcome them.

2 The unbearable scarcity of data

The basic technique is to obtain a large number of sentences along with the correct syntactic analysis for each one. We choose a selection of language features, and then count how many times each feature appears in the training corpus. For example, if *green jets* appears with *car* three times and with *buy* once, then we would correctly choose the first analysis for the sentence above. Naturally one can tweak and devise more sophisticated models, providing hours of entertainment.

Unfortunately there are a lot of words in the English language, and most of them appear quite infrequently. Even in an enormous training corpus, most words will only appear once if they appear at all. If we have not seen a word before, then our simple model faces a hard road in deciding between intepretations. Worse, enormous training corpora are very hard to come by since every sentence has to be hand-analysed by human annotators - a dauntingly boring task.

One way around this is to ignore words altogether, or at least group words into well-defined categories such as nouns, verbs and adjectives. (Carroll, 1993) has developed such a non-lexicalised parser and the results are surprisingly good. A fair number of potential ambiguities can be resolved just by looking at how nouns and verbs behave *in general* rather than considering the actual words themselves. Unfortunately there are still some stubborn kinds of ambiguity, such as prepositional phrases, that really need lexical information to resolve.

3 Bootstrapping for fun and profit

My research aims to improve the accuracy of this parser by adjusting the probabilities it assigns to each candidate parse. Instead of modifying the parser itself I have developed a separate statistical model that reranks its output. This model considers both lexical and non-lexical features. It can also consider the entire parse when assigning probabilities — the base parser is forced to make some decisions before it has finished analysing the sentence.

The problem of training data is solved by a neat trick: using the output of the base parser itself to train the new model. This bootstrapping approach allows us to create as much training data as we want, but of course the data is not as accurate as we would like. Worse, the base parser is most inaccurate in exactly those areas we want to improve. Nevertheless, I have managed to achieve a slight increase in accuracy using this technique. The key is to screen the data and only learn from unambiguous samples. First, each parse is broken down into

its constituents, such as the subject-verb relation (I'm writing), the verb-object (writing thesis) and modifiers (writing on table).

The various analyses differ in how they treat *table*, but they all agree that the predicate *writing* belongs with the subject *I'm*. Therefore I can keep that relation in my training data and eliminate the other two. The table below shows the results of the base parser, my reranking model and an oracle that magically chooses the best parse from the top 5 candidates. This is an encouraging, although not statistically significant result. The new model is still failing in many of the same areas as the base parser, in particular with prepositional phrases.

Base Parser	Reranking Model	Oracle
76.7	77.2	82.0

This approach opens up several avenues for research that are closed to people working with traditional human-annotated corpora. Firstly we can examine how the parsing performance improves with very large amounts of training data. A surprising result here is that large corpora are not necessary, since accuracy levels off once the corpus has 10,000 - 30,000 sentences. Secondly, we can train and test the parser on different kinds of language. A parser trained on scientific journals will learn different probabilities to one trained on prose and poetry, for example. This is an ongoing research area.

References

Carroll, J. (1993). Practical unification-based parsing of natural language. In *Technical report 314.* University of Cambridge. UK.

Affecting motivation considering student's goal orientation

Erika Annabel Martínez Mirón

eamm20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

This paper describes my interest to improve educational software by taking into account students' motivational states during the interaction, giving particular attention to their goal orientations.

On one hand, some motivational strategies (Keller, 1987; Malone & Lepper, 1987) propose that the software has to fulfill several requirements in order to have more possibilities to maintain the user's motivation. In Table 1 are shown the points in common between these strategies.

On the other hand, goal theory (Ames, 1992; Dweck & Elliot, 1988; Nicholls, 1984) classifies students into two main groups: students with mastery goals who aim to develop new skills and competencies, and students with performance goals who try to demonstrate competence or try to achieve at high levels of normative ability. In addition, it is suggested that more effort is expended by mastery- than performance-oriented students because the former think that by attempting more, they have more possibilities to achieve their goals; where the latter think that by spending more effort they show lack of ability (Dweck & Elliot, 1988).

So, should teachers ask for more effort from students who believe that spending more effort is an indication of low ability? It seems more reasonable to choose an indirect way to get the same final result: spend more effort. For instance, we believe that stressing the students' goals can motivate them to keep trying until they get the correct result. In relation to feedback, it is suggested that this should be frequent, clear, constructive and encouraging (Lepper, 1988), attributing success to personal effort (Keller, 1987). With respect to this, students with different goal orientations should be offered different feedback, taking into account that mastery-oriented students attribute the success to their effort, where performance-oriented students lay their success on ability.

To summarize, the main hypothesis for this research is that providing a goal context corresponding to student's goal orientation, their motivation could be affected positively. That is, students with mastery/performance goal orientation will be more motivated with the respective mastery-oriented/performance-oriented system. In order to prove it, the Ecolab system will be used. This is a learning environment (with domain in food chains and food webs) developed to explore the way in which a computerised tutor might offer collaborative support (Luckin & du Boulay, 1999).

References

Ames, C. A. (1992). Classrooms: Goals, structures, and student motivation. *Journal of Educational Psychology, 84*, 261-271.

Dweck, C. S., & Elliot, E. (1988). Goals: An approach to motivation and achievement. *Journal of Personality and Social Psychology, 54*, 5-12.

Keller, J. (1987). Strategies for stimulating the motivation to learn. *Performance and Instruction Journal, 20*, 1-7.

Lepper, M. R. (1988). Motivational considerations in the study of instruction. *Cognition and instruction, 5*(4), 289-309.

Luckin, R., & du Boulay, B. (1999). Ecolab: The development and evaluation of a vygotskian design framework. *International Journal of Artificial Intelligence in Education, 10*, 198–220.

Malone, T. W., & Lepper, M. R. (1987). Making learning fun: a taxonomy of instrinsic motivations for learning. In *Aptitude, learning, and instruction. volume 3: Conative and affective process analyses* (Vol. 3, p. 223-251). Hillsdale, New Jersey: Lawrence Erlbaum Associates, Inc. Publishers.

Nicholls, J. G. (1984). Conceptions of ability and achievement motivation. In *Research on motivation in education* (Vol. 1, p. 39-73). San Diego, CA: Academic Press.

Table 1: Points in common between Keller's and Malone and Lepper's motivational strategies

a) variability in audio and visual effects;
b) clear goals or an environment where the students can generate them for themselves;
c) instruction responsive to learner motives and values (meaningful goals);
d) appropriate metaphors or analogies;
e) challenging experiences (graded difficulty levels);
f) techniques to offer personal control (responsive learning environment, activities with moderately high levels of choice);
g) frequent, clear, constructive, and encouraging performance feedback;
h) verbal praise, real or symbolic rewards;

The postposed subject in locative inversion vs object in Setswana in Lexical Mapping Theory

Setumile Morapedi

sm39@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

The structural role of the postposed grammatical subject in sentences which have undergone locative inversion has been a topic of some controversy for some time, with Perez (Perez, 1983):141 calling it the subject of the verb while Bresnan and Kanerva (1989):15 refer to it as an unaccusative object, that is, the grammatical object that occupies the subject position in intransitive verbs, such as *nna* 'sit', *goroga* 'arrive'. Locative inversion refers to a situation in which the grammatical subject and the locative noun phrase exchange positions. In that case, the locative noun phrase occupies the sentence-initial position while the grammatical subject appears after the verb, as shown by the construction in (1b) below:

(1a)

Pinki	o	ntse	mo khoneng
1a-name	1a-sm	sit-PRES	loc 9-corner

'Pinkie is sitting in the corner'

(1b)

Mo khoneng	go	ntse	Pinki
18-loc corner	17- sm	sit-PRES	1a-name

'In the corner its sitting Pinki'

Sentence (1a) is the unmarked (neutral) version of (1b). In example (1b) the sentence initial locative noun phrase takes up a new impersonal concord *go* 'it' which links it with the verb. It is referred to as impersonal, in that it is like a pronoun it, which does not carry any specific meaning. In this paper, I show that there is a difference between the postposed grammatical subject and the object. This will be determined through the three most frequently used criteria (tests) by Bantu linguists, such as, (a) word order (b) capability to become a subject in passivisation and (c) capability of being expressed as an object marker (OM) within the verb (Hyman & Duranti, 1983):220. I will start by giving a background on Setswana and the arrangement of parts of a locative phrase and those of a sentence in Setswana. These are very important in that they will enhance the understanding of the locative inversion analysis.

2 Background in Setswana: morpohology and syntax

Setswana is a Bantu language spoken in Botswana and some parts of South Africa in Southern Africa. It is characterised by noun classes, numbering from 1 up to 18. These noun classes have prefixes with singular/plural pairs. However, the infinitival class (15), and three locatives fa (class 16), go (class 17), and mo (class 18) corresponding to the Proto-Bantu (language of origin), *pa, *ku, and *mu (of classes 16, 17 and 18) do not have plural forms, as in examples (2) below:

(2a)

Fa	fatshe
16-loc	16-ground

'On the ground'

(2b)

mo	godimo
18-loc	17-top

'On the top/above'

(2c)

mo	morago
18-loca	3-back

'At the back/behind'

Further the arrangement of words in Setswana is such that the subject comes before the verb, which itself is followed by the object whilst the adverbial, which usually expresses time and place, follows the verb. Some words in the sentence may be emphasised by the speaker, in that case, they are focused. However, the Setswana language may misplace some of the parts of a sentence by placing them either at the beginning of a sentence where it is followed by a comma or at the end of the sentence where it is separated from the sentence by a comma. These peripheral constructions are topicalised.

3 Postposed subject vs object

In this section, I show that the postposed subject fails to pass subjecthood criteria.

3.1 Word order

This refers to a situation in which the object occurs in the position immediately following the verb, as in (3) below. This object occurs in the short form present of Setswana while the long form with 'a' does not take it:

(3a)
Ngwana o ja *borotho*
1-Child 1-sm eat-PRES 14-bread
'The child is eating bread'

(4a)
Mosadi o neela *ngwana* dilekere
1-woman is give-PRES 1-child 10-sweets
'The woman gives the child sweets'

In the examples, (3) and (4) the italicized words occurring immediately after the verb are objects. The former is the patient while the latter is the benefactive. On the contrary, the postposed grammatical subject does not have any of these roles.

3.2 Object as subject

This is a situation in which the noun phrase assumes a subject role through passivisation, as in (5a) and (5b) below:

(5a)
Ngwana o neel-wa dilekere.
1-child 1-sm give-PASS 10-sweets
'The child is given the sweets'

(5b)
Dilekere di neel-wa ngwana
10-sweets 10-sm give-PASS 10-child
'The sweets are given the child'

(6a)
Mo khoneng go ntse Pinki
18-loc 9-corner 17-sm sit-PERF 1a-name
'In the corner it/is sitting Pinkie'

(6b)
*Pinkie 0 nts-wa mo khoneng
1a-name 1a-sm sit-PASS 18-loc 9-corner
'Pinkie is sat in the corner' lit.

In (5a) and (5b) *ngwana* 'child' and *dilekere* 'sweets' have assumed subject roles in the passivised constructions. Conversely, the postposed grammatical subject cannot become the subject of the passive constructions, as shown by the ungrammaticality of the sentence (6b) above:

3.3 Object marker prefixed to verb

Yet another criterion which encounters the same problem is the one in which the object is expressed by the object marker attached to the verb, as in (7) below:

(7a)
Noga, monna o a *e* bolya
9-snake 1-man 1-sm PRES 9-om kill
'The snake, the child is killing'

(7b)
Bana ba a *di* ja dijo
9-children 2-sm PRES 8-m eat 8-food
'The children are eating it, the food'

(8)
*Mo khoneng go ntse -o
18-loc 9-corner 17-sm sit -OM
'In the corner is sitting her'

On the contrary, the postposed grammatical subject cannot be expressed through the object marker, as shown by the ungrammaticality of example (8) above.

4 Theoretical Implications

The theory used in this paper is that of Lexical Mapping (Bresnan & Kanerva, 1989). In this theory, each grammatical function, such as subject or object is assigned semantic roles (C.Harford, 1990). Further, the semantic roles exhibit restricted or unrestricted grammatical properties [+/-r] and also if they can have object properties or not [+/-0]. These features show the extent to which these roles can occur and their limitations. For instance, Agents are classified as non-objects [-0], meaning that they can only function as subjects or as oblique (by preposition) when passivisation has occurred. The themes are assigned the feature [-r] (not restricted), this implies that they can function as objects that become the subject, in the case of unaccussative verbs or passivisation in transitive verbs. With regard to the locative inversion, it occurs in the environment in which the theme is the highest role (Demuth & Mmusi, 1997):2.

5 Conclusion

To sum up, I have argued that the postposed subject that occurs after the verb is the grammatical subject but not the object. The reasons I adduced are that, it does not pass the objecthood tests that are most frequently used by Bantu linguists. I have also given a theoretical explanation that locative inversion occurs when the theme is the grammatical subject. In that case, when it gets postposed it remains the grammatical subject.

References

Bresnan, J., & Kanerva, J. M. (1989). Locative inversion in Chichewa: A case of factoriza-

tion. *Linguistic Inquiry, 20*, 1-50.

C.Harford. (1990). Locative inversion in Chishona. *Current Approaches in African Linguistic, 7*, 137-144.

Demuth, K., & Mmusi, S. (1997). Presentational focus and thematic structure in comparative Bantu. *Journal of African Languages and Linguistics, 18*, 1-19.

Hyman, L. M., & Duranti, A. (1983). On object relation in Bantu. *Syntax and Semantics: Studies in Transitivity, 15*, 37-46.

Perez, C. H. (1983). Locative psuedo-subjects in Shona. *Journal of African languages and linguistics, 5*, 131-155.

Assumption free cognitive modelling

Anthony F. Morse

anthonfm@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

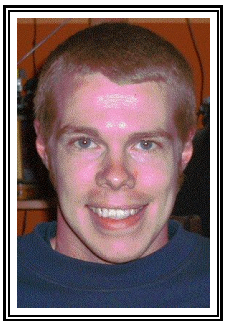

1 Introduction

Cognitive modelling has proved to be extremely useful in the development of psychological and, more recently, biological theories of cognition. However there remain some fundamental methodological problems which limit both the application and the explanatory power of these models. Typically, symbolic cognitive models contain explicit rule sets governing their behaviour, but provide no explanation as to where, how, or why these rule sets have come into existence. Conversely, non-symbolic cognitive models do not use rule sets but, typically, each model an aspect of cognition in a fully developed adult, often in small isolated tasks. Whilst the adult mind clearly is highly modular, there is a great deal of evidence that this is the result of ontogenetic development rather than pre-specified genetically determined modules. My research is concerned with the development of methodologies and tools for the construction of cognitive architectures via a developmental process of self-organisation. The title of this paper, 'assumption free cognitive modelling' is intended to reflect the idea that no assumptions are made concerning what knowledge the model will know and what rules will guide it's behaviour. Rather, the responsibility for these aspects of the model has been placed onto the various learning algorithms and dynamical systems involved. As such, the methods are intended to be as general purpose as possible and have a breadth of application to rival that of production systems.

2 Cognitive Modelling and Psychology

Following philosophical inquiry into systematicity and associative theories of mind, I have developed hybrid methods for the autonomous and unsupervised instantiation of Interactive Activation and Competition (IAC) networks (given appropriate inputs, see next section). Hand-wired IAC networks have been extensively used as cognitive models in the development of psychological theories - see (Young & Burton, 1999) for a summary - and it is hoped that the autonomous generation of these architectures will extend their psychological applicability to incorporate developmental theories. In addition to other cognitive properties, these networks can be used to predict changing inputs as a consequence of agent actions and environmental factors. Associating input features with appropriate homeostatic variables provides not only heuristics with which to compare the predicted consequences of various actions and select appropriately, but also provides goal oriented motivations (where the goal is to maintain various homeostatic variables).

3 Philosophy and Mathematics

At this point, the issue of what 'appropriate inputs' are must be addressed. Philosophically this is known as the Frame Problem, the problem of making salient those features relevant to a current problem whilst ignoring the rest. Note that this is different from choosing to ignore (Dennett, 1984). Typically the Frame Problem is not seen as problematic outside the symbolic paradigm, however, I claim that the difficulty in solving non-linearly separable problems is simply a reformulation of the Frame Problem in the non-symbolic paradigm. As stated in (Clark & Thornton, 1997), it is possible to trade representation against computation. That is to say, by changing the way a problem is presented, the computational requirements of the solution may also change. Equally, by altering the computational aspects of a solution the required representations may also change. This leads to a reformulated Frame Problem as the difficulty in matching representation and computation successfully to solve given problems. Further to this, the simple statistical learning methods incorporated in the Autonomous IAC architecture are known to be capable of solving all linearly separable problems. With this in mind, we can fix our computational

abilities and know that 'appropriate inputs' must present problems as linearly separable.

4 Computational Neuroscience and Dynamical Systems

Cortical Microcircuits, as examples of Liquid State Machines (LSM) (Maass, Natschlager, & Markram, 2002) are complex dynamic neural systems, which without training or design implement integration over time and kernel functions. These kernel functions project input data into a higher dimensional space where, it is claimed that given sufficient resources, all non-linearly separable problems will be transformed into linearly separable problems. Whilst the validity of this claim is unproven, these systems certainly do transform many non-linearly separable problems into linearly separable ones. Typically this is achieved by explicitly training linear readout units to interpret the sates of the network on sample data. A number of benchmark tests have been performed using Cortical Microcircuits with impressive results (Maass et al., 2002). My most recent work has been in the development of a very simple method, which autonomously trains readout units from a Cortical Microcircuit based on an analysis of the relationships (interference) between existing readout units. Essentially this technique identifies non-linear relationships between existing categories and trains new categories that will have a linear relationship to the pre- existing non-linear one. These categories/readout-units then provide an appropriate input for the autonomous generation of IAC architectures. By generating self-training signals, the entire architecture is rendered unsupervised (although reward signals are necessary). Psychologically, it may be interesting to investigate the extent to which the starting categories (which can also be autonomously generated) determine which categories or concepts can later be discovered.

5 The Future

Future developments of the architecture will include testing whether, having learnt about and planned actions in an environment, readout units can be autonomously trained to produce the same behavioural response directly from the Cortical Microcircuit. If so, this may provide a possible psychological explanation for the automation of behaviour, and possibly a radically different conception of the difference between long and short-term memory.

Future work with the architecture will hopefully demonstrate that it's applicability as a cognitive model has been significantly extended beyond that of current IAC architectures, and that it is practically useful in the development of psychological theories. Further to this, as an associationist theory of mind, this model captures not only psychological properties, but also psychotherapeutic properties, and may be usefully employed in the simulation of unethical situations such as the development of phobias.

References

Clark, A., & Thornton, C. (1997). Trading spaces: Computation, representation, and the limits of unsupervised learning. *Behavioural and Brain Sciences, 20(1).*

Dennett, D. C. (1984). Cognitive wheels: The frame problem of AI. *Minds, Machines and Evolution*, 129-151.

Maass, W., Natschlager, T., & Markram, H. (2002). Computational models for generic cortical microcircuits. *Computational Neuroscience: A Comprehensive Approach*, CRC-Press.

Morse, A. F. (2003). Autonomous generation of Burtons' IAC cognitive models. *in press.*

Young, A. W., & Burton, A. W. (1999). Simulating face recognition: Implications for modelling cognition. *Cognitive Neuropsychology, 16-1,* 1-48.

Baseline Confessions: Goats Lick Buttercups

Darren Pearce

darrenp@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Natural Language Processing is one of the many disciplines in which baselines are used to highlight the effect of utilising a particular source of information. To some extent, this corresponds to the notion of a control in many other research disciplines.

This paper discusses the result of a baseline technique in the evaluation of systems for extracting 'collocations' from the British National Corpus — a very large collection of textual material taken from a wide range of sources. Collocations can be defined informally as 'useful' phrases that tend to function as units in human language. For example:

amino acid	*big business*
county council	*disk drive*
essential element	*fossil fuel*
greenhouse gas	*home help*
inward investment	*jam jar*
long leg	*mass media*
neural network	*olive oil*
planning permission	*red rose*
social service	*table tennis*
vinho verde	*world war*

The data discussed in this paper results from doctoral research that develops new techniques to automatically extract collocations from large collections of text. These techniques all follow from the observation that substituting constituent words of collocations for synonyms tends not to result in commonly occurring phrases. For example, substituting *war* for a synonym such as *warfare* in the phrase *world war* results in *world warfare* which is comparatively rare.

Developing a sophisticated automatic approach to this task is not a straightforward process.[1] However, using the example above, a *naïve* approach could simply compare the frequencies of the phrases *world war* and *world warfare* directly. It is important to note that this comparison would *not* take into consideration the fact that *warfare* is a far less common word than *war*; *observed* frequencies must be compared to *expected* frequencies.

[1] See Pearce (2002b) for a brief discussion of some of the issues involved.

2 Experimental Details

2.1 Results

In order to comparatively evaluate the effectiveness of using the expected frequencies, a naïve baseline was developed that deliberately ignored this information. Figure 1 shows the top 15 phrases in which an adjective modifies a noun and the top 15 where a noun modifies another noun. The salacious nature of the results of this experiment was entirely unexpected.

	Adjective-Noun		**Noun-Noun**	
r	**Phrase**	f	**Phrase**	f
1	cheap seat	4	rear seat	79
2	wet grass	47	tail lot	1
3	stiff tail	4	heart seat	1
4	batty deal	2	bottom bunk	5
5	wet rot	20	bottom brace	2
6	great deal	1607	zip bottom	1
7	slick butt	1	rear zip	1
8	cracked pot	1	child seat	22
9	stiff bump	1	mass star	1
10	blind child	18	lot bull	1
11	awful lot	153	fish tail	3
12	complete mess	22	bottom line	142
13	complete prat	5	rear child	4
14	tight line	19	pile child	1
15	sharp butt	1	mint humbug	3

Figure 1: Extracted phrases using the naïve baseline. The top 15 phrases in which an adjective modifies a noun and the top 15 in which a noun modifies another noun are shown in rank order (r). Each phrase is accompanied by its corresponding occurrence frequency, f, in the BNC.

2.2 Analysis

The phrase that was first noticed when analysing the data was *slick butt* and so it is used here as a worked example. WordNet (Miller, 1990) lists 11 words that all mean *slick* in the sense of 'marked by skill in deception':[2]

crafty	*cunning*	*dodgy*	*foxy*
guileful	*knavish*	*slick*	*sly*
tricksy	*tricky*	*wily*	

and 24 words that all mean *butt* in the sense of 'the fleshy part of the human body that you sit on':

arse	*ass*	*backside*	*behind*
bottom	*bum*	*buns*	*butt*
buttocks	*can*	*derriere*	*fanny*
fundament	*hindquarters*	*keister*	*posterior*
prat	*rear*	*rump*	*seat*
stern	*tail*	*tooshie*	*tush*

resulting in $11 \times 24 = 264$ possible different phrases. The words that constitute such phrases may not always occur consecutively in a sentence so frequencies are obtained using word dependencies resulting from processing the BNC using the system described in Carroll, Minnen, and Briscoe (1998). Of these 264 phrases, the only one that occurs is *slick butt* and it occurs just once. This single occurrence is in fact an error in the processing of the sentence:

> *Get your butt on a bird, Slick, and let's make Bad Money.*

in which *Slick* was identified as an adjective that modified *butt* rather than a proper noun. Given the other 263 phrases did not occur at all, this single (erroneous) occurrence was given a very high score by this naïve baseline.

3 Conclusions

This technique was not expected to extract useful phrases; it was intended as a basis for comparison. What *was* unexpected was the salacious nature of the resulting highly-scored phrases. Variants of this baseline that also do not utilise expected frequencies extract phrases of a similar nature and worse.

The high scores for phrases involving the word *butt* or its synonyms is due largely to the fact that there a high number of these synonyms. In fact this is the largest synonym group in WordNet.[3] Other groups of synonyms that include slang terms also tend to be larger than average. A possible explanation for this is that the underlying concepts such words represent require a wide range of possible realisations that are suitable for different social contexts with different connotations. For example, the word *fundament* is used in very different contexts to the word *butt*.

4 Future Work

Forthcoming work builds on the comparative evaluation described in Pearce (2002a) and compares this technique to other, more sophisticated variants as well as proposing new ways in which to evaluate collocation extraction techniques.

Acknowledgements

The author wishes to thank Rochelle English for telling him that goats lick buttercups, and his supervisor, John Carroll, for permitting the publication of this paper.

References

Carroll, J., Minnen, G., & Briscoe, T. (1998). Can subcategorisation probabilities help a statistical parser. In *6th ACL/SIGDAT Workshop on Very Large Corpora* (pp. 118–126). Montreal, Canada.

Miller, G. A. (1990). WordNet: An on-line lexical database. *International Journal of Lexicography*.

Pearce, D. (2002a). A comparative evaluation of collocation extraction techniques. In *Third International Conference on Language Resources and Evaluation.* Las Palmas, Canary Islands, Spain.

Pearce, D. (2002b). Frame alerts and fire alarms. In J. Weeds (Ed.), *White House Papers* (Vol. 15). Herstmonceux Castle. (COGS, Sussex University, Cognitive Science Research Paper 552)

[2] Depending on the interpretation of the meaning of the phrase *slick butt*, the synonyms of *slick* listed here may be considered incorrect. Automatically deciding the 'correct' sense of any word — a task called 'sense tagging' — is difficult and the subject of much research in Natural Language Processing. The problems introduced into subsitution-based techniques for collocation extraction through the lack of sense information in the BNC are discussed in Pearce (2002b).

[3] The latest version of WordNet also includes the word *nates* in this group.

Be or not be? First language interference in second language learning

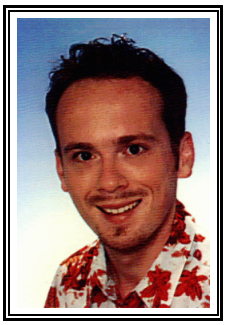

Sebastian Rasinger

sr26@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

It seems to be commonly agreed that when learning any second language (henceforth L2), one's first language (L1) influences the acquisition process. However, the extent of linguistic transfer, the "copying" of L1 features - particularly of grammatical characteristics - onto the L2 has yet not been explained unambiguously. Nicol (Nicol, 2001) argues that though they have automatized the retrieval and use of L2 rules, even highly proficient L2 speakers might - unconsciously - use production routines from their native language, leading to ungrammatical L2 utterances. This negative transfer may sometimes lead to drastic divergences from the rules of the L2 (Odlin, 1989). However, in contrast to negative transfer, positive transfer may facilitate (Second Language Acquisition) SLA. For example, similarities in the lexicon (cf. e.g. striking lexical similarities in Romance languages) may make it easier for learners to acquire vocabulary (Odlin, 1989). Similarly, analogies in L1 and L2 may provide learners with advantages in the acquisition of certain grammatical structures (Ellis, 1994). In this paper, I will exemplify the aspect of negative transfer with the acquisition of English copula constructions by native speakers of Bengali.

2 Copula constructions in English and Bengali

Copula verbs are parts of language which have "little or no independent meaning", and whose only function is to link certain syntactic elements (Crystal, 1999). In a sentence like

1) I am tired.

am constitutes the copula verb; its task is to link the subject *I* and the complement *tired*. The omission of the copula verb 'be' (realized by the inflected form 'am') leads to the ungrammatical utterance

2) I tired.

Nevertheless, the actual content of the utterance is (rather) unaffected, and its meaning can be conveyed to the interlocutor despite the apparent ungrammaticality. The main English copula is the verb to be, and its inflected forms (I am, you are etc.). Unlike English, the Bengali language does not possess copula verbs for subject-complement constructions:

3) ishkul bondo. [The] school closed. (Chalmers, 1996)

Example 3 is a grammatical sentence in Bengali, although it is missing a copula verb that links the subject *ishkul* (school) with its complement *bondo* (closed). Thus, it could be assumed that Bengali learners of English show rather high error rates, i.e. tend to produce incorrect sentences like example 2, because they copy their native language feature "missing copula for subject complement constructions" onto their second language English, which requires a copula verb. A negative transfer is likely to occur.

3 Methodology

Speech data was collected from 10 Bengali learners of English, who live in East London. Learners started learning English at different ages, are of several socioeconomic backgrounds, and have been living in England for between 2 and 30 years. Data was analyzed for the occurrence of copula constructions and the relative frequencies for correctly and incorrectly produced constructions calculated.

4 Results

Analysis of the influence of learners' ages, and their length of English learning/exposure - usually 2 reliable predictors for L2 outcome - has shown that neither age nor length do significantly predict learners' performance with respect to copula constructions: even young learners, who are usually said to acquire a second language more easily, as well as learners who have been learning English for

several years, show somewhat drastic error rates. On average, only 67 per cent of all constructions are produced correctly. On the basis that the two usual predictors (age and length) seem not to have a significant influence, and that the difference in English and Bengali copula constructions is structurally essential (not only on the ”visible” structure of the sentence, but also on the abstract level of mental sentence processing), it seems plausible to argue in favour of a rather strong negative L1 transfer - the copying of a distinct L1 feature (missing copula) onto the L2, producing incorrect English copula/subject-complement constructions. This phenomenon may also be facilitated by the lack of actual meaning of copula verbs.

References

Chalmers, R. (1996). In *Learning sylheti.* London: Centre for Bangladeshi Studies. Roehampton Institute.

Crystal, D. (1999). In *The penguin dictionary of language vol. 2.* Penguin Books, London, UK.

Ellis, R. (1994). In *The study of second language acquisition.* Oxford University Press, UK.

Nicol, J. L. (2001). One mind, two languages: bilingual language processing. In . Blackwell, Oxford, UK.

Odlin, T. (1989). Language transfer. cross-linguistic influence in language learning. In *The cambridge applied linguistics series.* Cambridge University Press, UK.

Assesing motivational issues in a Vygostkyan ITS

Genaro Rebolledo-Mendez

gr20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 abstract

An analysis of the influence of motivation on the Zone of Proximal Development (ZPD) is being carried out. To investigate the impact of motivational factors on the ZPD the inclusion of a motivation differentiator layer into the Ecolab (Luckin, 1998) is proposed. The mechanisms and values to measure the motivational states of the learners and scaffold the motivational strategies to adjust them to individual pupils are outlined. The main research question on which this proposal is centred is: How can motivational factors be incorporated within a Vygotskyan framework for an Intelligent Tutoring System (ITS) ?

2 Introduction

The main issue of this proposal is the relationship of the concept of the 'Zone of Proximal Development' (Vygotsky, 1978) to motivation. Although such a relationship is not directly addressed by Vygotsky, it is argued here that some motivational aspects such as effort or confidence are implicit in the ZPD. This relationship is made explicit in order to be able to include a motivational module into the Ecolab (Luckin, 1998). The Ecolab is an ITS designed within a Vygotskyan framework with which children aged 10 and 11 can investigate food chains and webs. In order to include motivational issues into the Ecolab, a relationship between metacognition and motivational variables is established and a mechanism to measure the motivational state of the pupil is proposed. This work extends on that of del Soldato (Gaunaurd et al., 1998) but expands it by describing a model of motivation in a Vygotskyan ITS while trying to establish a relationship between motivational factors and the ZPD.

3 Motivational variables and the ZPD

The theoretical basis of this work can be found in Motivational Theory (Brill & Helweg, Accessed 20/05/02), (Gaunaurd et al., 1998), (Feng & Liu, 2001) and Social Cognitive Development Theory (Luckin, 1998), (Vygotsky, 1978). Motivation theory in general and motivational instructional design in particular address ways of understanding motivation and its mechanisms. The inclusion of motivational modules into instructional systems has dealt with topics such as the diagnosis of the learner's motivational state to create a learner model. The work of Vygotsky is also reviewed in order to identify points in common with the concepts of motivation. The analysis of the Ecolab reveals important features that, together with an understanding of motivational variables and social cognitive theory, constitute the research proposal.

4 Research Proposal

What is the nature of a relationship between the Zone of Proximal Development (Vygotsky, 1978) and Keller's (Feng & Liu, 2001) Attention, Relevance, Confidence and Satisfaction (ARCS) model? In a Vygotskyan framework, metacognition indicates the degree of self-regulation and awareness of the learning process which is desirable in learners. This implies that high displays of four motivational variables: effort, independence, control and confidence are all characteristics of developed learners. In order to scaffold motivational issues within the ZPD, the more able partner should increase in the learner the effort and the feelings of independence, control and confidence during the learning situation. The modelling of motivational issues within Ecolab requires that effort, independence, control and confidence be measured. Another important aspect of this study consists of the reactions that the system will provide so as to scaffold the motivational state of the learner if the model detects a low state. Keller's ARCS model (Feng & Liu, 2001) provides a set of reactions that were taken into account to build the new motivational layer for the Ecolab, as well as aspects of narrative (Waraich, 2002).

5 Future work

Work for the future include the design and implementation of the motivational modeller. The design process will involve learners in a “Wizard of Oz” (Anderson, Au, Larsen, & Hansen, 1999) study so that the final product would engage learners within the framework of story which aims are the learning of ecological concepts embedded in the Ecolab. A future study would test this modeller, and if there exist an influence of motivation in the ZPD more research could be done to establish the effect of individual differences such as the learner’s goals (Sansone & Harackiewicz, 2000) in the learning process.

References

Dahlback, N., Jonasson, A., & Ahrenberg, L. (1993). Wizard of oz studies: Why and how? *Knowledge Based Systems, 6,* 258-266.

De Vicente, A., & Pain, H. (1998). Motivation diagnosis in its systems. In *Proceedings of the fourth international conference on intelligent tutoring systems.* Springer, San Antonio, TX.

Del Soldato, T., & B. du Boulay. (1995). Implementation of motivational tactics in tutoring systems. *International Journal of Artificial Intelligence in Education, 6,* 337-378.

Keller, J. M. (1987). Strategies for stimulating the motivation to learn,. *Performance and Instruction Journal, 8,* 1-7.

Luckin, R. (1998). Ecolab: Explorations in the zone of proximal development. In *Phd thesis.* School of Cognitive & Computer Sciences, University of Sussex, UK.

Sansone, C., & Harackiewicz, J. M. (2000). In *Intrinsic and extrinsic motivation: The search for optimal motivation and performance.* Academic Press, San Diego, CA.

Vygotsky, L. S. (1978). In *Mind in society : The development of higher psychological processes.* Harvard University Press, Cambridge, MA.

Waraich, A. (2002). In *Designing motivating narratives for interactive learning environments.* School of Education, Computer Based Learning Unit, University of Leeds, UK.

End-user and system interface; network breakdowns and a language mismatch

Jon Rimmer

jonr@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

The last decade has seen an explosion in terms of people's use of the Internet. The World Wide Web is no longer primarily used by experts - it has been globally adopted by all manner of people, from young children to grandparents.

Internet use has become ubiquitous and now accommodates an ever-increasing amount of functionality allowing novel interactions and wider forms of communication. However, the speed at which these technologies have been developed by experts, and embraced by novices, has left little time for considered replacement of the terminology of old expert systems with more people-friendly terms. One only has to read the content of many occurring 'error messages' when using the Web, to see that the language is often couched in very system-orientated terms, accompanied by threatening vocabulary delivered with an imperious tone. These messages offer very few clues to the average user as to what has happened, and how they can recover and resume the task they were undertaking before being so rudely interrupted.

The field of Human-Computer Interaction offers user-advocacy in order to inform the process of interaction design. I have undertaken a series of studies that have focused on people's use of the Internet, web and other networked technologies. These highlight that the language used within many interfaces is mismatched with users' knowledge and expectations. This inevitably leads to user frustration as they have little to go on and disappointing user experiences. Moreover, increased costs are involved when organisations are forced to provide helpdesk support to disentangle users' problems. These often stem from misunderstandings relating to how systems can be used effectively.

The results of these studies have led to the development of a variety of techniques to capture and analyse user language so that it can be better employed in the development of useful technology.

A computational approach to pulsatility in the neuroendocrine system

Enrico Rossoni

er28@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Mammals use their brains to perceive changes in the external environment, and appropriate responses are regulated via the nervous and the endocrine systems. In particular, the endocrine responses are orchestrated by the hypothalamus and the pituitary gland. The hypothalamus controls body temperature, water balance, blood pressure, food intake, energy balance, and much of the instinctive or reflex behaviours, including male and female sexual behaviour and maternal behaviour. Further, it governs the release of a host of hormones from the pituitary gland, including oxytocin, vasopressin, growth hormone, thyroid stimulating hormone, just to mention a few. For most of the pituitary hormones, secretion must be pulsatile in order to be biologically effective. Within this context, the case of the hormone oxytocin (OT) can be considered as an exemplar. Moreover, it is particularly appealing for theoretical modelling, since much is known about the electrophysiology of the OT neurons, and the response of the OT system is largely accessible to experimental investigation. Oxytocin controls milk let-down in response to suckling, and the progress of parturition by its action on the uterus. It is released from nerve terminals in the pituitary, in response to action potentials that originate in the OT cell bodies in the hypothalamus. Normally, OT cells discharge at low frequency (1-3 spikes/s) and asynchronously. During suckling, every 5 minutes or so, every oxytocin cell displays a burst of activity, lasting 1-3s, during which firing rate rises up to 50-100Hz (Lincoln & Wakerley, 1974). Notably, bursting activity is synchronised across the whole population of OT cells, resulting in a massive pulsatile hormone secretion from the pituitary. Despite the rather simple, low-dimensional, signal which is encoded by the OT system, a clear understanding of the mechanisms underlying its behaviour is still lacking. Indeed, by considering some of the experimental findings, the whole picture seems quite confused:

- Bursts are not simply a passive response to a pulsatile input, rather the output pattern appears to be a response to different levels of tonic input. Bursting behaviour is a 'specific' response to an appropriate input, whereas other modes of activation, such as hyperosmotic stimuli, do not induce any bursting.

- OT cells appear to be synchronised only during the bursts. Moreover, asynchronous bursting has never been observed, and bursting cannot occur in isolated OT cells. This would suggest that OT cells lack an intrinsic ability to generate a bursting behaviour.

- During suckling, oxytocin is also released centrally by dendrites (Pow & Morris, 1989), and acts on several targets within the hypothalamus, facilitating bursting (Leng, 1999). Reported effects include: modulation of both excitatory and inhibitory synaptic transmission (Kombian, 1997), direct excitation of OT cells (Yamashita, 1987), and facilitation of OT release from the dendrites (Lambert, 1994).

- Dendrites of OT cells are bundled together during lactation, so complex dendro-dendritic communications among OT cells are likely to occur, probably mediated by OT release.

We support the hypothesis that synchronised bursting may be an emergent behaviour, resulting from network interactions: OT neurones would appear to be organised as weakly 'pulse- coupled' oscillators, and positive feedback on cell excitability, via local OT release, could drive them toward bursting. In such a framework, the previous observation could be interpreted in a coherent way. A direct validation of such an hypothesis, however, is hard to achieve. We will examine whether the main experimental findings about the OT system are sufficient to provide an explanation of the synchronised bursting, by implementing them in a computational model. As a starting point, a model

of reduced complexity has been considered to describe OT neurons, and network topology. This approach is aimed to isolate the key variables that contribute to the observed dynamics, and to define the functional relationship among them, neglecting all those properties which have no major impact upon the selected behaviour.

References

Kombian, S. B. (1997). Dendritic released peptides act as retrograde modulators of afferent excitation in the supraoptic nucleus in vitro. *Neuron, 19*, 903-912.

Lambert, R. C. (1994). A rise in the intracellular Ca^{2+} concentration of isolated rat supraoptic cells in response to oxytocin. *J. Physiol, 478*, 275-288.

Leng, G. (1999). Physiological pathways regulating the activity of magnocellular neurosecretory cells. *Prog. Neurobiol, 57*, 625-655.

Lincoln, D. W., & Wakerley, J. B. (1974). Electrophysiological evidence for the activation of supraoptic neurones during the release of oxytocin. *E. J Physiol, 242*, 533-554.

Pow, D. V., & Morris, J. F. (1989). Dendrites of hypothalamic magnocellular neurons release neurohypophysial peptides by exocytosis. *Neuroscience, 32*, 435-439.

Yamashita, H. (1987). Oxytocin predominantly excites putative oxytocin neurons in the rat supraoptic nucleus in vitro. *Brain Res, 416*, 364-368.

The use of temporal models to investigate Dolphin Echolocation

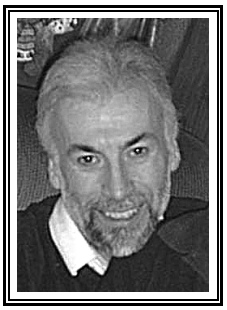

John Rowston

jwr20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Many marine mammals have, over a long period of time, evolved sophisticated sonar systems that allow them to freely navigate around and to have a large degree of control over their environment. One of the most successful of these mammals is the dolphin, which emits a series of high frequency 'clicks' lasting a very brief 50 millionths of a sec ($50\mu s$). These clicks cover a broad frequency spectrum (about $37kHz$) up to a maximum of around $150KHz$, with the higher frequencies tending to be used when there is danger of the signal being drowned out by background noise. It has been found that the discrimination abilities of the dolphins studied are sensitive to very small changes in target dimensions or elastic material composition. This sensitivity holds important implications for the design of future sonar systems. Not only mines made of hard to detect materials could be pinpointed but also 'fingerprints' of vessels, via material composition etc., could be held in databases. This would enhance the classification of targets and lead to detection of alien as opposed to friendly naval underwater objects.

2 Main Research Area

The main thrust of dolphin echolocation research has been in the area of Artificial Neural Networks (ANN'S), kind of artificial brains that have proven to be very efficient tools for pattern recognition tasks.

A network consists of a matrix of nodes in layers connected via weights (see figure 1). The nodes act like very simplified neurons with the weights acting as the synaptic strengths making connections between them. After training with several examples the connections become stronger or weaker depending on the training patterns presented, so that now the network is able to generalise the classification task to new patterns previously not seen.

Much of the work into dolphin echolocation has been carried out under the auspices of the Office of Naval Research based in San Diego. Moore et al correctly classified $90-93\%$ of echo trains, using a 'Gateway Integration Network' (GIN), which combined the information from multiple echoes from the same target. It was assumed that because the dolphin emitted a number of clicks per trial that it averaged or summed information from the spectral returns until it could confidently classify the target.

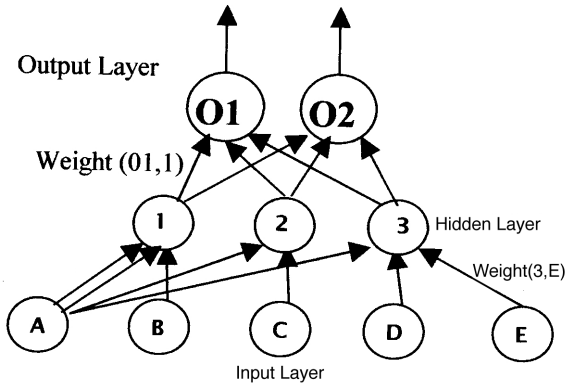

Figure 1: A typical 'Feed Forward' Neural Network. The hidden layer is so called, as it is an intermediate layer with no direct access to its outputs. (NB: not all connections are shown, in order to improve clarity).

(Anderson et al., 1999) used simulated dolphin clicks to classify echoes from 10 stainless steel hollow cylinders. They found that, using a combination of matched filter envelope detection, a gammatone filter bank, time integration and principal component analysis, they were able to classify different wall thickness to within $0.15mm$ with 99% accuracy.

(Gaunaurd et al., 1998) examined a large set of back-scattered echoes resulting from dolphin-emitted acoustic signals and concluded that certain features in the echo contain information about the material composition, size and filler characteristics of the targets. (Brill & Helweg, Accessed 20/05/02) tested the dolphin's ability to discrimi-

nate synthetic signals and reported their sensitivity to spectral differences spaced as close as $10\mu s$.

3 Temporal Model

Recently there has been much interest shown in a new form of network known as an 'Integrate and Fire Spiking Network' (IFSN). This type of network uses spikes instead of analogue values to perform pattern matching etc. It is thought to be much more biologically plausible in the way that individual nodes sum the incoming values (membrane currents), and fire a spike (action potential) if a set threshold is overcome (see figure 2). It is then the spiking rate or spike pattern that is used in the matching process. This type of process accords much more with the experimental data that is available from neurobiologists and neuroscientists working in the area of brain research and is much more able to cope with temporal data.

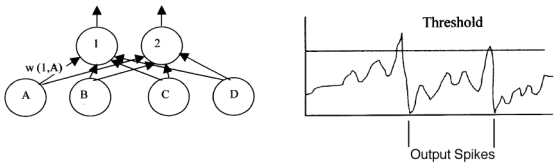

Figure 2: IFSN showing the input to neuron 1 (Summed Potential), as the summed output of the product from neurons A-D and their respective weights $w(x,x)$. Only when that input is driven over the threshold value, does it result in an action-potential spike output.

There is also recent research into the role of inhibition in the discrimination task (Feng & Liu, 2001) with the involvement of second order statistics and the variance proving to be of more importance than the mean firing rates alone.

4 Conclusion

The initial emphasis of this research centers on the use of spiking networks together with various signal-processing techniques in order to attempt replication of the discrimination abilities of the dolphin in a more biologically plausible manner. At the same time it is envisaged that more general discrimination tasks might well benefit from these approaches.

References

Anderson, L. N., Au, W., Larsen, J., & Hansen, L. K. (1999). Discrimination of cylinders with different wall thicknesses using neural networks and simulated dolphin sonar signals. In Y. H. Hu, J. Larsen, E. Wilson, & S. Douglas (Eds.), *Proceedings of the IEEE workshop on neural networks for signal processing* (p. 477-486). 1X Piscataway New Jersey IEEE.

Brill, R. L., & Helweg, D. A. (Accessed 20/05/02). Alternative mine detection and classification technologies SPAWARSYSCEN. (project proposal with some results). *www.cee.hw.ac.uk/~aramis/ONR/REPORTS/TSBRI*.

Feng, J., & Liu, F. (2001). *Discriminating between different input signals via single neuron activity* (Tech. Rep. No. 2001:07). University of Sussex, UK.

Gaunaurd, G. C., Brill, D., Huang, H., Moore, P. W. B., & Strifors, H. C. (1998). Signal processing of the echo signatures returned by submerged shells insonified by dolphin "clicks": Active classification. *Journal of Acoustical Society of America, 103(3)*, 1547-1557.

Learning for task based visual control

Kingsley Sage

khs20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Task based control

There is much current research interest in the development of frameworks for cognitive vision systems. Many of these frameworks draw inspiration from work on active (purposive) vision ((Aloimonos, Weiss, & Bandopadhay, 1987), (Bacsy, 1985), (Ballard, 1991), (Tsotsos, 1992) and (Ullman, 1984)). Task based visual control is one useful element of the cognitive vision paradigm and is concerned with the application of a task relevance control structure to guide low level vision processes. The benefit of this approach is computational efficiency both in terms of the low level processes and subsequent scene interpretation. Processing is limited to only that data that is considered task relevant. Task based visual control can therefore be thought of as data driven (bottom) up) processing limited in scope by task based (top-down) control.

2 Learning in the context of task based visual control

Bottom up non-cognitive approaches to vision require hand crafting of feature detectors and domain knowledge to build a system that is capable of recognising defined events of behaviours. In contrast, learning in the context of task based visual control is concerned with the modelling and recognition of activities involving highly structured and semantically rich behaviour.

For task based control, the high level representation structure would identify task relevant primitive objects and provide a basis for predicting interactions between those objects. So, for example, the representation would be sufficiently powerful to provide predictive cues for spatio- temporal tracking (i.e. primitive objects that move in a purposeful manner) and object interactions (e.g. for a task of ”making a cup of tea” it is likely that an empty hand moving in a purposeful manner away from a body torso might be about to pick up a cup, but rather less likely to attach itself to an electrical socket).

3 Integrating low level feature extraction with statistical behaviour analysis

A common theme throughout much previous work has been the separation of feature detection from the high level interpretation or behavioural analysis. Recent work by Frey and Jojic shows how the two may be combined into a single powerful learning procedure. (Frey & Jojic, 1999) describes the Transformed Mixture of Gaussians (TMG) model that combines established mixture modelling techniques with latent transformation variables that can represent a wide range of spatial transformations such as translation, rotation and shearing. This has the advantage over traditional mixture modelling in the extraction of statistically significant features is separated from their spatial localisation (which is described by the transformation variables). Jojic et al (Jojic, Petrovic, Frey, & Huang, 2000) have extended this generative transformed model using Hidden Markov Models (HMMs) to build a system capable of clustering unlabelled video segments and forming a video summary in an unsupervised manner.

Figure 1 below illustrates a TMG model in practice. The model was trained using a test set consisting of two 5*5 shapes (a square and a cross) superimposed over an 11*11 normally distributed background. The test set consisted of 200 images with the shapes translated horizontally and vertically at random. The model was then used to learn 2 TMG prototypes with translation invariance from a random initial configuration. Figure 1a shows a number of frames from the test set. Figure 1b shows the learned prototypes. The mean maps (top) show the mean values of the mixture model and the variance maps show how the significance of the mixture components to the model where dark pixels have low variance and are most significant.

These approaches are particularly relevant for task based control. Generative models can provide a natural predictive mechanism well suited to attentional control. For example, HMMs can provide ranked statistical estimates of the most likely spatial transformation of a feature in an ordered time

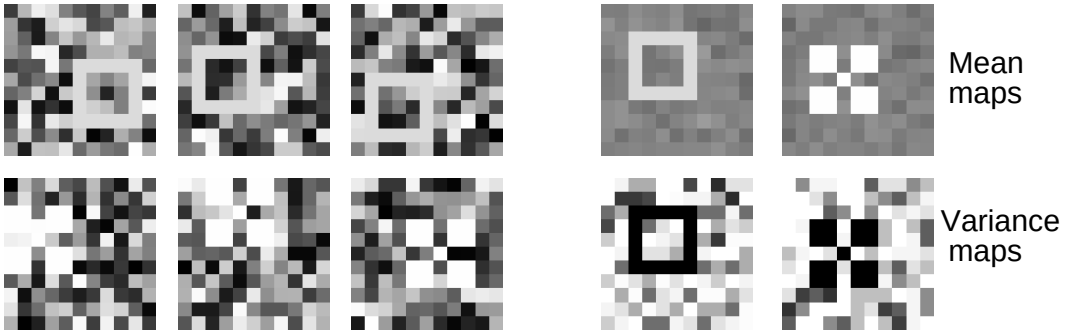

Example images from the training set Learned prototypes

Figure 1(a) Figure 1: Figure 1(b)

sequence. This knowledge can be used to focus computation to those areas of the scene that are most likely currently task relevant. This provides a cognitive model of ”perception guided by expectation”.

As far as task based control is concerned, the key to using models such as the TMG rests in a knowledge (as a function of time) of the index to the set of transformations. A further extension of the TMG (Jojic et al., 2000) to incorporate HMM modelling of transform selection as a function of time demonstrates this in practice. Work in progress includes investigation of higher and variable (using variable length Markov model) order temporal dynamics of index of transformations.

References

Aloimonos, Y., Weiss, I., & Bandopadhay, A. (1987). Active vision. In *IEEE international conference on computer vision* (p. 35-54). London, UK.

Bacsy, R. (1985). Active perception versus passive perception. In *IEEE workshop on computer vision.* Washington, US.

Ballard, D. (1991). Animate vision. *Artificial Intelligence, 41*, 57-86.

Frey, B., & Jojic, N. (1999). Estimating mixture models of images and inferring spatial transformations using the EM algorithm. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (p. 416-422).

Jojic, N., Petrovic, N., Frey, B. J., & Huang, T. S. (2000). Transformed hidden markov models: Estimating mixture models of images and inferring spatial transformations in video sequences. In *IEEE conference on computer vision and pattern recognition.*

Tsotsos, J. (1992). On the relative complexity of active vs. passive visual search. *International Journal of Computer Vision, 79*(2), 127-142.

Ullman, S. (1984). Visual routines. *Cognition, 18*, 97-159.

Working Memory in Adults with Asperger's Syndrome

Brenda Smith

brendas@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Asperger's Syndrome (AS) is a developmental disorder believed to be of neurological basis. The early symptoms of the disorder are similar to those typical of autism in that there is no imaginative play and a lack of physical communication (e.g. eye contact and pointing), but differs from autism in that there is no developmental delay in language acquisition. Whilst adults with AS are usually high functioning, they often maintain atypical body language and struggle to understand some subtleties of human interactions. Consequently, AS is considered to be a disorder on the autistic spectrum, the most debilitating aspect of which is the social impairments that make it difficult for this population to form meaningful relationships.

Most of the proposed cognitive explanations for autism concentrate on the deficiencies in cognitive functioning, but some symptoms of the disorder are more positive in nature. For example, some individuals possess exceptional rote memory, or display savant abilities in particular areas (music, mathematics or drawing). Explanations for these abilities range from simply reflecting less imaginative thought patterns e.g. (?), to more complicated neuropsychological explanations e.g. (?). Memory research may help establish whether differences in cognitive functioning are of a quantitative or qualitative nature.

For example, one reliable finding is that the free recall of semantically related nouns is significantly poorer in AS adults than in a control group (matched for age and IQ), whereas no such difference exists in the free recall of unrelated nouns (?). In the non-autistic population, free recall of semantically related words is higher than non-related. This phenomenon is often explained in terms of the 'Levels of Processing Model' (?), which posits that the greater elaboration involved at the encoding stage, the greater the chances of recall. Typically, semantically-linked tasks lead to greater recall of words than rhyming tasks, which in turn lead to greater recall than tasks involving structural questions (typeface etc.). However, results from a study by (?) appears to be at odds with this model of processing, at least with regard to individuals on the autistic spectrum. They found that in three separate encoding tasks on memory tests, a high functioning autistic group displayed enhanced phonological processing. They suggested was due to a low-level phonological bias in this population.

Further memory studies have identified selective impairments in some episodic memory skills in this population in comparison to a control group (?) and (?). Episodic memories are recollections of personal events, which can often be recalled in the context of time and location. Consequently, these findings could carry important significance with regard to the difficulty with communication in this population, and so establishing possible reasons for this deficit could help to build a support network to aid and improve social skills. One possible explanation for the differences in episodic memory functioning is the 'alternative processing' account (?). Rajaram distinguished between fluent processing (repetitive processing of stimuli of the same modality) and distinctive processing (processing which stimulates attention to meaning) and proposed that episodic memories are a product of the latter. If this is indeed the case, the type of memory or recall may well be influenced by the way that material is learned, or rehearsed, so again the lower-level processes may hold an important key.

The main aim of my research is to address possible reasons for the deficit in episodic memory functioning by comparing the underlying cognitive processes in AS and non-AS individuals. It is hoped that establishing potential differences in working memory processes will help to understand the relationship between the positive and negative symptoms of the disorder, as well as furthering the understanding of the social problems.

2 Experiment 1

12 adults with AS and 12 controls (matched for age and IQ) took part in a free-recall memory task on 3 types of word lists (categorised nouns, non-categorised nouns and rhyming words). During each trial, participants were asked to rehearse out loud so that rehearsals could be recorded and scored. Comparisons were made on rehearsal patterns and correctly recalled words. AS individuals displayed a slight tendency for a more repetitive style of verbal rehearsal compared to controls, but chi squared tests revealed that there was no significant association between the rehearsal style and the population (AS or Controls) for any of the different types of word lists. The AS group recalled fewer words than controls for both the categorised ($t = (22) - 2.190,\ p = 0.039$) and the rhyming ($t = (22) - 2.599,\ p = 0.016$) word lists, indicating that there may be less explicit awareness of phonological as well of semantic connections in this group. The latter finding is in contradiction with the suggestion of a phonological bias in this population (?). Further studies are planned to include visual as well as verbal stimuli, to and try and account for this apparent anomaly and further explore rehearsal style preferences in the Asperger population.

References

Bowler, D. M., Gardiner, J. M., & Grice, S. J. (2000). Episodic memory and remembering in adults with asperger's syndrome. *Journal of Autism and Developmental Disorders, 30(4)*, 295-304.

Bowler, D. M., Matthews, N. J., & Gardiner, J. M. (1997). Asperger's syndrome and memory: Similarity to autism but not amnesia. *Neuropsychologica, 35*, 65-70.

Craik, F., & Lockhart, R. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11*, 671-684.

Ericsson, & Faivre. (1988). What's exceptional about exceptional abilities. *The Exceptional Brain: Neuropsychology of talent and special abilities.*

Mottron, L., Morasse, K., & Bellville, S. (2001). A study of memory functioning in individuals with autism. *Journal of Child Psychology and Psychiatry, 42(2)*, 253-260.

Rajaram, S. (1996). Perceptual effects on remembering: recollective processes in picture recognition memory. *Journal of Experimental Psychology: Learning, Memory and Cognition, 22(2)*, 365-377.

Rimland, B., & Fein, D. (1988). Special talents of autistic savants. *The Exceptional Brain: Neuropsychology of talent and special abilities.*

Tracking a Beat: Cascades of Synchrony in Cortical Circuits

Sampsa Sojakka

sampsas@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Time is critical for life. Numerous ecologically relevant signals have a rich temporal structure that neural circuits must process in real time. Prevalent examples include auditory signal processing and movement detection in the visual domain. This presents two main requirements for models of cortical microcircuits:

1. *Spatiotemporal integration* - transient temporal data must be processed by a network of heterogeneous nodes (neurons) and retrieved to drive behaviour at timescales that are significantly slower than the timescales of neuronal performance (Harvey, 1997).

2. *Synchronisation* - such a network must be able to spontaneously establish and maintain coherence in the face of environmental noise without a centralised structure acting as a coordinator (Strogatz, Mirollo, & Matthews, 1992).

My research is aimed at investigating how cascades of synchrony in such integrative circuits arise and how they can be utilised to produce behaviour over a large range of timescales. This paper provides a brief overview of the motivations behind this work and discusses directions pursued by current research.

2 Integration

Autonomous robotics has traditionally centred on constructing models based on classical attractor neural networks or various feed-forward nets where information is encoded in the stable states of the system (Hopfield, 1982). Unfortunately the dependence on fixed states makes such network architectures brittle to lesions and synaptic decay and hence biologically unrealistic. Several enhancements have been proposed in the neuroscience literature, usually based on mechanisms modifying synaptic efficacy either through long-term potentiation akin to Hebbian learning (Hebb, 1949) or synaptic depression (Chialvo & Bak, 1999). While this has increased the robustness of controllers, a fundamental problem with time in such models still remains. Namely, since the timescales of activity of most classes of neurons are several orders of magnitude faster than the timescales of observed behaviour of the organism, it is increasingly unlikely that the activation of a single neuron is able to act as storage for a temporal pattern. The firing patterns simply decay too quickly to hold transient data long enough to drive behaviour.

The recent Liquid State Machine (LSM) model developed by Maass et al. has shown promise in circumventing the stable state problem (Maass, Natschlager, & Markam, 2002). This model hypothesises that the cortical microcolumn consists of stereotyped recurrent circuits of integrate-and-fire neurons connected randomly according to a few parameters. Such a network is not dependent on stable states since all events are transient (i.e. the network has fading memory, it forgets) but has been shown to have universal computational power. Essentially an LSM performs integration by projecting its input time series into a higher dimensional space so that a linear readout may suffice to make a classification. The choice of 'liquid' is not limited to spiking neural networks either, a point that was highlighted by recent work where an LSM was constructed using real water and used to solve complex pattern recognition tasks (Fernando & Sojakka, 2003).

3 Synchronisation

While the LSM model allows us to build systems capable of recognising temporal patterns at various timescales, it does not immediately suggest how such neuronal soups are able to maintain coherence. Classical models are of little use here since they are usually based on a central clock, a feature that is less ubiquitous in biological systems. However, examples of synchronisation in nature abound, ranging from the flashing of fireflies

to circadian pacemaker cells, lasers and superconducting Josephson junction arrays (for a review see (Strogatz & Stewart, 1993)). All of these cases of synchrony have been traced back to a phenomenon known as oscillator coupling.

The first mathematical model of coupled limit-cycle oscillators was developed by Winfree in (Winfree, 1980) and subsequently refined by Kuramoto to show that partially synchronised states appear as the diversity of nodes in the network decreases, eventually leading to perfect synchrony through phase locking (Kuramoto, 1984). In order to keep the mathematics tractable these models are based on idealised conditions where networks are fully connected, oscillators nearly identical and the coupling between them weak. Still, even with these restrictions networks display rich dynamics where synchronisation does not necessarily lead to periodic paths through state space. In other words rhythmic activity of the individual nodes may lead to complex non-rhythmic patterns of behaviour at a higher level (Ariaratnam & Strogatz, 2001).

This suggests that not only is synchrony vital for network coherence but could also provide a viable mechanism for spatiotemporal integration (Hopfield & Brody, 2000). While oscillators have been widely used in robotics to generate rhythmic behaviour, the applicability of transient synchrony as a mechanism for temporal pattern generation is still largely unexplored in robotics (Paolo, 2002). Main difficulties here centre on the complexity of the inherent dynamics of such circuits and the tracking of waves of synchrony through network lattices. Constraining network topology may prove a fruitful way to start (Strogatz & Watts, 1998).

4 Current Research

In order to construct a framework that can be used to study these questions, work is under way to simulate and build an active-dynamic walker based on networks of oscillator neurons. The aim here is to evolve walkers capable of traversing uneven terrain, a task that requires oscillator phase coherence at various timescales. This platform will be used to study how constraints in network topology affect robustness of the controller and what role synchrony plays in the production of behaviour.

References

Ariaratnam, J. T., & Strogatz, S. H. (2001). Phase diagram for the Winfree model of coupled nonlinear oscillators. *Physical Review Letters, 86*(19), 4278-4281.

Chialvo, D. R., & Bak, P. (1999). Learning from mistakes. *Neuroscience, 90*(4), 1137-1148.

Fernando, C., & Sojakka, S. (2003). Pattern recognition in a bucket. *To appear in ECAL '03.*

Harvey, I. (1997). Cognition is not computation; evolution is not optimisation. *Proceedings of the 7th International Conference on Artificial Neural Networks.*

Hebb, D. O. (1949). *The organization of behaviour.* Wiley, New York.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci., 79*, 3168-3171.

Hopfield, J. J., & Brody, C. D. (2000). What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. *Proc. Natl. Acad. Sci., 89*, 1282.

Kuramoto, Y. (1984). *Chemical oscillations, waves and turbulence.* Springer, Berlin.

Maass, W., Natschlager, T., & Markam, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation, 14*, 2531-2560.

Paolo, E. A. D. (2002). *Evolving robust robots using homeostatic oscillators* (Tech. Rep.). CSRP 548, University of Sussex.

Strogatz, S. H., Mirollo, R. E., & Matthews, P. C. (1992). Coupled nonlinear oscillators below the synchronization threshold: Relaxation by generalized Landau damping. *Physical Review Letters, 68*, 2730.

Strogatz, S. H., & Stewart, I. (1993). Coupled oscillators and biological synchronization. *Scientific American, 269*(6), 102-109.

Strogatz, S. H., & Watts, D. J. (1998). Collective dynamics of 'small-world' networks. *Nature, 393*, 440-442.

Winfree, A. T. (1980). *The geometry of biological time.* Springer, New York.

Cognitive Linguistic Analysis of Scientific Theory Formation in Genetics

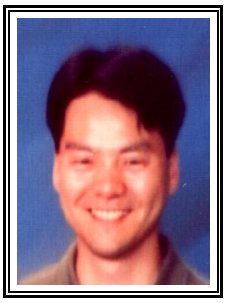

John Sung

jjs21@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

The importance of science in every facet of our society cannot be understated and the understanding of its development process would be crucial in improving our society. The cognitive linguistic approach is ideal for this task, because it accounts for cognitive processes by investigating linguistic phenomena, and because of the extensive scientific documentation available as data. The conceptual metaphor theory attempts to correlate human cognitive processes with linguistic phenomena. Thus, one may attempt to describe the scientific developmental process by examining the original phenomena, observations made, formulated theories and scientists' documentation of what they've sensed, perceived, and conceptualized in journals, books, etc. These documents are the linguistic evidence for the cognitive activities of scientists who generated those documents. This research will study the linguistic evidence of Genetics such as DNA IS A THREAD and DNA IS WRITTEN TEXT metaphors, and their developments to shed light on the scientific development process.

2 Cognitive Linguistics

Cognitive linguistics attempts to understand how humans conceptualize by examining the linguistic phenomena. Unlike the traditional linguistics where linguists start with literal language and treat figurative language as exceptions, Cognitive Linguistics start with figurative language. Furthermore, Cognitive Linguists maintain that most of our language usage are figurative in nature. For example, when we refer to time we tend to think of time as space or motion in space. In (1) below, the "hard times" are conceptualized as something that is behind us, even though we can not turn our heads, look behind us and see the past.

- (1) The hard times are behind us.

In (2) and (3) below, time is conceptualized as motion in space. We may either conceptualize time as moving and we're merely observing its motion, as in (2) or we are moving and time is some stationary reference point, as in (3).

- (2) Time is flying by.
- (3) We're coming upon Christmas.

The Conceptual Metaphor Theory (CMT), as presented in (Lakoff & Johnson, 1980) and (Lakoff & Johnson, 1999), proposes that the human conceptual system is metaphoric in nature. Furthermore, the linguistic data support the hypothesis that conceptual metaphors are systematic. The conceptual metaphor TIME ORIENTATION is used when we refer to past, present and future events as shown in (1). Furthermore, CMT proposes that these metaphors are not just singularities in our language, but are systematic in nature.

- (4) That's all behind us now.
- (5) We're looking ahead to the future.
- (6) He has a great future in front of him.

From the linguistic data (4) - (6) above, we may uncover the systematicity by describing the mapping of the TIME ORIENTATION conceptual metaphor as shown below. The "concept" on the left is mapped onto the concept on the right. So, the Current Location in the space domain represents The Present in the time domain.

- Current Location $\rightarrow$ The Present
- Space In Front $\rightarrow$ The Future
- Space In Back $\rightarrow$ The Past

In this manner, we map the space orientational concepts to the concepts of time. These linguistic data and other linguistic data indicate that the TIME ORIENTATION and other conceptual metaphors are systematic. Moreover, Cognitive Linguistics proposes that this systematicity reflects the human conceptualization.

3 Genetics

Genetics is a recently developed scientific field that has many moral and ethical issues. Also, its high potential to change our everyday life can not be disputed. Words such as Dolly the cloned sheep, human cloning, genetic therapy are well known by the public and the academics. This public awareness of Genetics and the extensive documentation of experiments, hypothesis and theories as the result of the scientific process, allows us to trace the conceptual development of Genetics by examining the scientific papers, books, etc. Furthermore, this extensive documentation of the scientific process allows us to use it as linguistic data. Two of the major conceptual metaphors used in Genetics that we'll be studying are DNA IS THREAD (7) - (10) and DNA IS WRITTEN TEXT (11) - (13) from (Weaver & Hedrick, 1999).

- (7) Tightly coiled threads of DNA.
- (8) Each strand of DNA consists of repeating nucleotide units
- (9) DNA takes the form of a highly regular double-stranded helix,
- (10) In order to cut and paste desired DNA fragments into vectors,
- (11) When genes are expressed, the genetic information (base sequence) on DNA is first transcribed
- (12) The genetic information of an organism can be stored in one or more distinct molecules
- (13) The change can be to insert a new nucleotide, to delete an existing one, or to change one nucleotide into another.

From this data, we can see a pattern in the way the two conceptual metaphors are being used. Whenever, the DNA is referred as thread, strand or fragment, adjectives and verbs that we might use for a physical thread is used, such as "tightly coiled", "takes the form", "cut", "paste", etc. This seems to imply that DNA IS THREAD metaphor is used when one is describing the structural aspect of the DNA. Similarly, the DNA IS WRITTEN TEXT seems to describe the informational content of the DNA, i.e. the particular order of the base sequences. Words such as "expressed", "genetic information", "transcribed", "insert", "delete", etc. seems to highlight the informational aspect of the DNA. Furthermore, these two metaphors are used in conjunction to provide a model of the DNA with a structural and informational properties. The words "cut" and "paste" in sentence fragment (10) could be interpreted as either cutting or pasting text or thread. Thus, many of the verbs used in one conceptual metaphor could have a valid interpretation in the other. Studying these types of linguistic phenomena is the main purpose of this research. Furthermore, we intend to understand the structure of conceptual metaphors being used in Genetics in depth and trace the historical development of these conceptual metaphors. This trace would be fueled by the extensive linguistic data available from the technical journals, text books, magazines, etc. The purpose of this endeavor is to understand how language is used in Genetics and science in general and its historical developmental processes for the betterment of the human race.

4 Conclusion

The cognitive linguistic theories, such as CMT, will be used to analyze the metaphoric systems and the conceptualization processes of scientific theory formation in Genetics. Because CMT allows us to analyze the linguistic phenomena, and gain insight into the human conceptual system and conceptual processes, it and other Cognitive Linguistic theories are perfect for investigating scientific theory formation in Genetics.

References

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by.* Chicago: University Chicago Press, US.

Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh.* New York: Basic Books.

Weaver, R. F., & Hedrick, P. W. (1999). *Genetics second edition.* Dubuque: Brown.

Vocal Repertoire Analysis using Artificial Neural Networks

Aisha Thorn

aishat@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

There are numerous reported cases of similar species and even separate groups of animals in the same species that occupy the same habitat, but nevertheless have quite different types of vocal repertoire. The social, environmental and behavioural reasons for these differences are indicative of the causes of vocal complexity and could ultimately reveal reasons for human language evolution.

Marler (Marler, 1970) performed extensive analysis on the vocalisations of two species of colobus with similar physiology, occupying the same forest habitat. He found the black and white colobus to have seven discrete call types and the red colobus to have a graded system surrounding three basic call types. It is likely to be different social structures that cause differences in repertoire structure. Black and white colobuses form small territorial groups and probably use discrete call-types for the ease of recognition it provides in their visually-deprived, between-group communication. Red colobuses form larger groups with emphasis on intra-group communication and can therefore afford to use graduated calls often accompanied by gesture.

There are currently no standard techniques for establishing and reporting the number of stereotypical calls in a repertoire, and the acoustic variation within and between call-types. I am developing neural network techniques to estimate repertoire size, where this is appropriate; to obtain a measure of the discreteness/gradedness of the repertoire; and to measure of the extent of the acoustic space that a call or a set of calls uses. The techniques are intended to be general enough to extend to any species, although the initial trials are being conducted on macaque vocalisations.

One such techanique is based on a self-organising map (SOM) (Kohonen, 1981) and has two separate stages. In the training phase, a corpus of calls from the chosen species is pre-processed using for example auto-correlation to extract source (glottal) characteristics, linear prediction to extract filter (vocal-tract) characteristics or an auditory filterbank. The spectra obtained from the pre-processed calls are presented as inputs to a SOM that updates its weights according to a dynamic warping algorithm (Rabiner, Rosenberg, & Levinson, 1978) ensuring invariance to slight frequency shifts. This initial phase is computationally expensive but only needs to be done once for each species.

In the testing phase, calls are pre-processed in the same manner and when presented to the network, each spectral frame is matched with the single most similar node and over the course of the call a trajectory is generated through the spectral space on the map (figure 1). This is recorded for later comparison with the trajectories of other calls from the same species. These comparisons are also made using the dynamic warping technique, to ensures invariance to temporal shifts or contortions. The result is a matrix of distance measures between each call and each other call, in a particular spectral domain. These can be interpreted as a measure of the gradedness of the repertoire and provide probability estimates of different repertoire sizes.

As well as providing these measures it is hoped that the techniques used will make it possible to synthesize sound signals which fall into a (natural) continuum of sounds that lie between two different calls useful for testing hypotheses about categorical perception in playback studies. The fact that the SOM is created using a corpus of only the species calls limits the acoustic space used to synthesize the calls and ensures that they make use of only the types of sounds that the animals can actually produce themselves.

It is plausible that at some level of abstraction the SOM mimics the structures and processes by which cortical structures in the brain learn and recognise patterns and although it remains uncertain exactly how these processes occur in the auditory system of mammals, there is evidence that the primary auditory cortex self-organises to form a map of frequency and spectral space that is continually modified by experience (Bakin & Wein-

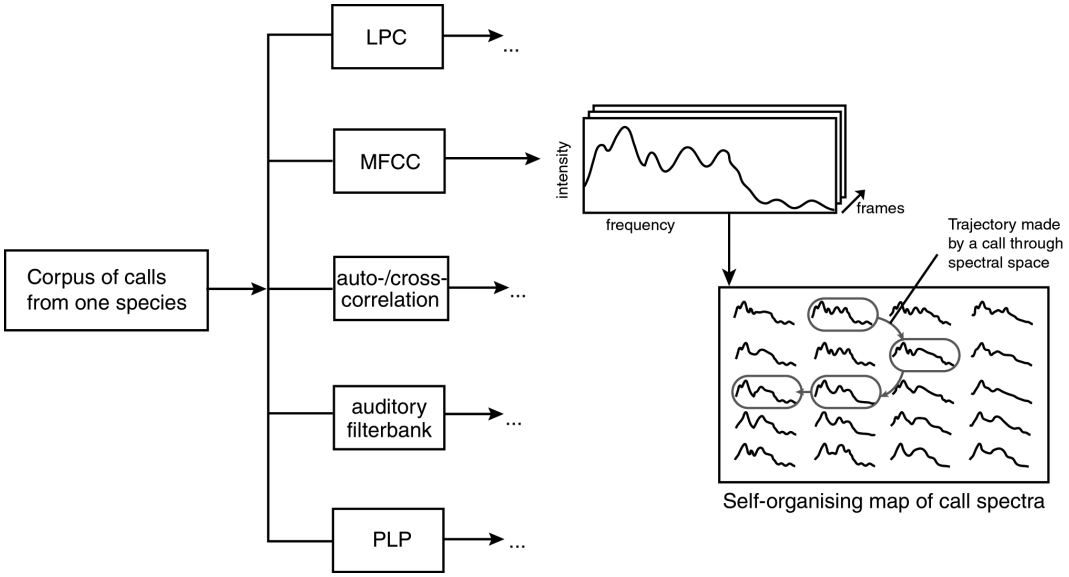

Figure 1: In phase 1 the corpus of calls is pre-processed and resulting spectra are presented as training inputs to the SOM in random order. In phase 2, after the SOM has been created, the test calls (not necessarily contained in the original corpus) are pre-processed in the same way and their spectra presented to the SOM in order. The trajectory formed by each call is recorded for later analysis and comparison with other calls' trajectories.

berger, 1990). A further intended modification to the technique is therefore to incorporate map modification during the testing phase, thus mimicking cortical plasticity.

References

Bakin, J. S., & Weinberger, N. M. (1990). Classical conditioning induces cs-specific receptive field plasticity in the auditory cortex of the guinea pig. *Brain Res.*, *536*, 271-286.

Kohonen, T. (1981). Automatic formation of topological maps of patterns in a self-organizing system. In E. Oja & O. Simula (Eds.), *Proceedings of 2SCIA, scand. conference on image analysis* (p. 214-220). Helsinki, Finland.

Marler, P. (1970). Vocalizations of east african monkeys. i. red colobus. *Folia Primatol.*, *13*, 18-91.

Rabiner, L. R., Rosenberg, A. E., & Levinson, S. E. (1978). Considerations in dynamic time warping algorithms for discrete word recognition. *IEEE Trans., ASSP-26 (6)*, 575-582.

Modelling Path Integration using an Evolutionary Robotics Approach

Robert Vickerstaff

robertvi@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Path integration (PI) is a navigational strategy widely used in animals, being particularly highly developed in the desert ant *Cataglyphis fortis.* Experiments show that this ant can forage on the flat, featureless saltpans it inhabits without the aid of land marks or pheromone trails, and return on a straight course to its nest entrance whenever it discovers food (Wehner, Gallizzi, Frei, & Vesely, 2002). The only method available to it is PI. This is a process whereby an animal continuously integrates its velocity (i.e. direction and speed) in order to calculate its current location. This can be used to return to the starting point after a circuitous journey and relies only on knowledge of the animal's orientation (compass direction) and speed (or distance travelled).

2 Background

Mathematically, the integration process required is easily expressed in either a polar or Cartesian coordinate system (Maurer & Seguinot, 1995). Very little is known about how this process is carried out in the nervous systems of animals, and this is likely to remain so for sometime given the difficulty of monitoring the brain of a moving animal. Another way to investigate how PI might be implemented in the brain is the use of neural networks - computer simulations of simple neuron-like elements interacting with each other in a network structure. Two hand-designed neural network architectures capable of carrying out PI have been proposed. One of these (Wittmann & Schwegler, 1995) carries out the integration process exactly without introducing any error in the animal's estimate of its location. The other (Hartmann & Wehner, 1995) uses an approximate method designed to mimic systematic PI errors known to occur in many species under certain conditions (Muller & Wehner, 1998). A key question raised here is whether animals' nervous systems are capable of performing exact PI, or whether, because of the energetic cost of a more complex brain, they have evolved to perform only an approximation of it, as is suggested by the presence of systematic errors. An alternative explanation for these errors is that they have an adaptive benefit by ensuring that the animal recrosses its outward path as it returns home, thereby increasing the chances of recognising familiar landmarks (Wittmann & Schwegler, 1995).

My work is aimed at producing alternative neural network models of PI, using an Evolutionary Robotics (ER) approach. In common with much work in ER, I am using a Genetic Algorithm (GA) (Goldberg, 1989) to evolve neural network controllers for a simulated agent that moves around a virtual arena, navigating by PI. The GA works on a population of candidate controllers, each of which is assessed for its ability to perform accurate PI. The most successful controllers are duplicated, mutated and used to replace less successful ones, and thereby, over time, the quality of the controllers can be improved.

3 Current work

I am working with an agent in a simulated two dimensional arena which must first visit a sequence of beacons placed at random locations in the arena. After visiting the last beacon it is required to return to its starting location using PI. The only sensory inputs available are two visual sensors, and compass and speed sensors. This is a difficult task since, when evolving the initial beacon approach behaviour, the network has no reason to pay attention to its speed and compass sensors. In order to solve this problem I have adopted a 'scaffolding' approach, first evolving the agent to signal its spatial coordinates using two of its neurons, before selecting for the ability to return to the nest. I will make full use of the flexibility of the GA by comparing the performance of different coordinate systems and styles of neural network.

References

Goldberg, D. E. (1989). In *Genetic algorithms in search, optimization, and machine learning.* Addison-Wesley, Reading, Mass, US.

Hartmann, G., & Wehner, R. (1995). The ant's path integration system: a neural architecture. *Biol. Cybern.*, **73**, 482-497.

Maurer, R., & Seguinot, V. (1995). What is modelling for? a critical review of the models of path integration. *J. Theor. Biol.*, **175**, 457-475.

Muller, M., & Wehner, R. (1998). Path integration in desert ants, *Cataglyphis fortis. Proc Natl Acad Sci*, **85**, 5287-5290.

Wehner, R., Gallizzi, K., Frei, C., & Vesely, M. (2002). Calibration processes in desert ant navigation: vector courses and systematic search. *J. Comp. Physiol.*, **188**, 683-693.

Wittmann, T., & Schwegler, H. (1995). Path integration - a network model. *Biol. Cybern.*, **73**, 569-575.

Automatically Acquiring Semi-Sense-Tagged Corpora

Xinglong Wang

xw20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

The problem of assigning semantical senses to the words in an open text, known as Word Sense Disambiguation (WSD), is central to many Natural Language Processing applications. According to previous research, supervised methods have been considered achieving better results. The supervised disambiguation process is based on the probability of occurrence of a particular sense in a given context. The context is determined by linguistic features such as parts of speech of surrounding words, keywords, syntactic relations, collocations, etc. Supervised methods consist usually of two phases:

1. a training phase, in which features are learned by using various algorithms on correctly sense-tagged corpora.

2. a testing phase, in which the features acquired in the previous step are used to determine the most probable sense for a particular word.

The disambiguation accuracy is strongly affected by the quantity and quality of the corpora used in the supervised disambiguation process. Unfortunately, large sense-tagged corpora are rare and manually producing them is extremely expensive and time-consuming. This problem is the so-called knowledge acquisition bottleneck.

In this short paper, I describe the work I intend to do, attempting to open the bottleneck, by using multilingual methodology to automatically acquire sense-tagged corpora from the World Wide Web.

2 Resources on the World Wide Web

The traditional way of manually creating sense-tagged corpora is slow and expensive. Not surprisingly, the huge amount of text available on the Internet has become a great potential resource for language research. A lot of work has been carried out on exploring it. For example, Mihalcea et al. (Mihalcea & Moldovan, 1999) presented a method that enables the automatic acquisition of sense-tagged corpora, based on the information found in WordNet (Miller, Beckwith, Fellbaum, Gross, & Miller, 1990), a Machine Readable Dictionary, and on the very large collection of texts gathered from the Internet using existing search engines. The idea is to firstly obtain monosemous synonyms or the gloss of a given word (W) for which a corpus is to be acquired, from the WordNet, and then using them to query a search engine, such as AltaVista, and finally to gather and refine retrieved snippets of text which should have similar contexts with W and could become sense-annotated corpora for W. Their work has shown that very large sense-tagged corpora can be automatically generated using the Web and even though the corpora might be noisy, still it is much easier and less time consuming to check an already existing tagged corpus for correctness, then to start tagging free text from scratch.

Another property of the Web is its multilinguality. Although there is no doubt that English is dominant on the Web, research by Grefenstette and Nioche (Grefenstette & Nioche, 2000) has shown that non-English languages are growing at a faster pace than English is. The increasing amount of text in other languages on the Web makes it a potential resource to do multilingual research and to acquire parallel text, which is also useful but rare.

3 Are Two Languages Better Than One?

Using one or more source languages as an aid to study a target language has become a recent trend in the NLP community. Such techniques take advantage of the fact that cross-language lexicalisations of the same concept tend to be consistent, but mappings between word forms and senses are different from one language to another.

Let's look at a simple application using the multilingual paradigm. If we have an aligned French-English corpus, we can easily acquire correct senses for ambiguous English words accord-

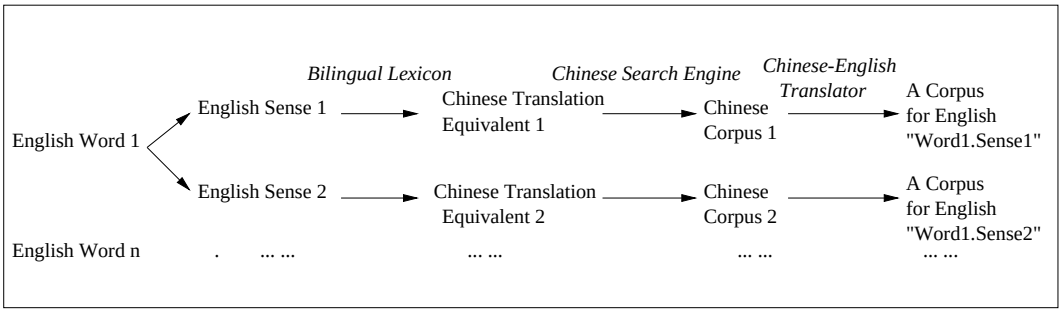

Figure 1: Abstract Process of Automatically Acquiring Semi-Sense-Tagged Corpora

ing to their French translation equivalents. For example, to sense-tag the English word *sentence*, we search where its French translations appear in the corpus. There are two translations: *peine* or *phrase*. Then we assign *sentence* the judicial sense if it is translated as *peine* and the grammatical sense otherwise.

It sounds like a good method. But unfortunately, parallel corpora are rare and aligned bilingual corpora are even rarer (knowledge acquisition bottleneck again!). Therefore, such a technique wouldn't work well in reality.

4 A Proposal

Unsupervised methods, which learn rules without using costly sense-tagged training data, offer an alternative way to open the knowledge acquisition bottleneck. In this section, I give a short description of my proposal for automatically creating semi-sense-tagged corpora. A short version of this algorithm includes the following steps: (see figure 1)

1. Translate each sense of ambiguous English words into Chinese using a bilingual dictionary.

2. Query a Chinese search engine using the Chinese translation equivalent of each English sense. Every snippet of text retrieved by the search engine should contain the query which can be viewed as a Chinese realisation of an English concept (an English sense).

3. Translate the Chinese text back to English either word by word using a bilingual dictionary or paragraph by paragraph using a machine translatoin software and then we should be able to get an English corpus for each English sense.

The corpora generated can be noisy and there is little chance to learn English syntactic rules from them. However, we believe a very large semi-sense-tagged corpus for each English sense will still be useful.

References

Grefenstette, G., & Nioche, J. (2000). Estimation of English and non-English language use on the WWW. In *Proc. riao (recherche d'information assistee par ordinateur).* Paris, France.

Mihalcea, R., & Moldovan, D. I. (1999). An automatic method for generating sense tagged corpora. In *Proc. of the* 16th *conference of the american association of artificial intelligence.*

Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., & Miller, K. J. (1990). Introduction to WordNet: An on-line lexical database. *Journal of Lexicography, 3(4)*, 235-244.

Children and adolescents use of self-presentation tactics: How does this relate to their peer relationships?

Dawn Watling

dawnwa@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

Self-presentational behaviour - behaviour designed to shape the impression that others form of the self (Goffman, 1959) - is an important aspect of social interaction in childhood and adolescence as well as adulthood. Self-presentational tactics have been classified into two general categories: assertive and defensive tactics (Lee, Quigley, Nesler, & Tedeschi, 1999). Assertive tactics (AT) are used to help individuals build the impression that they want their audience to hold; for instance, when children wish to be liked by their peers, they may flatter them or express agreement with their attitudes (a form of ingratiation). In contrast, defensive tactics (DT) are used when individuals believe that their desired identity is threatened; for instance, children may try to justify why they are about to do something that might be judged negatively by others (a form of disclaimer).

2 Theoretical background

Childhood is a time where children are concerned with gaining social approval and they develop a concern for social evaluation (Parker & Gottman, 1989). In contrast, adolescence has long since been discussed as a difficult time in one's life - trying to discover one's own identity, fit in with peers, and achieve independence from one's parents.

Research into adolescent behaviour has shown that this is a time where individuals are very interested in achieving and protecting their self-image (Berger, 2001). From the aforementioned research it is expected that children will use more AT than DT in an attempt get others to like them and form a positive impression of them. Additionally, it is expected that adolescents will use DT to a greater extent than AT in an attempt to maintain their desired identity.

3 Recent research

Recent research has examined children's and adolescent's self-rated usage of four different self-presentational tactics, two assertive and two defensive:

1. ingratiation (AT), which is used when an individual wishes to appear likeable;
2. self-promotion (AT), which is used when an individual wishes to appear competent in a particular area or with regard to particular skills;
3. excuse (DT), which is used after an individual did poorly on some task or performed some behaviour that might be seen negatively; and
4. disclaimer (DT), which is used before an individual believes that he or she will do poorly on some task or will perform some behaviour that might be seen as wrong.

The results of the research strongly suggest that there are between- and within-age group differences in children' use of tactics for managing the impressions they make on others. Furthermore, this variability is related to differences in peer relations among boys and girls. The research provides new insights into the types of social behaviour that predict successful peer interaction for boys and girls of different ages.

References

Berger, K. (2001). In *The developing person through childhood and adolescence.* Worth publishers, New York.

Goffman, E. (1959). In *The presentation of self in everyday life.* Penguin books, UK.

Lee, S., Quigley, B., Nesler, M., & Tedeschi, A. (1999). Development of a self-presentation

tactics scale. *Personality and Individual Differences, 26,* 701-722.

Parker, J., & Gottman, J. (1989). Social and emotional development in a relational context: Friendship interaction from early childhood to adolescenc. *Peer Relationships in Child Development,* 95-131.

Why Not to Throw Out Your Thesaurus

Julie Weeds

juliewe@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

Introduction A thesaurus groups together synonyms and other related concepts. For example, the words *courgette* and *zucchini* might be listed as synonyms and specified as types of *vegetable*, along with other vegetables such as *carrot*, *mushroom* and *aubergine*. Further, such relationships between words might be automatically derived from text by considering the contexts in which the words occur (Weeds, 2002). For example, we would expect all *food* words to appear as the object of the word *eat* in a large enough corpus of text.

In this paper, we turn to the question of why knowing these semantic relationships between words is useful. In everyday life, people may turn to a thesaurus when they can't quite think of the right word or when they want to vary the word they use for a particular concept. Here, however, we will present a brief survey of some other uses which computational linguists have found for thesauruses.

Text Simplification The PSET project (Carroll, Minnen, Canning, Devlin, & Tait, 1998) involved simplification of English text for aphasic readers by substituting rare words or expressions with their more common synonyms. For example, "I was quaffing amber brew at a drinking establishment" might become, "I was drinking beer at a pub".

Collocation Extraction Collocations can be a nightmare for foreign language learners. For example, how does one know that we say "fire alarm" rather than "flame alert" (Pearce, 2002) when the second surely means the same as the first? In order to automatically extract lists of collocations, Pearce (2002) substitutes each word in a candidate pair with its synonyms and counts the number of occurrences of each combination in a corpus. Word pairs which occur significantly more than would be expected (based on the relative frequencies of the use of each word for each concept) are deemed to be collocations.

Term Expansion Thesauruses can be used to broaden a search to return documents containing semantically similar words as well as exact matches. For example, recipes containing *zucchinis* could be returned when a search is made for *courgette* recipes. Similarly, a search for *hotels* in a certain area could be automatically expanded to include *hostels*, *B&Bs*, *motels* and *campsites*.

Spelling Correction Current spell-checkers are good at checking each word to see if it exists in the dictionary but not so good at spotting real-word spelling errors. For example, in the phrase "there was an extensive display of guns, canons and other military paraphernalia", a naive spell-checker would not spot that *canon* should be spelt *cannon* since *canon* exists in the dictionary (with meanings relating to *law* and *religion*). However, a more sophisticated spell checker e.g. (Budanitsky & Hirst, 2001) might note that *canon* is very similar in spelling to *cannon* and that *cannon* is closer in meaning than *canon* to *gun*.

Prepositional Phrase Attachment Ambiguity Resolution If we consider the sentence, "He shot the woman with a pistol," we probably imagine a situation where a man uses a pistol to shoot a woman. However, in, "He shot the woman with a Gucci handbag," we imagine a situation where a woman with a Gucci handbag gets shot by a man. Resolution of this problem requires knowledge of what types of things can be used to shoot with. Fortunately, these things, types of guns and other shooting implements, form a semantic class i.e. are found together in a thesaurus.

Compound Noun Interpretation In the interpretation of compound nouns, it is necessary to determine the implicit relationship between the two nouns. For example, in the sentence, "We are

concerned about terrorist activities in the Middle East," we are referring to the activities *of* terrorists, whereas in, "We are concerned about gun crime in the United Kingdom," we are referring to crime *with* guns. With the aid of a thesaurus, we can learn about or establish these relationships over semantic classes (such as *people* and *weapons*) rather than for each individual lexical item.

Conjunction Scope Reading the sentence, "The oldest boys and girls led the march," we infer that it is the oldest girls who accompany the oldest boys at the front of the march. However, in, "He collects old coins and computers," we probably imagine that he collects both old and new computers, it is just the coins that are old. This distinction can be made by looking at the similarity between the two nouns on either side of the conjunction. *Boys* and *girls* are semantically very similar whereas *coins* and *computers* are semantically further apart.

Associative Anaphora Resolution As discussed in Meyer & Dale, 2002, an associative anaphor is a definite referring expression used to refer to an entity not previously mentioned in the text. In the example, "A bus came round the corner. The driver had a mean look in her eye," we are likely to imagine that the *driver* in the second sentence refers to the driver of the *bus* in the first sentence. It may be that the associated noun entities have a direct relationship in the thesaurus, e.g. if the word *vehicle* was used to refer to a previously mentioned *bus*. Alternatively, they may have a relationship which can be established over semantic classes. For example, we can learn that vehicles have drivers and as *bus* is a type of *vehicle* in the thesaurus, it therefore has a driver.

Topic Identification Texts tend to be cohesive and thus will contain many words all relating to the main topic or theme of the text. Accordingly, if we identify which words in the text are semantically related using a thesaurus and then determine which set of related words is the largest or covers most of the text (Silber & McCoy, 2002), we might be able to identify the topic of the text. Further, such a technique often forms the first phase in the automatic summarization of a text.

Text Segmentation Also related to topic identification is the issue of text segmentation. Knowing that texts tend to be cohesive, natural boundaries in text can be found by considering where the topic shifts i.e. where the set of semantically related words which best covers the text changes.

Word Sense Disambiguation Many words have more than one meaning or sense. For example, the word "plane" has senses relating to a *flat surface*, a *tool*, a *tree* and an *aeroplane*. Knowing which sense of a word is intended is very important particularly in machine translation (where the target language may have multiple words for a single word in the source language). In the context of a sentence, such as, "The plane circled the airport for an hour," it is fairly obvious to a human reader which sense of *plane* is intended. The two clues in the sentence are the words *circled* and *airport*. *To circle* is a verb of movement and therefore its subject needs to be something which can move such as a vehicle. The relationship between *airport* and the *aeroplane* sense of plane is an associative one (similar to that discussed in the section on associative anaphor resolution.)

Conclusion By illustrating a number of different applications, we have demonstrated why it is useful to learn or store semantic relationships between words, or, in other words, why not to throw out your thesaurus.

References

Budanitsky, A., & Hirst, G. (2001). Semantic distance in WordNet: An experimental, application-oriented evaluation of five measures. In *NAACL-01*.

Carroll, J., Minnen, G., Canning, Y., Devlin, S., & Tait, J. (1998). Practical simplification of english text to assist aphasic readers. In *AAAI-98*.

Meyer, J., & Dale, R. (2002). Using the WordNet hierarchy for associative anaphora resolution. In *SemaNet02*.

Pearce, D. (2002). Flame alerts and fire alarms. In J. Weeds (Ed.), *The 15th whitehouse papers*.

Silber, H., & McCoy, K. (2002). Efficiently computed lexical chains as an intermediate representation for automatic text summarization. *Computational Linguistics, 28*(4).

Weeds, J. (2002). The meaning of taflig: Distributional similarity for rare words. In D. Pearce (Ed.), *The 14th whitehouse papers.*

On Autism, Social Cognition and Perception-Action

Hanneke De Jaegher

hdj20@cogs.susx.ac.uk

School of Cognitive and Computing Sciences

1 Introduction

> Ricky is strolling through the shop. In the candy section she halts. She sees the Mars-bars. Her hand sways out, she takes two bars and is about to slip them into her jacket pocket. Then she feels someone is looking at her. She turns and her eyes meet the gaze of an older girl standing a few meters away. She swiftly repairs her movement. She takes the candy bars to the counter, pays for them, but as she walks past the girl on her way out, she gives her a sinister look and sticks her tongue out.

What happens when we interact with others? How do we understand each other often without so many words? What is it that enables us to repair misunderstandings in a conversation (in the broadest sense of the word), when it goes wrong, as easily and smoothly as we often do?

Why these questions? Because there are people for whom interaction is far from easy. And because for a group of them, perception and thinking are different as well; individuals who have autism. What is specific about the difficulties with social interaction that people with autism and those who interact with them experience? And is there a connection between their communication difficulties and their distinctive ways of thinking?

2 What is Autism?

We all have our bad days, when we would have rather stayed in bed and not faced anyone. People with autism however, face a constant and continuous struggle with understanding others and the world organised by these others. Autism is a disorder characterised by difficulties in interpersonal interaction and communication, as well as in fluency and flexibility of thought, and by a limited range and scope of interests and imagination. People with autism come across as rigid and awkward in interaction. But not only in social situations, also when they have to do something as apparently simple as going for bread at the baker's, they get stuck on the complexity and range of the possibilities for solving this problem. Should I go now, or only when the clock strikes exactly four? Should I cross the street exactly in front of the door through which I go out, or in front of the baker, which is a few houses down the street, or should I walk on until I see a pedestrian crossing, which might be another 20 meters away, but which is legal? What if the baker does not have the bread I always have? What if the baker isn't there and someone else, whom I have never seen there before, is behind the counter? What if I don't have the exact amount of change?

3 Autistic Sociality

Most visibly to non-autistics however, people with autism are impaired in the ease of interpersonal goings-on. Explanations of the social difficulty have been given in terms of 'mind-reading' or 'mentalising'. People with autism are said to be impaired in, or even to lack, a 'theory of mind'. A theory of mind is a computational mechanism in the head that is responsible for mind-reading. It processes states in the form of statements, in order to come to their logical conclusion. Each person's mentalising mechanism computes what people are thinking, using as premises perceptions of their overt behaviour, knowledge of the world, and knowledge of social regularities. Theory of Mind-theory is an explanatory theory of social cognition and states that, whatever we do in interpersonal interactions, we do on the basis of our employment of this mechanism. Theory of Mind-proponents also put forward a course and timescale for the development of this mechanism in the human child. According to (Baron-Cohen, 1995), the

full- fledged mechanism does not come online until around 4 years of age, and after three preliminary mechanisms - the first one for detecting intentionality (the intentionality detector), the second for detecting eyes and what they are looking at (eye direction detector), and the third for sharing attention (shared attention mechanism) - have successively come into place. Theory of Mind-theory is also an explanatory theory of autism and as such it says that in people with autism one or more of the precursory devices does not come online or runs faulty, which results in an incapability or hampered capacity in persons with autism to read other people's minds.

However, Theory of Mind-theory is criticised on the following grounds: First, is it plausible to propose a proposition calculator in the head to account for our ease and fluency in social interaction? Second, do we need a proposition-calculator to be the fluent interactors that we are (most of the time at least)?

The answers to these criticisms hint at the scope of the puzzle of social interaction. First, such a calculator has not been found; it is neurologically not very probable. More to the point, it is not even likely that social cognition takes place mainly in the head. Gallagher (Gallagher, 2001) puts it astutely; Theory of Mind-theorists have a detached, Cartesian approach to social cognition, where what happens takes place in the remote mental realm [1]. That, however, is *not* where social cognition takes place.

The mechanism described above is also developmentally too simple. Even before the Theory of Mind-device "comes online" (around age 4), we are already social beings. Infants are active partners in their interactions with their mothers (Trevarthen, 1979); (Trevarthen & Aitken, 2001), that is, even before the Theory of Mind-mechanism is proposed to have come out. Infants are sources of im- and expressions, to themselves as much as to the persons interacting with them, see also (Hobson, 2002). According to (Gallagher, 2001), and I agree with him, even in adult life, our social capacities are not purely mediated by a mechanism in the head. He proposes that we look for the embodied practice, or primary intersubjectivity, which is the basis of our social talent. (Hobson, 2002) proposes a similar approach, but both authors fail to go deeper than a descriptive, psychological level.

[1]According to Gallagher, the same criticism applies to simulation theory, but there is no space here to go into this.

4 Capturing Intersubjectivity

What is necessary, after its description and localisation, is a method for probing the mechanism of social interaction. This is the system comprising of the partners in the conversation, i.e. the interactors, with their specific structure, [2] and the environment in which they interact. A crucial characteristic of this 'interpersonal choreography' is that it develops from infancy to adulthood, and changes according to who you interact with. Research that goes some way in the direction of unraveling this phenomenon by modeling turn-taking in simulated 'developing' robots, is done by (Di Paolo, 1999); (Di Paolo, 2000) and by (Ikegami & Iizuka, 2003). Further work is necessary to understand the mechanism and how it is connected with thinking. This can be done by comparing and integrating research of the sort referred to with investigations of autism.

References

Baron-Cohen, S. (1995). In *Mindblindness: An essay on autism and theory of mind. cambridge.* MIT Press, Cambridge, MA.

Di Paolo, E. A. (1999). On the evolutionary and behavioural dynamics of social coordination: Models and theoretical aspects. In *Unpublished DPhil thesis.* University of Sussex, UK.

Di Paolo, E. A. (2000). Behavioural coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behaviour, 8(1),* 25-46.

Gallagher, S. (2001). The practice of mind: theory, simulation or primary interaction? *Journal of Consciousness Studies, 8(5-7),* 83-108.

Hobson, R. P. (2002). In *The cradle of thought.* Macmillan, London, UK.

Ikegami, T., & Iizuka, H. (2003). Joint attention and dynamics repertoire in coupled dynamic recognizers. In *Aisb '03 convention: Cognition in machines and animals.* Aberystwyth, Wales, UK.

Trevarthen, C. (1979). In *Communication and cooperation in early infancy: a description*

[2]The structure of each interactor, i.e. their biological and historical make-up, everything that makes them the person they are and not another one, is a substantial part of this system, but not the only part, or even the main part.

of primary intersubjectivity before speech (p. 321-347). Cambridge University Press, UK.

Trevarthen, C., & Aitken, K. J. (2001). Infant intersubjectivity: research, theory, and clinical applications. *Journal of Child Psychology and Psychiatry, 42(1),* 3-48.