

ALife XI Conference Schedule

4th – 8th August 2008

Winchester, UK

Contents

Guide to the program	1
Guide to the venue	1
Overview of conference program	3
Program for Tuesday	4
Program for Wednesday	5
Program for Thursday	6
Program for Friday	7
Abstracts for all talks	8
Author index	78

Guide to the program

Welcome to Artificial Life XI. This conference marks a change from previous meetings in the series: we have been able to include many more talks — more than 180 in total — by switching to a multi-track format. You will therefore have a lot of decisions to make about which sessions you want to attend, and this program will give you the information you need to make the best choice.

On page 3, there is an overall plan of the conference, showing the timings for plenary sessions, parallel sessions, lunches, coffee breaks, the banquet, etc. This is your high-level map of how the conference will run. On the following pages you will find detailed programs covering the plenary and parallel sessions for each day. This should help you to plan your day, although note that in some cases the titles and lists of authors have been abbreviated in order to fit everything on the page.

Each day starts with a keynote lecture from 9am to 10am: given the new multi-track format, these plenary sessions are especially important as a way of bringing the whole community together. The keynote lecture is followed by a coffee break, and then the first of three 90-minute parallel sessions, each consisting of six individual tracks. Each of these sessions includes three talks, so each

talk slot is thirty minutes: twenty minutes for the presentation and ten minutes for questions. Lunch is from 12pm to 1:30pm, followed by two more parallel sessions separated by an afternoon coffee break.

If there's a session you really don't want to miss, it may be a good idea to arrive early. Some of the rooms are smaller than others and although we have done our best to match the most popular topics to the largest rooms, it's possible that some sessions will be over-subscribed. If you arrive late and a session is full, please try to have a second choice in mind.

To make this a complete guide to the conference, we have included the abstracts of all presentations in the final part of the document. In this section you can also get the complete titles and author listings for each paper, information about the authors' affiliations, and the time and room for the talk. Finally, there is an author index at the end. This will prove useful if you are trying to find a talk by a specific person.

Guide to the venue

Although the conference is being hosted by the University of Southampton, the venue is the King Alfred campus of the University of Winchester. As you can see from the map in Figure 1, most of the conference activities centre around the Stripe Theatre building. From Tuesday morning onwards, the registration desk can be found in the lobby, and the plenary talks are held in the Stripe Main Theatre.

Parallel sessions will be held in these six venues, all of them either within the Stripe building or a short walk away:

- Stripe Main Theatre (300-seat downstairs lecture theatre)
- Stripe Lecture Room (120-seat lecture room, upstairs)
- Stripe Studio 1 (off the lobby)
- St Edburga Room 4 (ground floor of the St Edburga building)

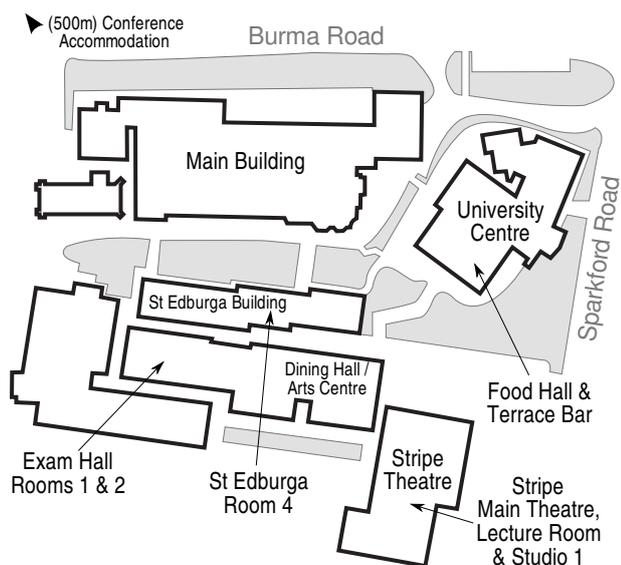


Figure 1: Map of the venue

- Exam Hall Rooms 1 and 2 (ground floor of the Arts Centre building — walk past the entrance to St Edburga to get to these rooms)

Coffee breaks will be held in Stripe Studio 2, also accessible from the lobby. Lunch is served in the food hall of the new University Centre building, which is on the hill above the Stripe. The Terrace Bar is also part of the University Centre building, and this will be open on Wednesday and Thursday evenings.

If you're looking for the route back to your accommodation, most people will be staying in the West Downs Student Village, which is a ten-minute walk to the north-west. (The larger-scale map on the back cover of this document should help.) The best way into central Winchester is to take the path through the cemetery: exit the Stripe and turn right, crossing Sparkford Road, and you're on your way.

Overview of Conference Program					
	Monday Aug 4	Tuesday Aug 5 Day 1	Wednesday Aug 6 Day 2	Thursday Aug 7 Day 3	Friday Aug 8 Day 4
8:00am–9:00am		<i>Registration</i>	<i>Registration</i>	<i>Registration</i>	<i>Registration</i>
9:00am–10:00am		Keynote: Ikegami	Keynote: Jablonka	Keynote: Ellington	Keynote: Schuster
10:00am–10:30am		<i>Coffee break</i>	<i>Coffee break</i>	<i>Coffee break</i>	<i>Coffee break</i>
10:30am–12:00pm		Cell Signalling Networks 1 Evolutionary Theory 1 Autonomous Robots Ecological Modelling Methods Spatial Organisation 1 Culture & Social Evolution 1	Open-Ended Evolution 1 Development 1 Networks 1 Communication 1 Models for Gaia Theory Amorphous & Soft Robotics	Synthetic Biology Information 1 Philosophical Issues 1 Self-* Computing Sensorimotor Connections 1 Fitness Landscapes	Artificial Consciousness Embodied Communication Dynamical Systems Analysis 2 Biochemical Modelling Entropy Information 4
12:00pm–1:30pm		<i>Lunch</i>	<i>Lunch</i>	<i>Lunch</i>	<i>Lunch</i>
1:30pm–3:00pm		Cell Signalling Networks 2 Evolutionary Theory 2 Neural Systems Methodologies Spatial Organisation 2 Culture & Social Evolution 2	Open-Ended Evolution 2 Development 2 Networks 2 Communication 2 Robots & Behaviour Major Transitions 1	Ecological Modelling 1 Information 2 Philosophical Issues 2 Fault Tolerance & Robustness Sensorimotor Connections 2 Wet ALife	Closing address
3:00pm–3:30pm		<i>Coffee break</i>	<i>Coffee break</i>	<i>Coffee break</i>	
3:30pm–5:00pm	<i>Registration opens, West Downs Student Village</i>	Sustainability & Homeostasis Evolutionary Theory 3 Markets Space in Neural Systems Swarm Systems Chemical & Molecular Models	Open-Ended Evolution 3 Development 3 Learning & Attention Multi-Agent Applications Cellular Automata Major Transitions 2	Ecological Modelling 2 Information 3 Philosophical Issues 3 Artificial Chemistries Embodiment Dynamical Systems Analysis 1	
5:00pm–5:30pm		<i>Break</i>	<i>Break</i>	<i>Terrace Bar</i>	
5:30pm–7:00pm	<i>Drinks reception, welcome barbecue & registration</i>	Sustainability Summit	Roke Manor Robot Demo		
7:00pm–7:30pm		<i>Free evening</i>	<i>Free evening</i>	<i>Walk to town centre</i>	
7:30pm–				<i>Conference banquet, Winchester Guildhall</i>	

Tuesday morning keynote: Takashi Ikegami, 9:00am–10:00am, Stripe Main Theatre

Tuesday morning session, 10:30am–12:00pm

Evolving Cell Signalling Networks <i>in silico</i> 1 Stripe Main Theatre	Evolutionary Theory 1 Stripe Lecture Room	Autonomous Robots Stripe Studio 1	Ecological Modelling Methodology St Edburga Room 4	Spatial Organisation 1 Exam Hall Room 1	Culture & Social Evolution 1 Exam Hall Room 2
Evolutionary dynamics in an artificial chemistry.... <i>Decraene et al.</i> Tracking the evolution of chemical computing networks. <i>Lenser et al.</i> Multiple functionalities of biochemical reaction networks. <i>Steijsaert et al.</i>	Group selection increased by coexistence dynamics.... <i>Powers et al.</i> Adapting parasite virulence to combat disengagement.... <i>Cartridge</i> Evolution of complexity in RNA-like replicator systems. <i>Takeuchi & Hogeweg</i>	Strategies for maintaining large robot communities. <i>English et al.</i> Adaptive multi-robot bucket brigade foraging. <i>Lein & Vaughan</i> Optimal robot recharging strategies for time discounted labour. <i>Wawerla & Vaughan</i>	Artificial-life ecosystems: what are they and what could they become? <i>Dorin et al.</i> Linking A-life with ecological and pedagogical research.... <i>Stafford et al.</i> What can artificial life offer ecology? <i>Noble et al.</i>	Exam Hall Room 1 Process-oriented modelling and simulation of space.... <i>Andrews et al.</i> Dynamical complexity of spatially embedded networks. <i>Buckley et al.</i> Cooperation in games on small world networks. <i>O'Riordan et al.</i>	Exam Hall Room 2 Conformist transmission and the evolution of cooperation. <i>Peña</i> Longevity of distinct cultures in a model of memetic drift.... <i>Matthews</i> Diet learning and the evolution of cultural inheritance. <i>van der Post et al.</i>

Tuesday early afternoon session, 1:30pm–3:00pm

Evolving Cell Signalling Networks <i>in silico</i> 2 Stripe Main Theatre	Evolutionary Theory 2 Stripe Lecture Room	Neural Systems Stripe Studio 1	Methodologies St Edburga Room 4	Spatial Organisation 2 Exam Hall Room 1	Culture & Social Evolution 2 Exam Hall Room 2
Mass conserving reaction networks.... <i>Liekens et al.</i> Effects of signalling on gene regulatory networks.... <i>Jenkins & Stekel</i> The evolution of evolvability in gene transcription networks. <i>Izquierdo & Fernando</i>	Initial modelling of the alternative phenotypes hypothesis. <i>Worgan & Mills</i> Species selection of aging for the sake of diversity. <i>Woodberry et al.</i> Adaptive units of selection.... <i>Mills & Watson</i>	Topological properties of evolved robot brains. <i>Adami et al.</i> Evolutionary advantages of neuromodulated plasticity.... <i>Soltoggio et al.</i> Optimal noise in spiking neural networks.... <i>Oros et al.</i>	Hierarchical components and entity-based modelling.... <i>Webster & Malcolm</i> Complex systems models: engineering simulations. <i>Polaek et al.</i> An autonomous agent approach to central planning.... <i>Greenwood</i>	Exam Hall Room 1 Modelling spatial market dynamics with retail agents. <i>Birkin & Heppenstall</i> Distributed gradient optimization.... <i>Litus & Vaughan</i> Reaction-diffusion spots as a model for autopoiesis. <i>Virgo & Harvey</i>	Exam Hall Room 2 Simulated trust: towards robust social learning. <i>Vanderelst et al.</i> The effects of payoff preferences on agent tolerance. <i>Howley & O'Riordan</i> Quantifying creative symbiosis.... <i>Pepper et al.</i>

Tuesday late afternoon session, 3:30pm–5:00pm

Sustainability & Homeostasis Stripe Main Theatre	Evolutionary Theory 3 Stripe Lecture Room	Markets Stripe Studio 1	Spatial Organisation in Neural Systems St Edburga Room 4	Swarm Systems Exam Hall Room 1	Chemical & Molecular Modelling Exam Hall Room 2
What can ALife offer socio-ecological sustainability? <i>Penn</i> Modelling resilience of agro-ecosystems.... <i>van Apeldoorn et al.</i> Homeostasis via chaos.... <i>Harvey</i>	The group selection debate and ALife.... <i>Powers & Watson</i> Altruism amongst spatial predator-prey animals. <i>Scogings & Hawick</i> How epigenetic evolution can guide genetic evolution. <i>Groom et al.</i>	Periodic and continuous markets.... <i>Chaggar et al.</i> CAT: a market design competition. <i>Gerding et al.</i> Local information and trading opportunities.... <i>Ladley & Bullock</i>	Time and space in neuronal networks.... <i>Wombie & Cohen</i> Evolving functional symmetry.... <i>Jones et al.</i> GasNet spatial embedding in a delayed-response task. <i>Vargas et al.</i>	Exam Hall Room 1 Self-organizing heterogeneous swarms.... <i>Sayama</i> On the ability of swarms to compute the 3-coloring of graphs. <i>Cases et al.</i> The spatiality of swarms... <i>von Mammen & Jacob</i>	Exam Hall Room 2 In silico evolution of chemotaxis. <i>Goldstein & Soyer</i> A model chemical memory in an evolved animat. <i>Dale</i> Protein folding with stochastic L-systems. <i>Danks et al.</i>

Tuesday evening: Sustainability summit, 5:30pm–7:00pm, Stripe Main Theatre

Wednesday morning keynote: Eva Jablonka, 9:00am–10:00am, Stripe Main Theatre

Wednesday morning session, 10:30am–12:00pm

Open-Ended Evolution 1	Development & Morphogenesis 1	Networks in Natural & Artificial Systems 1	From Communication to Language 1	Models for Gaia Theory	Amorphous & Soft Robotics
Stripe Main Theatre	Stripe Lecture Room	Stripe Studio 1	St Edburga Room 4	Exam Hall Room 1	Exam Hall Room 2
Learning drives the accumulation of complexity.... <i>Burtsev et al.</i>	Evolution and morphogenesis of multicellular organisms.... <i>Knabe et al.</i>	Group formation and social evolution.... <i>Geard & Bullock</i>	Limited utility of communication.... <i>Scheutz & Schermerhorn</i>	Environmental regulation by higher level selection.... <i>Williams & Lenton</i>	A minimal approach to modular assembly. <i>Studer & Harvey</i>
Passive and driven trends in the evolution of complexity. <i>Yaeger et al.</i>	Timing of critical periods in development. <i>Winks & Berthouze</i>	Developing a methodology for social network sampling. <i>Franks et al.</i>	How learning can guide evolution of communication. <i>Suzuki & Arifa</i>	Utility and experimental testability of the Gaia hypothesis. <i>Free et al.</i>	Mechanism as mind: tensegrities and caterpillars.... <i>Rieffel et al.</i>
Design strategies for open-ended evolution. <i>Caritani</i>	Adaptive evolution in a model evo-devo system.... <i>Psujek & Beer</i>	The connectivity of NK landscapes' basins.... <i>Vérel et al.</i>	Uncertainty and communication complexity in iterated games.... <i>Andras</i>	Canalization and environmental engineering. <i>Yamauchi & Okanoya</i>	"Soft" continuum robots.... <i>Cowan & Walker</i>

Wednesday early afternoon session, 1:30pm–3:00pm

Open-Ended Evolution 2	Development & Morphogenesis 2	Networks in Natural & Artificial Systems 2	From Communication to Language 2	Robots & Behaviour	Major Evolutionary Transitions 1
Stripe Main Theatre	Stripe Lecture Room	Stripe Studio 1	St Edburga Room 4	Exam Hall Room 1	Exam Hall Room 2
Exploiting open-endedness.... <i>Lehman & Stanley</i>	Evolution, development and environment.... <i>Tuffe</i>	Hawks and doves in an artificial society.... <i>Pestelacci & Tomassini</i>	Social coordination of behavior with public symbols.... <i>Türkmen & Zugic</i>	Communication channels for a human-robot dance.... <i>Bryden et al.</i>	Initiation of multicellularity in bacterial biofilms.... <i>Penn et al.</i>
Origins and maintenance of organismal complexity.... <i>Channon</i>	Adaptive growth processes: a model inspired by Pask's ear. <i>Virgo & Harvey</i>	A computational model of gene regulatory networks.... <i>Gonçalves & Costa</i>	Towards language acquisition in autonomous robots. <i>Tikhonoff et al.</i>	Behavior chaining.... <i>Bongard</i>	Does coexistence solve the prebiotic information problem? <i>Branciamore et al.</i>
Organizational requirements for 'open-ended evolution'. <i>Ruiz-Mirazo et al.</i>	Evo-devo in silico.... <i>Joachimczak & Wróbel</i>	RNA sequence-to-structure map for functional evolution.... <i>Ulrich et al.</i>	Animatronic model of a human tongue. <i>Hofe & Moore</i>	Dancing robots.... <i>Silverman et al.</i>	Can individual selection favour significant higher-level selection? <i>Watson et al.</i>

Wednesday late afternoon session, 3:30pm–5:00pm

Open-Ended Evolution 3	Development & Morphogenesis 3	Learning, Action Selection, & Attention	Multi-Agent Applications	Cellular Automata	Major Evolutionary Transitions 2
Stripe Main Theatre	Stripe Lecture Room	Stripe Studio 1	St Edburga Room 4	Exam Hall Room 1	Exam Hall Room 2
Enrichment of interaction rules in a string-based artificial chemistry. <i>Kelly et al.</i>	A developmental system based on sequential growth rules.... <i>Cohen et al.</i>	Basic principles of adaptive learning through variation and selection. <i>Burtsev</i>	KohonAnts: a self-organizing ant algorithm.... <i>Mora et al.</i>	Conceptual structure in cellular automata.... <i>Marques-Pita & Rocha</i>	The emergence of specialization. <i>Bersini</i>
The arrow of complexity hypothesis. <i>Bedau</i>	Programmable architectures that are complex and self-organized.... <i>Doursat</i>	Mammalian decisions. <i>Trimmer et al.</i>	Smart plankton.... <i>Anguitta et al.</i>	Homeodynamics in the game of life. <i>Suzuki & Ikegami</i>	How might group selection explain the major evolutionary transitions? <i>Bryden</i>
Nonconsumable resources facilitate complex evolution. <i>Kampis et al.</i>	Simple creature morphology and metabolism.... <i>Cussat-Blanc et al.</i>	Selective attention in artificial organisms. <i>Ruini & Parisi</i>	Stigmergy and evolved heterogeneity. <i>James et al.</i>	Evolving asynchronous cellular automata for density classification. <i>Jeanson</i>	

Wednesday evening: Roke Manor robot demonstration session, 5:30–7:00pm, Stripe Studio 2

Thursday morning keynote: Andrew Ellington, 9:00am–10:00am, Stripe Main Theatre

Thursday morning session, 10:30am–12:00pm

Synthetic Biology	Information in Complex Systems & ALife 1	Philosophical Issues 1	Self-* Computing	Sensorimotor Connections 1	Fitness Landscapes
Stripe Main Theatre A functional account of minimal cellular life. <i>Bedau Seth</i> Can we rebuild a cell? <i>Bryopsis</i> — an experimental model! <i>Vladimirescu</i>	Stripe Lecture Room Measuring emergence via nonlinear Granger causality. <i>Seth</i> Investigating emergence by coarse graining CAs... <i>Weeks et al.</i> Emergence of glider-like structures in a modular robotic system. <i>Lizier et al.</i>	Stripe Studio 1 Is the organism really a machine? <i>Nicholson</i> Autopoiesis, enactivism, and the extended mind. <i>Wheeler</i> Life in time: the missing temporal dimension in autopoiesis. <i>Di Paolo</i>	St Edburga Room 4 Autonomic co-operative service management. <i>Saffre & Shackleton</i> Adaptive spam detection inspired by the immune system. <i>Abi-Haidar & Rocha</i> Immune-inspired networked service delivery. <i>Gowans et al.</i>	Exam Hall Room 1 Arousal-based model of exploratory behaviour.... <i>Hollte & Cañamero</i> Consumer preferences from rudiments of visual processing.... <i>Tosh et al.</i> Morphodynamics and perceptual worlds.... <i>Almeida e Costa et al.</i>	Exam Hall Room 2 NKo: an extended NK model. <i>Hebbron et al.</i> Ruggedness and evolvability — an evolution's-eye view. <i>Barnett</i> Fitness landscapes and the maintenance of diversity.... <i>Paperin et al.</i>

Thursday early afternoon session, 1:30pm–3:00pm

Ecological Modelling 1	Information in Complex Systems & ALife 2	Philosophical Issues 2	Fault Tolerance & Robustness	Sensorimotor Connections 2	Wet ALife
Stripe Main Theatre Personalities in fish without genetic differences.... <i>Oosten & Hemelrijk</i> Agent-based model of behavioral rules in flocks of cowbirds.... <i>Smith</i> Communication channels in a honeybee colony.... <i>Thenijs et al.</i>	Stripe Lecture Room Information dynamics of phase transitions in RBNs.... <i>Lizier et al.</i> Control in a bio-hybrid robot device.... <i>Artmann et al.</i> Common concepts in agent groups.... <i>Möller & Polani</i>	Stripe Studio 1 When to ascribe life or mind to artificial systems? <i>Keijzer</i> Misrepresentations. <i>Harvey</i> Artificial mental life. <i>Barandiaran & Di Paolo</i>	St Edburga Room 4 Redefining robustness. <i>Silverman & Ikegami</i> Adaptive fault tolerance in wireless sensor networks. <i>Irons et al.</i>	Exam Hall Room 1 Evolving morphological and behavioral diversity.... <i>Pitchler & Cañamero</i> Anticipating future experience.... <i>Mirza et al.</i> Movement strategies for learning in visual recognition. <i>Bermudez et al.</i>	Exam Hall Room 2 Intelligent locomotion of eukaryotic cells. <i>Nishimura & Sasai</i> Computing by amoeba-based neurocomputer.... <i>Aono & Hara</i> Chemical basis for minimal cognition. <i>Hanczyc & Ikegami</i>

Thursday late afternoon session, 3:30pm–5:00pm

Ecological Modelling 2	Information in Complex Systems & ALife 3	Philosophical Issues 3	Artificial Chemistries	Embodiment	Dynamical Systems Analysis 1
Stripe Main Theatre Phenotype-based evolution of complex food webs. <i>de Back et al.</i> Evolving CSR strategies in virtual plant communities. <i>Bornhofen & Lattaud</i> Analysing honeybees' division of labour.... <i>Schnickl & Crailsheim</i>	Stripe Lecture Room How global structure is reflected in local structure.... <i>Anthony et al.</i> Modelling stigmergic gene transfer. <i>Polani et al.</i> Information and cooperativity in mobile robots.... <i>Der et al.</i>	Stripe Studio 1 Abstract organization and material parts in ALife and synthetic biology. <i>Etxeberria</i> Dimensions of adaptivity. <i>McGinn</i> Fitness transmission: a genealogic signature of adaptive evolution. <i>Miconi</i>	St Edburga Room 4 An artificial chemistry-based model of economies. <i>Stratman et al.</i> Towards the identification of the transition to life.... <i>Mendoza et al.</i> Autocatalytic replication of polymers revisited. <i>Blundell</i>	Exam Hall Room 1 Embodied time perception.... <i>Rohde & Di Paolo</i> Evolutionary robotics and the morphological turn.... <i>Almeida e Costa</i> Simulating active touch with a simple embodied agent. <i>Iizuka & Ikegami</i>	Exam Hall Room 2 Recurrent neural network for walking and chemotaxis.... <i>Izquierdo & Buhmann</i> Chemical organizations in living systems. <i>Dittrich et al.</i> Self-organising synchronisation in a robotic swarm. <i>Trianni & Nolfi</i>

Thursday evening: Conference banquet, 7:30pm onwards, Winchester Guildhall

Friday morning keynote: Peter Schuster, 9:00am–10:00am, Stripe Main Theatre

Friday morning session, 10:30am–12:00pm

Artificial Consciousness: From ALife to Mind		Embodied Communication		Dynamical Systems Analysis 2		Biochemical Modelling		Entropy		Information in Complex Systems & ALife 4	
Stripe Main Theatre	Stripe Lecture Room	Stripe Studio 1	St Edburga Room 4	Exam Hall Room 1	Exam Hall Room 2						
Cognitive integration and consciousness... <i>Shanahan & Connor</i>	Self-assembly in physical autonomous robots... <i>Tucl et al.</i>	Spatial prisoner's dilemma sensitive to asynchronism... <i>Grilo & Correia</i>	Emergent intracellular network states and cell fate decision... <i>Nene & Nagl</i>	An adaptive model of marine biogeochemistry in the Archean. <i>Clark & Williams</i>	Relative system–environment distinction... <i>Nozawa & Kondo</i>						
Reconfiguration process of subjective temporal order... <i>Nakajima & Ikegami</i>	Can body language shape body image? <i>Steels & Spranger</i>	Understanding robustness in random Boolean networks. <i>Willadsen et al.</i>	Artificial life as cancer research... <i>Bentley et al.</i>	Entropy production in an energy balance Daisyworld model. <i>Dyke</i>	Quantum entanglement and biological feedback... <i>Haruna</i>						
Towards mental life as it could be: a robot with imagination. <i>Marques et al.</i>	Referential communication in embodied agents... <i>Williams et al.</i>	'Psychoanalysis' of a minimal agent. <i>Manicka & Harvey</i>	Gene-function relationship in yeast and digital organisms... <i>Gerlee et al.</i>	Energy, entropy and work in computational ecosystems... <i>Jacyno & Bullock</i>	Fitness based identification of a robot structure. <i>Zagal et al.</i>						

Friday afternoon: Closing address, 1:30pm–2:30pm, Stripe Main Theatre

Abstracts for all talks

Adaptive spam detection inspired by the immune system

Alaa Abi-Haidar and Luis Rocha

Department of Informatics, Indiana University, Bloomington and Instituto Gulbenkian de Ciencia, Oeiras, Portugal

Thursday 10:30am–12:00pm, St Edburga Room 4, Self-* Computing.

This paper proposes a novel solution to spam detection inspired by a model of the adaptive immune system known as the cross-regulation model. We report on the testing of a preliminary algorithm on six e-mail corpora. We also compare our results with those obtained by the Naive Bayes classifier and another binary classification method we developed previously for biomedical text-mining applications. We obtained very encouraging results that we discuss and which can be further improved with development of this bioinspired model. We thus show that the cross-regulation model is promising as a bio-inspired algorithm for spam detection in particular, and binary classification in general. We also present evidence that a bio-inspired model is relevant for understanding immune regulation itself.

Topological properties of evolved robot brains

Christoph Adami¹, Nicolas Chaumont¹, Jeffrey Edlund² and Arend Hintze¹

¹Keck Graduate Institute

²California Institute of Technology

Tuesday 1:30pm–3:00pm, Stripe Studio 1, Neural Systems.

The topological structure of animal brains is likely to be interesting because the computational power of brains is thought to be almost entirely due to its wiring pattern and hierarchical organization. At the same time, this pattern is not at all well understood, and the information about the wiring pattern of the nematode *C. elegans*, for example, is unique in the literature. A promising direction for the study of network topology in the absence of detailed biological data is the Artificial Life approach, where functional networks are evolved that determine the survival of artificial organisms in an artificial chemistry and genetics. Recently, we used this approach to understand modularity in evolved artificial metabolic networks and developed new tools to dissect their topological and functional characteristics. Here, we apply some of these tools to the study of the brains of robots that have evolved to behave in a simulated world. The robots that are controlled by these brains are simulated versions of real robots (the ATRV Jr. of the iRobot Corporation) whose properties we tested in our laboratory. Both the robot and its environment are simulated in a three-dimensional world that implements realistic rigid body dynamics via the Open Dynamics Engine (ODE). As a consequence, evolved con-

trollers could in principle be transplanted onto the simulated robots' real-world counterparts.

Neural computational tissues (“brains”) are grown from genomes that implement neural network development and function based on a set of rules (“genes”) that are conditionally executed, that is, regulated, by a set of simulated proteins produced by the cells in the tissue. This system (“Simnoesis”) is based on the “Norgev” platform but was completely rewritten in order to be able to evolve complex tissues that process many temporally varying input signals. We evolve neural tissues on two-dimensional grids (of up to 15x15 neurons) that control a simulated ATRV Jr with 19 sensors (17 sonars, a compass, and a sensor relaying distance to goal), controlling two motors driven by two actuators for differential steering. The evolved tissues control complex robot behavior, such as wall-following, obstacle avoidance, and goal-finding, using a complex network structure reminiscent of the *C. elegans* connection graph. The fitness evaluation of a genome consists of growing the network, and evaluating the behavior of the robot in a 3D environment akin to the fitness evaluation in the work of Sims). Fitness evaluation and evolution via a Genetic Algorithm is implemented within the EVO software.

We analyze the properties of evolved neural networks using standard tools (such as edge-distribution, shortest-path length, and betweenness centrality), as well as new tools that reveal robustness and modularity via clustering methods and information theory. We find that the topological properties of evolved functional networks are very different from their randomized counterparts, and characterize the “rarity” of these networks with standard statistical tests. Finally, we compare the topological properties of our evolved networks to the connection graph of *C. elegans*.

Evolutionary robotics and the morphological turn: an epistemological perspective

Fernando Almeida e Costa

Centre for Computational Neuroscience and Robotics, University of Sussex

Thursday 3:30pm–5:00pm, Exam Hall Room 1, Embodiment.

What are the philosophical and epistemological implications of work in evolutionary robotics (ER) dealing with the evolution of morphologies and morphogenesis? So far, its theoretical consequences for cognitive science have not been fully fleshed out. Also, investigation in morphologically-based ER has not shown its affiliation in the long tradition of a *morphological thought*. Understanding theoretical implications and the phylum of thought of some line of research may be of great importance not only to the historians of science but to the future development of the research itself. I propose that the shift towards morphodynamics belongs to an old phylum of

cognitive and biological thought, both naturalist and structuralist in nature, and is a component of a broader *morphological turn*. Examples of its manifestations include early investigations in A.I., e.g. the studies on morphogenesis and cognitive structures by Alan Turing, the semiophysics firstly proposed by René Thom, or the morphological, non logico-combinatorial structuralism of Claude Lévi-Strauss in cultural anthropology. More recently, work in ER addressed not only the morphodynamics of an agent's body and environment but also the morphological properties of its "perceived world" (Almeida e Costa et al, 2008, Alife XI, these proceedings). By highlighting what is common to these apparently unrelated lines of research, light is shed on the entire framework of the morphological turn. Structuralism is normally seen as the continental current of thought that developed in the 60's and 70's and had its major tenets on linguistics or literary criticism, while being of a logico-combinatorial, algebraic and static nature. But there is another phylum of structuralist thought: one that can be traced back to the works of D'Arcy Thompson and even to Goethe. This other phylum thinks of structures as dynamical forms in development. It is a naturalist and non-formalist (in the sense of formal logic) approach that considers forms as morphodynamically self-organised wholes. The concept of "transformation" and its mathematical treatment is central to this perspective. Morphodynamically inspired robotics exploits all intrinsic and extrinsic physical properties available. This entails the denial that the "cognitive" properties are to be found at an algorithmic level that dominates the physical properties. Cognitive activity relies crucially on the agent's morphodynamics, actually implying the minimization of control at the algorithmic level. The functionalist principle of the irreducibility of the cognitive level to the physical medium is put aside. Thus, the *morphological turn* opens up the possibility of a non reductionist physics of meaning. The success of modern science, i.e. the physical mechanism that emerged in the XVII century was only possible due to the abandon of the dynamics of forms: the Aristotelian physics of qualities. This implied the impossibility of connecting the new "objectivity" with the qualities of the world as it is perceived. It is often pointed out that the recent embodied approach to cognition refuses the Cartesian mind-body dualism. It should be noted that it also refuses the divide between physics and qualitative form. On a supplementary note, this perspective is totally consistent with high-level and low-level cognitive abilities forming a continuum. The hypothesis of an evolutionary path leading from the emergence of particular human morphologies (feet and hands), hence the ability to walk, to the emergence of language, put forward by palaeontologist André Leroi-Gourhan, in the same structuralist vein afore mentioned, is entirely consistent with this orientation.

Morphodynamics and perceptual worlds: conceptual approach and an experiment in evolutionary robotics

Fernando Almeida e Costa¹, Ian Macinnes² and Inman Harvey¹

¹Centre for Computational Neuroscience and Robotics, University of Sussex

²RBS

Thursday 10:30am–12:00pm, Exam Hall Room 1, Sensorimotor Connections 1.

From New A.I. came the deep conceptual insight that cognition is a consequence of the opportunistic exploitation of all morphodynamical properties of an agent's body and environment, which acts to minimize control at the algorithmic level. These properties structure the agent's perceptual world. Here, we aim at those morphological properties that structure the environment "as perceived" by the agent. Jakob von Uexküll's functional circle hypothesis, provides a general framework to understand active perception as morphologically-based. This framework must be seen within a broader *morphological turn*; this occurred in the second half of the 20th century across various fields of research, and is morphologically-based, as opposed to information-theoretic based (Almeida e Costa, 2008, Alife XI, these proceedings). For some authors, developments in dynamical systems theory opened the possibility of a nonreductionist physics of perception and meaning (e.g., Petitot 2000, *Physique du Sens*, CNRS Éditions). A dynamicist approach regards organisms as being perturbed by and responding to cues they have been evolutionarily selected to respond to, rather than mirroring or extracting information from the outside world. The morphological structuring of the perceived environment is highly constrained by the particular morphologies of its body, and by the dynamics of those morphologies. To exploit this aspect for engineering and conceptual purposes, a particular method in evolutionary robotics is proposed, based on the functional circle hypothesis by Jakob von Uexküll (Macinnes et al., 2005 *Adaptive Behaviour*, Vol:14.2. p. 147). A functional circle is an abstract structure that describes the functional relationship between an organism, its "perceived world", and its environment. According to the functional circle hypothesis, a perceptual sign of an object (say, the smell of a mammal's butyric acid, captured by a tick) give rise to a perceptual cue, the *subjective experience* of that object in the organism's (the tick's) *Umwelt*: the word *Umwelt* was used by von Uexküll to describe the biologically evolved world of perceptions, as perceived by a particular organism/species, which results from the morphodynamical interaction with its environment. This leads to an effector cue which drives the animal to perform some action (say, fall down from the tree under which the mammal is passing), changing the organism's relationship to the object. After the action is

performed, the perceptual cue is gone and therefore that functional circle is extinguished but may lead to another (say, dealing with the fur, finding warm skin, then biting). The proposed method consists of changing the mutational operators to evolve functional circles instead of directly evolving sensorimotor loops. The agent's morphodynamics and perceptual world are co-evolved; evidence suggests this enables a closer coupling between body, controller, and environment. The evolving functional circle hypothesis predicts that adding multiple perceptual cues produces robots more adapted to their environment than they would be otherwise. A comparative analysis of the evolved robots suggests that this is the case. An explanation is suggested: the specific positions of the sensors using mutable locations together with body morphology define spatial and temporal relationships with the environment. Co-evolving the agent's morphology, locations of its sensors, and controllers, evolve these relationships as well which implies that we are evolving perceptual cues, and therefore evolving perceptual worlds.

Uncertainty and communication complexity in iterated cooperation games

Peter Andras

Newcastle University

Wednesday 10:30am–12:00pm, St Edburga Room 4, From Communication to Language 1.

Iterated cooperation games (e.g. Prisoner's Dilemma) are used to analyze the emergence and evolution of cooperation among selfish individuals. Uncertainty of outcomes of games is an important factor that influences the level of cooperation. Here we present an agent-based simulation that implements the uncertainty of outcomes together with the communication of intentions between agents. This simulation is used to analyze the relationship between uncertainty and the complexity of the language that the agents use to communicate about their intentions. The complexity of the language is measured in terms of variability of its usage among agents. The results show that more outcome uncertainty implies lower complexity of the agent language.

Investigating patterns for the process-oriented modelling and simulation of space in complex systems

Paul Andrews¹, Adam Sampson², John Markus Bjørndalen³, Susan Stepney¹, Jon Timmis¹, Douglas Warren² and Peter Welch²

¹University of York

²University of Kent

³University of Tromsø

Tuesday 10:30am–12:00pm, Exam Hall Room 1, Spatial Organisation 1.

Complex systems modelling and simulation is becoming increasingly important to numerous disciplines. The CoS-MoS project aims to produce a unified infrastructure for modelling and simulating all sorts of complex systems, making use of design patterns and the process-oriented programming model. We provide a description of CoS-MoS and present a case study into the modelling of space in complex systems. We describe how two models — absolute geometric space and relational network space — can be captured using process-oriented techniques, and how our models can be refactored to allow efficient, distributed simulation. We identify a number of design, implementation and refactoring patterns that can be applied to future complex systems modelling problems.

Smart plankton: a new generation of underwater wireless sensor network

Davide Anguita, Davide Brizzolara, Alessandro Ghio and Giancarlo Parodi

University of Genova

Wednesday 3:30pm–5:00pm, St Edburga Room 4, Multi-Agent Applications.

The aquatic world, which covers more than the 70% of the earth, has been largely unaffected by the WSN revolution (ignited by DARPA funded UC Berkeley “Smart Dust” project) due to the difficulty of transferring most of the knowhow, developed for terrestrial and aerial systems and devices, to their underwater counterparts. Nowadays underwater wireless networks are expensive (US\$ 10k or more), sparsely deployed (a few nodes, placed kilometers apart), typically communicating directly to a base-station or sometimes based on the use of underwater manned or unmanned vehicles. Our research is aimed to develop a new generation of UWSN (Underwater Wireless Sensor Network), called Smart Plankton, by drawing inspiration from marine biology and aquatic micro-organism such as zooplankton and phytoplankton. Our target is to develop a self-organizing network composed by a relatively large number of innovative nodes, equipped with sensors for monitoring, surveillance, underwater control and many others potential applications. Inspired to the rich inventory

of plankton adaptations our research is oriented to explore innovative solutions in following areas:

1. implementation of the single network node focusing on: a) the use of reconfigurable architecture, balancing the need of computation (sense, communicate, etc.) with the survivability constraints (energy foraging and storage), with algorithms recently developed for computational embedded intelligence (e.g. kernel methods for embedded and pervasive systems [D. Anguita, A. Ghio and S. Pisciotta, *Adaptive Hardware and Systems*, p.571, 2007]) for acquiring or improving intelligent behavior; b) a mobility system based on body thermal expansion of solids and liquids and compression under pressure, such as in sperm whale which uses spermaceti, a semi-liquid, waxy substance for movement and stability;

2. communication between nodes because, in comparison with ground-based sensor networks, mobile UWSNs cannot employ radio frequency (RF). The alternative and more innovative method of optical communication, suggested also by the natural world (e.g. quorum sensing through bioluminescence in plankton shoals [F.J.Jochem, *Marine Biology*, Vol. 135, p.721]), can allow the development of a high rate, low power, long life and low cost communication link among devices. At our Department, we are testing the use of LEDs and phototransistors for developing an underwater optical communication system (based on 802.11a protocol) considering that experimental tests have shown that the better wavelength lies around 420 nm (blue-violet wavelengths) and that the value changes in presence of turbidity;

3. energy scavenging in order to allow a long life to the network; energy can be generated by using electrochemically active bacteria [B.E.Logan and J.M.Regan, *Environmental Science and Technology*, 40, p. 5172] which have been recently discovered and have the property to oxidize organic matter and release the electron to an electrode. This has some definite advantages over the use of a chemical catalyst, as bacteria can sustain themselves and recover after inadvertent poisoning;

4. shoal intelligence in order to allow Smart Plankton to perform complex tasks by cooperation of the individuals; this approach can be considered as an application of Swarm Intelligence model [G. Beni and J. Wang, *Proceed. NATO Advanced Workshop on Robots and Biological Systems*, 1989] for dealing with the peculiarities of the harsh underwater environment.

On preferred states of agents: how global structure is reflected in local structure

Tom Anthony, Daniel Polani and Chrystopher Nehaniv

University of Hertfordshire

Thursday 3:30pm–5:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 3.

We investigate the correlation between the information theoretic measure of *empowerment* and the graph theoretic measure of *closeness centrality*, to better understand the structural conditions that must exist in a world for learning and adaptation. We examine both measures in both a simple gridworld scenario, represented as a graph, and on a scale-free graph. We show a strong correlation between the two measures, and discuss the strengths and weaknesses of both. We go on to show how the local measurement of empowerment can in many cases predict a measure for the global measurement of closeness centrality.

Combinatorial optimization and self-disciplined computing by amoeba-based neurocomputer

Masashi Aono and Masahiko Hara

RIKEN

Thursday 1:30pm–3:00pm, Exam Hall Room 2, Wet ALife.

Creating a biocomputer with its hardware incorporating biological materials, is it possible to implement some unique functions that are difficult for conventional digital computers to deal with? A living organism is a hierarchically-structured system in which a number of self-organization processes run simultaneously on its multiple levels with their characteristic spatiotemporal scales. Because a self-organization process at each level involves a certain kind of benefit optimization such as energy minimization and stability maximization, it would be sound to assume that an organism is a particular kind of concurrent computing system in which a number of computing processes to solve different benefit optimization problems run concurrently by sharing common computational resources such as the energy and substances. Despite a lack of predefined decision program, if these hierarchically-intervened optimization processes are capable of making a self-disciplined decision, for example, a decision to accept a loss in short-term benefits of its local part for the sake of long-term gains of its global body, the decision capability may be exploited for discovering some unprogrammed but reasonable optimization criteria when incorporated in a biocomputer.

With this expectation, we created a computing system incorporating an amoeboid unicellular organism, a true slime mold *Physarum polycephalum*, known to exhibit rich spatiotemporal oscillatory behavior and sophisticated computational capabilities. Introducing an optical feedback according to a recurrent neural network model, we

lead the amoeba's photosensitive branches to expand or shrink within a network-type chamber in search of a solution to the traveling salesman problem (TSP).

Here we demonstrate our system's high optimization capability of solving four-city TSP. Our system reaches and stabilizes an optimal solution, as the amoeba having photoavoidance changes its shape in search for the most stable configuration allowing the amoeba to maximize its body area while minimizing the risk of being illuminated.

Intriguingly, the maintained stabilizing mode of the solution, however, spontaneously switches to the destabilizing mode without any explicit external perturbation. Contrary to the photoavoidance, the amoeba starts to destabilize the once-reached solution by spontaneously expanding its branch under illumination, and restarts the solution-searching process. Consequently, our system finds multiple solutions by repeatedly switching between the stabilizing and destabilizing modes.

As long as the amoeba maintains the photoavoidance and stabilizes the solution without changing its shape, the amoeba is stuck in a stalemated situation eliminating any possibility of nutrient acquisition. However, the amoeba spontaneously takes a risk of being illuminated locally and temporally to restart its shape change. It may be possible to view this spontaneous behavior as implying biological systems' capability of self-disciplined decision to put their resources available at present into risky investments to target resource acquisitions in the future.

We speculate that the spontaneous destabilization occurs due to the existence of chaotic dynamics capable of amplifying tiny fluctuations in a microscopic level to affect the unstable shape change in a macroscopic level. Indeed, applying several nonlinear time series analysis to the amoeba's oscillatory movements, we obtained results suggesting that an individual amoeba might be characterized as a set of coupled chaotic oscillators.

Additionally, we present a new technique that we call "autonomous meta-problem solving." In this approach, our system not only can solve a given problem but also can find new problems and then determine solutions in a self-disciplined manner, by exploiting the amoeba's unique searching ability and spontaneous behavior.

Information-theoretic aspects of control in a bio-hybrid robot device

Stefan Artmann¹, Soichiro Tsuda² and Klaus-Peter Zauner²

¹Friedrich-Schiller-University

²Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Thursday 1:30pm–3:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 2.

Information processing in natural systems radically differs from current information technology. This difference is particularly apparent in the area of robotics, where both

organisms and artificial devices face a similar challenge: the need to act in real time in a complex environment and to do so with computing resources severely limited by their size and power consumption. The formidable gap between artificial and natural systems in terms of information processing capability motivates research into the biological modes of information processing. Such undertakings, however, are hampered by the fact that nature directly exploits the manifold physical characteristics of its computing substrates, while available theoretical tools in general ignore the underlying implementation. Here we sketch the concept of bounded computability in an attempt towards reconciling the information-theoretic perspective with the need to take the material basis of information processing into account. We do so in the context of *Physarum polycephalum* as a naturally evolved information processor and the use of this organism as an integral component of a robot controller.

Artificial mental life

Xabier E. Barandiaran¹ and Ezequiel Di Paolo²

¹Dept. of Logic and Philosophy of Science, University of the Basque Country, Spain

²Centre for Computational Neuroscience and Robotics, University of Sussex, UK

Thursday 1:30pm–3:00pm, Stripe Studio 1, Philosophical Issues 2.

Work in Artificial Life aimed at informing Artificial Intelligence (Steels & Brooks, 1994, *Artificial Life route to Artificial Intelligence*, Lawrence Erlbaum) has drawn inspiration from biology mainly at two levels: i) a bottom-up modelling approach conceiving cognition as the evolutionary complexification of adaptive behaviour, and ii) appeals to self-organization in the domain of behaviour and neural dynamics in analogy with self-organized chemical and biological processes. But little attention has been paid to the possibility of conceiving (and modelling) behaviour in terms of a self-maintaining organized unity in analogy with minimal forms of proto-cellular (or autopoietic) life. We propose that the behavioural counterpart of a network of self-sustaining chemical reactions should be a network of interactively maintained sensorimotor dissipative structures (*habits*) that emerge from the continuous reciprocal interaction between brain, body and world (and not, as in previous attempts, between molecular processes and neural processes, conceiving the nervous system as *operationally closed* Varela, 1979, *Principles of Biological Autonomy*, Elsevier).

Despite its popularity among pre-Darwinian biologists (such as Aristotle, Lamarck or Bichat), pragmatists and phenomenologists alike (Dewey, Merleau-Ponty) and among pre-computationalist psychologists (like James, Goldstein, Ivo Kohler or Piaget) the notion of habit has received little attention within Artificial Life. Habits possess key properties that make them extremely attractive for

modelling the organization of behaviour: a) the structure of habits can be traced back to a fully operational-dynamicist framework, b) they do not presuppose a distinction or a causal priority between perception and action, c) habits are inherently situated or enactive structures cutting across brain, body and environment, d) habits are plastic and malleable, e) habits provide a concrete sense of self-maintenance (they are both cause and effect of their occurrence) potentially implying an intrinsic and an interactive teleology and f) habits can be nested or composed at different scales. This opens up the possibility for an operational notion of what might be called *Mental Life* (Barandiaran, 2007, *The World, the Mind and the Body*, p. 49, Imprint Academic) as the continued formation of a web of habits through sensorimotor interactions whose cohesive self-maintenance constitutes the identity of a *cognitive* (as opposed to barely biological) agent and the world it thereby co-defines.

We use some recent evolutionary robotic models on preference and habit formation (Di Paolo & Iizuka, 2008, *Biosystems*, 91, p. 409) to illustrate and explore the theoretical and philosophical implications of taking sensorimotor habits as the building blocks of behavioural organization. This organization takes the form of an attractor landscape whose stability is homeodynamically maintained through sensorimotor coupling. *Mental Life* opens up a new object of modelling in its own right, closer to the Aristotelian notion of psyche (or even the Heideggerian notion of Dasein) than to the notion of information processing, adaptive problem solving or weak conceptions of autonomy in robotics. Artificial *Mental Life* involves a shift from building artificial systems that satisfy externally imposed norms (engineering or evolutionary) to systems capable of generating their own norms: those required to sustain their own behavioural organization. In turn, it can become a source of new research questions to investigate the dynamics of assimilation and accommodation into an existing organization, its shaping by social interactions and institutions, or mental disorders dealing with stability, stress, identity, etc.

Ruggedness and evolvability — an evolution’s-eye view

Lionel Barnett

Centre for Computational Neuroscience and Robotics, University of Sussex

Thursday 10:30am–12:00pm, Exam Hall Room 2, Fitness Landscapes.

In both artificial and biological evolution, autocorrelation is commonly cited as a statistic which speaks to the “ruggedness” — and by implication evolvability — of a fitness landscape. But while the standard definition of autocorrelation involves *uniform* sampling of genotypes, it is a truism that evolution most decidedly does *not* sample a landscape uniformly. This is of particular significance in

difficult artificial evolution problems, or indeed in natural evolution, where the vast majority of genotypes tend to be of poor or lethal fitness. On such landscapes uniform sampling is effectively biased towards precisely those (poor quality) genotypes which, from an evolutionary perspective, are of limited interest. To address this problem we suggest instead to take an “evolution’s-eye” view of autocorrelation: that is, we let evolution itself do the sampling.

How are we to go about this? We note first of all that autocorrelation may be considered naturally in terms of *mutation*. Indeed, the significance of autocorrelation to evolutionary dynamics lies precisely in the (statistical) relationship between the fitness of parents and their mutant offspring. We thus propose that a more cogent and useful statistic is just the correlation between parent/mutant fitnesses *as sampled over the ensemble of evolutionary histories*. We argue that this alternative autocorrelation is both conceptually compelling and also practicable, in the sense of being amenable to finite sampling.

We note that our new statistic is no longer “evolutionarily agnostic”; rather, it is tightly bound to the dynamics of a particular evolutionary scenario. This, however, we regard as a strength. We can imagine, for example, that the same fitness landscape might “appear smoother” to one evolutionary algorithm than to another, thus providing insight into the suitability of a particular evolutionary algorithm to a particular problem in artificial evolution.

We also demonstrate how autocorrelation may be derived from the *mutant fitness distribution* — a finer-grained statistic — and we introduce the notion of *linear regressive* fitness landscapes. We illustrate our ideas with generalised NK landscapes, which are particularly tractable to analysis.

A functional account of minimal cellular life

Mark Bedau

Reed College

Thursday 10:30am–12:00pm, Stripe Main Theatre, Synthetic Biology.

Life seems to be one of the most fundamental categories in nature. But how exactly do material objects that are living differ from those that are not? Is there any a fundamental difference? And if so, is it a dichotomy or a matter of degree?

One answer to these questions abstracts away from chemical details and instead concentrates on living system’s functional properties. In fact, the protocell research community more or less agrees that minimal cellular life forms are defined by chemically integrating three functionalities (Rasmussen et al., 2008, *Protocells*, p. 71). First, the system maintains an identity over time by localizing all its components, concentrating reagents and protecting key chemical reactions from molecular parasites and poisons. Second, it utilizes free energy from its environment to digest environmental resources in order to

maintain and repair itself, to grow, and ultimately to reproduce. Third, these processes are under the control of inheritable information that can be modified during reproduction. The three functionalities mutually enable and support each other. They are collectively autonomous in the sense that they are created and sustained by the operation of the whole functional triad itself, rather than by any external governing agency.

Why should we believe that minimal cellular life is a chemically integrated functional triad of container (C), metabolism (M), and genetic program (P)? The rough consensus in the protocell community lends CMP view some weight, but not enough to convince skeptics. Other functionalities often associated with life-like reproduction, autonomous behavior, and sensitivity to the environment can be explained by the functional triad, which lends it further support. Going even further, the CMP view can be explained as a consequence of a more fundamental view according to which the essence of life is open-ended evolution. Elsewhere I have defended this view on the grounds that it best explains life's familiar hallmarks (Bedau, 1996, *The Philosophy of Artificial Life*, p. 332, Oxford UP) and puzzles (Bedau, 1998, *Art. Life*, 4, p. 125).

One puzzle about life concerns whether the distinction between life and non-life is dichotomous or continuous. The functional triad view implies that there is an array of thousands of different possible kinds of functional organizations, and they all more or less match the paradigm organization used to define life. Instances of some other functional organizations would be pretty clearly alive, and instances of others would be pretty clearly not alive, and a gray zone of further possible functional organizations separates those two clear cases.

One could divide the gray zone with any number of bright lines purporting to separate those systems that are "really alive" from those that are not, but I recommend not doing this. Instead, I think there is no deeper fact of the matter about the life/nonlife distinction other than the graded array of functional organizations. Attempting to find a more precise "definition" of life would be to invent a categorical distinction that does not exist in nature.

The arrow of complexity hypothesis

Mark Bedau

Reed College

Wednesday 3:30pm–5:00pm, Stripe Main Theatre, Open-Ended Evolution 3.

The evolution of the biosphere exhibits a trend of increasing complexity of the most complex organisms. Even though we are uncertain about the proper way to measure complexity, it is hard to deny the trend that the earliest prokaryotic cells are simpler than the eukaryotic cells that arose from them, and these were simpler than the multicellular life forms that evolved from them, and so on. But this trend is controversial to interpret and explain, and even to

describe properly. Some think that the trend has for all intents and purposes already been explained. In contrast, I argue that the trend is not yet adequately explained but instead is a major remaining challenge in understanding the creativity of evolution.

Progress on this challenge is slowed in part because many people fail to realize that the explanation of life's complexity is still a mystery. Some people believe that natural selection given an infinite space of genetic possibilities will inevitably produce more and more complex adaptations. But soft artificial life models like *Tierra*, *Avida*, and *Echo* show conclusively that those mechanisms are in general insufficient to produce a trend of increasing complexity. The proof is simple: The models embody those mechanisms but they don't exhibit the requisite behavior. Mechanisms like natural selection in an infinite space of genetic possibilities might be necessary for explaining the trend, but they are not sufficient.

This implies that we need new concepts, theories, and models if it is to resolve the arrow of complexity hypothesis. Fortunately, soft artificial life models can be just the right tool for exploring answers to this question. But these models are not fool-proof. Some models beg the interesting questions, and others fail to produce the relevant behavior. So, proper use of these models requires care and experience. But in the right hands, they can provide a public, repeatable, and empirically grounded method for making incremental progress on the question of the creativity of evolution.

Artificial life as cancer research: embodied agent modelling of blood vessel growth in tumours

Katie Bentley, Paul Bates and Holger Gerhardt

Cancer Research UK

Friday 10:30am–12:00pm, St Edburga Room 4, Biochemical Modelling.

Tumours need to signal the growth of new blood vessels (angiogenesis) in order to obtain an oxygen supply and continue to grow. Angiogenesis in tumours, as opposed to normal tissue, generates abnormal, tortuous and leaky vessels. The vessels poor quality keeps oxygen levels low in the tumour, which keeps mutation rates high, and causes metastases to develop and spread across the body.

Angiogenesis as a process is a fascinating example of adaptive, environment-driven, morphogenesis of a spatial network; the most suitable and practical approach, and framework for simulation, was artificial life. Our multidisciplinary research aims to 1) understand the mechanisms of angiogenesis, 2) understand why the tumour environment causes abnormal vessels and 3) develop novel cancer therapies which could normalise tumour angiogenesis and thereby prevent metastases through reduction in hypoxia and increased genetic stability.

We have developed a multiscale agent-based model of a blood vessel interacting with its environment in order to investigate the effects that different environmental factors have on the initial stages of angiogenesis. The simulated endothelial cells in the vessel exist across multiple grid sites in a 3D gridded lattice. Each cell is comprised of many autonomous agents, representing sections of the cell membrane. In the first incarnation of the model, agents create new agents to change cell morphology. In the current version they move and are connected by springs, which realistically mimics membrane tension during cell migration.

Each agent communicates with its local environment, including other agents, to decide whether to activate receptors, release ligands and/or alter the cells local morphology. Overall a cell's behaviour and morphology emerge from the low-level interactions of its agents with the environment, which in turn then determines the vessel network morphology and development.

With this approach we have realistically modelled the initial stages of angiogenesis and made interesting predictions concerning abnormal endothelial cell fate determination in tumours, which are now being tested in the laboratory. The model is now being developed further to fully simulate cell migration and fusion of cells as the network develops.

Movement strategies for learning in visual recognition

Edgar Bermudez, Andrew Philippides and Anil Seth

University of Sussex

Thursday 1:30pm–3:00pm, Exam Hall Room 1, Sensorimotor Connections 2.

In this paper we study the role of movement strategies during learning in object recognition models. We show that a simple model, the RBF, can outperform a more complex hierarchical model, the HMAX, when rotation and scale invariance are provided by the training phase. Moreover, we assess the exploitation of temporal information by the RBF using optic flow. The results show that the RBF model can only exploit the temporal information using optic flow when the training and testing trajectories are the same. This work exemplifies the idea that the complexity of the neural mechanisms in object recognition can be understood not only in the brain but also in the interaction between brain, body and environment.

The emergence of specialization

Hugues Bersini

IRIDIA — ULB

Wednesday 3:30pm–5:00pm, Exam Hall Room 2, Major Evolutionary Transitions 2.

The emergence of specialization remains a true challenge. Suppose a world initially filled with specialized and generalist agents, and let's define these later as able to endorse the various competences characterizing the specialists and to endorse them as well as the specialists. In an evolutionary perspective, it is obvious to see why a specialist will always be less adapted than a generalist, which can indeed alternatively acts as many experts. The generalists will meet much more agents and much more situations to which they are adapted to and then cumulate much more payoff (unless the tasks done by generalists are systematically of pitiful quality). This is indeed a paradox to face in order to make sense of a world nonetheless full of specialists. This paper will discuss various ways, beyond the obvious possibility of unfavoring multi-specialization by paying a high cost, to allow specialists to survive the presence of generalists.

Modelling spatial market dynamics with retail agents

Mark Birkin and Alison Heppenstall

University of Leeds

Tuesday 1:30pm–3:00pm, Exam Hall Room 1, Spatial Organisation 2.

The power of agent-based modelling (ABM), when integrated with other AI-based and conventional approaches, can be greatly enhanced. The resulting hybrid systems offer a flexible modelling environment that exploits the benefits of the component methods. In particular, the ABM paradigm can be used to explore and understand systems that are governed by complex, non-linear relationships and self-organisation.

In earlier research, the authors have described an agent-based model of retail behaviour in which customer transactions are simulated using spatial interaction models (Heppenstall et al., 2005, *Trans. in GIS*, 9, p.35). The model has been used to simulate processes such as the diffusion of price changes through a retail network, and the interdependence of pricing behaviour between competing retail chains (Heppenstall et al., 2006, *J. Artif. Societies & Soc. Simulation*, 9).

In the research which is now presented, we draw an insight from an established method which explores the behaviour of retail provision when customer transactions are simulated by a spatial interaction model, but in which structural change is driven by a simple equilibrium-seeking mechanism. Established methods have provided

useful insights into retail patterns under equilibrium-seeking behaviour, but have done relatively little to enrich our understanding of the dynamic processes and decisions from which change arises. Through the combination of these approaches, it is suggested that a much richer model architecture is possible, in which interacting retail agents produce a spatially heterogeneous distribution of supply. This structure brings together coevolution in the economic and geographical variables (price and provision) through a dynamic model of competition amongst agents.

A series of numerical experiments are introduced to demonstrate how the use of agents can introduce more behaviour. The simulations are embedded in a real local retail environment. We evaluate the extent to which this work can be considered to present an improved understanding of this system.

Autocatalytic replication of polymers revisited

Ben Blundell

University of Sussex

Thursday 3:30pm–5:00pm, St Edburga Room 4, Artificial Chemistries.

A simple computational model for the emergence of autocatalytic sets as described in (Farmer et al., 1986) is reimplemented. Results are found to generally agree with the major theme in the original work: increasing the initial polymer variety in a toy chemical soup scenario increases the likelihood that a complex autocatalytic set will suddenly bootstrap itself into existence. Quantitatively, however, critical probabilities derived from this careful reimplementation are very much higher than those reported in the original work. A full resolution is not reached, but a theoretical argument supports the simulation results gained in this instance.

Behavior chaining: incremental behavioral integration for evolutionary robotics

Josh Bongard

University of Vermont

Wednesday 1:30pm–3:00pm, Exam Hall Room 1, Robots & Behaviour.

One of the open problems in autonomous robotics is how to consistently and scalably integrate new behaviors into a robot with an existing behavioral repertoire. In this work a new technique called Behavior Chaining (BC) is introduced, which allows for gradually expanding the behavioral repertoire of a dynamically behaving robot. The approach relies heavily on scaffolding: gradually restructuring the robot's environment such that selection pressure favors the incorporation of a new behavior. BC is demonstrated here on virtual robots that locomote toward and then lift objects. Finally, it is shown that behavioral trajectories may be unidirectional: it is easier for the robot to

learn lifting and then incorporate locomotion than it is to learn locomotion and then incorporate lifting.

Evolving CSR strategies in virtual plant communities

Stefan Bornhofen and Claude Lattaud

LIAP5, Paris Descartes

Thursday 3:30pm–5:00pm, Stripe Main Theatre, Ecological Modelling 2.

This paper introduces a functional-structural plant model based on Artificial Life concepts and reports studies on evolutionary dynamics in virtual plant communities. The characteristic of the present approach lies in plant evolution at both functional and structural levels. The conducted experiments focus on the emergence of different life history strategies in an environment with heterogeneous resource availability and disturbance frequency. It is found that, depending on the encountered conditions, the plants develop three major strategies classified as competitors, stress-tolerators and ruderals according to Grime's CSR theory. Most of the evolved characteristics comply with theoretical biology or field observations on natural plants.

Does coexistence solve the prebiotic information problem?

Sergio Branciamore¹, Walter de Back² and Enzo Gallori¹

¹Department of Evolutionary Biology, University of Florence, Italy

²Collegium Budapest, Institute for Advanced Study, Hungary

Wednesday 1:30pm–3:00pm, Exam Hall Room 2, Major Evolutionary Transitions 1.

The information problem in prebiotic evolution arises from constraints on the amount of information that can be maintained in Darwinian evolution. The error threshold limits the transmissible length of a template under high mutation load. Existing solutions to the error threshold assume that the primitive genome consisted of multiple coexisting unlinked templates. Such coexistence requires a mechanism of cooperation to counterweigh competitive exclusion. Although much attention is given to ecological coexistence properties of cooperation mechanisms, little is known about the information carrying capacity of these systems under high mutation rates. Template coexistence may escape the error threshold, but it simultaneously raises the new problem of maintaining cooperation. Cooperation is threatened by the production of parasites through mutation. This results in an additional constraint on the information content in a template ensemble, which we call the 'parasite threshold'. If the parasite threshold of cooperative system is lower than the error threshold, template coexistence does not solve the prebiotic information problem.

We study the information carrying potential in a spatial eco-evolutionary model, based on the metabolic model

(Czárán and Szathmáry, 2000). We define a surface in which each site is either empty or occupied by a molecule. A genome consists of d different templates of length $l=i/d$ nucleotides, where i is the information content of the genome. We assume that replication of a template is only possible when all d functional template types are present in the local neighborhood. Mutation in a functional template results in a nonfunctional parasitic copy, with probability m per nucleotide. Growth and decay rates of functional molecules and parasites are equal. Every reaction step is followed by diffusion of the molecules over the surface.

Genomes with varying number of fragments ($d=1,\dots,4$) are compared under high mutation load ($m=0.01$). We observed the parasite numbers and the maximum maintainable genomic information. Our results show that, surprisingly, the information storage capacity is highest for unfragmented, single replicator, genomes ($d=1$), despite their high production of parasite. While the number of parasites decreases with the number of templates, the vulnerability of fragmented genomes to parasitism sharply increases. Because the latter effect outweighs the former, the benefit of template coexistence is lost. When the different genomic strategies are put in direct competition with each other, fragmentation can out-compete the single template genome, but only in a situation with low mutational rate, and length-dependent growth rate. Close to the error threshold, however, fragmented genomes are competitively excluded by the single template strategy.

We conclude that template coexistence by itself does not solve the prebiotic information problem, because cooperative systems are limited by the 'parasite threshold'. We demonstrate that in the metabolic model, template coexistence does not increase information content and is excluded in direct competition. Although more realistic conditions concerning catalytic specificity, length-dependent neutrality and growth may refine these results, it is clear that limitations arising from cooperation must be taken into account in solving the information problem in prebiotic evolution.

How might group selection explain the major evolutionary transitions?

John Bryden

University of Leeds

Wednesday 3:30pm–5:00pm, Exam Hall Room 2, Major Evolutionary Transitions 2.

The central idea of group selection is that an individual will reduce its fitness so that the mean fitness of a group (of possibly non-related individuals) may increase. This is relevant to the major evolutionary transitions where an individual will cooperate by stopping reproduction on its own and reproduce instead as part of a group. Explaining the major evolutionary transitions should simply be a case of applying models of group selection. However,

group selection may only work when there is a small difference between the fitness of a cooperator and a defector (Traulsen et al., 2006, Proc. Nat. Acad. Sci. U.S.A., 103, 10952).

A new approach (Bryden, 2008, PhD thesis, University of Leeds) sheds fresh light on this topic however. We must carefully look at what we mean by fitness and what we mean by group.

Recent perspectives on fitness (Metz et al., 1992, Trends Ecol. Evol., 7, p.198) argue that fitness should be calculated over a range of environments. This contrasts with the Hamiltonian perspective of fitness (Hamilton, 1964, J. Theor. Biol., 7, p.1) which is the number of adult offspring. By calculating over a range of environments, some traits which prosper in some environments decay in other environments. Tools are available (e.g., Tuljapurkar, 1990, Proc. Nat. Acad. Sci. U.S.A., 87, p.1139; Bryden 2008) for modelling long-run growth rates over varied environments.

To apply this long-run perspective on fitness to the major evolutionary transitions, resource allocation strategy modelling has been done by Bryden (2005, ECAL, p.551; 2007, ECAL, p.645; 2008). The problem of the major evolutionary transitions is reformulated as a question as to whether an individual will invest resources in a higher reproductive process: a process of generating new offspring with two or more individuals having some genetic stake in, and contributing resources to, the new offspring. When an individual reproduces clonally, it will grow faster in favourable environments than those that contribute toward a higher reproductive process. However, reproduction can be risky and leave the fast reproducing lineage dangerously low on resources during unfavourable environments.

Analytic methods and computer simulations have shown how a strategy of collective reproduction (by sharing resources between several individuals and one offspring equally) can dominate a strategy of producing clonal offspring (Bryden 2007). When individuals are selfish and only contribute minimal resources, increases in the amplitude of environmental resource fluctuations becomes increasingly significant (Bryden 2008). The reason the shared strategy is successful is because the clonal strategy is very weak in harsh environments.

These results demonstrate that it is plausible that an individual may lower its Hamiltonian fitness (i.e., its reproductive output) to increase its long-run fitness by contributing to a higher reproductive process (such as those in the major evolutionary transitions). If I suggest a definition of a group as a lineage that is temporally spread out across several environmental eras, rather than all being present at the same time, we may then compare the long-run fitnesses of lineages to determine the most successful. In other words, the group that has the greatest long-run fitness is selected for. This theory calls for verification through scientific experiments and expansion through fur-

ther modelling of the major evolutionary transitions.

Building artificial personalities: expressive communication channels based on an interlingua for a human-robot dance

John Bryden, David Hogg, Sita Popat and Mick Wallis

University of Leeds

Wednesday 1:30pm–3:00pm, Exam Hall Room 1, Robots & Behaviour.

The development of artificial personalities requires that we develop a further understanding of how personality is communicated. This can be done through developing human-robot interaction (HRI). In this paper we report on the development of the SpiderCrab robot. This uses an interlingua based on Laban Movement Analysis (LMA) to intermediate a human-robot dance. Specifically, we developed metrics to analyse data in real time from a simple vision system and implemented a simple stochastic dancing algorithm on a custom built robot. The system was tested with professional dancers and members of the public and the results (formal and anecdotal) are presented herein.

Dynamical complexity of spatially embedded networks

Christopher Buckley¹, Lionel Barnett² and Seth Bullock¹

¹Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

²Centre for Computational Neuroscience and Robotics, University of Sussex

Tuesday 10:30am–12:00pm, Exam Hall Room 1, Spatial Organisation 1.

Living systems are embedded within physical space. While this embedding can be viewed as a restrictive constraint on structure, where the prohibitive costs of establishing or maintaining interactions over long distances mitigate against certain kinds of potential organisation, it can also be seen as an enabling factor, bringing about correlations, regularities and symmetries that can be exploited by evolution. Artificial life research on spatially embedded games, ecologies, networks, evolution, and agents has shown that projecting a well-mixed system into a low-dimensional medium and constraining interactions to be local can confer interesting properties (e.g., stability, honesty, robustness to parasites) that are otherwise absent or unstable. This paper explores the question: what is the contribution of spatial embedding to the dynamical complexity of networks.

In previous work some of us have developed a general framework for characterising the impact of spatial constraints on network topology (Barnett et al., Phys. Rev. E 76, 056115, 2007), and some of us have explored the dynamical complexity of spatially embedded artificial neural networks (Buckley & Bullock, ECAL 2007). Here we

combine these two threads to discover what graph theoretic properties of networks confer high dynamical complexity, and to explore the extent to which spatial embedding tends to encourage exactly these topological properties in networks that are random in other respects.

We first return to the original formulation of the dynamical complexity measure due to Tononi, Sporns and Edelman (PNAS 91, 5033, 1994) and correct an error in a widely used approximation of this measure. This correction impacts on intuitions about the structural and functional roots of dynamical complexity. However, we are able to rescue these intuitions by re-deriving the approximation for a continuous-time dynamical system rather than the discrete dynamical system used in the original formulism. This process emphasises some key differences between the dynamics of continuous and discrete dynamical systems.

We then go on to derive and extend a graph theoretical interpretation of dynamical complexity for the corrected discrete measure and the new continuous measure. This allows us to strengthen our understanding of the relationship between properties of spatially embedded structures and high complexity. In particular, we are able to concretise the notion that the structural contribution of spatial embedding to high dynamical complexity results from the introduction of cycles of connectivity at many structural scales. Furthermore, we are able to address a misconceived equivalence between the small world property and systems of high “dynamical complexity”. Specifically, we find that while systems of high dynamical complexity may possess the small world property, neither property is either necessary or sufficient for the other.

Basic principles of adaptive learning through variation and selection

Mikhail Burtsev

Keldysh Institute of Applied Mathematics of RAS

Wednesday 3:30pm–5:00pm, Stripe Studio 1, Learning, Action Selection, & Attention.

The evolutionary theory relies on the principles of variation and selection to explain adaptation. It is reasonable to fit these powerful principles to the learning theory. A number of selectionist approaches were proposed but found modest recognition so far. This theoretical paper attempts to review an application of basic ideas of the evolutionary adaptation to the lifetime learning. The analysis demonstrates that an adaptive value can be translated from the level of evolution to the level of individual through the innate repertoire of behaviors. This primary repertoire forms initial attractor for the behavioral dynamics. Learning starts when an environment offsets an organism from the existing attractor trajectory. Blind variations of behavior are generated until the return to the target attractor. These variations are retained to make up new branches of

basin of attraction. It is important that the existed behavioral trajectory should not be altered as the learning unfolds because it keeps knowledge about adaptations survived selection through the evolutionary and learning history.

Learning drives the accumulation of adaptive complexity in simulated evolution

Mikhail Burtsev¹, Konstantin Anokhin² and Patrick Bateson³

¹M.V.Keldysh Institute of Applied Mathematics, Moscow

²P.K.Anokhin Institute of Normal Physiology, Moscow

³Sub-Department of Animal Behaviour, University of Cambridge

Wednesday 10:30am–12:00pm, Stripe Main Theatre, Open-Ended Evolution 1.

Developmental plasticity and particularly learning enable organisms to cope with new environmental challenges. But if learning is costly, the same behavior could evolve through Darwinian modifications of development that substitute for the role of learning in the acquisition of that behavior, a hypothesis known as the Baldwin effect. Computer simulations have confirmed that learning accelerates evolutionary adaptation to a single problem posed by the environment. What has not been shown, however, is the way in which the driving force of learning can generate ever greater complexity in organization of evolved behavior, the one that has very small chance to appear in one step in the course of evolution. Here we report such consequences of the role of learning using a model fitted with sequentially appended adaptive systems.

A central component of our approach is that of a functional system. It asserts that each adaptive behavior is executed by a distributed system of phenotypic elements that cooperate towards organism's fitness in this particular task. Novel challenges lead to generation of new adaptive functional systems that can be added to the existing ones either by evolution of development or by learning. This allows the establishment of elaborate behavior patterns and results in increased complexity of organisms at the systems level. Selection assesses organisms by the adaptiveness of their functional systems and the more functions the individual possesses, the more competitive we assume it is. The greater complexity in our model implies more extensive repertoire of behaviors supported by greater amounts of equipment for monitoring and coping with the environment. From a biological standpoint an organism with the higher functional complexity will be better able to deal with a variety of challenges from the environment and therefore will be more likely to survive. If considered from the perspective of a single ecological challenge requiring just one functional system our model is similar to a classical single-peaked landscape simulation by Hinton & Nowlan. However, the main highlight of the model is its operation in a complex evolutionary landscape similar to the "Royal Staircase" fitness function of

van Nimwegen & Crutchfield, which allowed us to examine coordinated evolution of multiple functional systems under the impact of learning and developmental plasticity.

The results of simulations demonstrate that ability to learn dramatically accelerates the evolutionary accumulation of adaptive systems in model organisms with relatively low rates of mutation. The growth of complexity is mediated through a process of allelic substitutions that simulate emergence of evolutionary predispositions for learning of certain behaviors and simultaneously release organisms capacities for acquisition of next tasks. The effect of learning on evolutionary growth of complexity is even greater when the number of elements required for adaptive system is increased. These results suggest that as the difficulty of challenges from the environment become greater, so learning exerts an ever more powerful role in meeting those challenges and in opening up new avenues for subsequent genetic evolution of complex adaptations.

Design strategies for open-ended evolution

Peter Cariani

Harvard Medical School

Wednesday 10:30am–12:00pm, Stripe Main Theatre, Open-Ended Evolution 1.

Open-endedness is an important goal for designing systems that can autonomously find solutions to combinatorically-complex and ill-defined problems. We distinguish two modes of creating novelty: combinatoric (new combinations of existing primitives) and creative (new primitives). Although combinatoric systems may differ in numbers of possible combinations, their set of possibilities is closed. Creative systems, on the other hand, have open-sets of possibilities because of the partial- or ill-defined nature of the space of possible primitives. We discuss classes of adaptive and self-modifying cybernetic robotic devices in terms of these two kinds of processes. We consider such systems whose hardwares are constructed from genetically-directed pattern-grammars. Here although the space of accessible structures is closed, the space of functions is open. We conclude that genome sequence spaces and gene-product structure spaces are closed, whereas, being ill-defined, phenomic function-spaces are potentially open. Some general principles are proposed for designing and building systems with open-ended capabilities.

Dynamically adapting parasite virulence to combat coevolutionary disengagement

John Cartlidge

University of Central Lancashire

Tuesday 10:30am–12:00pm, Stripe Lecture Room, Evolutionary Theory 1.

Participating in an evolutionary arms-race, natural coevolutionary predator-prey and host-parasite systems often exhibit accelerated evolution. Competitive Coevolutionary Genetic Algorithms (CCGAs) attempt to harness this evolutionary acceleration by engaging (multiple) evolving populations in competitive self-play.

By evaluating individuals competitively, CCGAs afford the possibility of tackling problems that are ill-defined, open-ended and lacking in formalism. This offers CCGAs a potential advantage over more traditional Genetic Algorithms (GAs) when fitness evaluation is difficult to operationally define. Analogously, one imagines that it is much easier to formally define the rules of the game of tennis than it is to define tennis playing ability. In practice, however, defining an *appropriate* game is often non-trivial.

Competitive evaluation leaves CCGAs susceptible to some adverse evolutionary dynamics. One such hindrance is “disengagement”. This occurs when one coevolving population gets the upper hand and begins to easily outperform the other. Since it becomes impossible to discriminate between individuals according to ability, the selection gradient disappears and the coevolving populations begin to stagnate. The result is a stymied system that is left to flounder aimlessly.

To prevent disengagement, the author has previously introduced the “Reduced Virulence” technique (Cartlidge & Bullock, 2004, *Evol. Comp.*, 12, p.193). This technique helps avoid disengagement by reigning in a population that inherits an advantageous bias. Rather than reward individuals that maximally damage a competitor, Reduced Virulence favors individuals that give opponents a chance. Perhaps counter-intuitively, Reduced Virulence enables accelerated evolutionary progress by disadvantaging a population’s most successful individuals.

In this work, Reduced Virulence undergoes a rigorous sensitivity analysis in the Counting Ones domain (introduced by Watson & Pollack, 2001, *GECCO*, p.702, Morgan Kaufmann); an analytically tractable substrate designed to highlight the dynamics of coevolution. Following intuition, it is shown that for optimal performance, virulence should be increasingly reduced as the asymmetrical bias (and thus likelihood of disengagement) between coevolving populations increases. Interestingly, even when coevolution is unbiased, “Maximum Virulence” — equivalent to the canonical fitness evaluation of “reward all victories” — is shown not to be ideal. Thus, results suggest that (in the Counting Ones domain) when population sizes

are small, it is never the case that the canonical coevolutionary setup should be favored. The generality of this result, however, is an open question.

Utilizing this information, a novel “Dynamic Virulence” algorithm is introduced. This algorithm adapts population virulence over time as populations evolve. It is shown that Dynamic Virulence is able to cope with varying bias better than fixed virulence and allows the discovery of optimal solutions under a much wider range of conditions than any individual fixed virulence setting.

Finally, it is discussed how analyzing the role of virulence in artificial systems may allow us to better understand virulence in nature. For instance, perhaps there is potential for a “Reduced Virulence” approach to tackling infectious diseases. Rather than killing mosquitoes to eradicate malaria, one could alternatively encourage malaria-resistant strains that are better able to survive.

On the ability of swarms to compute the 3-coloring of graphs

Blanca Cases, Carmen Hernandez, Manuel Graña and Alicia d’Anjou

UPV/EHU

Tuesday 3:30pm–5:00pm, Exam Hall Room 1, Swarm Systems.

Swarm Intelligent Systems are computational models of the navigation of populations as a behavior emerging from locally controlled movements, based on decisions taken on local information. The increment of agents’ cognitive capabilities, endowing them with memory and with the ability of selecting the rules of movement depending on an internal state, gave place to Ant Colonies, Particle Systems Optimization, and allowed the application of Self-Organizing Particle Systems to heuristic problem solving. Our focus in this work is on the complex emergent behavior that arising from endowing the individuals with another elementary cognitive ability: the perception of the affinity of another individual. The individual perceives another individual as belonging to We or to Them. The first class is attractive while the second is repulsive. The first class is associated with amity, security and comfort while the second is interpreted as danger, enemies and things to avoid. This localized individual perception produces the emergence of teams and classes at a global level. In this work we formulate this behavior in terms of Designing for Self. We have found empirically that Self-Organizing Particle Systems perform the task of 3-coloring graphs with the same precision than the Brelaz coloring heuristic, which is the best greedy heuristic known for this purpose.

The effects of periodic and continuous market environments on the performance of trading agents

Satpal Singh Chaggar¹, Jason Noble¹ and Dave Cliff²

¹School of Electronics and Computer Science, University of Southampton

²University of Bristol

Tuesday 3:30pm–5:00pm, Stripe Studio 1, Markets.

Simulation experiments are conducted on simple continuous double auction (CDA) markets based on the experimental economics work of Vernon Smith. CDA models are usually periodic, with an experiment consisting of a sequence of discrete trading days. In our experiments we look at both periodic and continuous-time (always open) versions of the CDA. In this we build on the work of Cliff and Preist (2001) with human subjects, but we replace human traders with Zero Intelligence Plus (ZIP) trading agents, a minimal algorithm that can produce equilibrating market behaviour in CDA models. Our results indicate that continuous-time markets are similar to periodic markets in their ability to show equilibration dynamics. Secondly we show that although both models produce the same behaviour of price formation, they are different playing fields, as periodic markets are more efficient over time than their continuous-period counterparts. We also find, however, that the volume of trade in periodic markets is concentrated in the early period of each trading day, and the market is in this sense inefficient. We look at whether ZIP agents require different parameters for optimal behaviour in each market type, and find that this is indeed the case. Overall, our conclusions mirror earlier findings on the robustness of the CDA, but we stress that a continuous-time marketplace equilibrates in a different way to a periodic one.

A measure for natural selection's contribution to the origins and maintenance of organismal complexity

Alastair Channon

University of Keele

Wednesday 1:30pm–3:00pm, Stripe Main Theatre, Open-Ended Evolution 2.

The aim of Artificial Life research into Open-Ended Evolution is, initially at least, to develop artificial evolutionary systems in which new adaptive traits continue to evolve and the maximum complexity of organisms, ecosystems or behaviours continues to increase. The main proponents of this approach have presented systems that invoke natural or biotic selection, as opposed to artificial or abiotic selection, as the drive for both the generation of new adaptive traits and (potentially) a sustained increase in complexity. However, within both Biology and Artificial Life, doubts have been raised as to natural selection's role as the

drive for increasing complexity (Lynch, 2007, Proc. Natl. Acad. Sci., 104, p.8597; Miconi, 2008, Artif. Life, 14, in press), with the suggestion put forward that nonadaptive evolutionary forces (such as mutation, recombination and genetic drift) or mathematical/statistical constraints may be the primary drives, through either a passive increase in variance of complexity in the presence of a lower bound, or a constraint-driven drive toward complexity. The question therefore arises, how to determine natural selection's contribution to increases in complexity?

This work introduces first a measure for the phenotypic complexity of an individual, based on the class of components previously used in the population-level analysis of an Artificial Life system classified as exhibiting unbounded evolutionary dynamics (Channon, 2006, Genet. Program. and Evolvable Machines, 7, p.253): components that are approximately equivalent to the biological notion of a gene or coding DNA; and second a measure for the contribution made by selection to increases in complexity, based on the mechanism and methodology used in the development and application of component-normalised activity statistics to that system. These measures enable us to address the question posed above, about a fundamental aspect of evolution in general, in a way that would not be possible given the biological world alone. They also provide a mechanism for detecting increases in phenotypic complexity and attributing them to either adaptive or non-adaptive forces.

Results from the application of the measures to evolution in runs of the above system suggest that, according to these definitions, natural selection initially (but only briefly) opposes the level of increases in complexity (new active/coding DNA) that would be brought about by the nonadaptive forces alone, presumably because the new active/coding DNA would be nonadaptive; but that as evolution progresses, natural selection maintains and drives the increase in adaptive complexity with remarkable consistency: natural selection can (and does, in this system at least) provide a sustained drive toward increasing complexity.

An adaptive model of marine biogeochemistry in the Archaean

James Clark and Hywel Williams

University of East Anglia

Friday 10:30am–12:00pm, Exam Hall Room 1, Entropy.

Here we describe ongoing work that applies adaptive modelling techniques from artificial life to open questions in Earth system science. Understanding the Earth system in terms of global chemical cycling is of critical importance for the interpretation of Earth history and the prediction of climate change. Many of the key chemical reactions that facilitate biogeochemical cycling occur during the metabolism of organisms. The range of metabolic reactions present in the biota has changed significantly over

evolutionary timescales, often with dramatic effects on the Earth system. Yet models of biogeochemistry have traditionally only included static representations of the biotic components of the major nutrient cycles. More recently, some models have begun to include a number of functional types of organism, each with different prescribed biogeochemical properties. However, none of these models address the dynamic adaptation of the biota over time, which can change the way in which different species interact and alter their effects on biogeochemical cycling. Here we present a new adaptive individual-based model of the marine ecosystem in the Archaean period of Earth history. We use a simplified version of the major physical chemical processes and metabolic functions that are known to have existed at this time; the marine ecosystem during the Archaean is in any case simpler than the modern ecosystem, being solely based on microbial life. We specify a number of microbial guilds containing species that have similar metabolic reactions and biogeochemical functions, e.g., photosynthesisers, chemoheterotrophs, etc. Within each guild, multiple species may coexist and compete on the basis of various physiological traits. Individual microbes each have a genotype that specifies various metabolic traits and thus determines their growth rate as a function of their environment. We consider each model individual to represent an aggregation of many genetically similar real world individuals; this assumption allows us to study phenomena on a global scale. Successful microbes (i.e., those that are well suited to their environment) grow and reproduce, while unsuccessful microbes starve and die. Mutation can occur during reproduction, allowing the creation of new species. Competition for nutrients drives ongoing adaptation that dynamically changes the chemical environment in a coevolutionary loop of interaction. Our model considers a vertical column separated into three compartments representing the deep and surface ocean layers and the atmosphere. We find that diverse self-sustaining ecosystems emerge over time without being prescribed, and that the distribution of nutrients and organisms between the three compartments is qualitatively similar to that believed of the Archaean. In particular, photosynthesisers tend to dominate the sunlit surface ocean, fixing inorganic carbon (CO_2) into organic forms that support other populations, while the deep ocean ecology is dominated by methanogens, which are able to survive in the dark and anoxic deep ocean conditions. Adaptation of nutrient uptake and light sensitivity traits creates species within each guild that are optimally suited to their environment. Recent work has looked at evolutionary trade-offs and the possibility of a biological trigger for the Great Oxidation event. This work represents the first step in a greater program of study that will seek to model evolutionary adaptation in marine biogeochemistry.

Measuring the robustness of a developmental system based on sequential growth rules

Michael Cohen¹, Mark Miodownik² and Buzz Baum¹

¹UCL

²Kings College London

Wednesday 3:30pm–5:00pm, Stripe Lecture Room, Development & Morphogenesis 3.

Understanding how complex structures emerge from localised interactions in a robust way is essential to unravelling the mechanisms that underlie developmental processes in both biological and artificial systems. This study investigates the effects of genome complexity on robustness using a simple, evolved developmental system in which cellular automata (CA) rules are applied in sequence in order to generate a 1D pattern of cells. The system employs a 1D two state CA with 128 distinct nearest neighbour update rules. Each developmental run is initiated with a single cell. The cell update rules adopted by every cell at each time-step are allowed to change sequentially at different times according to the instructions contained in a ‘genome’. In order to generate a set of productive developmental programs for this analysis, a genetic algorithm was used to select for individuals whose cell states, after a fixed number of time steps, match a set of predefined target patterns. This was repeated for genomes of different sizes. The robustness of evolved and randomized CA patterns were compared by systematically applying single cell state perturbations during pattern development. This analysis revealed that in these evolved systems genome size has a positive effect on robustness by freeing the system to generate patterns using a relatively unbiased set of rules, which have very different individual properties. In contrast, smaller genomes are frequently forced to rely on complex patterning rules to generate complex patterns, which amplify damage and hence reduce their robustness. In addition, pattern size (the number of cells) was found to be a major factor in the measured robustness in this system. This is because the cumulative damage induced by developmental perturbations does not scale with pattern size. As a result, increasing pattern size reduces the percentage damage following perturbations and improves overall robustness. In conclusion, we have shown that pattern robustness is an additive effect of the ability of individual rules to propagate and heal defects resulting from environmental perturbation in this simple CA system, and is potentially increased by increasing pattern size and genome size. These results have implications for our understanding of robustness in biological and artificial systems.

“Soft” continuum robots: the interaction of continuous and discrete elements

Lara Cowan and Ian Walker

Clemson University

Wednesday 10:30am–12:00pm, Exam Hall Room 2, Amorphous & Soft Robotics.

In this paper, we examine key issues underlying the design and operation of “soft” robots featuring continuous body (“continuum”) elements. We contrast continuum and continuum-like robots created to date with their counterparts in the natural world. It is observed that natural continuum locomotors or manipulators almost invariably rely on hard/discrete elements (in their structure and/or operation) in their interactions with their environment. Implications for the successful operation and deployment of continuum robots are identified and discussed.

From single cell to simple creature morphology and metabolism

Sylvain Cussat-Blanc, Herve Luga and Yves Duthen

IRIT

Wednesday 3:30pm–5:00pm, Stripe Lecture Room, Development & Morphogenesis 3.

In order to produce diversity in virtual creatures to populate virtual worlds, different techniques exist. Some of these use blocks or sticks. In this morphological approach, blocks and sticks can be considered as organs, which means body parts able to perform different functions. Another approach, artificial embryogenesis, consists in developing organisms from a single cell. In this paper, we propose a bridge between these two approaches: a model that will create creatures with a particular morphology and which is organized in organs. The creature development will start from a single cell. In this paper, we propose a unique model able to produce organisms that perform a specific function and to produce organisms with a user-defined morphology.

A model chemical memory in an evolved animat

Kyran Dale

University of Sussex

Tuesday 3:30pm–5:00pm, Exam Hall Room 2, Chemical & Molecular Modelling.

This paper describes work carried out to investigate whether a classic simulated reaction-diffusion (RD) system could be used to control a ‘minimally cognitive’ animat during the course of a simple memory test. This test required the animat to remember an arbitrary signal and adapt its behaviour as a result. A further requirement was that the effects of the first signal be reversed by a second signal, returning the animat to its default behaviour. In this way the two signals combined to form a behavioural switch, regulated by a memory-trace preserved in the ho-

mogeneous chemical substrate. The reaction-diffusion system chosen was that first described by Gray and Scott (Gray-Scott) and the minimally cognitive behavior of a class introduced by Beer et. al, involving the fixation and avoidance of a falling circular object by a whiskered animat. The parameters of this RD-controller were evolved using an evolutionary, or genetic, algorithm (GA).

Protein folding with stochastic L-systems

Gemma Danks, Susan Stepney and Leo Caves

University of York

Tuesday 3:30pm–5:00pm, Exam Hall Room 2, Chemical & Molecular Modelling.

Protein molecules adopt a specific global 3D structure in order to carry out their biological function. To achieve this native state a newly formed protein molecule has to fold. The folding process and the final fold are both determined by the sequence of amino acids making up the protein chain. It is not currently possible to predict the conformation of the native state from the amino acid sequence alone and the protein folding process is still not fully understood. We are using L-systems, sets of rewriting rules, to model the folding of protein-like structures. Models of protein folding vary in complexity and the amount of prior knowledge they contain on existing native protein structures. In a previous paper we presented a method of using open L-systems to model the folding of protein-like structures using physics-based rewriting rules. Here we present an L-systems model of protein folding that uses knowledge-based rewriting rules and stochastic L-systems.

Phenotype-based evolution of complex food webs

Walter de Back¹, Sergio Branciamore² and George Kampis¹

¹Collegium Budapest, Institute for Advanced Study, Hungary

²Department of Evolutionary Biology, University of Florence, Italy

Thursday 3:30pm–5:00pm, Stripe Main Theatre, Ecological Modelling 2.

Explaining the relation between structure and dynamic of food webs remains one of the most daunting challenges in ecological theory. The trophic structure is a fundamental property of ecological communities that relates their diversity to productivity and stability. Understanding the generic properties of food webs requires insight in the way these networks are assembled. Although ecological communities are often augmented by invasion rather than adaptive speciation, the diversity in an ecological network can ultimately only be explained from an evolutionary perspective. Recently, several aggregate population models for the evolution of food webs have been proposed. Here, we present an individual-based model in which we show

the evolutionary emergence of complex trophic networks from the bottom-up.

We use a spatial individual-based model with a non-linear mapping between genotype and phenotype. An individual is specified by a genotype that determines its phenotype in a redundant (“many-to-one”) and epistatic (“one-to-many”) fashion. All behavioral/ecological properties of the individual (e.g. reproduction/mortality rate, auto/heterotrophy, prey preference, etc.) are derived from specific aspects of the phenotype, such that trade-offs in ecological function are inherently introduced. Autotrophs only consume abiotic resources, while heterotrophs can consume individuals in their spatial neighborhood. The outcome of the consumption interactions depends on prey preference and the distance between both phenotypes. Mortality and (asexual) reproduction are based on energy that decreases linearly, and is increased by consumption. Mutations occur by substitutions in the genotype with a low probability per locus.

The system is inoculated with a single genotype coding for a random autotroph. After the appearance of heterotrophs, coevolution between auto- and heterotrophs causes phenotypic diversification. The number of species (i.e. individuals with same phenotype) varies around 40-60 for the given system size (+/- 75000 individuals), of which approximately half are heterotrophs. The evolved species show a large variation in life history and consumption patterns, balancing various trade-offs. The structure and composition of the network can continually change, although evolutionary stability increases over time. We assess the ecological stability of the evolved food webs by canceling mutation. After a partial collapse, the truncated food web (with typically 10-15 species) persists over an indefinitely long period, during which it shows chaotic population dynamics.

We demonstrate the evolutionary emergence of ecological stable food webs in an individual-based model. The genotype-phenotype mapping provides efficiency and robustness in the exploration of phenotype space, while spatial interactions stabilizes population dynamics and preserves diversity. By encapsulating the ecological properties and their trade-offs in the phenotypes, rather than defining them as global state variables, allows for the niche differentiation by which food webs are assembled.

Unexpected evolutionary dynamics in a string based artificial chemistry

James Decraene, George G. Mitchell and Barry McMullin
Dublin City University

Tuesday 10:30am–12:00pm, Stripe Main Theatre, Evolving Cell Signalling Networks *in silico* 1.

This work investigates closure in Cell Signaling Networks, which is one research area within the ESIGNET project. To assist this research, we employ the Molecular Classifier Systems(MCS.b), a string-based Artificial Chemistry based on Holland’s broadcast language. In this paper we present a series of experiments which use the MCS.b, these experiments focus on the emergence and evolution of self-maintaining molecular organizations. Such experiments naturally relate to similar studies conducted in artificial chemistries such as Tierra, Alchemy and Alpha-Universes. However, our results indicate counter intuitive outcomes as opposed to those suggested in the literature. Each of these “unexpected” evolutionary dynamics (including an elongation catastrophe phenomenon) are examined and explained both informally and formally. We also demonstrate how the elongation catastrophe can be prevented using a multi-level selectional model of the MCS.b (which acts both at the molecular and cellular level). This work provides complementary insights into the understanding of evolutionary dynamics in minimal artificial chemistries.

Predictive information and emergent cooperativity in a chain of mobile robots

Ralf Der¹, Frank Güttler² and Nihat Ay¹

¹MPI MIS Leipzig, Germany

²University Leipzig

Thursday 3:30pm–5:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 3.

Measures of complexity are of immediate interest for the field of autonomous robots both as a means to classify the behavior and as an objective function for the autonomous development of robot behavior. In the present paper we consider predictive information in sensor space as a measure for the behavioral complexity of a chain of two-wheel robots which are passively coupled and controlled by a closed-loop reactive controller for each of the individual robots. When in a maze with many obstacles, the predictive information (approximated by the mutual information in the time step) of the sensor values of an individual robot is found to have a clear maximum for a controller which realizes the spontaneous cooperation of the robots in the chain so that large areas of the maze can be visited.

Life in time: the missing temporal dimension in autopoiesis

Ezequiel Di Paolo

University of Sussex

Thursday 10:30am–12:00pm, Stripe Studio 1, Philosophical Issues 1.

There is a widespread view in the artificial life community that life is not so much about materiality but about organization. However, one of the favourite candidate theories that explain what this organization should be like has an ambivalent position with respect to materiality. This is the theory of autopoiesis (Maturana and Varela, 1980, *Autopoiesis and cognition*, D. Reidel). Accordingly, autopoiesis is the self-production of a unity in the domain of processes of material construction, transformation and destruction. The identity of the living system, in this view, is sustained in its form in spite (or rather because) of its material flux.

This notion seems to imply an inherent temporal dimension to autopoiesis. It is, intuitively, a dynamical concept. However, because of its insistence on a set-theoretic focus on conservation of autopoiesis, the theory says little about the temporality of life (leaving relevant phenomena such as stress, fatigue, pathologies and development untouched).

Efforts have been made to bring thermodynamic material constraints into the theory of autopoiesis (Ruiz-Mirazo and Moreno, 2004, *Artificial Life*, 10, p. 235). The implications are not trivial, nor are they foreseen in the theory as postulated in the original literature and interpreted purely in formal terms. One immediate consequence of realising autopoietic organizations in dissipative physical structures far from equilibrium is that there is an obvious time arrow introduced into the process of life: the arrow of thermodynamics. Bare autopoiesis, paradoxically, does not present us with a similar time arrow (a time-reversal thought experiment leads to this conclusion).

But this temporality belongs to the nature of dissipative processes and is in some sense only inherited by life because such material processes constitute it.

I shall argue that living systems enjoy a different kind of temporality, given by their own interactive and teleological organization. This temporality is richer and different from that of the background time's arrow — it is a consequence of expanding the theory of bare autopoiesis with the notion of adaptivity, (Di Paolo, 2005, *Phenomenology and the Cognitive Sciences*, 4, p. 429). This temporality is characterized by intentional direction, minimal granularity, rhythmicity, and the presence of historical transitions (in behaviour and development). It belongs to the organization of an adaptive autopoietic system under precarious circumstances.

It is conceivable that this temporality could emerge even if the system were not subject to thermodynamical con-

straints since it is a consequence of the higher order relations between interactive and constitutive aspects of self-maintenance. It is also conceivable for aspects of the temporality of life to contradict the temporality of the thermodynamics time's arrow, since the temporality of life is inherently related to its intentionality. Even at a minimal level, a living system may retroactively alter the virtual possibilities of the past through its sense-making activity in the present.

Chemical organizations in living systems

Peter Dittrich¹, Florian Centler², Christoph Kaleta³, Naoki Matsumaru³ and Pietro Speroni di Fenizio⁴

¹Friedrich Schiller University Jena

²UFZ Leipzig

³FSU Jena

⁴DCU Dublin

Thursday 3:30pm–5:00pm, Exam Hall Room 2, Dynamical Systems Analysis 1.

Complex dynamical reaction networks consisting of many components that interact and produce each other are difficult to understand, especially, when new components may appear over time. In this talk, I outline a theory, which has been inspired by artificial chemistry research, to deal with such systems. It has been successfully applied to regulated metabolic networks, virus-immunsystem dynamics, chemical information processing, chemical evolution, and planetary atmosphere photochemistries. I will show how the approach can be used to predict growth phenotypes and to evaluate the quality of large bio-models. The theory consists of two parts. The first part introduces the concept of a chemical organization as a closed and self-maintaining set of components. This concept allows to map a complex (reaction) network to its set of organizations. The theory provides a new view on the system's "organizational structure", which is fundamentally different from a pathway-oriented view. The second part of the approach connects dynamics with the set of organizations, providing a link to classical dynamical systems theory, e.g., by mapping a movement of the system in state space to a movement in the set of organizations. It is shown that every dynamically stable state must be an instance of an organization.

Artificial-life ecosystems: what are they and what could they become?

Alan Dorin¹, Kevin Korb¹ and Volker Grimm²

¹Monash University

²Center for Environmental Research, Leipzig-Halle

Tuesday 10:30am–12:00pm, St Edburga Room 4, Ecological Modelling Methodology.

This paper summarises the history of the terms ecology and ecosystem, before examining their application in the early and recent literature of A-Life agent-based software simulation. It investigates trends in A-Life that have led

to a predominance of simulations incorporating artificial evolution acting on generic agents, but lacking a level of detail that would allow the emergence of phenomena relating to the transfer and transformation of energy and matter between the virtual abiotic environment and biota. Implications of these characteristics for the relevance of A-Life's virtual ecosystem models to Ecology are discussed. We argue a position that the inclusion of low-level representations of energetics, matter and evolution, in concert with pattern-oriented modelling techniques from Ecology for model validation, will improve the relevance of A-Life models to Ecology. We also suggest two methods that may allow us to meet this goal: artificial evolution can be employed as a mechanism for automating pattern-oriented ecological modelling from the level of individual species up to that of the ecosystem, or it may be employed to explore general principles of ecosystem behaviour over evolutionary time periods.

Programmable architectures that are complex and self-organized: from morphogenesis to engineering

Rene Doursat

Complex Systems Institute, CNRS & Ecole Polytechnique, Paris
Wednesday 3:30pm–5:00pm, Stripe Lecture Room, Development & Morphogenesis 3.

Outside biological and social systems, natural pattern formation is essentially “simple” and random, whereas complicated structures are the product of human design. So far, the only self-organized (undesigned) *and* complex morphologies that we know are biological organisms and some agent societies. Can we export their principles of decentralization, self-repair and evolution to our machines, networks and other artificial constructions? In particular, can an “embryomorphic” engineering approach inspired by evo-devo solve the paradoxical challenge of planning autonomous systems? In this work, I wish to better understand and reproduce complex morphogenesis by investigating and combining its three fundamental ingredients: *self-assembly* and *pattern formation* under *genetic regulation*. The model I propose can be equivalently construed as (a) *moving* cellular automata, in which cell rearrangement is influenced by the pattern they form, or (b) *heterogeneous* collective motion, in which swarm agents differentiate into patterns according to their location. It offers a theoretical framework for exploring the causal and programmable link from genotype to phenotype.

Entropy production in an energy balance Daisyworld model

James Dyke

University of Sussex

Friday 10:30am–12:00pm, Exam Hall Room 1, Entropy.

Daisyworld is a simple mathematical model of a planetary system that exhibits self-regulation due to the nature of feedback between life and its environment. A two-box Daisyworld is developed that shares a number of features with energy balance climate models. Such climate models have been used to explore the hypothesis that non-equilibrium, dissipative systems such as planetary atmospheres are in a state of maximum entropy production with respect to the latitudinal flux of heat. When values for heat diffusion in the two-box Daisyworld are selected in order to maximize this rate of entropy production, the viability range of the daisies is maximized. Consequently planetary temperature is regulated over the widest possible range of solar forcing.

Strategies for maintaining large robot communities

Stephen English, Jeffrey Gough, Alexis Johnson, Robert Spanton, Joanna Sun, Richard Crowder and Klaus-Peter Zauner

School of Electronics and Computer Science, University of Southampton
Tuesday 10:30am–12:00pm, Stripe Studio 1, Autonomous Robots.

The confluence of progress in micro-actuators, power sources, and mixed-signal microelectronics have recently moved swarm robotics and robot communities from simulation to reality. Swarms of 20 to 100 robots are in use already, implementations with several hundred robots are practicable, and communities exceeding a thousand robots are certainly conceivable. Such large robotic communities provide platforms for numerous exciting research directions including collaborative swarms and self-reconfiguring structures.

Maintaining hundreds of robots, however, poses significant practical challenges. The literature on strategies for maintaining software and hardware in large robot communities is sparse, even if applicable concepts from wireless sensor-networks are included. Crucial for the viability of any such strategy is its impact on cost per robot.

To provide a realistic setting we introduce a robot platform designed to be fabricated in full on standard printed circuit board (PCB) assembly lines. In this context we introduce a framework for on-line testing and calibration based on code pieces, termed plasmids, that migrate among the micro-controllers of the robots. The proposed approach allows the robots access to a larger library of code than what could be stored locally.

The robot consists of a single PCB that doubles as chassis and contains no custom mechanical components. Inexpensive motors (mass produced to vibrate mobile phones) are directly soldered to the circuit board and used in direct drive. Our prototypes use a 200 mAh rechargeable lithium polymer battery giving the robot over an hour of autonomy while moving at a speedy 1 m/s. An MSP430F2131 microcontroller controls the robot and communicates with neighbouring robots via infrared light. The simplicity of the design allows the entire robot to be assembled with low-cost PCB manufacturing techniques and is well suited for small-scale mass production of several hundred robots.

While this design significantly reduces the current cost barrier to obtaining a robot swarm, it also shifts the attention to the practical problem of maintaining hundreds of robots. Recharging batteries, sieving out robots with worn tyres or accidental damage is one aspect. A second aspect is testing and calibration. It can not be performed in the PCB assembly process and cost considerations prevent proprioceptive sensor. Collaboration among robots to verify performance and provide feedback (e.g., drift direction during a run and return) provide a scalable alternative. A third aspect is the maintenance of software in the robot community.

Our plasmid framework addresses all three aspects with a design that is lightweight enough to run on the microcontrollers. Pieces of code and associated attributes (version, target number for redistribution, lifetime, conditions for transmission) are propagated among robots that meet. For example, the code may perform a test on the robot and require to be forwarded to four other robots that have not encountered this test, before it is deleted. Such test plasmids traverse the robot community which, in its collective memory, contains and executes more code than would fit within the program memory of a single device.

Abstract organization and material parts in ALife and synthetic biology

Arantza Etxeberria

University of the Basque Country

Thursday 3:30pm–5:00pm, Stripe Studio 1, Philosophical Issues 3.

Two decades after the initial steps of ALife another field, Synthetic Biology (SynBio), defends the need of synthetic methodologies in the life sciences. Both have used similar arguments: the epistemological principle that to understand how something works, we must know how to build it, and the instrumental goal of producing useful things. But they come from different traditions and constitute different scientific communities; probably ALife's closest ancestor is cybernetics, whereas SynBio's is molecular biology.

In the case of ALife, the main goal is to study living organization. For Langton (1996 *Artificial Life*, MIT

Press.), synthesis makes possible to explore life-as-it-could-be in order to understand the necessary and the contingent of living organization in principle in a materiality different from the carbon-based. For others, the stress is in characterising life as an autonomous organization. ALife, like SynBio, is very diverse, and although it has predominantly pursued to create "life-like behaviours within computers and other artificial media", occasionally ALife models have been developed in vitro with biochemical components. Also, although the main goal was to construct new forms of life, in practice the field has produced more models (with scientific purposes) and tools (with instrumental purposes) than real instantiations.

In SynBio, however, the goal is not organisation, but design. The field intends to construct engineered organisms (biofacts) out of the components of existing life, by changing specific parts. Like ALife, SynBio is very diverse; O'Malley et al. (2008, *BioEssays*, 30, p. 57) have distinguished three different approaches: DNA-based device construction, genome-driven cell engineering, and protocell creation. Because it is difficult to see what they have in common, we may consider that the first is the most characteristic so far, or at least we could say that it is the one sharing less with ALife (whereas there is certainly some overlapping in the case of the third).

This paper intends to analyse some epistemological similarities and differences of both fields, especially in what concerns their views on living organization and the importance of materiality. In what concerns the first, in ALife living organization has been considered as an invariant to be found/constructed, whereas in SynBio, the goal is to engineer or to create life (sometimes close to the field of the origins of life). One idea has been that "nature is imperfect and should and can be revised and improved" (cited in Morange, 2008, Unpublished manuscript on Synthetic Biology). In what concerns the second, ALife aims to understand parts, composition and function as emerging properties, thus avoids fixed parts and aims at construction; in contrast, SynBio uses existing parts to change the design of life (one major effort is to build an open-access library of presynthesized biological parts and devices, the *Registry of Standard Biological Parts*), aiming at intervention.

Developing a methodology for social network sampling

Daniel W. Franks¹, Richard James², Jason Noble³ and Graeme Ruxton⁴

¹University of York

²University of Bath

³University of Southampton

⁴University of Glasgow

Wednesday 10:30am–12:00pm, Stripe Studio 1, Networks in Natural & Artificial Systems 1.

Researchers are increasingly turning to network theory to describe and understand the social nature of animal populations. To make use of the statistical tools of network theory, ecologists need to gather relational data, typically by sampling the social relations of a population of animals over a given time-period. Due to effort constraints and the practical difficulty involved in tracking animals, these sampled relational data are almost always a subset of the actual network. Measurements of the sample – such as average path length, clustering, and assortativity – are assumed to be informative as to the structure of the real-world social network. However, this assumption is problematic. Due to artefacts of the sampling process, the various network measures taken on the sample may be biased estimators of the true values. For example, just as we would get a biased estimate of mean human height by selecting for a sample those people who stood out in a crowd, we will get a biased estimate of a measure like mean connectivity if we sample individuals who are socially prominent.

This problem can only be solved by developing a qualitative theory of network sampling, answering questions such as what proportion of the whole network needs to be sampled before a given level of accuracy is achieved, and what sampling procedures are least biased? To develop such a theory, we need to be able to generate networks from which to sample. Ideally, we need to perform a systematic study of sampling protocols on different known network structures. But currently available data on animal social networks are unsuitable as these networks were themselves sampled.

The simulation methods of artificial life provide the way forward. We have developed a computational tool for generating artificial social networks that have user-defined distributions for network properties (such as the number of nodes, and the density) and for key the measures of interest to ecologists (such as the average degree, average path length, clustering, betweenness, and assortativity). This tool allows us to perform the required systematic analyses of the biases inherent in different sampling regimes (e.g., snowball sampling) applied to different network structures. We will present details of this system, and show we are using it to develop robust sampling methods for social network data. We see the system as the first in a series of works that will allow us to develop a

qualitative theory of social network sampling to aid ecologists, and eventually social scientists, in their social network data collection.

Utility and experimental testability of the Gaia hypothesis

Andrew Free, Amanda McNeil, Olivia Mozley and Rosalind Allen

University of Edinburgh

Wednesday 10:30am–12:00pm, Exam Hall Room 1, Models for Gaia Theory.

Gaia theory describes the life-environment system of the Earth as stable and self-regulating. It has remained at the fringes of mainstream biological science owing to historical definition problems and its apparent incompatibility with individual-level natural selection. However, various bodies of ecological and evolutionary research suggest ways in which the biosphere might tend towards stability and self-regulation. Here we review this research, relate the results to a plausible and informative formulation of ‘Gaia theory’, and ask how the theory extends the perspectives offered by these disciplines.

We then address the question of how Gaia theory might be tested experimentally. Such tests require the (reasonable) assumptions that life, where it evolves, will exploit essentially all thermodynamically-feasible forms of metabolism, and that Gaian regulation should be possible with a purely microbial biosphere. The biosphere is a closed system driven by solar radiation, and we describe here a laboratory microcosm which is an appropriate analogue of such a system. We then describe our preliminary experimental results from characterisation of this system, and discuss how we will use advanced molecular techniques employed by modern microbial ecology, in conjunction with computer simulations of inter-species interactions, to study the system and answer questions of relevance to Gaia theory. We also describe how this combined experimental-simulation approach can be applied to many questions of evolutionary and ecological interest which lie within the research areas bounded by Gaia.

Group formation and social evolution: a computational model

Nicholas Geard and Seth Bullock

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Wednesday 10:30am–12:00pm, Stripe Studio 1, Networks in Natural & Artificial Systems 1.

The tendency to organise into groups is a fundamental property of human nature. Despite this, many models of social network evolution consider the emergence of community structure as a side effect of other processes, rather than as a mechanism driving topological change. We present a model of social network evolution in which the

group formation process forms the basis of the rewiring mechanism. Exploring the behaviour of our model, we find that rewiring on the basis of group membership reorganises the network structure in a way that, while initially facilitating the growth of groups, ultimately inhibits it.

CAT: a market design competition

Enrico Gerding¹, Peter McBurney², Jinzhong Niu³ and Simon Parsons³

¹University of Southampton

²University of Liverpool

³City University of New York

Tuesday 3:30pm–5:00pm, Stripe Studio 1, Markets.

In the past decade the market paradigm has been adopted as the *prima facie* candidate for achieving coordination between self-interested, autonomous software agents. In this context, the area of mechanism design has been increasingly applied in order to achieve desired global outcomes based on local interactions, by developing rules of interaction that align the individual needs with social desiderata. Crucially, most of the current literature considers these mechanisms or markets in isolation and ignores any competition between them. In practice, however, we see similar markets competing, including Internet auction sites such as eBay and Amazon, and stock exchanges from around the world. Similarly, we argue that competition needs to be considered within multi-agent systems when designing mechanisms for them.

Against this background, we have developed an international market-design competition called CAT (short for *catactactics*, the science of exchanges) where each entrant to the competition sets the rules of interaction, as well as a fee structure. Each market consists of a double auction exchange where, similar to a stock exchange, the role of the market is to match buyers with sellers. Our competition framework includes the set of trading agents who participate in these markets, and who will choose the market which has proven the most profitable. Given this, the goal of the market designer is to attract profitable traders and make profits by charging appropriate fees. This can be done by changing the market rules and fees dynamically as a response to changing market condition. Overall, the objective of the competition is to see whether dynamic markets outperform static ones, and to study the types of markets that emerge from such complex interactions.

The CAT market design competition is an international competition and is part of the trading-agent competition (TAC). The first competition was successfully held in conjunction with AAI 2007, and this year's competition is going to be held in conjunction with AAI 2008.

The gene-function relationship in the metabolism of yeast and digital organisms

Philip Gerlee¹, Torbjorn Lundh², Bing Zhang³ and Alexander Anderson⁴

¹Niels Bohr Institute

²Chalmers and University of Gothenburg

³Vanderbilt University

⁴University of Dundee

Friday 10:30am–12:00pm, St Edburga Room 4, Biochemical Modelling.

Many natural and artificial systems form structures which can best be viewed as networks consisting of a set of nodes and links connecting the nodes. This perspective has been helpful in elucidating the organisation of a variety of systems ranging from power-grids and the internet to protein interactions, gene regulation and cell metabolism. Many of these networks exhibit a scale-free degree distribution and therefore deviate from the classical description of complex networks which predicts a Poisson degree distribution, which for degrees larger than the average degree scales as an exponential distribution.

We have studied the metabolic gene-function network in yeast and digital organisms from the artificial life platform Avida. The gene-function network is a bipartite network in which a link exists between a gene and a function (pathway) if that function depends on that gene, and can also be viewed as a decomposition of the more traditional functional gene networks, where two genes are linked if they share any function. We show that the gene-function network exhibits two distinct degree distributions: the gene degree distribution is scale-free while the pathway distribution is exponential. This is true for both yeast and digital organisms which suggests that this is a general property of evolving systems. One possible explanation for this structure is that in the network the genes acquire new links according to preferential attachment while the pathways receive new links independent of their degree.

This hypothesis was tested in Avida by tracking the evolution of the gene-function network in repeated simulations and measuring the rate of link attachment. Here the single commands takes the role of genes and the functions are the evolved boolean functions for which the organisms are rewarded. The results show indeed that the a gene is more likely to become involved in new functions (i.e. increase its degree) the more links it already has. The link attachment of the functions on the other hand occurred independent of the degree. In real cells it is known that gene duplication is the main mechanism by which new genes are created. If the two genes would retain similar functionality then we would expect pathways which involve many genes to increase their degree. This is contradicted by the exponential degree distribution and the observations in Avida and suggests that the rate of the gene divergence in yeast must be high. The duplication of pathways/functions

could on the other hand explain the scale-free distribution of the genes, and this mechanism has already been observed in *Avida*. Measuring the overlap between different pathways in terms of the genes which constitute them, showed that this also is a likely mechanism in yeast evolution. In conclusion we have presented a new way of analysing the gene-function dependence which sheds new light on the evolution of genes and functionality, and suggests that function duplication could be an important mechanism in evolution.

In silico evolution of chemotaxis

Richard Goldstein¹ and Orkun Soyer²

¹NIMR

²CoSbi

Tuesday 3:30pm–5:00pm, Exam Hall Room 2, Chemical & Molecular Modelling.

Bacteria are able to sense their environment and move towards sources of attractant and away from sources of repellent through the process of chemotaxis. The best understood model for chemotaxis comes from *Escherichia coli*, where the biochemical pathways have been extensively studied. In *E. coli*, the bacterium switches between swimming and tumbling based on the changes in the local concentration of attractant. It is unclear, however, how similar the behaviour and biochemical mechanisms are for other organisms. Work is proceeding on evaluating the chemotactic behaviour of a number of different bacteria, indicating substantial differences with *E. coli*. Even in *E. coli*, the fact that bacteria ‘guttled’ of most of the chemotactic machinery still displays effective chemotaxis indicates that there are still unanswered questions even in this organism. Finally, there are issues about how this particular strategy and implementation have evolved. Was this only one of a number of possible strategies? How did the strategy depend upon the environment, the properties of biochemical networks, or on the evolutionary process? Would similar strategies result under different conditions?

Ideally we would like to take non-chemotactic bacteria and evolve them to perform chemotaxis under a variety of different conditions, a daunting and lengthy experiment. In contrast, we can do this easily in a virtual world. In addition, we can keep complete records of the evolutionary process as it occurs. We create a population of digital organisms that move in a virtual world of attractant. The organisms contain a rudimentary set of biochemical elements, that is, sensors of the external attractant concentration, a reversible motor that can cause the bacteria to tumble or swim, and a set of proteins that can be activated, and while activated, have the potential to activate or deactivate other proteins. We then allow the biochemical network to evolve, selecting those digital bacteria that better co-localise with the attractant to reproduce for the next generation. This allows us to combine molecular-level evolutionary dynamics with phenotype-level selec-

tion of a relevant fitness parameter in an exactly specified environment. We find that these digital organisms quickly are able to display effective chemotaxis. Interestingly, the dominant mechanism is one that is very different from that observed in *E. coli*, but is similar to that observed in other bacteria as well as in gutted *E. coli*. The required network can be extremely simple, and can be fulfilled by coupling the bacteria’s metabolic network to the regulatory network, suggesting an explanation for the behaviour of gutted *E. coli* as well as suggesting a possible evolutionary route to how chemotaxis first arose.

A computational model of gene regulatory networks and its topological properties

Ângela Gonçalves¹ and Ernesto Costa²

¹European Space Agency

²University of Coimbra

Wednesday 1:30pm–3:00pm, Stripe Studio 1, Networks in Natural & Artificial Systems 2.

A new model for Gene Regulatory Networks (GRN) is proposed. The model is potentially more biologically sound than other approaches, and is based on the idea of an artificial genome from which several products like genes, mRNA, miRNA, non-coding RNA, and proteins are extracted. These products are connected giving rise to a heterogeneous directed graph. The topology of the obtained networks is studied using degree distributions. We make some considerations about the biological meaning of the outcomes of these simulations.

Immune-inspired networked service delivery

David Gowans, Paul Marrow and Richard Tateson

BT Group plc

Thursday 10:30am–12:00pm, St Edburga Room 4, Self-* Computing.

The expansion in distributed computing capabilities has led to a need to deliver services at different locations in networks, where demand is unpredictable. Giving all nodes in a distributed system the potential to deliver any service is likely to be a waste of resources, while extreme specialization on the part of individual nodes is likely to incur large overheads in terms of messaging and service transmission across the network. Other Computer Science researchers have drawn inspiration from one of real life’s most adaptable distributed systems, the vertebrate immune system, to attempt to solve problems in a variety of different application areas (de Castro et al., 2003, *Soft Comp.*, 7, 526). We describe an immune system-inspired method for service management in a distributed network services scenario. Each node in the network runs our system and the actions of all these local instances mesh together to provide overall service management.

The self/non-self theory of activation of the vertebrate immune system suggests how it can generate a response

to non-self entities, protecting the organism concerned. Our system uses an analogous activation sequence system to respond to requests arriving unpredictably at different nodes in a network, and enables the efficient delivery of responses to the requesting node without requiring complete specialization of capabilities in each node. In our system we have focused on the network of stimulatory interactions amongst antigens, antibodies, T-Cells and B-Cells. Requests for services are represented by antigens which interact with elements (nominally representing T- and B-Cells) in a two stage activation process to produce fully activated B-Cells. The fully activated B-Cells are monitored by Service Runners to indicate the level of demand in the system. The activation of B-Cells also releases antibodies which act as adverts for services. We believe that reproducing the logic of this interaction network, along with biologically plausible parameters for longevity and diffusion rates of the various cell types, give us an artificial system which can adapt to service demand patterns in the same way that the immune system adapts to patterns of antigenic challenge. A key issue is the level of complexity in our system. It aims to be useful, with some of the advantages of the analogous natural system, without attempting to model it slavishly. Our simulations show that an adequate level of complexity was chosen, as simpler incarnations lost some desirable properties, whereas more complex implementations would have led the system more towards modelling the immune system than focusing on service delivery.

In the simple network simulations described here some of the major issues for distributed systems are successfully counteracted. Unevenly distributed demand (which would otherwise result in excess demand on particular nodes) is balanced across suitable processing nodes. We see a memory effect whereby nodes can build up a more rapid response to requests based on their history of responses, allowing nodes to become specialists at dealing with particular request types. The diffusion of 'cells' deals with the issue of locating a suitable node to satisfy a request, even across several network hops. The large number of cells in the network originating from many different nodes, provides a level of fault tolerance, directing requests to alternative nodes in the event of node failure. We look forward to testing the system in larger, and more realistic, distributed network services scenarios.

From market to non-market: an autonomous agent approach to central planning

Dan Greenwood

University of Westminster

Tuesday 1:30pm–3:00pm, St Edburga Room 4, Methodologies.

In the longstanding debate in political economy about the feasibility of socialism, the Austrian school of economists have argued that markets are an indispensable means of evaluating goods, hence a prerequisite for productive efficiency. From an Austrian perspective, the prices generated by markets neatly encapsulate, in terms of a single numerical unit, highly complex information about the relative levels of demand and supply for different goods. Furthermore, it is emphasised that prices enable both consumers and producers to discover new economic knowledge about more efficient ways of attaining their particular ends. The Austrians contend that, such is the complexity of economic choices facing producers, no adequate level of economic efficiency could be achieved in the absence of markets, even on the assumption that a set of production objectives has been agreed. This problem of productive calculation is referred to as the 'economic calculation problem' for socialism.

Socialist advocates of a non-market economy have yet to provide a satisfactory response to this Austrian argument for the indispensability of markets. Some have sought to develop computational solutions to the economic calculation problem using techniques such as linear programming. Yet the computational models proposed are strongly influenced by the equilibrium model of neo-classical economics. From an Austrian perspective, these models overlook the essence of the calculation problem by assuming the availability of knowledge which can only be acquired through the market process itself.

The debate in political economy about the feasibility of a computational solution to the problem of non-market calculation has not yet considered the recent emergence of agent-based systems and their applications to resource allocation problems. Agent-based simulations of market exchange offer a promising approach to fulfilling the dynamic functions of knowledge encapsulation and discovery in a decentralised way, as the Austrians show to be performed by markets. Further research is needed in order to develop an agent-based approach to the economic calculation problem. Given that the macro-level objectives of agent-based systems can be easily engineered, it is suggested that such an approach holds the potential to become a desirable alternative to the real markets that the Austrians favour.

What makes the spatial prisoner's dilemma game sensitive to asynchronism?

Carlos Grilo and Luís Correia

LabMag, Dep. Informática, Faculdade de Ciências da Universidade de Lisboa

Friday 10:30am–12:00pm, Stripe Studio 1, Dynamical Systems Analysis 2.

We investigate aspects that control the Spatial Prisoner's Dilemma game sensitivity to the synchrony rate of the model. Based on simulations done with the generalized proportional and the replicator dynamics transition rules, we conclude that the sensitivity of the game to the synchrony rate depends almost exclusively on the transition rule used to model the strategy update by the agents. We then identify the features of these transition rules that are responsible for the sensitivity of the game. The results show that the Spatial Prisoner's Dilemma game becomes more and more sensitive for noise levels above a given noise threshold. Below this threshold, the game is robust to the noise level and its robustness even grows, compared to the imitate the best strategy, if a small amount of noise is present in the strategy update process.

How epigenetic evolution can guide genetic evolution

Graeme Groom, Rob Mills and Richard Watson

School of Electronics and Computer Science, University of Southampton
Tuesday 3:30pm–5:00pm, Stripe Lecture Room, Evolutionary Theory 3.

The expression level of a gene in future generations can be modified both by genetic mutations and by the attachment of methyl groups to the DNA. Since the DNA methylation pattern along a genome is inherited, methylation patterns constitute a significant epigenetic inheritance mechanism that is subject to evolution by natural selection. The variation rate of methylation patterns is generally higher than that of DNA which suggests that evolution of methylation patterns might be more rapid than that of genetic evolution. But, common consequences of methylation, such as reduced expression of methylated genes, could also be produced by genetic changes and these would have higher heritability. The question we address in this work is how the evolution of epigenetic methylation-dependent phenotypes might interact with the evolution of genetic DNA-determined phenotypes.

There is no biological mechanism known to directly transfer methyl groups into equivalent DNA changes. However, in principle an indirect mechanism could cause evolved methylation patterns to enable the subsequent evolution of equivalent genetic patterns in a manner analogous to the Baldwin effect (Baldwin, *Am. Nat.*, 30:441-451, 1896; Jablonka et al, *TREE*, 13:206-210, 1998). The Baldwin effect describes how non-heritable acquired char-

acteristics can influence the evolution of equivalent genetic characteristics without any direct Lamarckian inheritance of acquired characters. This occurs because the ability to acquire or learn a new behaviour changes the selective pressures acting on genetic changes. Specifically, genetic changes that support this behaviour, e.g. by reducing learning time by making a small part of the behaviour genetically innate, may be selected for when the learning mechanism is present even though these same genetic changes may not be selected for when the learning mechanism is absent. Over generations, the modified selection pressures so produced can cause genetic assimilation of a phenotype that was previously acquired, even to the extent of making the acquisition mechanism subsequently redundant. Thus a learned behaviour can guide the evolution of an equivalent innate behaviour (Hinton & Nowlan, *Complex Systems*, 1: 495-502, 1987).

In the Baldwin effect a rapid mechanism of lifetime adaptation guides the relatively slow genetic evolution of the same behaviour. By analogy, Jablonka et al have suggested that “genetic adaptations may be guided by heritable induced or learnt phenotypic adaptations”. Here we hypothesise that “inherited epigenetic variations may be able to ‘hold’ an adapted state for long enough to allow similar genetic variations to catch up”, as they put it, even if the epigenetic variations are not induced or learnt but simply evolved by natural selection on methylation patterns. We assume that an individual may only express one phenotype in its lifetime, but that a given genome will persist relatively unchanged on a timescale that allows its methylome to adapt by natural selection. Thus, in contrast to the Baldwin effect, in this case two mechanisms of evolution by natural selection are coupled — one acting at a different variation rate from the other. We present a simple model to illustrate how a rapidly evolving methylome can guide a slowly evolving but highly-heritable genome. This is used to show that methylome evolution can enable genetic evolution to cross fitness valleys that would otherwise require multiple genetic changes that were each selected against. This finding suggests that the relatively rapid evolution of methylation patterns can produce novel phenotypes that are subsequently genetically assimilated in DNA evolution without direct transfer or appeal to induced phenotypes. This can enable the genetic evolution of new phenotypes that would not be found by genetic evolution alone, even if methylation is not significant in the ultimate phenotype.

Chemical basis for minimal cognition

Martin M. Hanczyc¹ and Takashi Ikegami²

¹ University of Southern Denmark

²University of Tokyo

Thursday 1:30pm–3:00pm, Exam Hall Room 2, Wet ALife.

We have developed a simple chemical system capable of self-movement in order to study the chemical-molecular origins of movement, perception and cognition. The system consists simply of an oil droplet in an aqueous environment. The aqueous phase contains a surfactant that modulates the interfacial tension between the drop of oil and its environment. We embed a chemical reaction in the oil phase that reacts with water when an oily precursor comes in contact with the water phase at the liquid-liquid interface. This reaction not only powers the droplet to move in the aqueous phase but also allows for sustained movement. The direction of the movement is governed by a self-generated pH gradient that surrounds the droplet. In addition this self-generated gradient can be overridden by an externally imposed pH gradient, and therefore the direction of droplet motion may be controlled. Also we noticed that convection flow is generated inside the oil droplet to cause the movement, which was also confirmed by simulating the fluid dynamics integrated with chemical reactions (Matsuno et al., 2007, ACAL 07, Springer, p.179, Springer). We can observe that the droplet senses the gradient in the environment (either internally generated or externally imposed) and moves predictably within the gradient as a form of primitive chemotaxis (Hanczyc, M., et al., 2007, J. Am. Chem. Soc., 129, p. 9386).

By creating a pH gradient and concomitant convection flow, the droplet behaves as if it can perceive the environment. We believe that the geometry of the interface shape can control sensitivity to the environment (Ikegami et al., 2008, BioSys., 91, p.388). This geometry-induced fluctuation is the source of fluctuation of motion, which we think is tightly linked with the idea of biological autonomy. There is empirical evidence to support the above ideas. Some form of internal bias is necessary for breaking symmetry to cause self-movement and the bias may be the result of perception of the environment.

Such simple oil droplet systems show autonomy in the sense that the droplets move in response to the self-generated pH and the environmental gradient. In our modeling, we demonstrated that a computational autopoietic cell could move by continuously self-repairing the membrane, but in this case failed to show any gradient-climbing behavior (Suzuki et al., 2008, Artificial Life, in press). This may be due to the fact that the autopoietic cell can only survive in the narrow range of environments that support a certain substrate density. Compared with that autopoietic cell model, our oil droplets are more stable and they strive to maintain their boundary structures. We hypothesize that the pH gradient around the droplet results

in an unbalanced interfacial tension at the interface. The droplet then responds by motion in order to maintain a balanced interfacial tension. Once the tension forces around the droplet are balanced the droplet would stop moving. In this way, we contend that a kind of homeostasis is a basis for self-movement. Different from the mere physical-chemical process, any life system preserves its own identity and consistency with respect to the environment. This homeostasis, rooted on the sensory motor couplings, will organize minimal cognition (see also, Ikegami, T. et al., 2008, BioSys., 91, p.388]

Being arranged in advance: quantum entanglement and biological feedback

Taichi Haruna

Kobe University

Friday 10:30am–12:00pm, Exam Hall Room 2, Information in Complex Systems & ALife 4.

The categorical semantics of quantum protocols proposed by Abramsky and Coecke reveals that a prearranged quantum entanglement brings a strange quantum information flow in the quantum teleportation protocol. Their formal argument leads us to the distinction between an information flow sequence and a causal sequence on the same event. If this distinction is applied to information processing biological networks, we can claim that a prearranged biological feedback can play the same role as the quantum entanglement on the emergence of a specific local structure of networks. The aim of this paper is to provide a first step toward formal arguments on changes in biology without the external time parameter.

Homeostasis via chaos: implementing the unisector as a dynamical system

Inman Harvey

University of Sussex

Tuesday 3:30pm–5:00pm, Stripe Main Theatre, Sustainability & Homeostasis.

Ashby's Homeostat (Ashby, 1952, Design for a Brain, Chapman and Hall) was a demonstration of how an extended form of homeostasis, defined by him as ultrastability, could be achieved with a relatively simple mechanism. Homeostasis refers to the process whereby an organism, or a machine, actively maintains certain 'essential variables' (EVs) within the critical bounds of viability. The simplest form is negative feedback, but a higher order of homeostasis can sometimes be observed when an EV, on approaching a critical value, triggers one or more periods of *positive* feedback that reorder the dynamics until a new stable equilibrium (based once again on negative feedback) is found. This ultrastability can be viewed as an interaction between two coupled dynamical systems (DS): the primary DS comprises the EVs, and their direct, parameterised interactions; the secondary DS only kicks in tem-

porarily when the EVs of the first are threatened, and then it alters the parameters of the first DS until some equilibrium is found that no longer threatens the EVs. Hence this is a form of selection between multiple possible steady states.

In Ashby's Homeostat, the secondary DS was implemented by the 'Uniselector'. Under normal circumstances it maintained a fixed set of parameters for the first DS. When it was triggered, it picked a different set of parameters (in practice drawn from a lookup table of random numbers), and continued doing so until the triggering factor ceased. In Evolutionary Robotics one common method for designing an artificial 'nervous system', coming from the DS perspective on cognition, is to evolve the parameters (weights, biases and time constants) for a Continuous Time Recurrent Neural Network (CTRNN; Beer, 2006, *Neur. Comp.* 18(12), p. 3009). One way of implementing an Ashbyan ultrastability mechanism would be to incorporate the Uniselector as an add-on to the CTRNN. An alternative approach proposed here is to incorporate the Uniselector-effects within the CTRNN, rather than as a separate add-on.

We require a very large number of different attractors (corresponding to different sets of random numbers in the Uniselector); and a trigger mechanism that initiates random or chaotic jumps to a new attractor. This can be done with a core of just 3 interconnected variables, equivalent to 3 nodes of a CTRNN if we extend the class of transfer functions at each node to include sine waves as well as sigmoids. Drawing on a result of Thomas (Kaufman et al., 2003, *C. R. Biologies* 326, p. 205), we show how this can be implemented; we can switch between chaotic 'search' and settling into one amongst many possible attractors. These attractors are cyclic or strange, but can be used to set parameters for the remaining part of the CTRNN that comprises the 'primary DS'. There remain practical issues, somewhat glossed over by Ashby, in orchestrating how long is spent 'evaluating' each attractor visited before abandoning it for another one.

This approach demonstrates the possibility of composing a Homeostat entirely of such an (extended) CTRNN, with the Uniselector-substitute as a distinct hand-designed sub-circuit or module. Further evolution can maintain the desired ultrastable characteristics, whilst relaxing these architectural constraints of modularity.

Misrepresentations

Inman Harvey

University of Sussex

Thursday 1:30pm–3:00pm, Stripe Studio 1, Philosophical Issues 2.

The concept of "representations", and particularly "internal representations", can be controversial in Cognitive Science and AI. It is suggested here that much time-wasting confusion could be avoided if participants in such controversies came to recognize the variety of different senses, often incompatible, in which such terms are used. A hypothesis is presented as to why there is so much reluctance to recognize this. Once such fruitless controversies are swept aside through linguistic hygiene, there remain interesting real problems, which are eminently appropriate for being tackled by an Artificial Life methodology.

NK α : non-uniform epistatic interactions in an extended NK model

Tom Hebborn¹, Seth Bullock¹ and Dave Cliff²

¹Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

²University of Bristol

Thursday 10:30am–12:00pm, Exam Hall Room 2, Fitness Landscapes.

Kauffman's seminal NK model was introduced to relate the properties of fitness landscapes to the extent and nature of epistasis between genes. The original model considered genomes in which the fitness contribution of each of N genes was influenced by the value of K other genes located either at random or from the immediately neighbouring loci on the genome. Both schemes ensure that (on average) every gene is as influential as any other. More recently, the epistatic connectivity between genes in natural genomes has begun to be mapped. The topologies of these genetic networks are neither random nor regular, but exhibit interesting structural properties. The model presented here extends the NK model to consider epistatic network topologies derived from a preferential attachment scheme which tends to ensure that some genes are more influential than others. We explore the consequences of this topology for the properties of the associated fitness landscapes.

Why should you care? An arousal-based model of exploratory behavior for autonomous robots

Antoine Hiolle and Lola Cañamero

University of Hertfordshire

Thursday 10:30am–12:00pm, Exam Hall Room 1, Sensorimotor Connections 1.

The question of how autonomous robots could be part of our everyday life is of a growing interest. We present here an experiment in which an autonomous robot explores its environment and tries to familiarize itself with the features available using a neural-network-based architecture. The lack of stability of its learning structures increases the arousal level of the robot, pushing the robot to look for comfort from its caretaker to reduce the arousal. In this paper, we studied how the behavior of the caretaker influences the course of the robot exploration and learning experience by providing certain amount of comfort during this exploration. We then draw some conclusions on how to use this architecture together with related work, to enhance the adaptivity of autonomous robots development.

Animatronic model of a human tongue

Robin Hofe and Roger K. Moore

Department of Computer Science, University of Sheffield

Wednesday 1:30pm–3:00pm, St Edburga Room 4, From Communication to Language 2.

Contemporary speech technology struggles with the variation that occurs in natural speech. In order to explain such variation, researchers frequently refer to processes of energy optimisation that govern speech production. Although there is plenty of qualitative evidence, quantitative data about speech energetics are still sparse because it is difficult to acquire them from human subjects. To overcome this problem, an animatronic model of a human tongue and vocal tract (AnTon) was designed. Human anatomy provided the guideline for its construction; functional considerations were made only when an approximation using available technology proved impossible or infeasible. Thus, the behaviour of the model derives from, and is grounded in, its anatomy. The tongue model presented here was developed using these 'biomimetic' principles. The human tongue consists almost completely of interwoven muscle fibres whose topology allows for complex movements. The incompressibility of muscle tissue is an important prerequisite, a property it shares with water; such structures are therefore called 'muscular hydrostats'. The soft silicon that forms the artificial tongue body approximates incompressibility. Muscles are represented by filaments that run along paths resembling real muscle fibre orientation. Wherever muscle fibres follow a curved path, regularly spaced glass beads prevent filaments from cutting into the silicon. The filaments connect to meshes that are embedded into the tongue body and serve both as an

attachment point and to distribute force evenly. The current tongue model comprises four of the main tongue muscles, represented by eleven filaments that are attached to servo motors. It is connected to a movable jaw and a hyoid bone; the latter is a horseshoe-shaped bone that supports the tongue root and is situated directly above the Adam's apple. AnTon is able to imitate a range of oral gestures and will be used for sound production as soon as the anterior part of the vocal tract is completed. Apart from studying speech energetics, AnTon has the potential to become a general tool for speech research, education, and speech therapy. A video is available on the AnTon project's website: <http://www.dcs.shef.ac.uk/robin/anton/anton.html>.

The effects of payoff preferences on agent tolerance

Enda Howley and Colm O'Riordan

National University of Ireland, Galway

Tuesday 1:30pm–3:00pm, Exam Hall Room 2, Culture & Social Evolution 2.

An objective of multi-agent systems is to build robust intelligent systems capable of existing in complex environments. These environments are often open, noisy and subject to rapid, unpredictable changes. This paper will explore how agents can bias their interactions and choices in these complex environments. Existing research has investigated how agents can bias their interactions based on factors such as similarity, trust or reputation. Unfortunately, much of this research has ignored how agents are influenced by their preferences for certain game payoffs. This paper will show that individual payoff preferences have a significant effect on the behaviors that emerge within an agent environment. We argue that agents must not only determine with whom to interact, but also the levels of benefit or risk these interactions should represent. This paper presents a series of game theoretic simulations examining the effects of agent payoff preferences within an evolutionary setting. Our experiments show that these factors promote tolerance and reciprocity throughout the population. We provide an experimental benchmark using an almost identical game environment where payoffs are not considered by agents. Furthermore, we also present simulations involving noise, thereby demonstrating the ability of these more tolerant agents to cope with uncertainty in their environment.

Simulating active touch with a simple embodied agent

Hiroyuki Iizuka¹ and Takashi Ikegami²

¹Osaka University

²Tokyo University

Thursday 3:30pm–5:00pm, Exam Hall Room 1, Embodiment.

In the present paper, we simulate the differentiation of sensory- and motor-like interfaces, and their dynamics in an agent of novel architecture. This paper reconsiders the boundary between an agent and its environment. In the standard framework, we tend to assume that the physical boundary between the agent and its environment exists due to the independent physical devices of its sensory and motor interfaces. The boundary between agent and environment thus appears as a static and rigid boundary. We contest this view by creating a simple simulation model in which this boundary can vary dynamically. Thus we argue that an agent's perception is an outgrowth of the complex interference between the efferent and afferent copies of its action patterns.

The agent consists of a body with two straight arms that move freely from -90 to 90 degrees. Those arms are controlled by a continuous-time neural network. The neurons connected to the arms sometimes receive body states from the arms (the afferent copy) and then will move the arm in order to explore the environment (efferent copy). The role of sensory system and motor system are thus switching temporally, and there is no explicit sensory-motor flow.

There are some neuro-psychological experiments supporting the above dynamic view. For example, Yamamoto and Kitazawa (2001) demonstrated that with their arm-crossing experiment in which the perceived temporal ordering of haptic stimuli on a subject's hands is reversed when successive stimuli are temporally very close together. This experiment implies that we may be able to explain our perception by something based on body image rather than just by reactions to sensory input.

The agent in this simulation is trained via a standard genetic algorithm. The agent was intended to learn to distinguish between fans with differing numbers of wings which the agent can touch and manipulate using its arms. From this experiment we draw two primary conclusions:

- 1) The two arms become differentiated into sensor-like and motor-like arms, and the agent can successfully distinguish different fans by touching them

- 2) The agent's active and passive touching of the fans is driven by the stability of dynamic attractors associated with those touching behaviours. By placing time delays into the neural connection from the arm neurons to the body neurons, we notice that motor-like arms appear more fragile than sensor-like arms, and the attractor associated with active touch is more fragile than that associated with passive touch.

Our simulation demonstrated that the differentiation of sensory and motor functions emerged by mixing free and constrained arm motion conditions. When there are no obstacles, the arm moves freely and the neural network activations maintain coherence. This coherence could be taken as a sensation of body image. When there are obstacles present, however, this coherence is lost, and this interruption is transmitted as sensory information. By investigating the coherence/decoherence events of these internal neural dynamics, we will argue that the interference between efferent and afferent sensory copies drives sensation.

Adaptive fault tolerance in wireless sensor networks

David Irons¹, Fabrice Saffre² and Chris Cannings¹

¹University of Sheffield

²British Telecom

Thursday 1:30pm–3:00pm, St Edburga Room 4, Fault Tolerance & Robustness.

Systems including wireless sensor networks and amorphous computing involve scattering a large number of nodes (sensors / processors) over a geographical area and then creating local connections to carry out some global objective. For example, monitoring, positioning and reporting back on changes in environmental features; such as heat, light, sound and / or motion. Geographical, cost or size constraints may also imply that the nodes need to be dispersed semi-randomly, implying a great deal of uncertainty about their relative positions and the exact size of the area being covered.

Various research challenges exist in these systems, such as creating strategies for self organisation, energy management, fault tolerance, maintaining security and adapting to cope with geographical features (e.g. walls, hills, buildings). Moreover, these challenges become increasingly important as technological advances lead to smaller and smaller nodes, making them (potentially) more error prone and difficult to position accurately. In this study, we introduce a fault tolerance / energy management strategy that allows the system to robustly self re-organise, when nodes break or run out of power. We then test this strategy on different geographical landscapes.

In order to model these sensors networks, we use random geometric graphs. Here, n nodes are randomly placed on a surface and edges can form between nodes if they are within distance R of one another. Additionally, we assume that each node has a randomly assigned lifespan corresponding to time taken for the node to run out of power / break down.

We then test the following strategy. Suppose each node can take one of two states, 'active' or 'hibernated', with the hibernated state consuming less power. Additionally, suppose any active node is capable of sending all other

nodes within distance $r \leq R$ into hibernation. Now, beginning with every node in hibernation, the system evolves by individual hibernated nodes (asynchronously) trying to become active (i.e. turning active if there are no existing active nodes within radius r). Whenever the lifespan of an active node is exceeded, it disappears from the system, allowing the opportunity for hibernated nodes to take its place.

In order to score the integrity of the system, we compare the evolution of such networks to an equivalent system where all nodes are active. In particular, we compare various statistical properties of the network at regular time intervals; including (1) the proportion of active nodes in the largest component, (2) diameter and (3) area covered by the largest component. Such performance measures were chosen since these systems may need to adequately monitor a geographical area (over a period of time), with network connectivity being essential for communication of data / results. We find that, for a large range of parameters, the new strategy presented here is more robust, in that the integrity of the system is maintained over a longer period of time. In particular, robustness increases as the initial number of nodes (n) increases. We believe that as nodes become cheaper to mass produce, having a large number is a feasible strategy for these systems.

As well as testing our strategy on flat two dimensional surfaces, we also extend our approach to more complicated three dimensional geographical landscapes with the additional constraint that nodes can only communicate if they are in line of sight of one another.

Analysis of a dynamical recurrent neural network evolved for two qualitatively different tasks: walking and chemotaxis

Eduardo J. Izquierdo¹ and Thomas Buhrmann²

¹Centre for Systems Biology, University of Birmingham

²Natural Motion Ltd

Thursday 3:30pm–5:00pm, Exam Hall Room 2, Dynamical Systems Analysis 1.

Living organisms perform a broad range of different behaviours during their lifetime. It is important that these be coordinated such as to perform the appropriate one at the right time. This paper extends previous work on evolving dynamical recurrent neural networks by synthesizing a single circuit that performs two qualitatively different behaviours: orientation to sensory stimuli and legged locomotion. We demonstrate that small fully interconnected networks can solve these two tasks without providing *a priori* structural modules, explicit neural learning mechanisms, or an external signal for when to switch between them. Dynamical systems analysis of the best-adapted circuit explains the agent's ability to switch between the two behaviours from the interactions of the circuit's neural dynamics, its body and environment.

The evolution of evolvability in gene transcription networks

Eduardo J. Izquierdo¹ and Chrisantha T. Fernando²

¹Centre for Systems Biology, University of Birmingham

²Mathematical Biology, National Institute for Medical Research

Tuesday 1:30pm–3:00pm, Stripe Main Theatre, Evolving Cell Signalling Networks *in silico* 2.

We present a case of a genotype-phenotype map, that when evolved in variable environments optimizes its genetic representation to structure phenotypic variability properties, allowing rapid adaptation to novel environments. How genetic representations evolved is a relatively neglected topic in evolutionary theory. Furthermore, the “black art” of genetic algorithms depends on the practitioner to choose a representation that captures problem structure. Nature has achieved remarkably efficient heuristic search mechanisms without top-down design. We propose that an important example of this, ubiquitous in biology is the structuring of the phenotypic variability properties of gene networks. By studying a simple model of gene networks in which topology is a function of interactions between transcription factor proteins and transcription factor binding sites (TFBS), we show that transcription factor binding matrices (TFBM) evolve to positively constrain phenotypic variability in response to transcription factor binding sequence mutations.

Energy, entropy and work in computational ecosystems: a thermodynamic account

Mariusz Jacyno and Seth Bullock

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Friday 10:30am–12:00pm, Exam Hall Room 1, Entropy.

Recently, computer scientists have begun to build computational ecosystems in which multiple autonomous agents interact locally to achieve globally efficient organised behaviour. Here we present a thermodynamic interpretation of these systems. We highlight the difference between the regular use of terms such as energy and work, and their use within a thermodynamic framework. We explore the way in which this perspective might influence the design and management of such systems.

Solving the division of labour problem using stigmergy and evolved heterogeneity

Emyr James, Jason Noble and Richard Watson

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Wednesday 3:30pm–5:00pm, St Edburga Room 4, Multi-Agent Applications.

Evolving cooperative teams is a research area with applications in the fields of robotics and software agents. Progress on this problem could also help us to understand

the evolution of cooperation in natural systems such as the social insects. The overarching question is how cooperative teams should be represented in order to promote efficient evolutionary search. More specifically, what should serve as our basic unit of selection — the individual or the team? — and how can the division-of-labour problem be solved? In order to answer these questions we have taken a benchmark problem from the genetic programming (GP) literature, the artificial ant problem, and extended it so that teams of ants must cooperate to complete the task.

In this model, the ants are centrally placed in a bounded grid with each square containing food. The goal of the team is to harvest all the food in the environment in as few moves as possible. In the initial version of the problem, the members of the team are all clones, each having exactly the same GP controller program. Many solutions will have poor performance as the team members will all behave in the same way, and will therefore fail to cover the grid efficiently. To perform better, the ants must evolve to take advantage of stigmergic interactions to break the symmetry of the problem and clear the world of food efficiently. This division of labour through stigmergy is indeed what is seen to evolve during the simulations.

A further extension is made by assigning each member of the team an identity tag, and adding the ability to execute different subtrees of the cloned controller based on this tag. When these operations are allowed, higher fitnesses are achieved than with the purely stigmergic situation above.

During evolution, selection acts at the team level. We can therefore view the members of the team as being equivalent to cells in a multicellular organism. The identity branching operation is analogous to cell differentiation within this abstract organism. Using this scheme, the degree of differentiation is not specified a-priori and is controllable through evolution. This allows the full continuum from purely homogeneous teams to entirely heterogeneous teams to be expressed. There is also the potential to use this method as a way of measuring the degree to which a task demands heterogeneous solutions.

The relative importance of stigmergy and innate heterogeneity in achieving the necessary division of labour were compared with a third experimental manipulation. The ability to influence each other stigmergically was removed by placing each ant in its own world and tallying the pieces of food consumed by the team as a whole. In this scenario, the most efficient way to tackle the problem is for the team to evolve complete heterogeneity.

We conclude that the division-of-labour problem in the evolution of cooperative teams can be solved by both stigmergic communication and innate heterogeneity. Furthermore, the technique of allowing the level of heterogeneity of the team to be open to selection shows promise for future work.

Evolving asynchronous cellular automata for density classification

Francis Jeanson

University of Sussex

Wednesday 3:30pm–5:00pm, Exam Hall Room 1, Cellular Automata.

This paper presents the comparative results of applying the same genetic algorithm (GA) for the evolution of both synchronous and asynchronous cellular automata (CA) for the computationally emergent task of density classification. The present results indicate not only that asynchronous CA evolve more quickly and consistently than their synchronous counterparts, but also that the best performing asynchronous CA find equally good solutions on average to the density classification task in fewer computational steps than synchronous CA.

Effects of signalling on the evolution of gene regulatory networks

Dafyd Jenkins and Dov Stekel

Centre for Systems Biology, School of Biosciences, University of Birmingham

Tuesday 1:30pm–3:00pm, Stripe Main Theatre, Evolving Cell Signalling Networks *in silico* 2.

We investigate whether observed transcription network structures and network motifs are a byproduct of the mechanisms by which DNA strands evolve, or if they are fundamental to the function of the network. We explore this with an evolutionary model with stochastic Boolean network simulation. Structurally distinct regulation strategies are observed in populations evolved with and without internal energy signalling. However, food signalling is not used in either population in the case when the food supply itself is constant. Parallels between the evolved networks and CRP-cAMP regulation in *Escherichia coli* and the endosymbiont *Buchnera aphidicola* are presented and discussed. Comparing the evolved networks with neutrally evolved populations indicates that networks evolve to lose most regulatory activity, due to loss of binding sites and transcription factor activity, including losing global regulation mechanisms.

Evo-devo in silico: a model of a gene network regulating multicellular development in 3D space with artificial physics

Michal Joachimczak and Borys Wróbel

Institute of Oceanology, Polish Academy of Sciences

Wednesday 1:30pm–3:00pm, Stripe Lecture Room, Development & Morphogenesis 2.

We present a model of multicellular development controlled by a gene network in which the connectivity is determined by the proximity of sequences in N-dimensional

space. Thus the sequences of individual genes can be visualised as points in space which approach or move away from one another as the genomes evolve. The genotype-phenotype (morphology) mapping in our model is indirect, relies on artificial physics, and allows cell adhesion and free movement in 3D space. Cell differentiation is allowed by positional information provided by factors that diffuse in this space, and the differential gene expression in each cell determines the cell fate (such as division, death, growth and movement). We apply a genetic algorithm to find genotypes that can direct morphogenesis of non-trivial asymmetrical shapes. We then investigate the mechanism of such developmental process and the features of gene regulatory network that direct the embryogenesis.

Evolving functional symmetry in a three dimensional model of an elongated organism

Ben Jones¹, Yaochu Jin², Bernhard Sendhoff² and Xin Yao¹

¹School of Computer Science, University of Birmingham, UK

²Honda Research Institute Europe

Tuesday 3:30pm–5:00pm, St Edburga Room 4, Spatial Organisation in Neural Systems.

In evolutionary–developmental biology, it is well established that neural organization is coupled to a given organism’s body-plan. Many theories attempt to underpin this coupling and the transitions involved during the organism’s evolution, for example the transition from radial to bilateral symmetry. Before theoretically tackling these transitions however, we felt it essential to first address, in this paper, precisely why bilateral symmetry might be advantageous for a simple eel-like agent. We find that neural architectures affording the best motor-coordinated behavior (architectures that allow directional swimming of the agent), will readily emerge in a way that is functionally–bilaterally symmetric, suggesting therefore, that symmetrical emergence for long elongated creatures can be essential if it needs to travel over some distance.

Nonconsumable resources facilitate complex evolution

George Kampis, Laszlo Gulyas and Walter de Back

Collegium Budapest

Wednesday 3:30pm–5:00pm, Stripe Main Theatre, Open-Ended Evolution 3.

Open ended evolution (OEE) is a problem closely related to the “arrow of complexity”, that is, the temporal emergence of complex and complicated structures and organisms in a progressive, never-ending fashion. A straightforward, and biologically proven means to approach this is to integrate organisms in complex environments, and to make them parts of an evolving food web where various

aspects of complexity may increase. Our recent work, reported in a series of papers, addresses OEE via the emergence of food webs and niche differentiation in a simple agent based population that starts from a seed such as a single species of producers (consuming nonreplicating resources and forming the basic trophic level of an ecosystem).

A notorious problem of such systems (as exemplified by recent work like DOVE, Webworld, etc) is the questionable ecological stability of the evolving trophic structure. To stabilize even a 3-species system by parameter tuning is an egregious task, let alone in systems of higher complexity and with the frequent introduction of newer species. A typical problem is that either a newly evolved consumer cannot grow (thus cannot become part of the system) or depletes the resources until both die out. Consequently, few systems currently can handle this problem. To breed restraint e.g. by selecting for moderate predators is a slow process which already presupposes the coexistence of species in a food web. But how do we get there?

Density dependent feedback (such as functional response) could be a viable solution at this point, but that typically requires complicated agents or an alien hand. An additional mechanism is to consider genotype-phenotype maps with tradeoffs in consumption, replication, and other items of life history. This is a road we have reported in another paper submitted to this conference.

Here we introduce and study a different idea, that of nonconsumable resources (NCRs), an idea which is usually underestimated in the evolutionary modeling of food webs and OEE. NCRs are factors required for the life history but not destroyed by the agents; NCRs can range from space to nesting ground to other abiotic factors to sexual partners to environmentally inherited properties to other species entering mutualism or (mild) parasitism with the given species. In natural systems NCRs are found everywhere and they pose important, yet little understood feedbacks on otherwise destabilizing dynamics. Using our agent ecosystem we show that phenotype to phenotype interactions in a consistently agent based (“fully embedded”) system naturally introduce NCRs dynamically. This, in turn, helps stabilizing the emerging ecosystems and permits complexification steps otherwise impossible due to kinetic instabilities.

This work is part of our development of the FATINT system and is part of a research project reported in “Feedback Self-Organization”, Springer, to be published in 2009. This work was supported by the EC grant QosCos-Grid IST FP6 #033883. The authors thank Collegium Budapest for their hospitality. L.G. acknowledges the partial support of the GVOP-3.2.2-2004.07-005/3.0 (ELTE Informatics Cooperative Research and Education Center) grant of the Hungarian Government.

Intentional and causal urges: when to ascribe life or mind to artificial systems?

Fred Keijzer

University of Groningen

Thursday 1:30pm–3:00pm, Stripe Studio 1, Philosophical Issues 2.

How do we decide whether or when Artificial Intelligence systems are cognitive or Artificial Life models are alive? Despite the many different ways in which these decisions are made and specific criteria are used, in both cases the final verdict has a strong intuitive component: The Turing Test is a way to harness the decision in a particular way; but why would making a specific impression on a human observer be a good criterion for deciding whether an AI system is cognitive? Similarly, why would a CA model of growth be considered a possible form of life? In both cases, many of us have tendencies to interpret mind-like or life-like phenomena as being actual example of mind or life. Still, many of us also have tendencies to question the step from looking like mind or life to actually being a mind or alive. The result has been a wavering attitude in which both tendencies coexist, while a more definite view on this issue remains elusive. Research on the automatic tendency to detect causality and animacy provides a way to make sense of this wavering attitude. The upshot of this research is that we have an inbuilt intentional urge: a prereflective tendency to categorize certain objects or displays as alive and even intentional. We also have a causal urge: a prereflective tendency to interpret events with specific visual features in causal terms. Both urges seem to be mutually exclusive. These automatic reactions are subsequently elaborated and scrutinized in a more reflective mode, which can lead to either the acceptance or rejection of the first automatic reaction. This research can have important and wide-reaching implications for discussions on the status of artificial cases of life and mind. Foremost, it would explain the wavering attitude when it comes to formulating criteria for making decisions either way. Strong intuitive judgements and reflection do not necessarily coincide. It seems clear that we should be more critical of our own intuitive judgements in this area. At the same time, it is less evident how exactly we can come to a more critical and ultimately sounder judgement on these matters. At present, it seems best just to start out with trying to sort out the options and issues. For example, a way out would be by differentiating between design research in which AI and ALife would fall and descriptive and explanatory research where empirical cognitive science and biology belong. Another important issue would be whether the intuitive mutual exclusiveness of causal and intentional interpretations is really to be trusted and should be taken to imply a serious dichotomy in the natural world. In this talk, I will start out with the problem of ascribing life and mind, introduce the literature on the

intentional and causal urges, and focus on the possible implications for artificial life and mind.

Enrichment of interaction rules in a string-based artificial chemistry

Ciaran Kelly, Barry McMullin and Darragh O'Brien

Dublin City University

Wednesday 3:30pm–5:00pm, Stripe Main Theatre, Open-Ended Evolution 3.

In this paper, we discuss our reasoning and progress in adding a mapping between information and enzymatic function to our Molecular Classifier System (MCS). MCS takes a bottom-up approach to building artificial biochemical networks. Unlike Holland's LCS system, which it is loosely based on, MCS has no overt demarcation between *rules* and *messages*. In our previous work, we explored a version of this Artificial Chemistry which had an impoverished interaction scheme. While this system did present some interesting results, it had very limited potential for evolving greater complexity. We present here a mechanism for enriching the reaction rules used in our Artificial Chemistry. This mechanism is similar to the folding of RNA to an enzymatically active form. To date, we have examined in detail the evolutionary trajectories of single reactors populated with this modified Artificial Chemistry and the results of this work are presented here.

Evolution and morphogenesis of differentiated multicellular organisms: autonomously generated diffusion gradients for positional information

Johannes F. Knabe, Maria J. Schilstra and Christopher L. Nehaniv

University of Hertfordshire

Wednesday 10:30am–12:00pm, Stripe Lecture Room, Development & Morphogenesis 1.

Development is the powerful process involving a genome in the transformation from one egg cell to a multicellular organism with many cell types. The dividing cells manage to organize and assign themselves special, differentiated roles in a reliable manner, creating a spatio-temporal pattern and division of labor. This despite the fact that little positional information may be available to them initially to guide this patterning. Inspired by a model of developmental biologist L. Wolpert, we simulate this situation in an evolutionary setting where individuals have to grow into "French flag" patterns. The cells in our model exist in a 2-layer Potts Model physical environment. Controlled by continuous genetic regulatory networks, identical for all cells of one individual, the cells can individually differ in parameters including target volume, shape, orientation, and diffusion. Intercellular communication is possible via secretion and sensing of diffusing morphogens.

Evolved individuals growing from a single cell can develop the French flag pattern by setting up and maintaining asymmetric morphogen gradients – a behavior predicted by several theoretical models.

The effects of local information and trading opportunities in a network constrained economy

Daniel Ladley¹ and Seth Bullock²

¹University of Leeds

²Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Tuesday 3:30pm–5:00pm, Stripe Studio 1, Markets.

Work within the field of artificial life has a history of exploring the ways in which locally constrained interactions between the elements of a system can give rise to organised behaviour at the level of the ensemble. Here we study the effect of constraining co-operative, competitive and communicative interactions within a market by embedding it within a network. We are particularly interested in how these different kinds of interaction are influenced by the structure of the market network. The paper aims to examine the effect of limited trading opportunities and information availability on the behaviour of individuals and of the market as a whole. It examines how a trader's ability to make profit is influenced by their location within a trade network and how trader strategy must be adapted to cope with this constraint. To this end we employ an agent-based model of trader interaction in which the actions of each trader are governed by individual behavioural rules. Traders are situated on the nodes of a network and interact with potential trading partners through the ties. The networks considered in this work are constructed via preferential attachment schemes resulting in networks both with and without positive assortedness. The behavioural rules of the traders are optimised for their respective locations within the networks through the use of a hill-climbing algorithm. It is demonstrated that a trader's ability to profit and to identify the equilibrium price is positively correlated with its degree of connectivity within the market. Better connected traders are able to exploit their market position at the expense of other market participants. When the effects of constraining trade and information are separated it is demonstrated that when traders differ in their number of potential trading partners, well-connected traders are found to benefit from aggressive trading behaviour. A higher number of potential trading partners allows these traders to demand better terms as there is a higher chance of another trader being willing to trade with them. Where information propagation is constrained by the topology of the trade network, connectedness affects the nature of the strategies employed. Better connected traders attempt to learn more quickly, taking in as much information as possible at the start of the market in order to exploit possible trading opportunities. Less well connected traders learn more slowly and average over time to

avoid being exploited by better connected individuals. We also demonstrate that traders are unable to exploit second order information and trade effects connected to the network. We show that it is not possible for traders to modulate their price or the way in which they weight information based on the connectedness of the potential trading partner/information source to make higher profits. When this situation is permitted all traders adopt strategies such that none benefit from the additional abilities.

Exploiting open-endedness to solve problems through the search for novelty

Joel Lehman and Kenneth Stanley

University of Central Florida

Wednesday 1:30pm–3:00pm, Stripe Main Theatre, Open-Ended Evolution 2.

This paper establishes a link between the challenge of solving highly ambitious problems in machine learning and the goal of reproducing the dynamics of open-ended evolution in artificial life. A major problem with the objective function in machine learning is that through deception it may actually prevent the objective from being reached. In a similar way, selection in evolution may sometimes act to discourage increasing complexity. This paper proposes a single idea that both overcomes the obstacle of deception and suggests a simple new approach to open-ended evolution: Instead of either explicitly seeking an objective or modeling a domain to capture the open-endedness of natural evolution, the idea is to simply search for novelty. Even in an objective-based problem, such novelty search ignores the objective and searches for behavioral novelty. Yet because many points in the search space collapse to the same point in behavior space, it turns out that the search for novelty is computationally feasible. Furthermore, because there are only so many simple behaviors, the search for novelty leads to increasing complexity. In fact, on the way up the ladder of complexity, the search is likely to encounter at least one solution. In this way, by decoupling the idea of open-ended search from only artificial life worlds, the raw search for novelty can be applied to real world problems. Counterintuitively, in the deceptive maze navigation task in this paper, novelty search significantly outperforms objective-based search, suggesting a surprising new approach to machine learning.

Adaptive multi-robot bucket brigade foraging

Adam Lein and Richard Vaughan

Simon Fraser University

Tuesday 10:30am–12:00pm, Stripe Studio 1, Autonomous Robots.

Bucket brigade foraging improves upon homogeneous foraging by reducing spatial interference between robots, which occurs when robots are forced to work in the same space, and must spend time avoiding one another instead of carrying out useful work. Bucket brigade foraging algorithms restrict the motion of each robot to at most some fixed distance from its starting location. We reproduce the performance of known bucket brigade foragers, and then present a new controller in which robots adapt the size of their foraging area in response to interference with other robots, improving overall performance. This approach also has the potential to cope with nonuniform resource distributions.

Tracking the evolution of chemical computing networks

Thorsten Lenser, Naoki Matsumaru, Thomas Hinze and Peter Dittrich

Friedrich Schiller University Jena

Tuesday 10:30am–12:00pm, Stripe Main Theatre, Evolving Cell Signalling Networks *in silico* 1.

How do chemical reaction networks that process information evolve? This is not only a fundamental question in the study of the origin of life, but also in diverse fields like molecular computing, synthetic biology, and systems biology. Here, we study the evolution of chemical flip-flops by means of chemical organisation theory. Additionally, we compare evolved circuits with manually constructed ones. We found that evolution selects for an organisational structure that is related to function. That is, the resulting computation can be explained as a transition between organisations. Furthermore, an evolutionary process can be tracked as a change of the organisational structure, which provides a fundamentally different view than looking at the structural changes of the reaction networks. In our experiments, 90% of evolutionary improvement coincide with a change in the organisational structure. We conclude that our approach provides a novel and useful perspective to study evolution of chemical information processing systems

Simulated evolution of mass conserving reaction networks

Anthony Liekens, Huub ten Eikelder, Marvin Steijaert and Peter Hilbers

Technische Universiteit Eindhoven

Tuesday 1:30pm–3:00pm, Stripe Main Theatre, Evolving Cell Signalling Networks *in silico* 2.

With the rise of systems biology, the systematic analysis and construction of behavioral mechanisms in both natural and artificial biochemical networks has become a vital part of understanding and predicting the inner workings of intracellular signaling networks. As a modeling platform, artificial chemistries are commonly adopted to study and construct artificial reaction network motifs that exhibit complex computational behaviors. Here, we present a genetic algorithm to evolve networks that can compute elementary mathematical functions by transforming initial *input* molecules into the steady state concentrations of *output molecules*. More specifically, the proposed algorithm implicitly guarantees mass conservation through an atom based description of the molecules and reaction networks. We discuss the adopted approach for the artificial evolution of these chemical networks, evolve networks to compute the square root function. Finally, we provide an extensive deterministic and stochastic analysis of a core square root network motif present in these resulting networks, confirming that the motif is indeed capable of computing the square root function.

Distributed gradient optimization with embodied approximation

Yaroslav Litus and Richard Vaughan

Simon Fraser University

Tuesday 1:30pm–3:00pm, Exam Hall Room 1, Spatial Organisation 2.

We present an informal description of a general approach for developing decentralized distributed gradient descent optimization algorithms for teams of embodied agents that need to rearrange their configuration over space and/or time, into some optimal and initially unknown configuration. Our approach relies on using embodiment and spatial embeddedness as a surrogate for computational resources, permitting the reduction or elimination of communication or shared memory for conventional parallel computation. Intermediate stages of the gradient descent process are manifested by the locations of the robots, instead of being represented symbolically. Each point in the space-time evolution of the system can be considered an approximation of the solution, which is refined by the agents' motion in response to sensor measurements. For each agent, motion is approximately in the direction of the local anti-gradient of the global cost function. We illustrate this ap-

proach by giving solutions to two non-trivial realistic optimization tasks from the robotics domain.

We suggest that embodied approximations can be used by living distributed systems to find affordable solutions to the optimization tasks they face.

Emergence of glider-like structures in a modular robotic system

Joseph Lizier¹, Mikhail Prokopenko², Ivan Tanev³ and Albert Zomaya⁴

¹CSIRO; The University of Sydney

²CSIRO

³Doshisha University

⁴The University of Sydney

Thursday 10:30am–12:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 1.

Information-driven evolutionary design has been proposed as an efficient method for designing self-organized multi-agent systems. Information transfer is an important candidate fitness function here. It is known to be an important component of distributed computation in many complex systems, and indeed it has been suggested that maximization of information transfer can give rise to interesting behavior and induce necessary structure in the system. In this paper, we report the first known application of a direct measure of information transfer, transfer entropy, as a fitness function to evolve a self-organized multi-agent system. The system evolved here is a simulated snake-like modular robot. In the most fit snakebot in the final generation, we observe coherent traveling information transfer structures. These are analogous to gliders in cellular automata, which have been demonstrated to represent the coherent transfer of information across space and time, and play an important role in facilitating distributed computation. These observations provide evidence that using information transfer to drive evolutionary design can produce useful structure in the underlying system.

The information dynamics of phase transitions in random Boolean networks

Joseph Lizier¹, Mikhail Prokopenko² and Albert Zomaya³

¹CSIRO; The University of Sydney

²CSIRO

³The University of Sydney

Thursday 1:30pm–3:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 2.

Random Boolean Networks (RBNs) are discrete dynamical systems which have been used to model Gene Regulatory Networks. We investigate the well-known phase transition between ordered and chaotic behavior in RBNs from the perspective of the distributed computation conducted by their nodes. We use a recently published framework to characterize the distributed computation in terms of its underlying information dynamics: information storage, information transfer and information modification.

We find maximizations in information storage and coherent information transfer on either side of the critical point, allowing us to explain the phase transition in RBNs in terms of the intrinsic distributed computations they are undertaking.

‘Psychoanalysis’ of a minimal agent

Santosh Manicka and Inman Harvey

Centre for Computational Neuroscience and Robotics, University of Sussex

Friday 10:30am–12:00pm, Stripe Studio 1, Dynamical Systems Analysis 2.

The dowry problem (also known as the secretary problem) is studied with minimal cognitive agents, being a problem that needs memory and judgment. A sequence of values, drawn from an unknown range, is presented; the agent has only one chance to pick a single value as they are presented, and should try to maximize the value chosen. In extension of previous work (Tuci et al. 2002), Continuous Time Recurrent Neural Networks (CTRNN) are evolved to solve the problem, and then their strategies are analyzed by relating mechanisms to behavior. Strategies similar to the known optimal strategy are observed, and it is noted that significantly different strategies can be generated by very different mechanisms that perform equally well.

Towards mental life as it could be: a robot with imagination

Hugo Marques¹, Owen Holland¹, Rob Knight² and Richard Newcombe³

¹University of Essex

²The Robot Studio

³Imperial College

Friday 10:30am–12:00pm, Stripe Main Theatre, Artificial Consciousness: From ALife to Mind.

“Artificial life is the study of man-made systems that exhibit behaviors characteristic of natural living systems.” So wrote Chris Langton two decades ago, and most people within the discipline still seem happy with his definition. But from the human perspective, the behaviour of our own living system is only half the story, because what matters to most of us is not our behaviour, but our experience — our mental life. Perhaps it is now time for artificial life to embrace the challenge of creating man-made systems that have experiences ‘characteristic of natural living systems’ and giving an account of ‘mental life as it could be’.

This talk will describe some of our recent work within the area known as machine consciousness. Several modern theories of consciousness, notably those of Damasio and Metzinger, stress the importance of embodiment for consciousness, and identify a body-centred self-structure as being a key element in the origin and support of phenomenal experience. In the CRONOS project, we reasoned that human consciousness — the only kind of which we have knowledge — must require both a human-like

body, and a human-like body model, and so we set out to construct suitable entities. The physical robot, CRONOS (Holland and Knight 2006), is a new kind of robot: it is anthropomorphic, faithfully copying the human skeleton, and equipped with appropriately placed elastic muscles and tendons. The body model, SIMNOS (Gamez et al. 2006), is an accurate physics based simulation that behaves sufficiently similarly to CRONOS to be controlled by the same control system.

For a self-model to have evolved, it must confer some functional advantage on its host. We are investigating the hypothesis that one useful function might be the ability to support imagination by predicting the outcomes of possible actions through internal simulation, more or less along the lines suggested by Hesslow (2002). Simulating the body and its controller is only the first step; it is also necessary to create an internal representation of the real world that interacts with the body sufficiently accurately to have predictive value. A physics based internal representation meets this requirement, but the problem of creating such a representation goes well beyond the traditional AI/computer vision challenge of achieving correct geometry, since the objects in the modelled world must also have the correct physical properties. We will report on the use of SLAM techniques and cross-modal integration for dealing with this problem.

In order to obtain functional benefit from imagination, it is not enough merely to have predictively valid models of the body and of the world; it is also necessary to embed the models within a suitable architecture to identify when imagination is required, to generate appropriate candidate actions, to evaluate the success, failure, and costs of the actions, to prevent candidate actions from being acted out by the real robot, to select the preferred action, and to execute the preferred action. We have developed a task-dependent taxonomy of these architectures, and have implemented and tested the ones that seem most important (Marques and Holland 2008). For the first time, this talk will describe and demonstrate an anthropomorphic robot using an architecture for embodied imagination, including a physics based internal model of itself and the world, to select and execute an appropriate action; we believe this may be a significant step towards building a robot with a form of mental life.

Conceptual structure in cellular automata: the density classification task

Manuel Marques-Pita and Luis Rocha

Indiana University

Wednesday 3:30pm–5:00pm, Exam Hall Room 1, Cellular Automata.

The notion of conceptual structure in CA rules that perform the density classification task (DCT) was introduced by Marques-Pita et al. (2006). Here we investigate the role of process-symmetry in CAs that solve the DCT, in

particular the idea of conceptual similarity, which defines a novel search space for CA rules. We report on two new best-known process symmetric rules for the DCT. We further discuss how our results are relevant to understand, control, and design the collective computation performed by other networks of automata, such as those used to model living systems.

The longevity of distinct cultures in an agent-based model of memetic drift

Jamie Matthews

University of Sussex

Tuesday 10:30am–12:00pm, Exam Hall Room 2, Culture & Social Evolution 1.

When a distinct cultural region forms, its rate of absorption into the surrounding culture may be an important variable to take into account when attempting to minimise conflict. This paper describes a re-implementation of Axelrod's agent-based model of cultural dissemination, and uses it to investigate how random drift influences the longevity of distinct regions. Cultural regions are found to be surprisingly resistant to such memetic drift.

Dimensions of adaptivity

Marek McGann

MIC, University of Limerick

Thursday 3:30pm–5:00pm, Stripe Studio 1, Philosophical Issues 3.

Artificial Life offers strong tools for making coherent and disciplined our understanding of the relationships between several fundamental concepts within Cognitive Science more generally such as cognition, value, function, learning and others. One concept which may act as an anchor for this set of ideas is that of adaptivity, in its intraorganismic sense of a system being able to change to maintain itself in the face environmental challenges. A-Life, along with many other areas in Cognitive Science, examines different forms of adaptivity, though rarely in an explicit fashion. We therefore miss many opportunities to analyse the concept of adaptivity itself, and what it might tell us about the relationship between life and mind, learning and value, and other issues within this constellation of notions. The present paper will attempt to outline some of the general characteristics of adaptivity, with the particular aim of identifying what distinguishes different forms of adaptivity (for example, homeostatic regulation and operant learning) and how such dimensions might be used to help organise and direct our thinking on the matter in future. Crucial elements include the timescale over which the adaptive mechanism operates (e.g. achievement of reward in an operant learning task versus strategic play in chess), the inertia of those mechanisms (e.g. the tolerance parameters of a homeostatic mechanism) and integrative capacity of the mechanisms (basically, pattern

recognition). It may be possible for these three dimensions to give us a coherent account of adaptivity and how it varies. This in turn would open new avenues of research into the relationship between cognition and value, and how that relationship changes through the operation of such adaptive mechanisms. The account proposed differs from the likes of Dennett's "Tower of Generate and Test" (1996, *Kinds of Minds*, Weidenfeld & Nicolson) and similar models as it is an attempt at an analysis of adaptivity per se, rather than the kinds of mechanism in a given organism that might produce different forms of competence or adaptive response.

The proposed account might therefore also offer ways in which we could codify the concept of mediacy in the interaction between an agent and its environment. In this, the framework might fit with other theorising on the matter such as Hans Jonas's (1966 *The Phenomenon of Life*, Greenwood Press) arguments concerning the increasing mediacy of the interaction between animals and their environments in evolution, the concept of the "recession of the stimulus" in Edwin Holt's (1915 *Some Broader Aspects of Freudian Ethics* p.134, Holt Company) description of learning and the variety of forms of cognition identified by Merlin Donald (1991 *Origins of the Modern Mind*, Harvard University Press) in his account of cognitive evolution.

An artificial chemistry towards the identification of the transition to life

Ivan Alfredo Mendoza, Patricia Muethé, Jonatan Gómez and Marcela Ewert-Sarmiento

National University of Colombia

Thursday 3:30pm–5:00pm, St Edburga Room 4, Artificial Chemistries.

A simplified model of a natural system in the form of an Artificial Chemistry is presented in order to explain the origin of life. It is intended to serve as a tool for modelling certain scenarios where conditions that could lead to the formation of the compounds, thought to have a major role in the emergence of life, are met. The "molecules" (first element of an Artificial Chemistry) in this model take the form of two-dimensional atoms modelled after those most commonly found in living creatures (with a justification for their inclusion). These atoms have specific shapes (and places for bonds) for metals, nonmetals and metalloids. Their last energy level valences, electronegativities, radii (based on the Helium radius along with all other distances in the system) and weights are also included in the model. The rules were designed with two main purposes, to move the atoms around the reactor and to make and destroy bonds. In this way, molecules are formed, moved around and eventually they give form to new molecular species too. Atoms are initially placed at random over a two dimensional space and also given an initial random velocity and acceleration. Their movement is based then

in one of two alternatives, since no forces can be modelled to move the atoms in this level: the first approach is to give them velocities that change depending on their relationship (bonding potential) with others in the neighbourhood. The second alternative is to randomly accelerate them but keeping a total constant value for this acceleration, proportional to the atomic weight. For bonding purposes a novel approach is introduced. It is based on a combined strategy that involves the electronegativity difference between two atomic species (when free valences are available) and the "gross affinity", a magnitude given for two specific atomic species that was obtained from counting real bonds from a database of bioorganic molecules. This gross affinity favours the formation of compounds that can be familiar in the context of the origin of life, so the future work for this model is based on the fine-tuning of initial conditions (quantities, placing and environment division) and the modification of the affinity to make the model to move towards specific desired results. Also, the emergence of a prevalence of more complex structures and behaviours could be achieved by tracking down the processes that eventually led to them, boosting these processes with new rules can be a form of mimicking natural selection. The preliminary results are optimistic in terms of the production of basic molecules and substrates and compounds as the ones found in experiments like the ones conducted by Miller and Urey.

Fitness transmission: a genealogic signature of adaptive evolution

Thomas Miconi

University of Birmingham

Thursday 3:30pm–5:00pm, Stripe Studio 1, Philosophical Issues 3.

We introduce fitness transmission as a simple statistical signature of adaptive evolution within a system. Fitness transmission is the correlation between the fitness of parents and children, where fitness is evaluated after the number of grandchildren, suitably normalised. This measure is a direct calculation based on a genealogical record, rather than on genetic or phenotypic observation. We point out that the Bedau-Packard statistics of evolutionary activity cannot be used as a reliable system-wide signature of adaptive evolution, because they can produce positive signals when applied to certain "random", non-evolutionary systems. We apply fitness transmission to simple evolutionary algorithms (as well as neutral equivalents) and demonstrate its capacity to accurately detect the presence or absence of Darwinian evolution.

Adaptive units of selection can evolve complexes that are provably unevolvable under fixed units of selection

Rob Mills and Richard Watson

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Tuesday 1:30pm–3:00pm, Stripe Lecture Room, Evolutionary Theory 2.

Symbiosis, the collaboration of multiple organisms from different species, is widespread amongst prokaryotes. As symbiotic associations become more absolute, an inseparable union can result (symbiogenesis). This union would typically be considered a new unit of selection once a set of symbionts became reproductively inseparable. However we consider a macroevolutionary model within which new units of selection are formed gradually from maturing symbiotic relationships, without the need for sybiogenic events to be complete. We find that these adaptive units of selection can evolve complexes that are provably unevolvable under fixed units of selection.

We perform experiments in several structured rugged landscapes (including spin glasses) in which there only few global optima amongst many local optima. Importantly, the number of initial conditions that must be sampled to find one of these global optima by individual selection is exponentially large (i.e., their basins of attraction are small).

Selection on individual species only has the ability to evolve to locally optimal configurations. Selecting on groups at the level of entire ecosystems requires an exponentially large number of ecosystems to be created to find optima of globally maximal fitness. Since this level of selection cannot utilise gradient information, global optima are only guaranteed to be found by enumerating all possible ecosystem configurations.

We model an ecosystem, in which species have variable length genotypes. Every species has the potential to develop a symbiotic association with any other species. Initially all species specify a single gene and have no associations. The ecosystem is sub-divided into several semi-independent demes, each evolving to local attractors via individual selection. Symbiotic associations between the entities that are present in the local attractor discovered by each deme are modified to reflect their frequency of co-occurrence in the attractors across all of the demes: when entities are repeatedly found to be successful with the same partners, the symbiotic relationship is strengthened. Strong associations between entities biases the likelihood of their future co-occurrence. These two phases are repeated, resulting in coalitions of increasing size, stability, and co-dependence. Selection starts at the level of the individual, when local optima are discovered. However, the unit of selection adapts as composites gradually develop. Some composites are favoured at the expense of

other composites, leading to the discovery of higher fitness optima. These composites consist of co-adapted groups that are locally optimal under individual selection, and not random configurations of large sets of entities. Thus, although the groups can contain many entities, the number of competing complexes is small: the configuration space is reduced as symbiotic relationships are strengthened.

We demonstrate the adaptive significance of these processes by applying them to provably difficult optimisation problems. In general, optimisation problems frequently have some structure that can be exploited, but only with an appropriate mechanism that can recognise that structure. When locally optimal configurations contain some information that is congruent with globally optimal solutions, such as when local optima are created from the frustration of a large number of low-order fitness interactions, the process described above can provide automatic problem decomposition.

When composite entities evolve through the intensification of symbiotic coalitions, the units of selection adapt causing the configuration space dimensionality to be collapsed. We show that this can result in complexes that are unevolvable when the units of selection are static.

Anticipating future experience using grounded sensorimotor informational relationships

Naeem Assif Mirza, Chrystopher L. Nehaniv, Kerstin Dautenhahn and René te Boekhorst

University of Hertfordshire

Thursday 1:30pm–3:00pm, Exam Hall Room 1, Sensorimotor Connections 2.

Operational definitions and applications of the sensorimotor experience of an artificial embodied organism are presented along with a mathematical metric for distance between experiences based on Shannon information. We describe a simple robotic experiment that illustrates how an artificial embodied agent can use its own history of experience combined with the experience metric to predict future experience. Present sensorimotor experience is used to find the most similar past experience using the geometry of its growing and changing experience metric space. This is then used to ground the ontogeny of autonomous prospective capability in interacting with the environment, e.g. to anticipate forthcoming changes in environment based on temporally extended past experiences.

Common concepts in agent groups, symmetries, and conformity in a simple environment

Marco Möller¹ and Daniel Polani²

¹Theory of Complex Systems Group, Institute of Solid State Physics, Technical University of Darmstadt, Germany

²Adaptive Systems Group, School of Computer Science, University of Hertfordshire, UK

Thursday 1:30pm–3:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 2.

We analyze representations of the world attained through an infomax principle by agents acting in a simple environment. The representations obtained by different agents in general differ to some extent from each other in different instances. This gives rise to ambiguities in how the environment is represented by the different agents. We now develop an information-theoretic formalism able to extract a “common conceptualization” of the world for a group of agents. It turns out that the common conceptualization intuitively seems to capture much higher regularities or symmetries of the environment than the individual representations.

Abstract We formalize the notion of identifying symmetries in the environment — both with respect to “extrinsic” operations on the environment, as well as with respect to “intrinsic” operations, i.e. the reconfiguration of the agent’s embodiment. In particular, using the latter formalism, we can re-wire an agent to conform to the highly symmetric common conceptualization to a much higher degree than an unrefined agent; and that without having to re-optimize the agent from scratch. In other words, we can “re-educate” an agent to conform to the de-individualized “newspeak” of the agent group with comparatively little effort.

KohonAnts: a self-organizing ant algorithm for clustering and pattern classification

Antonio M. Mora¹, Carlos M. Fernandes², Juan Julian Merelo¹, Vitorino Ramos², Juan Lu s Jim nez Laredo¹ and Agostinho Claudio Rosa²

¹Dpto. ATC. Universidad de Granada

²LASEEB-ISR/IST. University of Lisbon

Wednesday 3:30pm–5:00pm, St Edburga Room 4, Multi-Agent Applications.

In this paper we introduce a new ant-based method that takes advantage of the cooperative self-organization of Ant Colony Systems to create a naturally inspired clustering and pattern recognition method. The approach considers each data item as an ant, which moves inside a grid changing the cells it goes through, in a fashion similar to Kohonen’s Self-Organizing Maps. The resulting algorithm is conceptually more simple, takes less free parameters than other ant-based clustering algorithms, and, after some parameter tuning, yields very good results on some benchmark problems.

Considering the reconfiguration process of subjective temporal order by using recurrent neural networks

Kohei Nakajima and Takashi Ikegami

Department of General Systems Sciences, The Graduate School of Arts and Sciences, University of Tokyo

Friday 10:30am–12:00pm, Stripe Main Theatre, Artificial Consciousness: From ALife to Mind.

Starting from the famous conception of B. Libet, how the brain orders successive events has been a matter of intense debate in neuroscience. S. Yamamoto and S. Kitazawa (S. Yamamoto et al., 2001, *Nature Neuroscience*, 4, p.759) revealed that subjective temporal order of successive taps to hands, or to the tips of sticks held in each hand, are easily reversed just by crossing the arms. What is astonishing is that, especially in the crossing case, when the tapping interval was less than 0.3 second, the mistake rate grew to 100%. This result could be taken as general functionality of the real brain’s time-space reconfiguration process. In this session, we try to explain these phenomena and to characterize the real brain reconfiguration process of subjective temporal order, from dynamical systems perspectives. To construct the computational model, we adopted the following settings. We prepared an agent containing two arms and corresponding pairs of input nodes, consisting of proprioceptive and exteroceptive inputs, and output nodes to make the agent answer which hand had received the stimuli first. Proprioceptive nodes detect the location of the arms and exteroceptive nodes detect the stimulus applied to the hand. For the agent’s internal architectures, we adopted discrete time recurrent neural networks with plasticity and trained the network by using a genetic algorithm depending on the tasks. To prepare the same settings as S. Yamamoto and S. Kitazawa’s experiment, one calculation of internal networks was defined as 0.01 second, and one experiment takes 3 minutes by iterating 60 cycles of time period consisting of “a stimulus application interval”, “an agent responding interval” and “a resting interval”. Each interval was defined as 2 seconds, 2 seconds and 1 second, respectively. The agent receives the stimuli in “a stimulus application interval”, has to respond to the stimuli in “an agent responding interval”, and rests in “a resting interval”. Other actions are not allowed for the agent. Stimuli are designed to create alternative time delay between both hands. For the agent’s task, we set two: Task1 involved changing the locations of both arms for the entire parameterized region, and applying the stimuli to only one hand, after which the agent had to answer which hand had received the stimuli. Task2 involved changing the locations of both arms for a limited parameterized region, non-hand-crossing region, and applying the stimuli to both hands, after which the agent had to answer which hand had received the stimuli. For the experiment, we crossed the agent’s arms, applied the stimuli to both hands and then analyzed

the response and the agent's internal dynamics. Finally, determining the result, we discuss the implications of this model to real brain reconfiguration processes of empirical time and space.

Emergent intracellular network states and cell fate decision: a dynamically integrated study of the epidermal growth factor receptor network with an agent based model

Nuno Nene and Sylvia Nagl

University College London Cancer Institute, Centre for Mathematics and Physics in Life Sciences and Experimental Biology (CoMPLEX)

Friday 10:30am–12:00pm, St Edburga Room 4, Biochemical Modelling.

In biological networks (e.g. protein-protein) component interactions are highly nonlinear. Theoretical models based on ordinary differential equation (ODE) have dominated the simulation of intracellular networks in the literature of theoretical cell biology. This approach is parameter driven and therefore motivated by the data currently available for well studied signalling systems such as the epidermal growth factor receptor (EGFR). Although these models are incredibly detailed they become unmanageable with the increase of network size and hence lose their predictive power. Simplified techniques such as discrete time Boolean networks which strip completely the dynamical system of kinetic parameters and evaluate qualitative network behaviour have been an alternative option for large networks. The network studied in this work involves the extended EGFR signalling pathway and its validated nuclear targets [Oda, K., et al. *Mol Syst Biol* 1:2005.0010] for which only a small part has been modelled using ODEs. The epidermal growth factor receptor (EGFR) has been implicated in the regulation of cell proliferation, survival and differentiation via activation of signalling pathways. Overexpression or constitutive activation of EGFR has been associated with in-vitro tumorigenic transformation and linked for example to non-small cell lung carcinoma (NSCLC), breast and colon cancers. Small molecule kinase inhibitors of the EGFR have been developed and two of them, Gefitinib and Erlotinib, have already been licensed for clinical use in NSCLC. These drugs have been found to have positive impact — reduction in cell proliferation and induction of apoptosis — in patients with mutated EGFR. Nevertheless, the action of the inhibitors has many non-specific interactions [Fabian, M, et al., (2005), *Nat. Biotechnol*, 23, 329]. Both pro-cancerous and anti-cancerous kinases, in addition to EGFR, are inhibited exerting therefore additional effects on the cell. To study the possible consequences of an EGFR signalling network perturbation and devise strategic methods for cell fate decision control, through the administration of drugs, an Agent Based Modelling (ABM) approach was used. Cell phenotype is identified with particular network emergent dynamical states. An extension

of a class of continuous time Boolean networks was analysed under the ABM paradigm. Each agent, seen as a control gate, represents a 'molecular species' (signalling protein or transcription factor) with a normalized concentration value but with only ON/OFF output states determined by a threshold value. Its concentration evolution is represented by a piece-wise linear differential equation with time delays. The time delays stand for diffusion associated effects or other processes that escape the usual assumption of a well mixed system in cell simulations. Higher autonomy of each node is secured by the existence of an internal noise term associated with each time delay. Agents representing transcription factor nodes exhibit higher time delays which stand for the differences in time scale of signal transduction and gene transcription processes. Node targeting is performed by node removal, by modulation of the internal noise term or partial inhibition according to the binding assay data [Fabian et al (2005)]. This modelling approach is the simplest possible one retaining important properties of biological systems such as distributed control, asynchrony and noise.

Is the organism really a machine?

Daniel Nicholson

University of Exeter

Thursday 10:30am–12:00pm, Stripe Studio 1, Philosophical Issues 1.

Despite all its successes, modern biological science has done remarkably little to tackle the fundamental question that lies at the very heart of biology: what is the nature of the living organism? Contemporary biologists (and philosophers of biology, for that matter) seldom ask this question openly and explicitly. The reason is simple: they already presuppose the answer. The organism is a machine.

One could conceivably construe the history of biology since the seventeenth century as the story of the success of the Cartesian notion of the *bte*-machine. Although the dissatisfaction with this mechanistic conception is almost as old as the idea itself, most of those who found themselves in disagreement were labelled as vitalists and marginalised from the scientific discussion. Today, the organism-machine analogy is dominant in virtually every branch of biological science that studies the organism. This conception of life promotes the view that biology is a subsidiary branch of physical science within which the theories and methods of physics, chemistry, and engineering can be fruitfully applied. In this way, the organism-machine analogy serves to not only justify, but also actively encourage, a number of epistemic attitudes, such as strict methodological reductionism, that have dominated the study of life since the mid-twentieth century.

In this paper I question the conceptual coherence of the organism-machine analogy, and in so doing, I challenge some of the central presuppositions underlying con-

temporary biological research. Apart from promoting a misleading view of what the living organism is and how it behaves, I argue that the Cartesian notion of the bte-machine has actually more in common with a Creationist/ID theorist's conception of life than it does with a well-informed evolutionary-developmental account of the organism. This unhappy marriage of Cartesian mechanicism on the one hand, and Neo-Darwinism on the other, has led to a number of tensions which play out both at the theoretical and practical level concerning what the organism is and how it should be studied.

As a symptom of the pervasiveness of the organism-machine analogy in biological thinking, several research programmes have emerged in recent years which aspire to provide the ultimate vindication of the organism-machine analogy. Central among them is Synthetic Biology, although a number of lines of research in A-Life also appear to share this objective. Whilst not denying the legitimacy and usefulness of these new fields, I argue that the faith that has been bestowed upon these disciplines regarding their potential to substantially advance our biological understanding of life is likely to be misplaced.

Finally, I draw from the long-standing anti-mechanistic tradition in biology to propose an alternative, organisation-based conception of the organism that sidesteps the problems of the organism-machine analogy, eliminates some of the deep-rooted conceptual tensions generated by it, and provides an understanding of the organism that is more in accordance with its actual nature.

Intelligent locomotion of eukaryotic cells

Shin Nishimura and Masaki Sasai

Nagoya University

Thursday 1:30pm–3:00pm, Exam Hall Room 2, Wet ALife.

Eukaryotic cells such as cellular slime molds (*Dicystostelium descoideum*) and other animal cells are thought to share unified mechanisms for locomotion. In this work, emergence of “intelligent” behaviors of cells is discussed by developing a simple computational model of locomotion. The model describes changes in cell shape on the two-dimensional plane by considering a cell membrane, actin filaments embedded in the membrane, and an intracellular control factor called “cortical factor”. Actin filaments polymerized on the membrane push it outward to change the cell shape, whereas cortical factor suppresses polymerization of actin filaments. Cortical factor is conveyed from the leading edge to the rear of a moving cell by the intracellular flow of cortex and accumulates at the rear of the cell. This flow of cortical factor leads to the spontaneous locomotion of cells by amplifying the initial fluctuation in cell movement: If a fluctuating cell slightly moves into a direction, cortical factor begins to accumulate at the rear of the moving cell, which suppresses actin polymerization there and further promotes cellular locomotion in

the initial direction. This positive feedback mechanism reproduces the experimentally observed amoeboid-like and keratocyte-like locomotion and cytokinesis B- and C-like cytofission depending on the kinetic rate and the threshold value for actin polymerization in the model, where the amoeboid-like locomotion is a repeat of stop-and-go motion, and the cell usually changes its moving direction after the stopping phase, while the keratocyte-like locomotion maintains a moving direction for long duration. Cytokinesis B-like cytofission divides a cell into two parts, and a cell is torn into several pieces in cytokinesis C-like cytofission. Based on this model of eukaryotic cells, emergence of intelligent behaviors in locomotion is demonstrated. We assume that the reception of external chemical signal suppresses the activity of cortical factor, leading to chemotaxis of a cell toward the source of the chemical signal. We consider that there exist obstacles intercepting a cell on its way to the source of signal. When signal permeates through some obstacles to attract the cell (i.e. traps), the simulated cell falls into a trap at first but it suddenly escapes from the trap to find a way to get it around. The cell finds this detour because the distribution pattern of cortical factor is flushed while the cell is trapped and the occasional fluctuation in cortical factor amplifies locomotion against the gradient of the external signal. In this way, the feedback loop between cell movement and cortical factor is a key mechanism of the emergent behavior to find a detour. We also discuss efficient food finding and cells' hunting moving bacteria. Cognitive locomotion of the model eukaryotic cells is the result of the fluctuating dynamics in which the interaction with the environment and the internal chemical reactions are coupled through the feedback between cell movement and cortical factor.

What can artificial life offer ecology?

Jason Noble¹, Donna Clarke² and Rob Mills¹

¹Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

²School of Biological Sciences, University of Southampton

Tuesday 10:30am–12:00pm, St Edburga Room 4, Ecological Modelling Methodology.

Artificial life is the simulation and synthesis of living systems, and ALife models show how interactions between simple entities give rise to complex effects. Ecology is the study of the distribution and abundance of organisms, and ecological modelling involves fitting a linear model to a large data set and using that model to identify key causal factors at work in a complex ecosystem. We are interested in whether the individual-based modelling approach of ALife can be usefully employed in ecology.

ALife models are “opaque thought experiments” (Di Paolo et al., 2000, Proc. ALife VII, p.497). They show that a phenomenon can arise from a given set of assumptions in cases where the implication is not clear from intuition alone: e.g., that spatial structure in a population can

lead to altruistic behaviour. This type of modelling can be useful to ecology by showing the plausibility of a novel concept or process, which in turn suggests new natural experiments and new forms of data to collect. However, we argue that ALife models can go beyond this “proof of concept” role and serve as a direct account of data in the same way that statistical models do.

We focus on a typical problem from ecology: the effect of clearing powerline corridors through a forest on the local wildlife populations (Clarke et al., 2006, *Wildlife Research*, 33, p.615). The real data set in this case is complex and, of course, we don’t know the true effects that underlie it. We therefore generated a fictional data set that reflects aspects of the original problem while allowing complete control over the simulated environment. The idea is to construct a test case for looking at the relative success of different modelling approaches. We know the true picture because we generated the data, but which modelling approach will get closer to the truth? The fitting of generalized linear models as is conventional in ecology, or the use of individual-based simulations as in ALife?

Statistical models are fitted using some variant of the method of maximum likelihood: given the data, which of the models in the family we’re considering (e.g., a linear regression) makes the observed data most plausible? When dealing with simulations, however, it is difficult to establish that one model is a better fit to data than another. Simulations have many parameters, it may be difficult to determine a level of granularity at which the simulation output is supposed to “match” the data, and there will be no analytically tractable likelihood function. These problems are solved by the method of indirect inference (Gouriéroux et al., 1993, *J. Applied Econometrics*, 8, p.S85) in which an auxiliary model is fitted to both the real data and to the output from competing simulation models. The best simulation model is the one producing the closest match to the data in terms of fitted parameter values in the auxiliary model.

Using indirect inference with our fictional data set we demonstrate that ALife simulation models can be fitted to realistic ecological data, that they can out-compete standard statistical approaches, and that they can thus be used in ecology for more than just conceptual exploration.

Information-theoretic characterization of relative and fluctuating system–environment distinction

Takayuki Nozawa and Toshiyuki Kondo

Tokyo University of Agriculture and Technology

Friday 10:30am–12:00pm, Exam Hall Room 2, Information in Complex Systems & ALife 4.

Defining a system in distinction from its environment is a fundamental but elusive problem in artificial life as well as in real-world complex systems. While many notions of closure gives a qualitative and absolute criteria

for the system–environment distinction, the concept of “informational closure” proposed by Bertschinger et al. (Bertschinger et al., 2006, *Proc. GWAL-7*, p.9, IOS Press) gives gradual and relative evaluation of closure (or closedness). There, a system is tentatively defined in distinction to its environment, and the validity of the definition is judged according to how causally closed the system is, being quantified by information flow (transfer entropy) from the environment into the system. This quantitative approach for the characterization of closedness is expected to bring rich description of “relative” systems on a wide range of dynamical models.

In this study we proceed one step further in the direction of relativizing closure: for the evaluation of closedness we also utilize information-theoretic measures, such as the transfer entropy and difference of Boltzmann-type and KS-type entropies, but instead of evaluating closedness of a system with its elements fixed in time, we evaluate the closedness for the system’s specific states which are dissociated from the history of interaction with the environment. This dissociation excludes from the system–environment correlation the components which are realized by the system modeling or controlling the environment. Therefore, the measures evaluate solely how a state can prevent the invasion of uncertainty from the environment. This setting can be effective in describing partial closures which appear transiently and fluctuate in uniformly structured degrees of freedom or in a directed flow of information processing, while the original setting of informational closure would be more efficient when the meanings of a system and its environment are clear, their boundary is fixed, and advanced notions such as cognition, learning, self-reference, etc. are of immediate interest.

We apply the method to discrete dynamical networks and a cellular automata model which simulates physico-chemical self-organization (molecular aggregation). The spectrum of closedness is shown to depend on the dynamical properties of each model. (The investigation of the spectra has some similarity with the exploration of characteristic structures in the phase spaces of chaotic systems.) We will also discuss how reversibility of the models and introducing dissipative irreversibility, that is, disregarding information flow into the environment as a heat-bath, can influence the evaluation of the closedness.

Personalities in fish without genetic differences: a model

Johanneke Oosten and Charlotte Hemelrijk

Centre for Ecological and Evolutionary studies, University of Groningen
Thursday 1:30pm–3:00pm, Stripe Main Theatre, Ecological Modelling 1.

Differences in personality between individuals are usually attributed to genetic differences, and seldom to differences in experience. In this study, we investigated to what degree personalities may result from differences in experience, that are caused by self-organisation of behaviour and chance. We use an individual-based model to generate such an explanation for the experimental findings on personalities in perch (Magnhagen & Staffan., 2005, *Behav Ecol Sociobiol*, 57, p.295). In this study, small groups of individuals could either hide in vegetation, or visit an open area that contained food, but was near a predator. Individuals were attributed a personality, based on the time that they spend in the open area, and on how fast they fed there.

In our model, we mirror this experiment, with artificial individuals that are genetically identical. We test whether personalities arise as a consequence of three mechanisms: habituation, social facilitation, and competition. To this end, we study three models: a model of only habituation, a second one of habituation and social facilitation, and a third that includes all three mechanisms.

The first model focuses on habituation. Artificial individuals habituate by increasing their (initially low) tendency to enter the open area after each successful foraging event. Although this self-reinforcing effect led to personality differences, these differences disappeared as soon as all individuals were habituated.

In the second model, we added social facilitation. This implied that individuals tended to visit the open area more, if it was already occupied by group members. Social facilitation led to a positive correlation between personality of an individual and that of its group members. Because of this, groups more often arose, that consisted of a single personality type. Both observations resemble the empirical data. However, in this model, personality differences also disappeared because all individuals habituated.

In the last model, we added competition. This was represented by a ‘residence effect’: upon arrival in the open area individuals captured less prey, if group members had arrived before them. Here, personality differences appeared to be stable over long periods of time.

In sum, our models show that differences in speed of habituation may give rise to personality differences, and that these differences are reduced by social facilitation, and maintained by competition. We conclude that it is valuable to consider learning and social interactions as an explanation for the origination of personality in animals.

that can reside in either a protected area or an open area with food. However, all individuals are genetically identical, and start with a low tendency to reside in the open area. Similar to the empirical experiment, they can reside in either a protected area, or an open area with food.

Emergence of cooperation in N-player games on small world networks

Colm O’Riordan, Alan Cunningham and Humphrey Sorensen

National University of Ireland, Galway

Tuesday 10:30am–12:00pm, Exam Hall Room 1, Spatial Organisation 1.

The emergence of cooperation in social dilemmas has been addressed in a number of fields. In this paper, we illustrate how robust cooperation can emerge among a population of agents participating in a N-player dilemma when the agents are spatially arranged on a graph exhibiting small world properties. We present a graph structure with a high level of community structure, small diameter and a variance in the node degree distribution. We show that with simple learning rules, robust cooperation emerges. We also show that a population of agents whose interactions are constrained by such a graph can adapt to dramatic environmental changes.

Optimal noise in spiking neural networks for the detection of chemicals by simulated agents

Nicolas Oros, Volker Steuber, Neil Davey, Lola Cañamero and Rod Adams

University of Hertfordshire

Tuesday 1:30pm–3:00pm, Stripe Studio 1, Neural Systems.

We created a spiking neural controller for an agent that could use two different types of information encoding strategies depending on the level of chemical concentration present in the environment. The first goal of this research was to create a simulated agent that could react and stay within a region where there were two different overlapping chemicals having uniform concentrations. We used a spiking neural network to control the agent that encodes its sensory information as temporal coincidences when the level of chemical concentration is low, and as firing rates at high level of concentration. With this architecture, we could study synchronization of firing in a simple manner and see its effect on the agent’s behaviour. The next experiment we did was to use a more realistic model by having an environment composed of concentration gradients and by adding input current noise to all neurons. We used a realistic model of diffusive noise and showed that it could improve the agent’s behaviour if used within a certain range. Therefore, an agent with neuronal noise was better able to stay within the chemical concentration than an agent without.

Holey fitness landscapes and the maintenance of evolutionary diversity

Gregory Paperin, Suzanne Sadedin, David Green and Alan Dorin

Monash University

Thursday 10:30am–12:00pm, Exam Hall Room 2, Fitness Landscapes.

Analytical models show that high-dimensional fitness landscapes form “holey” rather than “rugged” topographies, but the implications of this finding for biological and artificial life systems remain largely unexplored. One of the reasons for this gap can be attributed to serious difficulties in the implementation of individual-based holey fitness landscape (HFL) models. Here, we introduce a method for simulating HFLs in spatially explicit individual-based models that overcomes these difficulties. We examine how the HFL changes predictions for the maintenance of genetic diversity in the face of migration. Previous models suggest that ecologically-based reproductive isolation will rapidly collapse under migration. Our results indicate that an underlying HFL can often maintain diversity in this situation. Hybrid species emerge frequently when HFL genetics are simulated, but are usually doomed to extinction because of small population sizes. However, hybridisation can also lead to novel adaptations and potentially the exploitation of new ecological niches. More generally, the results imply that HFL genetics should not be neglected in studies of adaptation and diversity.

Conformist transmission and the evolution of cooperation

Jorge Peña

Institute of Applied Mathematics — University of Lausanne

Tuesday 10:30am–12:00pm, Exam Hall Room 2, Culture & Social Evolution 1.

We study the effects of conformist transmission on the evolutionary dynamics of the Prisoner’s Dilemma, the Snowdrift and the Stag Hunt games in both well-mixed and spatially structured populations. The addition of conformism introduces a transformation of the payoff matrix that favours the stability of pure equilibria and reduces the basin of attraction of risk dominant equilibria. When both conformism and local interactions are present, the system can exhibit higher levels of cooperation than those obtained in the absence of any of the two mechanisms.

What can artificial life offer the development of methodologies in the field of socio-ecological sustainability?

Alexandra Penn

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Tuesday 3:30pm–5:00pm, Stripe Main Theatre, Sustainability & Homeostasis.

Issues of sustainability involve the dynamics and interactions of multiple complex adaptive systems at a variety of scales: climatic; ecological; economic; technological; political and social. Some of the most pressing challenges for society are inherently concerned with gaining a better ability to understand and manage the interacting systems upon which we rely. While there is widespread acknowledgement that the science of complex adaptive systems can provide key tools to address these challenges, there is little consensus on how to develop and apply these tools effectively. New approaches are therefore needed before effective policy-making can be informed by well-founded scientific modelling. Concepts and terminology such as complexity, complex adaptive systems, whole systems thinking, non-linear dynamics, co-evolution, autopoiesis, and self-organisation enjoy common currency in movements such as resilience thinking, sustainable systems approaches and permaculture (amongst others), which share a common ancestry with A-Life in systems theory. The use of such ideas as metaphors to guide thinking is valuable up to a point. For example, the concept of non-linear response to change, including at the extreme system tipping points, is an important understanding which must certainly guide policy in areas such as climate change. A large array of complex systems metaphors are also used as sources for design and management heuristics. However, there are presently enormous methodological leaps to be made before their full potential usefulness can be realised, and the availability of clear quantitative or qualitative measures and methodologies connecting theory with practice is extremely limited. A clear opportunity exists for the field of Artificial Life to contribute in this domain at this key time. In this talk I will give an overview of the current use of complex and dynamical systems concepts within the sustainability movement and associated challenges. I will detail practical tools being developed to measure the qualitative or quantitative behaviour (or health) of dynamical systems such as ecosystems, and discuss how we can move from a metaphorical understanding of such systems as complex, dynamical or adaptive, towards strategic intervention in or interaction with them with the goal of sustainability in mind.

I will focus on what I consider to be three key areas in which A-Life methodology can contribute: 1) The use of modelling to predict the gross behaviour of systems, with a particular emphasis on the incorporation of evolutionary

processes, network dynamics, and agent-based modelling into current resilience approaches; 2) The development of quantitative indicators of systems' "health" with regard to their ability to self-maintain; 3) The development of tools for management or steering of complex systems undergoing rapid change, including the potential for "engineering" or "programming" self-organisation of complex adaptive systems for increased resilience, robustness and "sustainability".

The aim of this talk is to initiate dialogue between theoreticians and practitioners towards practical use of A-Life methodologies in frontline sustainability.

Mechanisms for the initiation of multicellularity in bacterial biofilms

Alexandra Penn¹, Richard Watson¹, Simon Powers¹, Jeremy Webb², Alex Kraaijeveld², Tim Conibear² and Zoe Bigg²

¹Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

²School of Biological Sciences, University of Southampton

Wednesday 1:30pm–3:00pm, Exam Hall Room 2, Major Evolutionary Transitions 1.

Major transitions in evolution create the prerequisite features which allow natural selection to occur at a new level of organisation. Heredity, variation and reproduction must all be produced at a new higher level in order for a new evolutionary unit to arise. The origin of these features at new levels of organisation has been problematic for both theoretical biology and in artificial life models — but several partial theories exist. The necessary features may arise under the action of adaptive processes on existing units, and/or with the potential support of self-organisation of some kind. Limited mechanisms of heredity, potentially including ecological inheritance of constructed niches, may play an important role in bootstrapping the early stages of a transition to higher-level selection on new units. But, in short, theories for the precise routes by which new biological individuals might arise remain mostly speculative.

Our programme of theoretical modelling work has been focussing on simple individually-adaptable characters that may be involved in initiating higher-level units of selection. It is known that two key determinants of the efficacy of higher-level selection, between-group variation and heredity of group characters (e.g. the species composition of a group), are significantly affected by the modification of some simple parameters such as initial group size in an aggregation and/or the size of dispersal propagules. Our models have demonstrated that if simple features such as group structure parameters, group size and dispersal modes can be affected by characters that are under individual adaptive control then conditions that effect significant higher-level selection can be selected for despite individual self-interests. These features are simple

enough that we can begin an empirical programme to investigate and manipulate the relevant variables.

In this talk we describe an ongoing experimental programme to investigate parameters affecting the levels of selection using bacterial biofilms. The majority of bacteria spend most of their life cycle in single- or multi-species biofilms, complex collective structures formed when bacteria attach to surfaces, and in this form they display an extraordinary repertoire of coordinated behaviours and interactions. Bacterial biofilms have numerous high-impact application areas including bio-engineering, bioremediation and medicine where controlling the adaptation and co-adaptation of bacteria is vital. Reproduction in biofilms may be either via shearing off of groups of cells or by the production of individual motile cells. Despite this, such groups are also able to disaggregate into individual cells which reproduce in a planktonic phase. Accordingly, these organisms, in this case *Pseudomonas aeruginosa* (a common opportunistic pathogen problematic in cystic fibrosis), provide an excellent model system to address questions concerning the transition to multicellularity. They are fast-growing and experimentally tractable, allowing us to perform multi-generational evolutionary experiments over a relatively short timeframe, and they naturally exhibit physical characteristics, such as micro-colony formation, that implement group structure. They possess individual characters, such as siderophore production, that can be knocked-out to produce clearly identifiable 'cheats' and 'cooperators'. Crucially, there are easily identifiable individual characters, such as extra-cellular matrix production, that clearly affect the grouping parameters that are of interest, such as group size and propagule dispersal. We present our experimental methodology for manipulating these characters and thereby parameters that affect the strength of group selection; a vital first step in tackling the investigation of evolutionary transitions in real organisms.

Quantifying creative symbiosis: a lexical analysis of the evolution of technology as reflected in patent records

Noah Pepper, Charles Francis, Mark Bedau, Devin Chalmers and Bryan Head

Reed College

Tuesday 1:30pm–3:00pm, Exam Hall Room 2, Culture & Social Evolution 2.

The similarities and differences between adaptive dynamics in biological and cultural evolution is an important and controversial open question about evolutionary processes in the real world. One way to address this issue is by studying empirical data from biological and cultural evolution. Technology is itself an important part of culture, and one that is rather amenable to empirical investigation, and an excellent window on the evolution of technology is

to study patterns in the citations among patent records (Jaffe & Trajtenberg, 2005, *Patents, Citations, and Innovations*, MIT Press). Each patent must describe a novel concept and must cite previous related works and each patent is filed at a distinct point in time. This provides us with a time series corpus of formally written text that records all patented technological innovations, an ideal platform for studying the creativity of one form of cultural evolution.

Bedau and Skusa (2002, *ALife VIII*, p. 431, MIT Press) analyzed the dynamics of adaptation in the evolution of technology, by looking at patent citations and measuring evolutionary activity statistics over time (Bedau & Packard, 1991, *ALife II*, p. 431, Addison-Wesley; Bedau et al., 1998, *ALife VI*, p. 233, MIT Press). We here present a novel, complementary method of analyzing technological evolution using the textual content of patent records. Our analysis takes linguistic tokens as the unit of cultural adaptation and measures their occurrence and relations using several linguistic tools. This can reveal latent connections between conceptual or technological innovations.

We have analyzed a corpus containing thirty years of patents (over 4 million) with the WORDSPACE model for quantification of word relatedness within a corpus (Widdows, 2004, *Geometry and Meaning, CSLI*). This produces a fuzzy set of correlated terms based on concurrence within a text. The change in relatedness of n-grams over time provides a movie of a part of cultural evolution. We also analyze regressions on time series frequency counts of n-grams and groups of n-grams. (N-grams are linguistic tokens n words in length which represent technologies or concepts in the corpus.) Patterns in n-gram frequencies provide another window into the evolution of culture.

These two tools reveal and map the cross-temporal relationships between technologies. One can see classes of technologies share in significant evolutionary success, and then eventually decline (for example, the recent technology bubble). The tools can reveal when different technologies have symbiotic relationships (Margulis & Fester, 1991, *Symbiosis as a Source of Evolutionary Innovation*). These results help illuminate the nature of cultural change and open-evolution, and in particular, whether there is a fundamental difference in the evolution of cultural and biological systems.

Hawks and doves in an artificial dynamically structured society

Enea Pestelacci and Marco Tomassini

Information System Institute, HEC, University of Lausanne, Switzerland
Wednesday 1:30pm–3:00pm, Stripe Studio 1, Networks in Natural & Artificial Systems 2.

Using a dynamical network model of society, we show that cooperation is the norm in the Hawks-Doves game when individuals are allowed to break ties with undesirable neighbors and to make new acquaintances in their

extended neighborhood. This is an interesting result, as standard theory for mixing populations prescribes that a certain fraction of defectors must always exist at equilibrium. We discuss the empirical network structure reasons that allow cooperators to thrive in the population.

Evolving morphological and behavioral diversity without predefined behavior primitives

Peter-Paul Pichler and Lola Cañamero

University of Hertfordshire

Thursday 1:30pm–3:00pm, Exam Hall Room 1, Sensorimotor Connections 2.

Virtual ecosystems, where natural selection is used to evolve complex agent behavior, are often preferred to traditional genetic algorithms because the absence of an explicitly defined fitness allows for a less constrained evolutionary process. However, these model ecosystems typically pre-specify a discrete set of possible action primitives the agents can perform. We think that this also constrains the evolutionary process with the modellers' preconceptions of what possible solutions could be. Therefore, we propose an ecosystem model to evolve complete agents where all higher-level behavior results strictly from the interplay between extremely simple components and where no 'behavior primitives' are defined. On the basis of four distinct survival strategies we show that such primitives are not necessary to evolve behavioral diversity even in a simple and homogeneous environment.

Complex systems models: engineering simulations

Fiona Polack¹, Tim Hoverd¹, Adam Sampson², Susan Stepney¹ and Jon Timmis²

¹Computer Science, University of York

²Computing Laboratory, University of Kent

²Computer Science and Electronics Departments, University of York

Tuesday 1:30pm–3:00pm, St Edburga Room 4, Methodologies.

As part of research towards the CoSMoS unified infrastructure for modelling and simulating complex systems, we review uses of definitional and descriptive models in natural science and computing, and existing integrated platforms. From these, we identify requirements for engineering models of complex systems, and consider how some of the requirements could be met, using state-of-the-art model management and a mobile, process-oriented computing paradigm.

Modelling stigmergic gene transfer

Daniel Polani¹, Mikhail Prokopenko² and Matthew Chadwick²

¹University of Hertfordshire

²CSIRO ICT Centre

Thursday 3:30pm–5:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 3.

We consider an information-theoretic model studying the conditions when a separation between the dynamics of a 'proto-cell' and its proto-symbolic representation becomes beneficial in terms of preserving the proto-cell's information in a noisy environment. In particular, we are interested in understanding the behaviour at the "error threshold" level which, in our case, turns out to be a whole "error interval". We separate the phenomena into a "waste" and a "loss" component; the "waste" measures "packaging" information which envelops the proto-cell's information, but itself does not contain any information of interest, the "loss" measures how much of the proto-symbolically encoded information is actually lost. We observe that transitions in the waste/loss functions correspond to the boundaries of the "error interval". Secondly, we study whether and how different protocells can share such information via a joint code, even if they have slightly different individual dynamics. Implications for the emergence of biological genetic code are discussed.

The efficacy of group selection is increased by coexistence dynamics within groups

Simon Powers, Alexandra Penn and Richard Watson

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Tuesday 10:30am–12:00pm, Stripe Lecture Room, Evolutionary Theory 1.

Selection on the level of loosely associated groups has been suggested as a route towards the evolution of cooperation between individuals and the subsequent formation of higher-level biological entities. Such group selection explanations remain problematic, however, due to the narrow range of parameters under which they can overturn within-group selection that favours selfish behaviour. In principle, individual selection could act on such parameters so as to strengthen the force of between-group selection and hence increase cooperation and individual fitness, as illustrated in our previous work. However, such a process cannot operate in parameter regions where group selection effects are totally absent, since there would be no selective gradient to follow. One key parameter, which when increased often rapidly causes group selection effects to tend to zero, is initial group size, for when groups are formed randomly then even moderately sized groups lack significant variance in their composition. However, the consequent restriction of any group selection effect to small sized groups is derived from models that assume

selfish types will competitively exclude their more cooperative counterparts at within-group equilibrium. In such cases, diversity in the migrant pool can tend to zero and accordingly variance in group composition cannot be generated. In contrast, we show that if within-group dynamics lead to a stable coexistence of selfish and cooperative types, then the range of group sizes showing some effect of group selection is much larger.

The group selection debate and ALife: weak altruism, strong altruism, and inclusive fitness

Simon Powers and Richard Watson

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Tuesday 3:30pm–5:00pm, Stripe Lecture Room, Evolutionary Theory 3.

Models of the evolution of social behaviour are often framed in terms of either multi-level selection or inclusive individual fitness theory. Although both of these descriptions correctly predict changes in gene frequency (where group fitness is defined as the average individual fitness of the group members), it is still a hotly contested issue as to which provides a faithful description of the underlying causal processes at work. Furthermore, the type of model analysis used reflects the philosophical bias of the author. It is important for ALife researchers to be aware of this issue when evaluating or presenting models of social evolution, for many authors simply claim as a matter of fact that their model works via multi-level or (inclusive) individual selection, without acknowledging the alternative perspective.

In this talk, two particular areas of ongoing contention between multi-level and individual selectionists will be illustrated, using examples from the ALife literature. The first of these concerns the evolution of weakly altruistic traits. These are behaviours that provide a whole-group benefit at some cost to the actor. Crucially, however, the cost to the actor is more than offset by its share of the group benefit, such that the lifetime number of offspring of the actor is increased. In a recent paper West et al. (2007, *J. Evol. Biol.*, 20, p.415) have advocated that the evolution of such traits can be adequately explained in terms of direct fitness benefit, thus avoiding the need to invoke selection at the group level. However, this explanation hides the fact that weak altruists suffer a relative fitness disadvantage within every group. Indeed, the local attractor within any one group is the extinction of weak altruists. Therefore, the behaviour cannot spread unless groups compete and groups with more weak altruists are fitter than those with less. While the individualist methodology correctly predicts if the behaviour will evolve, it obscures the mechanistic explanation. This suggests that models couching the evolution of social behaviour in terms of individual benefit should be analysed to determine whether

group structure is playing any causal role in the evolutionary dynamics.

The second issue to be addressed by this talk concerns the evolution of strong altruism, i.e., behaviours where there is a reduction in the lifetime number of offspring of the actor. For such behaviours to evolve there must be a correlation in interactions, such that the recipients of an altruist's help tend to be altruists themselves. This correlation frequently occurs in nature through the limited dispersal of kin, and is usually modelled by inclusive fitness equations that contain no notion of group fitness. However, the underlying mechanism is that kin groups with more altruists outcompete those with less. Once this is realised, it becomes apparent that other assortative group formation mechanisms can in principle produce the same effect. Appealing to kinship is therefore simply invoking one kind of assortative grouping.

This talk will further elaborate on these points, including definitions of a group, and consider claims about the strength of group selection.

A strong local dependency of developmental bias orients adaptive evolution in a model evo-devo system

Sean Psujek¹ and Randall Beer²

¹Case Western Reserve University

²Indiana University

Wednesday 10:30am–12:00pm, Stripe Lecture Room, Development & Morphogenesis 1.

The differential production of phenotypes, that is, developmental bias, has been presumed to have a strong influence on the direction of evolutionary modification. In a computational evo-devo system generating neural connectivity, we demonstrate an orienting role for developmental bias during adaptive evolution. The differences in phenotypic transitions taken during evolution are, to a large degree, due to phenotypic accessibility being strongly dependent on the interactions of the developmental process as directed by a particular genotype. We define developmental bias as the differential production of phenotypes given uniform genetic variation. In our gene-based developmental system, we approximate the range of phenotypic variation possible by creating random genotypes and determining the phenotypes they produce. The resultant pattern of phenotypic variation indicates an intrinsic bias (global bias) in the developmental system. We then determine the accessibility of phenotypes from a given genotype (local bias) by determining the phenotypes generated through each and every single-base substitution. We find that the local bias patterns vary strongly with the genotype, even among phenotypically-neutral genotypes. These patterns also differ from the global bias pattern indicating local biases depend more on the dynamics of the developmental process than on the overall mechanisms of the developmental system.

During evolutionary simulations toward a target phenotype, the local dependency of bias dictates the phenotypic transformations that occur. For example, in two simulations at the generation preceding an increase in the populations' best fitness (the populations have approximately the same average fitness), the target phenotype is produced in one population but not in the other. The average of the local bias patterns for all individuals in the second population (population bias) shows the target phenotype is completely inaccessible through mutagenesis of the population. Other fitness-increasing phenotypic transitions show a similar result; the particular phenotype produced is dictated by the phenotypic variants accessible from the population. This results in multiple phenotypic pathways to the target phenotype across simulations.

Because local bias has such a strong dependency on the dynamics of the developmental process as determined by the regulatory structure of the genotype, bias patterns often change dramatically during evolution through the accumulation of mutations (neutral or otherwise). Phenotypic variants that are possible, as indicated by local bias patterns, occasionally are not able to be generated in subsequent generations. More importantly, phenotypes previously inaccessible often become available after multiple rounds of mutation. Mutations change the developmental context in which subsequent mutations operate. In one example, a mutation previously selectively-neutral eventually becomes a beneficial mutation, resulting in a change to a higher-fitness phenotype.

These results indicate that developmental bias has a strong influence on the direction of evolutionary modification. More generally, there are features of the genotype-to-phenotype and phenotype-to-fitness mappings that affect evolvability, the capacity to vary in phenotypic availability over time.

Mechanism as mind: what tensegrities and caterpillars can teach us about soft robotics

John Rieffel¹, Barry Trimmer¹ and Hod Lipson²

¹Tufts University

²Cornell University

Wednesday 10:30am–12:00pm, Exam Hall Room 2, Amorphous & Soft Robotics.

Recent interest is being applied to the possibility of soft robots with few if any rigid parts, able to substantially deform themselves in order to flow around, and even through objects. In order to accomplish these goals in an efficient and affordable manner, space and power will be at a premium, and so soft robots will most likely be both under-actuated and under-controlled. One promising approach to solving the ensuing actuation and control problems lies in attempting to embody a portion of both tasks within the dynamics of the structure itself. Such "morphological computation" is known to exist throughout the

biological world, from the behavior of cellular cytoskeletons up to the tendinous network of the human hand. Here we present two examples of morphological computation — one from biology, the *manduca sexta* caterpillar, and one from engineering, a large irregular modular tensegrity tower — and explore how ideas from these realms can be applied toward locomotion and control of a highly articulate, under-controlled, soft robot.

Evolutionary robotics models in the interdisciplinary study of embodied time perception

Marieke Rohde and Ezequiel Di Paolo

Centre for Computational Neuroscience and Robotics, University of Sussex

Thursday 3:30pm–5:00pm, Exam Hall Room 1, Embodiment.

A recent study by Cunningham et al. (Cunningham et al., 2001, *Psychological Science*, 12, p. 532) has shown that human subjects adapt to delayed visual feedback in a visuomotor task both behaviourally and experientially, i.e., the behaviour is altered in such a way that successful performance on the task relies on the presence of a visual delay (negative after-effect) and that the experience of simultaneity is re-adjusted to incorporate the visual delay. This adaptation effect is similar to those observed in experiments with visual displacements, but contrasts with earlier experiments with sensory delays, in which no such adaptation occurred (e.g., Smith and Smith, 1962, *Perception and Motion*, Saunders).

This discrepancy (i.e., adaptation in some situations but not in others) suggests that adaptation to sensory delays does not proceed automatically, on the basis of statistical properties of sensory inputs, but is contingent on the performed behaviour and the associated sensorimotor dynamics. Artificial Life and Evolutionary Robotics simulation models are proven tools in the study of non-linear sensorimotor dynamics, which are difficult to understand intuitively. In particular, our earlier work (Di Paolo et al., 2008, *New Ideas in Psychology*, forthcoming; Rohde, 2008, PhD Thesis, University of Sussex; Rohde and Di Paolo, 2007, *ECAL 2007*, p. 193, Springer) argues and demonstrates how Evolutionary Robotics simulation models can contribute to the scientific study of human sensorimotor adaptation.

In a combined experimental and evolutionary robotics modelling study, we have tested the (unconfirmed) hypothesis put forward by Cunningham et al. (Cunningham et al., 2001, *Psychological Science*, 12, p. 532) that adaptation to sensory delays occurs if there is time-pressure on the task (Rohde, 2008, PhD Thesis, University of Sussex; Rohde and Di Paolo, 2007, *ECAL 2007*, p. 193, Springer). On the basis of data analysis of both the artificial model agents' and the experimental subjects' sensorimotor recordings we revised our hypothesis: We now believe that, apart from time pressure, the task needs to

feature a systematic link between present motion and future sensation over a longer time span in order to make the task predictable.

This new hypothesis will be tested using a combined Evolutionary Robotics modelling and experimental psychophysics approach proposed and applied in (Rohde, 2008, PhD Thesis, University of Sussex) that aims at formalising and explaining the sensorimotor invariances associated with perceptual experience of time and simultaneity. We argue that the contingent relation between function and underlying mechanisms inherent in Evolutionary Robotics simulations helps to identify general dynamical principles and fundamental sensorimotor invariances across viable solutions. In this aspect, the approach taken is more general and less biased, even though also less transparent than related approaches like robotic forward model learning (e.g., Tani, 1996, *IEEE Trans. SMC (B)*, 26, p. 421). This novel methodological framework, which is characterised by a close match between simulation model and minimalist empirical experiment, can be applied to other problems of perceptual experience and opens up new powerful avenues for interdisciplinary research that uses Artificial Life methods to study of human perception and cognition in the closed sensorimotor loop.

Selective attention in artificial organisms

Fabio Ruini¹ and Domenico Parisi²

¹University of Plymouth, School of Computing, Communications and Electronics

²Institute of Cognitive Sciences and Technologies, National Research Council, Rome

Wednesday 3:30pm–5:00pm, Stripe Studio 1, Learning, Action Selection, & Attention.

Behaviour has two levels, which we will call strategic and tactical. At the strategic level the organism has to choose which particular activity to pursue at any given time (e.g., looking for food, finding a sexual partner, escaping from dangers, sleeping). At the tactical level it must implement the particular sequence of actions that makes it possible to achieve the goal of the chosen activity. The strategic level has important consequences for attention. Since an organism receives many different stimuli at the same time, it has to selectively attend to the ones that are relevant to the current activity while ignoring those which are not.

A population of organisms lives in an environment with randomly distributed food elements in which a predator appears from time to time. The organism's behaviour is controlled by a neural network with input units encoding the location of the nearest food and other input units encoding the location of the predator when it is present. Both sets of input units are connected to a single layer of internal units, which in turn are connected to the output units that control the organism's movements. The connection weights of the neural architecture are evolved using a genetic algorithm where the organism's fitness depends on

both the number of food elements eaten and the organism's ability to avoid being reached by the predator. We contrast two populations of organisms, one with the basic neural architecture and the other one with an architecture which includes an additional set of units which receive connections from the input units encoding the predator's location and send connections to the internal layer. Both populations are able to evolve the appropriate behaviour which consists in looking for food when the predator is absent and flying away from the predator when it is present, ignoring food. However, the population with the additional units reaches higher levels of fitness compared with the population with the simpler architecture. Two additional control simulations show that higher levels of fitness are not obtained if we simply increase the number of internal units or if we connect the additional units directly to the output layer.

To better understand why the additional units yield a better performance we compared the activation patterns of the internal units when the locations of both food and predator are encoded in the input units and when only the latter is encoded and there is no food. This comparison shows that the contribution of the additional units consists in making the activation patterns more similar in the two conditions, with the activation patterns becoming even more similar as the predator comes closer to the organism. In other words, the additional units allow the organism's nervous system to better filter out the information from food when the predator is present and, therefore, they might be considered as functionally equivalent to the modulatory influence of subcortical structures on frontal cortex in real organisms.

Organizational requirements for 'open-ended evolution'

Kepa Ruiz-Mirazo, Jon Umerez and Alvaro Moreno

Dept. Logic and Philosophy of Science, University of the Basque Country, Spain

Wednesday 1:30pm–3:00pm, Stripe Main Theatre, Open-Ended Evolution 2.

In this contribution we will review different conceptions of open-ended evolution and propose our own (Ruiz-Mirazo et al., 2008, *Biol. & Philos.*, 23, p.67). Then, we will consider what are the general conditions that would allow such an evolutionary process to take place, with a specific focus on the type of organization that the systems involved should have. It will be argued that a strong 'dynamic decoupling' is necessary, making possible the long-term maintenance of those systems, in which the individual (self-constructing) and collective (ecological and historical) spheres become deeply intertwined. Particular attention in the discussion will be given to bottleneck cases, like a hypothetical prebiotic 'RNA-world', which –according to our account– would not meet all the requirements. We will also reason why the evolution of a

prokaryote world is already open-ended, even if the transition to higher levels of complexity (eukaryotes, multicellular organisms, cognitive agents,) would imply further organizational bottlenecks and the fulfilment of additional conditions.

"Embryo": an autonomic co-operative service management framework

Fabrice Saffre and Mark Shackleton

BT Group plc

Thursday 10:30am–12:00pm, St Edburga Room 4, Self-* Computing.

In this paper, we present "Embryo", a fully decentralized service management framework inspired by morphogenesis and capable of installing components and modifying the topology of a peer-to-peer (P2P) interaction overlay network so as to meet the needs of the majority of all participating peers. Co-operation is an emergent property of the self-organisation process, which is underpinned by purely "selfish" decision-making based on incomplete information gathered through gossiping (local messaging). We provide a detailed description of the local reasoning loop governing the behavior of individual peers, as well as Monte Carlo simulation results that demonstrate the system's ability to converge to a stable state in which most peers have direct access to all the components they require via one of their first neighbors.

Self-organizing heterogeneous swarms designed through evolutionary methods

Hiroki Sayama

Department of Bioengineering, Binghamton University, State University of New York

Tuesday 3:30pm–5:00pm, Exam Hall Room 1, Swarm Systems.

Self-propelled particle swarm models are computational models of many particles capable of autonomous acceleration and local kinetic interaction. Their dynamics have been extensively studied in physics, theoretical biology, and computational science communities because of their useful implications for the understanding of collective behavior of various autonomous agents (e.g., bacteria, fish, birds, pedestrians) as well as their potential of application to practical problem solving.

Earlier studies mostly focused on homogeneous swarms, assuming that the same (or quantitatively similar) set of kinetic rules uniformly apply to all the particles. Some literature also assumed intra-specific variations among particles (such as in body size or velocity) but none of them systematically considered interactions between kinetically distinct types of particles. In real biological/ecological systems, however, there are cases where multiple distinct types of organisms interact to form non-trivial patterns. In a herd of animals, for example, males and females, or parents and offspring, occupy different

ecological positions and therefore adopt distinct behavioral rules. A unique formation may arise within the herd from interactions between those different types of organisms. Such self-organization of heterogeneous swarms could also be useful for engineering design purposes.

We therefore extend our scope to heterogeneous self-propelled particle swarm systems in which more than one type of particles can co-exist and interact with each other in the same space. Our model, "Swarm Chemistry" (Sayama, ECAL 2007, p.675, Springer), assumes self-propelled particles moving in a two-dimensional infinite continuous space. Each particle can perceive only the local center of mass and the average velocity vector of other particles within its local perception range, and change its velocity in discrete time steps according to kinetic rules similar to those of Reynolds' Boids (Reynolds, 1987, *Computer Graphics*, 21(4), p.25). Each particle is assigned with its own kinetic parameter settings that specify preferred speed, local perception range, and strength of each kinetic rule. Particles that share the same set of kinetic parameter settings are considered of the same type.

Using this model, we computationally studied what kind of patterns/motions could emerge out of the mixtures of multiple types of particles. In the first experiments testing the effects of two-type interactions, we found that heterogeneous particle swarms usually undergo spontaneous mutual segregation, often leading to the formation of multilayer structures. Driven by their own endogenous self-propulsion forces, the aggregates of particles may additionally show more dynamic macroscopic behaviors, including oscillation, rotation, and linear or even chaotic motion. Moreover, to explore the possibilities of more than two type interactions, we developed an interactive simulation tool with which a human experimenter can select, perturb, mix, and mutate heterogeneous swarms using an interactive evolutionary method. The second experiments using this interactive tool and human participants further revealed unexpected possibilities of more complex, mechanical, and/or even biological-looking structures and behaviors when several different types are mixed appropriately. Specifications of those patterns were indirectly and implicitly woven into a list of different kinetic parameter settings and their proportions, which would be hard to obtain through conventional design methods but can be obtained heuristically through evolutionary design methods. These results suggest a novel direction of understanding and engineering collective behavior of physical agents, such as distributed robotic systems.

The interactive simulation tool is implemented in Java and available at <http://bingweb.binghamton.edu/sayama/SwarmChemistry/>. Readers are invited to participate in the ongoing exploratory efforts of this project.

The limited utility of communication in simple organisms

Matthias Scheutz and Paul Schermerhorn

Indiana University

Wednesday 10:30am–12:00pm, St Edburga Room 4, From Communication to Language 1.

Many forms of communication have evolved in the animal kingdom for different purposes. In this paper we investigate the limits of communication for simple reactive organisms and show that communication has only limited benefits in biologically inspired foraging tasks and can even have detrimental effects in certain environments. Based on these results, we argue that simple agents with simple architectures need very special environmental conditions for communication to benefit them and thus to evolve.

Analysing honeybees' division of labour in broodcare by a multi-agent model

Thomas Schmickl and Karl Crailsheim

University of Graz, Dept. for Zoology

Thursday 3:30pm–5:00pm, Stripe Main Theatre, Ecological Modelling 2.

We describe a multi-agent model of a honeybee colony and show several applications of the model that simulate experiments that have been performed with real honeybees. Our special emphasis was on the decentralized, self-organized regulation of brood nursing, which we successfully simulated: We found that brood manipulations, food-deprivation experiments and colony-size manipulations can be explained by the mechanisms we implemented into our model described here. Our agents can perform various tasks (foraging, storing, nursing). The model is spatially resolved, and contains a designated broodnest area as well as a designated honey/nectar storage area. All bees (and larvae) consume nectar/honey at a task-specific rate, allowing us to track the flow of nectar through the colony. Several kinds of stimuli, which are important for division of labour, were modelled in detail: dances, contact stimuli and chemical signals.

Altruism amongst spatial predator-prey animats

Chris Scogings and Ken Hawick

Massey University

Tuesday 3:30pm–5:00pm, Stripe Lecture Room, Evolutionary Theory 3.

Understanding the emergence or suppression of altruism is an important step towards understanding real-life many-agent systems. We explore the relative survival traits of spatial animats in our predator-prey model and find some quantifiable emergent advantages of altruistic behaviour on the part of individual animats.

Measuring emergence via nonlinear Granger causality

Anil Seth

University of Sussex

Thursday 10:30am–12:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 1.

The concept of emergence is central to artificial life and complexity science, yet quantitative, intuitive, and easy-to-apply measures of emergence are surprisingly lacking. Here, I introduce a just such a measure, G-emergence, which operationalizes the notion that an emergent process is both dependent upon and autonomous from its underlying causal factors. G-emergence is based on a nonlinear time series analysis adapted from ‘Granger causality’ and it provides a measure not only of emergence but also of apparent ‘downward causation’. I illustrate the measure by application to a canonical example of emergence, an agent-based simulation of bird flocking, and I discuss its potential impact on perhaps the most challenging of all scientific problems involving emergence: consciousness.

Modeling the neural basis of cognitive integration and consciousness

Murray Shanahan and Dustin Connor

Imperial College London

Friday 10:30am–12:00pm, Stripe Main Theatre, Artificial Consciousness: From ALife to Mind.

This paper presents a number of models whose aim is to establish a computational basis for the hypothesis that conscious information processing in the brain is mediated by a mechanism of global broadcast. A possible role for this putative “global neuronal workspace” in achieving cognitive integration is mooted in the context of modular theories of mind, and an argument is advanced for its likely emergence within the sort of small-world brain network favoured by evolution. The paper concludes with some speculation on the relationship between life and consciousness as it could be.

Redefining robustness

Eric Silverman and Takashi Ikegami

University of Tokyo

Thursday 1:30pm–3:00pm, St Edburga Room 4, Fault Tolerance & Robustness.

The concept of robustness in the context of most work in Alife and complex systems implies that the results of a given model remain consistent despite unexpected variation. For example, homeostatic coupling between an animat and the environment is one possible simple form of robustness which we have demonstrated in simple simulation models (Ikegami et al., 2008, *BioSystems*, 91, p. 388), in which we defined robustness as a dynamic that

sustains variability. However, these models still were situated in simple virtual environments. We are currently using this simulation as a basis for developing a real-world robot experiment with “virtual” sound scape setups. Developing this link between real and simulated methodologies has led us to an examination of robustness in a broader sense.

We argue that robustness as commonly defined in Alife is no longer adequate for producing real insight into the functions of biological life. Robustness in one methodology or virtual world does not imply robustness in another, and likewise does not imply that we can develop a robust explanation of the behaviour of interest.

Robustness analysis as a concept is credited to Richard Levins who was the first to truly address pragmatic concerns important to biological modellers (1966, *Conceptual Issues in Evolutionary Biology*, p. 18, MIT Press). Levins argued that the construction of robust theorems from models involves studying similar but conceptually different models of the same phenomena and attempting to discern the common structures between them. Levins’ pragmatic concerns about modelling illuminate similar trade-offs made by modellers in the Alife community, leading some models to become mired in modelling for its own sake, creating simulations with little relation to the natural world (Silverman et al., 2008, forthcoming).

With this perspective in mind, a reevaluation of the concept of robustness within Alife is needed. While Alife can contribute to the search for common structures in biological systems which can drive behaviour, producing robust theorems about those behaviours also involves confirming that such structures are instantiated in the system of interest (Weisberg, 2005, *Phil. Sci.*, 73, p. 730). A unified framework under which to search for common structures is central to these concerns. Without a clear common relationship between conceptually related models, performing Levinsian robustness analysis becomes an impossible task.

Thus, we argue that finding robust theorems in Alife which demonstrate common structures is made difficult by the lack of common environments between models. A more critical analysis of what constitutes a useful environment for simulation and robotics is needed, and without such analysis, our concept of robustness falls short of Levins’ requirements for developing true robust theorems about the natural world. In essence, crafting a robust explanation of a behaviour using a model requires a robust demonstration of that behaviour through a suitable combination of modeling and experimentation. We contend that combining simulation and robotics with an approach using common methodologies and related environments as described above will allow us to develop a new definition of robustness in Alife.

Dancing robots: form, environment and context in human-robot interaction

Eric Silverman, Takashi Ikegami, Yuta Ogai and Jean-Julien Aucouturier

University of Tokyo

Wednesday 1:30pm–3:00pm, Exam Hall Room 1, Robots & Behaviour.

Most dancing robots to date have used patterns of preprogrammed motions or hard-coded interaction rules to produce this behaviour. In a departure from this approach, recently we used a form of embodied chaotic itinerancy (Ikegami, 2007, *J. Consc. Studies*, 14, p.111; Kaneko et al., 2003, *Chaos*, 13(3), p. 926) to generate motor movements for a robot in real time. We used the robot’s sensors to analyse audio input, processing it at regular time intervals to find the appropriate tempo, and used this information to send input pulses to a FitzHugh-Nagumo neural network model (Aucouturier et al., 2007, *Proc. of the 14th ICONIP*, Springer-Verlag). The dynamic properties of these neurons allow for interesting chaotic behaviour, as some inputs will produce entrained periodic states, while others produce chaotic or aperiodic responses.

The output neurons of this network drove the motors of our chosen robotics platform Miuro, a simple two-wheeled vehicle robot manufactured by ZMP (Tokyo, Japan). Despite the deterministic nature of the mechanisms driving Miuro, the resultant motions are heavily dependent on the music being played, and thus the robot displays complex transitions between quasiperiodic states of motion. The robot is able to demonstrate both synchronization and autonomy in its reactions to the music.

Currently we are organizing a workshop together with art students to invent a new type of environment for robots. In particular, we are aiming to generate “natural” sound environments in which mobile robots can generate complex and interesting dancing patterns. The workshop will also investigate the role of physical form in driving interaction between autonomous robots and human observers. The robot will interact with human observers in the same method as above, through sound, but our investigation of radical and novel physical forms for the robot will allow us to investigate new varieties of agent-environment couplings.

Ikegami and colleague Keiichiro Shibuya have started a series of sound installations, each of which uses ideas of Artificial Life and complex systems science to make unique soundscapes (<http://sacral.c.u-tokyo.ac.jp/index.php?Third%20Term%20Music>). Using robots, we can further develop this enlightening cooperation between science and art, which we think is a promising future avenue of artificial life study. Our collaboration with people from Art University will also encompass this issue, examining the concept of open-endedness in arti-

cial life studies and developing new methods for generating sound arts.

In both cases, our future work hinges on developing a new understanding of the robot-environment-human relationship. Through analysis of our work thus far, and discussion of the multitude of conceptual issues we have investigated with the artistic community, we will demonstrate new ways in which to examine these relationships. The interplay between form and function, between observer and performer, and between context and action will all influence the development of both the robot’s morphology and its control structure.

Evolving an agent-based model to probe behavioural rules in flocks of cowbirds

V. Anne Smith

University of St Andrews

Thursday 1:30pm–3:00pm, Stripe Main Theatre, Ecological Modelling 1.

Flocks of brown-headed cowbirds, *Molothrus ater*, self-organize social environments, which have strong impacts on social learning and behavior. To understand the rules by which such self-organization occurs, I develop an agent-based model of cowbird social association and evolve it to match observed patterns of association measured from real birds. The behavioral rules evolved in the model provide insight into the type of rules real birds use to organize their social environment. The evolved models successfully predicted both association patterns and additional related movement variables measured from a new flock of birds.

Evolutionary advantages of neuromodulated plasticity in dynamic, reward-based scenarios

Andrea Soltoggio¹, John Bullinaria¹, Claudio Mattiussi², Peter Dürri² and Dario Floreano²

¹University of Birmingham

²EPFL

Tuesday 1:30pm–3:00pm, Stripe Studio 1, Neural Systems.

Neuromodulation is considered a key factor for learning and memory in biological neural networks. Similarly, artificial neural networks could benefit from modulatory dynamics when facing certain types of learning problem. Here we test this hypothesis by introducing modulatory neurons to enhance or dampen neural plasticity at target neural nodes. Simulated evolution is employed to design neural control networks for T-maze learning problems, using both standard and modulatory neurons. The results show that experiments where modulatory neurons are enabled achieve better learning in comparison to those where modulatory neurons are disabled. We conclude that modulatory neurons evolve autonomously in the proposed learning tasks, allowing for increased learning and memory capabilities.

The virtual rocky shore — linking A-life with ecological and pedagogical research

Richard Stafford¹, Mark S. Davies² and Gray A. Williams³

¹University of Gloucestershire

²University of Sunderland

³University of Hong Kong

Tuesday 10:30am–12:00pm, St Edburga Room 4, Ecological Modelling Methodology.

The design and analysis of manipulative experiments is a key skill for undergraduate ecologists to learn. Despite this, students rarely conduct a true manipulative experiment with appropriate control treatments in the field, because of practical constraints such as time. Meaningful ecological experiments typically require months to obtain results, as well as constant maintenance and attention; and often require the use of large areas of undisturbed habitat. One solution to this problem is to allow students access to a virtual ecosystem in which they can conduct a possibly unlimited range of experiments which will quickly provide realistic data for subsequent statistical analysis and interpretation.

Currently, virtual ecosystems fall into two basic categories: those that are vast oversimplifications of real ecosystems, displaying simplistic and pre-programmed behaviours; and those that bear little resemblance to real ecosystems, made entirely of interacting digital organisms. While the latter are of more interest to the A-Life community, they are not user-friendly to biology students who are generally not computer literate beyond the basics of word processing, spreadsheets and internet technologies.

“The Virtual Rocky Shore”, is grounded in a variety of A-life techniques, including agent-based modelling, self-organisation, evolutionary algorithms and cellular automata. The present version of The Virtual Rocky Shore is based on the high intertidal region, a simple consumer / resource ecosystem consisting of grazing snails and a photosynthetic mat or biofilm of lichens, diatoms and bacteria. Although the system is simple, research underpinning the system’s models has demonstrated that these intertidal snails show many similar behavioural rules to those displayed by classic A-Life inspirations such as ants. The intertidal snails, for example, exhibit self-organisation as a result of trail following. The dynamics of the photosynthetic components of the virtual shore are also suitable to being modelled in space and time by use of cellular automata and computer-based optimisation processes including evolutionary algorithms.

Current development of The Virtual Rocky Shore using a user-friendly interface has already provided novel insights into the functioning and evolution of intertidal communities; with many of these insights backed up by empirically derived data from real shores and published in peer reviewed journals.

The first implementation of the Virtual Rocky Shore allows experiments to be designed and analysed in a matter of minutes, rather than the many months traditionally required, facilitating the active and potentially deep experiential learning of experimental design by students. The results obtained from the simulation are also similar, and result in comparable statistical analysis, to those obtained from experiments on real shores.

The full potential of The Virtual Rocky Shore, however, lies in its expansion to cover the mid and lower shore systems. The complexity found at these shore levels will allow many opportunities for further research at the interface of ecology and computer science, as well as the development of a wide range of potential experiments beyond simple grazer / biofilm interactions.

Can body language shape body image?

Luc Steels¹ and Michael Spranger²

¹University of Brussels (VUB), Sony CSL Paris

²Sony CSL Paris

Friday 10:30am–12:00pm, Stripe Lecture Room, Embodied Communication.

One of the central themes in autonomous robot research concerns the question how visual images of body movements by others can be interpreted and related to one’s own body movements and to language describing these body movements. The discovery of mirror neurons has shown that there are brain circuits which become active both in the perception and the re-enactment of bodily gestures, although it is so far unclear how these circuits can form, i.e. how neurons become mirror neurons. We report here further progress with our robot experiments in which a group of autonomous robots play language games in order to coordinate their visual, motor and cognitive body image. We have shown that the right kind of semi-otic dynamics can lead to the self-organisation of a successful communication system with which robots can ask each other to perform certain actions. The main contribution of this paper is to show that if the robot has the capacity to ‘imagine’ the behavior of his own body through self-simulation, he is better able to guess what action corresponds to a visual image produced by another robot and thus guess the meaning of an unknown word. This leads to a significant speed-up in the way individual agents are able to coordinate visual categories, motor behaviors and language.

Multiple functionalities of biochemical reaction networks

Marvin Steijaert, Anthony Liekens, Huub ten Eikelder and Peter Hilbers

Eindhoven University of Technology

Tuesday 10:30am–12:00pm, Stripe Main Theatre, Evolving Cell Signalling Networks *in silico* 1.

We consider a biological cell as a highly interconnected network of chemical reactions, which is constituted of a large number of semi-autonomous functional modules. Depending on the global state of the network, the separate functional modules may display qualitatively different behavior. As an example, we study a conceptual network of phosphorylation cycles, for which the steady-state concentration of an output compound depends on the concentrations of two input enzymes. We show that the input-output relation depends on the expression of the proteins in the network. Hence changes in protein expression, due to changes in the global regulatory network of the cell, can change the functionality of the module. In this specific example, that means switching between the functionality of various logical gates.

An artificial chemistry-based model of economies

Bas Straatman¹, Roger White² and Wolfgang Banzhaf³

¹Energy and Environmental Systems Group, ISEEE, University of Calgary

²Department of Geography, Memorial University, St. John's

³Department of Computer Science, Memorial University, St. John's

Thursday 3:30pm–5:00pm, St Edburga Room 4, Artificial Chemistries.

Economies can be modelled using Artificial Chemistry approaches. In this contribution we discuss the development of such a model starting from the well-known von Neumann's technology matrices. Skills and technologies that allow the transformation of raw materials into products are introduced in a form akin to chemical reactions. The dynamic flow of materials in such a system is simulated and connected through an agent-based market mechanism that assigns value to raw materials, labour, and products. Starting from a fixed set of raw materials, energy and labor, we observe the appearance of new products, the use of consumables and the general increase in complexity of such a system. Real evolutionary dynamics including waves of innovation can be demonstrated.

A minimal approach to modular assembly

Gregory Studer and Inman Harvey

University of Sussex

Wednesday 10:30am–12:00pm, Exam Hall Room 2, Amorphous & Soft Robotics.

In this work, we propose a new research direction into minimal assembling agents. Our goal is to use very simple, inflexible assembling units to form complex and flexible assemblies (or meta-modules), guided by global environmental signals. Instead of the focus in modular robotics and self-assembly on creating maximally flexible and programmable assembling units (Yim et al., 2007, IEEE Rob. Aut. Mag., p.43), we suggest a different, complementary approach in which assembled structures maintain or enhance the range of assembly behaviors atomic agents are capable of. Replacing the idea of complex autonomous modules which are able to build arbitrary structures, like cells building organisms, we are beginning to simulate robotic platforms which themselves have rather limited assembly behavior but stochastically form structures or meta-modules with more complex interactions, like proteins built from interactions of a few amino acids. This is inspired by stochastic assembly results in the real world at any scale (Krishnan et al., 2007, Proc. ASME IMECE, ASME.org), (Winfrey et al., 1998, Nature, 394, p.539), with an emphasis on understanding how function develops in these semi-controllable environments.

As a proof-of-concept and to gain intuition into how such units might look, we use a microbial genetic algorithm (MGA) to evolve the logic placed on simulated assembling agents. The agents are modeled as very simple units containing male (M) and female (F) assembly ports, as well as an input sensor, each of which may be in one of two states: enabled or not-enabled. Logic (in the form of Petri Nets) is generated by the MGA and identical copies placed on each agent, which are then allowed to assemble into chains in a well-mixed stochastic environment. Limited communication can occur between assembled agents' M and F ports. Instead of a traditional fitness function, however, where we might evaluate a logic as highly fit if it performs a particular assembly task, our fitness function rewards logics that maintain assembly behavior as the units assemble. In particular, we reward logic that maintains pairing behavior in response to a "start" signal. First, we enable the input sensor on all the agents, which may then form assembled structures including pairs which add to the fitness. If there are pairs, we then send a second "start" signal, and pairs may form of the pair structures themselves, and so on until no more pairing occurs. Higher level pairing was rewarded more than lower level pairing.

By limiting the complexity of the generated logics, and comparing the maximum fitness given these limits, there appears to be a lower complexity bound for our particular

assembling units to maintain their assembly behavior as they grow orders of magnitude in size. This demonstrates that our initial proposal of designing simple assembling units which build functional assemblies themselves is feasible, at least in some cases. The successful controllers generated are interesting in that they function similarly alone or when linked together in groups of any number of agents: the *behavior* scales. In future work, we hope to expand this result and demonstrate assembly controller designs which generate more complex assembly (and other) behavior as they grow. Our eventual goal is to discover designs for very simple, inflexible units which create programmable and controllable meta-modules in response to global environmental signals.

Homeodynamics in the game of life

Keisuke Suzuki and Takashi Ikegami

University of Tokyo

Wednesday 3:30pm–5:00pm, Exam Hall Room 1, Cellular Automata.

In this paper, we study the emergence of homeodynamics and adaptation on top of the game of life layer with the other cellular automata layer. Homeodynamics here is defined as a space-time dynamics that maintains the state-1 density in the life game space. A genetic algorithm is used to evolve the rules of the second layer cellular space to control the pattern of the game of life which is the first layer. We discovered that pattern generators emerged in this system to control the first layer pattern. One such generator creates cloud patterns for the initial high density environment, but keeps quiet for the initial lower density environment. Homeodynamics sustained by the pattern generators is discussed by comparing with Ashby's ultra-stability and the daisy world simulations.

How learning can guide evolution of communication

Reiji Suzuki and Takaya Arita

Graduate School of Information Science, Nagoya University

Wednesday 10:30am–12:00pm, St Edburga Room 4, From Communication to Language 1.

The Baldwin effect is known as a possible scenario of the genetic acquisition process of a learned trait without the Lamarckian mechanism. However, it is still controversial how learning can facilitate evolution in dynamically changing environments caused by internal factors. Our purpose is to clarify whether and how leaning can facilitate evolution in dynamic environments which arise from communicative interactions among individuals. We constructed a simple computational model for the evolution of communication ability and its phenotypic plasticity. In the model, the levels of adaptive communication, which correspond to the expected fitness value when the communication results in success, of signalling and receiving

processes are determined by different sets of traits under the assumption of the correlation between their fitness and the effects of epistatic interactions among traits. A communication is successful only when the levels of the signaller and the receiver are the same, and the individuals try to improve their communication levels through the learning process in which the values of plastic traits can be modified from their genetically determined values. The evolutionary experiments clearly showed that the Baldwin effect repeatedly occurred and facilitated the adaptive evolution of communication in this type of dynamic environments.

Evolution of complexity in RNA-like replicator systems

Nobuto Takeuchi and Paulien Hogeweg

Theoretical Biology and Bioinformatics, Utrecht University

Tuesday 10:30am–12:00pm, Stripe Lecture Room, Evolutionary Theory 1.

Background: The evolution of complexity is among the most important questions in biology. The evolution of complexity is often observed as the increase of genetic information or that of the organizational complexity of a system. It is well recognized that the formation of biological organization—be it of molecules or ecosystems—is ultimately instructed by the genetic information, whereas it is also true that the genetic information is functional only in the context of the organization. Therefore, to obtain a more complete picture of the evolution of complexity, we must study the evolution of both information and organization.

Results: Here we investigate the evolution of complexity in a simulated RNA-like replicator system. The simplicity of the system allows us to explicitly model the genotype-phenotype-interaction mapping of individual replicators, whereby we avoid preconceiving the functionality of genotypes (information) or the ecological organization of replicators in the model. In particular, the model assumes that interactions among replicators—to replicate or to be replicated—depend on their secondary structures and base-pair matching. The results showed that a population of replicators, originally consisting of one genotype, evolves to form a complex ecosystem of up to four species. During this diversification, the species evolve through acquiring unique genotypes with distinct ecological functionality. The analysis of this diversification reveals that parasitic replicators, which have been thought to destabilize the replicator's diversity, actually promote the evolution of diversity through generating a novel "niche" for catalytic replicators. This also makes the current replicator system extremely stable upon the evolution of parasites. The results also show that the stability of the system crucially depends on the spatial pattern formation of replicators. Finally, the evolutionary dynamics is shown to significantly depend on the mutation rate.

Conclusions: The interdependence of information and organization can play an important role for the evolution of complexity. Namely, the emergent ecosystem supplies a context in which a novel phenotype gains functionality. Realizing such a phenotype, novel genotypes can evolve, which, in turn, results in the evolution of more complex ecological organization. Hence, the evolutionary feedback between information and organization, and thereby the evolution of complexity.

[The original article is published as Takeuchi & Hogeweg, 2008, *Biology Direct*, 3:11]

How to know without having been there? — Investigating communication channels in the nectar collecting system of a honeybee colony

Ronald Thenius, Thomas Schmickl and Karl Crailsheim

Department for Zoology, Karl-Franzens-University Graz

Thursday 1:30pm–3:00pm, Stripe Main Theatre, Ecological Modelling 1.

Honeybees collect nectar from flowering plants in the environment to accommodate their energetic demands. In a honeybee colony a temporal caste, “foragers”, is collecting nectar. These foragers bring their harvest into the colony, where they unload their nectar loads to one or more specialised “storer bees”, another temporal caste in the colony, responsible for the next step of nectar processing. Natural selection has shaped the foraging-related processes of honeybees, like the communication between foragers via dances, in a way that a colony can react to changing environmental conditions in an adaptive way. To investigate this complex dynamic social system and the information and nectar channels we developed a multi-agent model of the nectar flow inside and outside of a honeybee colony. This model allows us to investigate the nectar collection process and nectar processing pathways on the colony level, as well as from the point of view of a single bee during the foraging trip and during the nectar proceeding inside the colony. The simulation includes near-natural environmental factors, like scattered nectar sources with variable distances between flowers and a near-natural model of the honeybee metabolism. The inside of the colony (the so called “dance floor”) was simulated as two one-dimensional transfer zones for foragers and storers, what enabled us to simulate the unloading-procedure in a highly abstract and defined manner. Our model predicts that a cohort of foragers, collecting nectar from a single nectar source, is able to detect changes in quality (e.g., the nectar flow) in other food sources they have never visited, by analysing side-effects of the nectar processing system of the colony: We identified two novel pathways of forager-to-forager communication by analysing the results predicted by our model. Foragers can gain information about changes in the nectar flow in the environment via two ways: Firstly, foragers can detect changes in their mean waiting time for unloadings, which

are performed by the storer bees. Secondly, the foragers can detect changes in the number of experienced multiple unloadings after returning from a foraging trip. The amount and quality of information available to the single forager about the environmental situation is increased, what enables the forager to modulate its individual decisions. The sum of this modulated forager decisions can lead to an optimisation of the foraging behaviour in an unsteady environment. This way, two distinct groups of foragers, that forage on different nectar sources and that never communicated directly, can share information via a third cohort of worker bees. We show that the communication channels within this noisy social network allow the colony to perform collective information processing. Simulation runs with fluctuations in the environmental nectar flow revealed, that the honeybee foraging system is even more adaptive (by exploiting the before mentioned communication channels) than was previously thought.

Towards language acquisition in autonomous robots

Vadim Tikhonoff¹, Angelo Cangelosi¹, Jun Tani² and Giorgio Metta³

¹University of Plymouth, UK

²RIKEN Brain Science Institute, Japan

³Italian Institute of Technology

Wednesday 1:30pm–3:00pm, St Edburga Room 4, From Communication to Language 2.

In this paper we present a novel cognitive robotic model to study language acquisition in autonomous robots through the grounding of words in sensorimotor representations. The aim of this new model is to extend previous work on language grounding in simulated cognitive agents to the new robotic platform iCub. The iCub is an emerging open platform for cognitive robotic research that will allow research groups to exploit a common hardware and software infrastructure in order to advance knowledge of natural and artificial cognitive systems. The language learning model is based on the use of artificial neural networks controllers. The model is based on a series of interconnected modules to gather and integrate visual and linguistic information for a language comprehension task. The model comprises of a vision module, a sound perception and feature extraction module and a language integration and recognition network. The vision acquisition module takes input from the robot’s cameras and applies approximation techniques for the purpose of detecting shapes, size and colour features of individual objects. The classification of a spoken word is based on the sequence of the most activated neurons of a self organizing map (SOM) with a 10 x 10 topological 2D grid. The SOM model has been trained on 112 English words and 544 syllable utterances both from two different speakers, for determining the ability of the system to distinguish between all words. The language integration module is based on a Recurrent

Neural Network with Parametric Biases (RNNPB). This network is particularly suitable for online learning of behaviour in robots. Two experiments were carried out to test the language learning model. The first consists in the recognition and classification of the speech signals as an imitation task without the integration of the vision module. This experiment has been based on the use of 20 words. Each training patterns (words) consists of a sequence of x/y coordinates of the SOM map. During the interaction phase of RNNPB training, the system learns to imitate the SOM word feature outputs pattern by predicting their next pattern. The network successfully learns to recognize spoken words with a final mean square error of the output nodes of 0.082. The second experiment consists of the integration of the vision and speech modules for learning and grounding of the names of objects. This experiment uses as input stimuli the combination of the features extracted from the visual module and the SOM output patterns. The output units predict the SOM sequence for the object name shown in the picture. The final square error of the output nodes was 0.003 over all the learning results. The model was able to categorize and name two objects which share some features (e.g. shape) but differ in other dimensions (e.g. colour). This preliminary work demonstrates the successful integration of a SOM network to classify spoken words with the RNNPB network capable of on-line learning and naming of visual objects. This model is being extended to include the learning of motor responses to be associated to the visual input of different objects and the capability to combine groups of words to describe visual scenes involving multiple objects. Although the current model primarily focuses on the naming aspects of language, our future plans include work on linguistic and communication capabilities.

Complex consumer preferences from rudiments of visual processing

Colin Tosh¹, Jens Krause¹ and Graeme Ruxton²

¹IICB, University of Leeds

²University of Glasgow

Thursday 10:30am–12:00pm, Exam Hall Room 1, Sensorimotor Connections 1.

Humans and many other consumer animals such as predators tend to select similar salient, pop-out, resource items from a visual scene. This observation hints at a common, simple mechanism underlying visual attention and driving the evolution of visual apparatus. Using simple artificial neural networks, we demonstrate that when information degrades early in a neural apparatus, and this degradation is compensated for in higher layers, many of the distinctive behaviours of consumer organisms emerge. These include preference for odd-looking resource items, resources that are spatially isolated, and resources that are on the edge of groups. We also observe evolution of a primitive visual fovea. While visual attention is struc-

turally and mechanistically complex in humans, the fundamental mechanisms driving the evolution of visual apparatus across different animal species may be simpler.

Self-organising synchronisation in a robotic swarm

Vito Trianni and Stefano Nolfi

Institute of Cognitive Sciences and Technology, CNR, Italy

Thursday 3:30pm–5:00pm, Exam Hall Room 2, Dynamical Systems Analysis 1.

Synchrony is a pervasive phenomenon: examples of synchronous behaviours can be found in the inanimate world as well as among living organisms (Strogatz, 2003, Sync, Hyperion Press). The synchronisation behaviours observed in Nature can be a powerful source of inspiration for the design of swarm robotic systems, where emphasis is given to the emergence of coherent group behaviours from simple individual rules. Much work takes inspiration from the self-organised behaviour of fireflies or similar chorusing behaviours. Here, we present a study of self-organising synchronisation in a group of robots based on minimal behavioural and communication strategies. We follow the basic idea that if an individual displays a periodic behaviour, it can synchronise with other (nearly) identical individuals by temporarily modifying its behaviour in order to reduce the phase difference with the rest of the group. In other robotic studies, synchronisation is based on the entrainment of the individual internal dynamics through some form of communication (see for instance Wishmann et al., Adaptive Behaviour, 14(2), p.113). In this paper, instead, we do not postulate the need of internal dynamics. Rather, the period and the phase of the individual behaviour are defined by the sensory-motor coordination of the robot, that is, by the dynamical interactions with the environment that result from the robot embodiment. We show that such dynamical interactions can be exploited for synchronisation, allowing to keep a minimal complexity of both the behavioural and the communication level. In order to define a robot controller able to exploit such dynamical agent-environment interactions, we use artificial evolution (Nolfi and Floreano, 2000, Evolutionary Robotics, MIT Press). The obtained results are analysed under a self-organising perspective, evaluating their scalability to large groups of robots.

The main contribution of this work consists in the analysis of the evolved behaviours, which is brought forth exploiting a dynamical systems approach: We introduce a dynamical system model of the robots interacting with the environment and among each other. This model offers us the possibility to deeply understand the evolved behaviours, both at the individual and collective level, by uncovering the mechanisms that artificial evolution synthesised to maximise the user-defined utility function. Moreover, we show how the developed model can be used to predict the ability of the evolved behaviour to efficiently

scale with the group size. We believe that such predictions are of fundamental importance to quickly select or discard obtained solutions without performing a time-demanding scalability analysis, as well as to engineer swarm robotic systems that present the desired properties.

Mammalian decisions

Pete Trimmer, John McNamara, Alasdair Houston, James Marshall, Liz Paul, Rafal Bogacz and Mike Mendl

University of Bristol

Wednesday 3:30pm–5:00pm, Stripe Studio 1, Learning, Action Selection, & Attention.

We study a simple scenario in which optimal decision-making requires a trade-off between speed and accuracy. Our analysis is set in the context of a mammal deciding whether to forage or invoke anti-predator action following an ambiguous cue. We assume that the brain has two systems which can make decisions. Thalamic decisions are fast but are less accurate than cortical decisions (which take longer).

We idealise the analysis by assuming that: 1) Thalamic decisions are made immediately, based upon a single piece of information (represented using Signal Detection Theory). 2) Cortical decisions are made by gathering information continuously until a confidence-threshold is reached (represented by applying a single-boundary version of the Sequential Probability Ratio Test to Brownian Motion with drift).

Following the analysis of each process in isolation, we examine how such decision systems might best be combined and used in the brain, discussing results in the context of information flow and the phylogeny of the mental architecture. We find that in some circumstances, if one system is weakened, the other system can largely compensate, thereby producing similar overall performance but with a different likelihood of response and decision timing.

The work may help to open areas of research on several topics, such as selective attention, mental stress and how learning affects decision-making (and vice-versa).

Self-assembly in physical autonomous robots: the evolutionary robotics approach

Elio Tuci¹, Christos Ampatzis¹, Vito Trianni², Anders Christensen³ and Marco Dorigo¹

¹IRIDIA-ULB

²ISTC-CNR

³DCTI-ISCTE

Friday 10:30am–12:00pm, Stripe Lecture Room, Embodied Communication.

This research work illustrates the details of a methodological approach to the design of homogeneous neuro-controllers for self-assembly in physical autonomous robots in which no assumptions are made concerning how agents allocate roles. Artificial evolution is used to set

the parameters of a dynamical neural network that when ported on two physical robots allows them to coordinate their actions in order to decide who will grip whom. The neural network directly controls the state of all the actuators. To the best of our knowledge, this work is the first example in which physical robots manage to self-assemble without relying on a priori injected morphological or behavioural heterogeneities. The results shed a light on the minimal requirements necessary to achieve self-assembly in autonomous robots.

Evolution, development and environment toward adaptation through phenotypic plasticity and exploitation of external information

Gunnar Tufte

Norwegian University of Science and Technology

Wednesday 1:30pm–3:00pm, Stripe Lecture Room, Development & Morphogenesis 2.

Biological organisms have an inherent ability to respond to environmental changes. The response can emerge as organisms that can develop into structural and behavioural different phenotypes. To achieve such properties in an artificial developmental setting external environmental information is included in the gene regulation of the developmental model. This implies interplay between evolution, development and the environment. An experimental approach is taken to investigate this interplay. The test case chosen is evolution of robustness to environmental fluctuations. Development models with and without environmental information included in the gene regulation are compared. Further, the developing organisms of the two models are exposed to environmental fluctuations for a more extensive investigation. The results indicate that including external information in the gene regulation can be favourable and exploitable, particularly for organisms developing in a dynamic environment.

Modelling the social coordination of behavior with public symbols

Ulas Türkmén and Radomir Zugic

University of Osnabrück

Wednesday 1:30pm–3:00pm, St Edburga Room 4, From Communication to Language 2.

In the recent decade, there has been an upsurge in models of language evolution. These models are based on the idea that multi-agent dynamics coupled with certain simple capabilities can lead to basic linguistic behavior. The general organization of these models, however, betrays an idea common to the cognitivist paradigm, namely that the linguistic symbols used by humans are labels for inner representations. The experiments in which these models are tested consist of two separate stages, which are (1) generating labels for sensory data and (2) using these labels for communication. The models also reflect

this division, and consequently contain separate modules (1) for creating representations and (2) for creating and transmitting vocalizations that correspond to these representations. An alternative to this division between perception/categorization and symbolic representation is using situated representations. The main idea behind situated representations is that the symbols human beings use in communication serve not only to carry meaning, but also to coordinate their embodied interaction. Inherent in this view is the social shaping of embodied activity through linguistic symbols. Traditionally, human cognition is divided into contrasting classes of high- and low-level processes, which then share inner representations among them in the mediation of thinking and action. Situated representations implicate that the relationship between high- and low-level processes, traditionally posited to be linear, rather is a dialectic one, where the social being of the agent affects, but is at the same time formed by, its embodied activity.

In order to model situated representations with multiple agents, the microworlds methodology used in the early days of artificial intelligence is ideal. In such a setup, it is possible to study in a task environment the necessary components of symbolic intelligence, namely social situatedness and ecological relevance. In the case of the model presented here, a robotic approach was chosen in order not to abstract away from the sensory-motor aspects of intelligence. Our approach to categorization is called “categorization without categories”, which refers to avoiding internal representations which do not have a linguistic function. In order to implement it, we have used an exemplar-based mechanism which relied on a simple similarity measure and the storing of whole sets of sensory data.

The principle idea of our experimental setup is providing the robot with an environment with a number of choices. In the first part of the experiment, the agents learn picking one of the choices, using solely the possibilities of the environment as cues. The second part starts with teaching one of the agents making a certain choice. Afterwards, a language game is played in which this agent instructs the other one. By developing a mechanism to instruct and to be guided by instructions, one agent can profit from the learning experience of the other and can directly choose the best out of a number of possible behavioural alternatives without having to go through the same learning experience again.

Using the RNA sequence-to-structure map for functional evolution of ribozyme catalyzed artificial metabolisms

Alexander Ullrich, Christoph Flamm and Lukas Endler

University of Vienna

Wednesday 1:30pm–3:00pm, Stripe Studio 1, Networks in Natural & Artificial Systems 2.

We introduce a novel genotype-phenotype mapping based on the relation between RNA sequence and its secondary structure for the use in evolutionary studies. The inspiration for this particular mapping emerged from the modeling of RNA enzymes within a simulation framework for the evolution of metabolic reaction networks. In our simulation we allow individuals, containing a genome and a metabolism, to evolve. The genome contains a number of RNA genes which then give rise to RNA enzymes acting on metabolites and thus shaping the metabolic network. Individuals are selected based on measures of this network and new individuals with mutated genomes are created. The use of our mapping allows not only for a more realistic study of the evolution of the entire system, but also enables us to observe the behavior of our enzymes itself and therefore possibly gain some insights about the evolution of catalytic molecules in general.

Enzymes typically have an active site where only few amino acids or bases determine its catalytic function and the remaining structure has mostly stabilization function. Accordingly, we extract structural and sequence information only from a restricted part of the fold. We decided to focus on the longest loop of the folded RNA. The idea for mapping the extracted information to a specific chemical reaction was encouraged by the fact that many enzymes catalyze a reaction by stabilizing its transition state. Recent work on hairpin ribozymes and other catalytic RNA support that as a common strategy for RNA enzymes. Given the definition of Fujita’s imaginary transition structures (ITS), we developed a unique index for all possible pericyclic chemical reactions, describing the constitution of the reaction’s transition state. Every RNA molecule is assigned such an reaction ID based on the information from its fold. The length of the longest loop specifies the number of involved atoms and the sequence within the loop determines the atom types. The bond types are derived from structural characteristics of the loop, such as the length and position of contained stems. Thus, a mapping from RNA sequence (genotype) to a chemical reaction (phenotype) is produced.

For many years it is known that neutral mutations have a considerable influence on the evolution in molecular systems. The folding of RNA sequences to secondary structures with its many-to-one property represents a mapping entailing considerable redundancy. Various extensive studies concerning RNA folding in the context of neutral theory yielded insights about properties of the structure

space and the mapping itself. We intend to get a better understanding of some of these properties and especially of the evolution of RNA-molecules as well as their effect on the evolution of the entire molecular system.

Besides using the mapping in several simulation runs which yielded realistic metabolic networks and connectivities, we performed several statistical tests commonly used in neutral theory, such as the number of visited phenotypes and the average discovery rate during a random neutral walk. We compared it with results of approaches using cellular automata, random boolean networks and other mappings based on RNA folding. It exceeds all non-RNA mappings in extent and connectivity of the underlying neutral network. Further, it has a significantly higher evolvability and innovation rate than the rest. Especially interesting is the highly innovative starting phase in RNA-based mappings.

Modelling resilience of agro-ecosystems: the co-evolution of regimes

Dirk van Apeldoorn, Kasper Kok, Marthijn Sonneveld and Tom Veldkamp

Wageningen University, Land Dynamics

Tuesday 3:30pm–5:00pm, Stripe Main Theatre, Sustainability & Homeostasis.

Understanding how natural resources co-evolve with management practices is fundamental to all questions concerning their sustainable management. In the search for new pathways of sustainable development we seek to understand how comparable initial conditions in agro-ecosystems lead to fundamentally different resource conditions. Agro-ecosystems are characterized by a high diversity, path dependence, self organization, cross-scale interactions and non-linear feedbacks. The resilience perspective claims it offers insight in these complex systems attributes. For exploring these dynamics we used a summary model approach. A summary model approach aggregates processes based on detailed knowledge, while allowing for integration of multiple scales and the identification of thresholds.

We analysed an area characterised by dairy farming in the Netherlands, within which two different farming systems can be found. One system is characterised by the modernisation paradigm with benefits of scale, intensification and specialisation, while the other system is characterised by low external input and less intensive farming. For a long time it was thought that the latter, alternative, system represented the laggards of the adoption rate of innovations. Scientific attention to the development of this region was attracted by the persistence of the alternative farming system, the increasingly valued effect on the landscape by these farms, and their mismatch with “modern” environmental regulations.

We developed a summary model that integrates the soil, feed, and animal compartments of the farming system.

The model enables us to simulate the effects of farm management decisions on the key natural resources soil carbon and nitrogen. When characteristic farm management of both the intensive and the alternative farming systems are used as model input, the systems evolve to different regimes. I.e., the different management systems lead to natural resources that respond differently to external drivers.

The two management systems evolved to different natural resource regimes and traverses to the other regime are slow, highly non-linear and involve large costs. Policies aiming at social-ecological regime change are currently not aware of these dynamics. We are now discussing with farmers strategies for regime change employing the non-linear dynamics in their system. With researches in the area we are formulating new hypothesis on system functioning. Meanwhile at the governance level we show how the environmental regulations, although initially successful, will actually drive the systems to an unwanted regime.

Simulated trust: towards robust social learning

Dieter Vanderelst, Rene Ahn and Emilia Barakova

Eindhoven University of Technology

Tuesday 1:30pm–3:00pm, Exam Hall Room 2, Culture & Social Evolution 2.

Social learning is a potentially powerful learning mechanism to use in artificial multi-agent systems. However, findings about how animals use social learning show that it is also possibly detrimental. In social learning agents act based on second-hand information that might not be trustworthy. This can lead to the spread of maladaptive behavior. Animals employ a number of strategies to selectively use social learning only when appropriate. This suggests that artificial agents could learn more successfully if they are able to strike the appropriate balance between social and individual learning.

In this paper, we propose a simple mechanism that regulates the extent to which agents rely on social learning. Our agents can vary the amount of Trust they have in others. The Trust is not determined by the performance of others but depends on the agents’ own rating of the demonstrations. The effectiveness of this mechanism is examined through a series of simulations. We first show that there are various circumstances under which the performance of multi-agents systems is indeed seriously hampered when agents rely on indiscriminate social learning. We then investigate how agents that incorporate the proposed Trust-mechanism fare under the same circumstances. Our simulations indicate that the mechanism is quite effective in regulating the extent to which agents rely on social learning. It causes considerable improvements in the learning rate, and can, under some circumstances, even improve the eventual performance of the agents. Finally, some possible extensions of the proposed mechanism are being discussed.

Diet learning and the evolution of cultural inheritance

Daniel van der Post, Bas Ursem and Paulien Hogeweg

Utrecht University

Tuesday 10:30am–12:00pm, Exam Hall Room 2, Culture & Social Evolution 1.

Social learning can give rise to cultural inheritance which forms an additional inheritance system next to genetic inheritance. Its evolution can be seen as a major transition in evolution. Using a spatial individual-based model we study the evolution of social learning and therewith the emergence of culture. We focus on diet learning in group foragers as a context in which cultural inheritance could have evolved. We model a rich environment in which foragers learn what to eat and focus on how environmental complexity can structure behavioural opportunities and lead to self-organizing processes. Our results show that social influences on learning arise as obligate side-effects of grouping. In patchy environments this can give rise to both traditional inheritance and cumulative cultural processes. Cultural phenomena therefore arise “for free” as soon as individuals learn by trial-and-error in groups. This shows the role of self-organizing processes in generating novelty in evolution. These self-organized processes set the context in which more sophisticated forms of social learning can evolve. By including copying behaviour in our model, we studied its adaptive influence and evolution. Results show that copying is not a fixed strategy and its adaptive value depends on resource distributions in the environment. On the one hand copying leads to collective problem solving within lifetimes. On the other hand it generates cumulative cultural diet optimization over lifetimes. Preliminary results of evolutionary simulations show that copying behaviour evolves because it allows for these adaptive processes. However copying also tends to reduce variation in groups and thus reduces the efficacy of natural selection. We conclude that self-organization plays a large role in the transition to cultural inheritance by means of generating obligate social influences on learning as side-effects of grouping. Moreover, this self-organized baseline affects the evolution of cognitively more sophisticated forms of social learning.

A study of GasNet spatial embedding in a delayed-response task

Patricia A. Vargas, Ezequiel Di Paolo and Phil Husbands

University of Sussex

Tuesday 3:30pm–5:00pm, St Edburga Room 4, Spatial Organisation in Neural Systems.

GasNet artificial neural networks can be used as complex neurocontrollers involving virtual chemical neuromodulation as well as synaptic interaction. The aim of this paper is to further explore the role of space in GasNet models on a delayed-response robot task. Comparative results demonstrate that the use of spatial constraints is not a prerequisite for a good performance of the original model in terms of speed of evolution.

The connectivity of NK landscapes’ basins: a network analysis

Sebastien Verel¹, Gabriela Ochoa² and Marco Tomassini³

¹University of Nice Sophia-Antipolis

²University of Nottingham

³University of Lausanne

Wednesday 10:30am–12:00pm, Stripe Studio 1, Networks in Natural & Artificial Systems 1.

We propose a network characterization of combinatorial fitness landscapes by adapting the notion of ‘inherent networks’ proposed for energy surfaces (Doye, 2002). We use the well-known family of NK landscapes as an example. In our case the inherent network is the graph where the vertices represent the local maxima in the landscape, and the edges account for the transition probabilities between their corresponding basins of attraction. We exhaustively extracted such networks on representative small NK landscape instances, and performed a statistical characterization of their properties. We found that most of these network properties can be related to the search difficulty on the underlying NK landscapes with varying values of K .

Adaptive growth processes: a model inspired by Pask’s ear

Nathaniel Virgo and Inman Harvey

University of Sussex

Wednesday 1:30pm–3:00pm, Stripe Lecture Room, Development & Morphogenesis 2.

We introduce the notion of an “adaptive growth process” in order to explain an experimental result from the 1950s in which a complex mechanism capable of distinguishing between two sounds emerges from a homogeneous chemical solution. We present a very simple computational model which exhibits an adaptive growth process. Adaptive growth processes could have practical applications in adaptive control systems and may also play a role in biological development.

Reaction-diffusion spots as a model for autopoiesis

Nathaniel Virgo and Inman Harvey

University of Sussex

Tuesday 1:30pm–3:00pm, Exam Hall Room 1, Spatial Organisation 2.

We analyse pattern formation in reaction-diffusion systems from an autopoietic point of view, emphasising the commonalities between living organisms and a certain class of so-called dissipative structures, namely those (such as spot patterns or hurricanes) in which there are more-or-less clearly defined unities, or individuals, which arise from the system's dynamics.

Previous authors have used cellular automata as a basis for studying the emergence of autonomous agent-like structures, but the continuous nature of reaction-diffusion systems gives them a substantial advantage over discrete cellular automata as it enables systems to be perturbed by an arbitrarily small amount. Since reaction-diffusion systems are simulations of physical/chemical systems the resulting model agents must obey the relevant thermodynamic constraints, an aspect of living systems that has generated a lot of recent discussion in the autopoietic literature.

The Gray-Scott model is perhaps the simplest reaction-diffusion system that can create complex patterns; it models a single type of autocatalyst feeding on a 'food' chemical that is continually added to the system; both are able to diffuse on a two-dimensional surface. One of the patterns that can be formed consists of blurred but individuated "spots" of autocatalyst separated by regions in which the autocatalyst is absent. We take a single spot as the basis for our model agent.

With the autopoietic description in mind we perform three experiments. Firstly, we put these spots into situations where there is a spatial gradient of the food molecule and find that they tend to move along it, usually away from areas where the level of food is too low for their survival. The relationship between constitution and behaviour is fundamental to the autopoietic theory, and this result opens the possibility of studying the interface between the two empirically.

Secondly we vary the rules of the system, allowing a different set of chemical reactions, which can result in agents with a more complex anatomy than just a single spot, and even a very limited form of heredity.

Finally we find that individuated spots are very likely to arise when there is a negative feedback between the whole system's activity and its overall supply of food. This situation is common in natural systems, and our result suggests a direction for further research into the conditions under which individuated unities are likely to occur in general.

Can we rebuild a cell? Bryopsis — an experimental model!

Alexandru Vladimirescu

National Institute for Research and Development for Microbiology and Immunology, Bucharest, Romania

Thursday 10:30am–12:00pm, Stripe Main Theatre, Synthetic Biology.

Generally speaking a violent mechanical treatment applied on a living cell or a unicellular organism destroys its structural and functional integrity and leads to death; any tentative to re-build the destroyed cell from the remaining cellular fractions turned into failure. There are few exceptions; one of them is the coenocytic seaweed *Bryopsis*. This widespread seaweed has the ability to be restored from its cellular fractions beginning with the spontaneous aggregation of cytoplasm and organelles (in the presence of seawater) and continuing with the formation of a temporary polysaccharide membrane surrounding the cytoplasm aggregates, formation of a lipid-based membrane and restoration of the cell wall; the result is a cell that has the ability to grow and form a new *Bryopsis* thallus. In the experimental approach to re-build the coenocytic alga *Bryopsis*, in early events after the mechanical destruction, cytoplasm and organelles can be mixed with biological (*E. coli* living cells carrying the *gfp* gene for the Green Fluorescent Protein) or inorganic particles (Fe_3O_4 nanoparticles for Ferrofluids) so that the new particles would be incorporated in the re-constructed *Bryopsis* protoplast/cell. The behavior of *Bryopsis* protoplasm and particles will be presented and discussed using Optical and Transmission Electron Microscopy investigations. The preliminary experimental data support the belief that the reconstruction of a designed *Bryopsis* cell, including artificial or natural foreign elements, will become a reality in the nearest future.

The spatiality of swarms — quantitative analysis of dynamic interaction networks

Sebastian von Mammen and Christian Jacob

University of Calgary

Tuesday 3:30pm–5:00pm, Exam Hall Room 1, Swarm Systems.

Many mathematical models, which try to capture emergent phenomena, are based on state transitions that depend on neighborhood relationships. Cellular Automata (CA) and Random Boolean Networks (RBN) are examples of such models, where connectivity patterns determine the flow of signals among interconnected units. Whereas neighborhoods in CA and RBNs remain static, the focus of our investigations are artificial swarms that act in three-dimensional space, where neighborhood relationships among the swarming agents change over time. In fact, it is through the dynamically changing neighbors that determine a swarm system's overall behavior. In this pa-

per we explore neighborhood dynamics of swarms and ask the question how each agents' time-dependent perception of its neighbors relates to specific flocking formations. We give examples of 'neighborhood functions' for choreographed swarming behaviors, such as line and figure-eight formations. We also evolve control parameters for swarm agents such that they approximate specific neighborhood functions that trigger switching and oscillations.

Can individual selection favour significant higher-level selection?

Richard Watson, Rob Mills, Alexandra Penn and Simon Powers

Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Wednesday 1:30pm–3:00pm, Exam Hall Room 2, Major Evolutionary Transitions 1.

How do new evolutionary units, supporting higher levels of functional organisation, arise from existing evolutionary units? The adaptive transformation of co-adapted species into new units, as in the major evolutionary transitions, is centrally implicated in the evolution of complexity but has proved very problematic for current evolutionary theory and understandably elusive in ALife. We investigate the evolution of new evolutionary units via individual adaptation in a multi-species ecosystem by modelling symbiotic associations that cause interaction probabilities to deviate from a freely mixed condition. It is well known that assortative grouping supports group selection in a well-defined sense, thus it is no surprise that enabling such associations will introduce some group selection effects. However, what form will this take when the control of such grouping is under individual adaptation? We tackle this by comparing the ecosystem attractors of the initial freely mixed system to those of the same system given the evolved association probabilities. In general we find that self-interested adaptation of associations tends to only reinforce species combinations that were already stable before the associations — which seems rather uninteresting. However, if the species densities of the ecosystem are occasionally perturbed whilst associations are developing this causes the system to visit different attractors and allows multiple, possibly incompatible, associations to be selected for in different contexts. Under these conditions, even when the attractors of the final system already existed as attractors in the freely-mixed system, competition between different combinations of species enlarges basins of attraction that lead to fit combinations at the expense of those that lead to less fit combinations. Thus, after the associations have evolved, a fit combination of species may be favoured in the niche that is constructed by the action of its association preferences, even if each species involved would be individually unfit if the system were freely mixed.

These findings show that evolved higher-level selection can have significant effects even when the new units result from the self-interest of the constituent sub-units. They also suggest that evolved complexes observed naturally may appear to be merely the result of individual selection because they are supported by individual self-interest, but in fact the reason that this complex persists and not some other is due to competition among species combinations.

Nonetheless, in small systems these mechanisms do not produce higher levels of complexity than those which occurred without evolved associations because the configurations that result were already visited in the initial freely-mixed system. However, we find that in large complex ecosystems with many local attractors, evolved associations naturally generalise over the relatively few attractors that are visited, enlarging attractors for fit species combinations even before they are visited. An idealisation of these processes has been shown to be far superior to conventional evolutionary algorithms on a fairly general class of difficult optimisation problems. This self-modification of ecosystem attractors therefore illustrates a mechanism that produces high-fitness biological complexes despite the fact that their evolution would seem highly implausible given the very small size of the basin of attraction that leads to this configuration under selection on the original units.

Optimal robot recharging strategies for time discounted labour

Jens Wawerla and Richard Vaughan

Simon Fraser University

Tuesday 10:30am–12:00pm, Stripe Studio 1, Autonomous Robots.

Energy is defined as the potential to perform work: every system that does some work must possess the required energy in advance. An interesting class of systems, including animals and recharging robots, has to actively choose when to obtain energy and when to dissipate energy in work. If working and collecting energy are mutually exclusive, as is common in many animal and robot scenarios, the system faces an essential two-phase action selection problem: (i) how much energy should be accumulated before starting work; (ii) at what remaining energy level should the agent switch back to feeding/recharging? This paper presents an abstract general model of a energy-managing agent that does time-discounted work. Analyzing the model, we find solutions to both questions that optimize the value of the work done. This result is validated empirically by simulated robot experiments that agree closely with the model.

Hierarchical components and entity-based modelling in artificial life

Matt Webster and Grant Malcolm

University of Liverpool

Tuesday 1:30pm–3:00pm, St Edburga Room 4, Methodologies.

We present notions of entity and entity-based model that are applicable to artificial life. We illustrate these notions by giving an abstraction of Langton's loops: loop-like structures that reproduce in a cellular automaton. Our abstraction takes as entities persistent configurations of the cellular automaton, and shows how these entities may be combined to form more complex entities. The resulting entity-based model of Langton's loops describes the functionality and interrelationships of these components, abstracting from their actual realisation in a cellular automaton. As well as providing a basis for the study of ecologies of interacting entities in artificial life, our approach provides a useful intermediate level of abstraction that can relate top-down and bottom-up approaches to the study of life-like systems.

Investigating emergence by coarse graining elementary cellular automata

Andrew Weeks, Fiona Polack and Susan Stepney

Computer Science, University of York

Thursday 10:30am–12:00pm, Stripe Lecture Room, Information in Complex Systems & ALife 1.

We extend coarse graining of cellular automata to investigate aspects of emergence. From the total coarse graining approach introduced by Israeli and Goldenfeld (Coarse-graining of cellular automata, emergence, and the predictability of complex systems, *Phys. Rev. E*, 2006) we devise partial coarse graining, and show qualitative differences in the results of total and partial coarse graining. Mutual information is used to show objectively how coarse grainings are related to the identification of emergent structure. We show that some valid coarse grainings have high mutual information, and are thus good at identifying and predicting emergent structures. We also show that the mapping from lower to emergent levels crucially affects the quality emergence.

Autopoiesis, enactivism, and the extended mind

Michael Wheeler

Department of Philosophy, University of Stirling

Thursday 10:30am–12:00pm, Stripe Studio 1, Philosophical Issues 1.

Two strands of recent embodied theorizing about cognition that are commonly held to be in harmony are actually in tension. This tension arises, in part, from the different ways in which the two positions in question — the extended mind hypothesis (EMH) and enactivism — conceive of the relationship between life and mind.

The history of enactivist ideas, plus a recent presentation of the view by one of its architects, Evan Thompson, suggest that the autopoietic theory of Maturana and Varela is a non-negotiable component of enactivism. An autopoietic system is a self-organizing autonomous system that, through its own endogenous activity, produces and maintains a physical boundary that distinguishes that system as a material unity in the space in which it exists. According to Maturana and Varela: (i) any living system is an autopoietic system; (ii) any autopoietic system is a living system; (iii) cognition is viability-maintaining activity in a domain of interactions defined by an autopoietic system's organization; (iv) enaction is the process by which significance is brought forth through the viable structural coupling of an autopoietic system with its environment.

Against this background, one striking claim made by autopoiesis theorists is that living is cognition. A natural way of hearing this claim (one that Maturana and Varela themselves often seem to recommend) is as asserting that the living system is identical with the cognitive system. If we add to this identity assertion the independently plausible thought that the living system (the organism) will be bounded by its skin, the implication is that, for the enactivist, the cognitive system is bounded by the skin. However, according to EMH it is possible for things and processes located beyond the skin sometimes to count as the proper parts of a cognitive system, which means that the boundary of the cognitive system may sometimes extend beyond the skin. So, it seems, the enactivist cannot endorse EMH.

The enactivist might reply that I have painted an impoverished picture of the relationship between life and cognition, as she understands it. Varela, in later work, depicted cognition as a process of sense-making. Di Paolo has argued that to explain sense-making, raw autopoiesis (autopoiesis as described above) must be supplemented with a capacity for adaptivity, itself established on the basis of an autopoietic organization. As I understand him, Di Paolo holds that being a raw autopoietic system is necessary but not sufficient for being a cognitive system (for realizing sense-making). However, since being a raw autopoietic system remains necessary and sufficient for being a living system, being a cognitive system remains sufficient for being a living system. So if our reconstructed enactivist did try to sanction EMH, she would be claiming (a) that an extended cognitive system is an autopoietic system, and (b) that an extended cognitive system is itself (it does not merely contain) a living system. Claim (a) is debatable and claim (b) violates our highly plausible thought that living systems don't extend. The enactivist still cannot endorse EMH. Enacted minds are not extended minds.

Understanding robustness in random Boolean networks

Kai Willadsen¹, Jochen Triesch¹ and Janet Wiles²

¹Frankfurt Institute for Advanced Studies

²University of Queensland

Friday 10:30am–12:00pm, Stripe Studio 1, Dynamical Systems Analysis 2.

Long used as a framework for abstract modelling of genetic regulatory networks, the Random Boolean Network model possesses interesting robustness-related behaviour. We introduce coherency, a new measure of robustness based on a system's state space, and defined as the probability of switching between attraction basins due to perturbation. We show that this measure has both upper and random-case bounds, and that these bounds are based on the size of individual attractor basins within the system. A mechanism for calculating these bounds is introduced, and the bounds are then used to define structural coherency, a measure of robustness attributable to system structure. Using these measures, we show that the decrease in coherency that occurs in the Random Boolean Network as its connectivity increases is related to a loss of structure in the system's state space.

Environmental regulation by higher level selection in a simulated network of microbial ecosystems

Hywel Williams and Tim Lenton

University of East Anglia

Wednesday 10:30am–12:00pm, Exam Hall Room 1, Models for Gaia Theory.

The Earth possesses a number of regulatory feedback mechanisms involving life. In the absence of a population of competing biospheres it has proved hard to find a robust evolutionary mechanism that would generate environmental regulation. It has been suggested that regulation must require altruistic environmental alterations by organisms and would therefore be evolutionarily unstable. This need not be the case if organisms alter the environment as a selectively neutral by-product of their metabolism, as in the majority of biogeochemical reactions, but the question then arises: why should the combined by-product effects of the biota have a stabilising, rather than destabilising, influence on the environment? In certain conditions selection acting above the level of the individual can be an effective adaptive force. Here we present an evolutionary simulation model in which environmental regulation involving higher level selection robustly emerges in a network of interconnected microbial ecosystems. The Flask model simulates an evolving microbial community suspended in flasks of liquid with prescribed inputs of nutrients. The system is seeded with a clonal population of 'microbes' that are subject to mutation on genetic loci that

determine their nutrient uptake patterns, release patterns, and their effects on, and response to, other environmental variables. Nutrient recycling loops robustly emerge from local adaptation, but populations are vulnerable to crashes caused by 'rebel' mutants which push abiotic conditions away from habitability. In previous work we have demonstrated a community-level response to artificial selection at the level of a single flask. Here we show that spatial structure in a network of interconnected flasks creates conditions for a limited form of higher level natural selection to act on the collective environment-altering properties of local communities. Local communities that improve their environmental conditions achieve larger populations and are better colonisers of available space, while local communities that degrade their environment shrink and become susceptible to invasion. The spread of environment-improving communities alters the global environment towards the optimal conditions for growth and tends to regulate against external perturbations. This work suggests a new mechanism for environmental regulation that is consistent with evolutionary theory. Interestingly, the system appears to be ultrastable — a term originally introduced in cybernetics by W. Ross Ashby — in that its stability requires the maintenance of key variables within bounds. We speculate that the biosphere may also be ultrastable.

Evolving referential communication in embodied dynamical agents

Paul Williams, Randall Beer and Michael Gasser

Indiana University

Friday 10:30am–12:00pm, Stripe Lecture Room, Embodied Communication.

This paper presents results from three experiments which investigate the evolution of referential communication in embodied dynamical agents. Agents, interacting with only simple sensors and motors, are evolved in a task which requires one agent to communicate the locations of spatially distant targets to another agent. The results from these experiments demonstrate a variety of successful communication strategies, providing a first step towards understanding the emergence of referential communication in terms of coordinated behavioral interactions.

Timing of critical periods in development

Oliver Winks and Luc Berthouze

University of Sussex

Wednesday 10:30am–12:00pm, Stripe Lecture Room, Development & Morphogenesis 1.

Critical periods are specific periods in the development of a living organism during which there is an increased sensitivity to external perturbations. Such perturbations result in a developmental trajectory significantly different from what is considered the norm. This study is concerned with the question of whether the presence and timing of a critical period can be predicted from the developmental profile without perturbation. We frame this question in the context of Waddington's epigenetic landscapes (Waddington, 1943, *Am. Midl. Nat.*, 30, p. 811) and put forth the hypothesis that bifurcations are more likely to take place when the system is undergoing rapid developmental changes, i.e., critical periods will occur when the rate of change is greatest.

To test this hypothesis, we developed a simulation of the early stages of embryonic development, specifically, the development of cellular structures and cell differentiation. The model was formed of two components, the genetic component and the cellular component. The genetic component simulated gene expression and genetic regulation where artificial transcription factors and proteins were synthesised that excited and inhibited genes. The cellular component simulated several cell functions that were controlled by proteins and that made it possible to grow cellular structures composed of cells of different types. The model existed in 3D space allowing for complex 3D structures to emerge over a fixed time period through the dynamics of differential gene expression and cellular functions. The amount of energy available to the system was kept constant so that energy consumption was a limiting factor that stopped the physically impossible scenario of infinite growth. Artificial evolution was used to create genomes capable of growing into organisms of a specific structure, and genome size was varied to allow for organisms to develop differently into their final structure. The presence of critical periods was tested by systematically depriving each developing organism of varying amounts of an extra-cellular signalling protein at different times between runs and by locating variations in fitness of the organism after development of more than two standard deviations. All organisms were found to exhibit critical periods, and the timing of these critical periods was found to correlate strongly with greatest rates of change in the energy profile of the organism developing without perturbation. Interestingly, these periods of change were linked to discontinuities in the consumption profiles of various signalling proteins, an observation that is consistent with a recent finding in developmental biology that morphogenetic variables are not monotonous in time (Cherdantsev et al.,

2005, *Ontogenez*, 36(3), p. 211). The ability to predict the critical periods of a developing system has broad implications not only in the clinical domain – in particular, the study of teratogens (Wilson, 1973, *Environment and Birth Defects*, Academic Press) – but also in the study of artificial developmental and adaptive systems.

Time and space in neuronal networks: the effects of spatial organisation on network behaviour

Stephen Womble and Netta Cohen

University of Leeds

Tuesday 3:30pm–5:00pm, St Edburga Room 4, Spatial Organisation in Neural Systems.

Complex networks in nature are often characterised by nontrivial spatial organisation. In such networks, the topology and spatial organisation are often tightly coupled to the activity on the network. In particular, spatial topologies frequently influence time delays through the network. In many real world networks, from transport systems to brains, increased distance (or wiring length) increases time delays in physical or information flow. Here we investigate this role of space on neural network dynamics.

Time delays are well recognised to influence dynamics in neural networks. It would not be surprising, therefore, if neural networks exploited spatial organisation to tune their time delays and obtain desired patterns of activity. However when we model such networks, either to study or to apply their properties to engineering problems, it is important to distinguish between two cases: either the effects of spatial organisation in the network are merely temporal, in which case the complex and computationally expensive spatial organisation can be succinctly abstracted out, or — more interestingly — spatial organisation may lead to behaviour that cannot be captured solely by temporal effects.

We investigated this question in biologically constrained spiking neural networks operating close to a critical bifurcation between stationary behaviour and population-wide oscillatory behaviour. In particular, we used a network of leaky-integrate-and-fire neurons and conductance based synapses with alpha-function kernels, operating in the so-called balanced regime, where each neuron receives similar excitatory and inhibitory drive. Specifically, we compare the transition from stationary to oscillatory network behaviour as a function of spatial organisation. The two modes of behaviour are found in biological neuronal networks and are important for qualitatively different information coding schemes (e.g. rate coding in the stationary regime and phase or temporal coding in the oscillatory regime).

First, we demonstrate that time delays play an important role in this type of network. We found that the region of parameter space for which stationary behaviour is found increases with shorter time delays. We compared this performance to that of an otherwise identical

network but with explicit spatial organisation. In our networks, spatial organisation was implemented by clustering inhibitory cells in the midst of a homogeneous population of excitatory cells, in keeping with dendritic length scales of inhibitory cells in some cortical areas. In addition to introducing a spatial patterning, the central clustering of inhibitory cells also reduces the mean inhibitory-to-excitatory time delays in the network. Thus, a purely temporal effect of clustering would yield similar results to shortening the time delays in the system. In fact, with this spatial network model we found a robust reversal of the network behaviour compared to that of the non-spatial network: the shorter the time delays, the more oscillatory the network becomes. This counter-intuitive result may suggest that the gradient in time delays imposed by spatial organisation may dominate the effect of reducing mean time delays in these networks. In fact, the role of space could be further enhanced (or suppressed) by introducing spatially graded connection probabilities (e.g. with closer cells more likely to be connected).

In conclusion, we have described a novel spatial effect. While the spatial organisation described is plausible, it and the corresponding effect on network dynamics would be interesting predictions to test in cortical networks.

Species selection of aging for the sake of diversity

Owen Woodberry, Kevin Korb and Ann Nicholson

Monash University

Tuesday 1:30pm–3:00pm, Stripe Lecture Room, Evolutionary Theory 2.

In their paper, “Punctuated Equilibria: An Alternative to Phyletic Gradualism”, Eldredge and Gould (1972) argue that most evolution occurs during geologically rapid speciation events, with species exhibiting stasis the vast majority of the time. Gould (2002, 2007) demonstrates that an implication of Punctuated Equilibrium is that selectionist theory is expanded to the level of species — defining species as basic units of macroevolution. In our paper, we demonstrate the evolution of aging rates in a species selection scenario. We have developed an ALife simulation environment with mating governed by evolving compatibility signatures, resulting in the formation of reproductively isolated subpopulations (i.e., species). Given a co-evolving parasite population, heterogeneity in a host subpopulation is beneficial for the health of that subpopulation. This can result in group selection pressure at the species level for the evolution of altruistic traits, such as a faster aging rate.

Initial modelling of the alternative phenotypes hypothesis

Simon Worgan¹ and Rob Mills²

¹University of Sheffield

²Science and Engineering of Natural Systems group, School of Electronics and Computer Science, University of Southampton

Tuesday 1:30pm–3:00pm, Stripe Lecture Room, Evolutionary Theory 2.

The alternative phenotype hypothesis contends that multiple phenotypes exist in a single genotype and are expressed by environmental or genetic cues. It further states that these multiple phenotypes will be maintained and improved in a population where the environment is unstable, in spite of the increased cost of this plasticity. In this work we propose a simple computational model to investigate the conditions under which alternative phenotypes become beneficial, and persist over evolutionary timescales. We find that the environment must vary to realise this hypothesis, and that these adaptations not only provide a fitness benefit in highly unstable environments but also continue to arise despite increasing stability and a corresponding gradual decline in fitness.

Passive and driven trends in the evolution of complexity

Larry Yaeger¹, Virgil Griffith² and Olaf Sporns¹

¹Indiana University

²California Institute of Technology

Wednesday 10:30am–12:00pm, Stripe Main Theatre, Open-Ended Evolution 1.

The nature and source of evolutionary trends in complexity is difficult to assess from the fossil record, and the driven vs passive nature of such trends has been debated for decades. There are also questions about how effectively artificial life software can evolve increasing levels of complexity. We extend our previous work demonstrating an evolutionary increase in an information theoretic measure of neural complexity in an artificial life system (Polyworld), and introduce a new technique for distinguishing driven from passive trends in complexity. Our experiments show that evolution can and does select for complexity increases in a driven fashion, in some circumstances, but under other conditions it can also select for complexity stability. It is suggested that the evolution of complexity is entirely driven just not in a single direction at the scale of species. This leaves open the question of evolutionary trends at larger scales.

Canalization and environmental engineering

Hajime Yamauchi and Kazuo Okanoya

RIKEN Brain Science Institute

Wednesday 10:30am–12:00pm, Exam Hall Room 1, Models for Gaia Theory.

Canalization is an umbrella term a leading C. H. Waddington introduced in 1960's. The term describes the ability of a population to buffer disturbances in phenotypic development. There are two major types of canalization based on the sources of disturbances: genetic and environmental. Genetic canalization describes the ability to express the same phenotype from different genotypes, while environmental canalization is to stabilize development under different environmental conditions.

With a graph of norms of reaction, evolutionary process of canalization is expressed as a squashing process of the pattern of phenotypic expression across a range of environments. It is also known that learning (or plasticity, in general) has a similar effect: it would normalize the end state of phenotypic development. However, it has not been thought of a case of canalization as it also enables a genotype to express two distinctive phenotypes (as the term "plasticity" implies).

In this work, we will challenge this view, and examine how learning would potentially provide a third type of canalization. Specifically, this work will show if phenotypic traits are partially inherited via cultural transmission (i.e., learning), and if the function of the trait is to do with social conformities among individuals in a population, it would possibly stabilize the learning environment itself so as for different genotypes to be able to express the same phenotype. This effectively narrows the range of trait variation not by (evolutionarily) modifying the reaction norm itself, but by manipulating the variation of environment the majority of the population would encounter during its learning period. As such, this type of canalization should be categorized neither genetic nor environmental canalization.

We use a multi-agent model which simulates evolution of language. In the model, the population consists of 200 agents allocated on a horizontal space. Each agent acquires her linguistic knowledge with inputs provided from adults' linguistic activities (i.e., communicative activities). Learnability of given linguistic knowledge is sensitive to 1. learner's genetic information, 2. learning ability, and 3. consistency of inputs across different adults. The fitness of an agent is measured by the number of successful communication with her neighbor peers.

With the condition that there is no evolution on the learning ability, the result shows that initially, selection works on genes so that agents' genetic information conforms with the linguistic knowledge dominating the population, as this eases the burden of learning—a classic example of the Baldwin effect. However, as doing so,

the population more and more converges into a single linguistic community. This paves the way to smoothing the learning environment (i.e., the collective state of learning inputs) from which later generations learn. Consequently, the environment is now confined itself to provide coherent inputs. Under this circumstance, the population can tolerate a certain degree of genetic disturbances as they can be absorbed by the plasticity which created such an environment itself.

As noted above, this type of canalization is primarily not a process of genetic evolution. Instead, it is an environmental engineering process which modifies the frame of norms of reaction itself.

Fitness based identification of a robot structure

Juan Cristobal Zagal¹, Javier Ruiz-del-Solar² and Adrian Galo Palacios¹

¹Centro de Neurociencia de Valparaiso, Universidad de Valparaiso

²Departamento de Ingenieria Electrica, Universidad de Chile

Friday 10:30am–12:00pm, Exam Hall Room 2, Information in Complex Systems & ALife 4.

Embodiment theory suggests that recurrent processes of sensorimotor activity give rise to cognitive structures. In the case of robots, internal sensorimotor activity generated with physics simulators can be exploited to expand the historical domain of action, however pre-engineered simulations are limited by the reality gap problem. Alternatively simulation might be inferred and self-constructed out of data collected during robot functioning. Fundamental to this line of research is defining a distance function to assess the potential of candidate robot simulations to reproduce real world activity. In this paper we study the characteristics of a distance function based on behavioral fitness measurements. We show how this function can be applied for the generation of behaviors using an algorithm that co-evolves a robot and its simulation. The experiments show how the monotonicity of the function increases with the amount of behaviors being tested in reality and with the genotypic diversity of corresponding robot controllers. Moreover it allows for the accurate identification of behavior-relevant parameters contained in the simulation. The metric shows an advantage, when compared to other metrics, for assessing the quality of simulators over long time scales of robot behavioral evaluation.

Author index

- Abi-Haidar, Alaa, 8
Adami, Christoph, 8
Adams, Rod, 51
Ahn, Rene, 69
Allen, Rosalind, 28
Almeida e Costa, Fernando, 8, 9
Ampatzis, Christos, 67
Anderson, Alexander, 29
Andras, Peter, 10
Andrews, Paul, 10
Anguita, Davide, 10
Anokhin, Konstantin, 19
Anthony, Tom, 11
Aono, Masashi, 11
Arita, Takaya, 64
Artmann, Stefan, 12
Aucouturier, Jean-Julien, 61
Ay, Nihat, 24
- Banzhaf, Wolfgang, 63
Barakova, Emilia, 69
Barandiaran, Xabier E., 12
Barnett, Lionel, 13, 18
Bateson, Patrick, 19
Bates, Paul, 14
Baum, Buzz, 22
Bedau, Mark, 13, 14, 53
Beer, Randall, 56, 74
Bentley, Katie, 14
Bermudez, Edgar, 15
Bersini, Hugues, 15
Berthouze, Luc, 75
Bigg, Zoe, 53
Birkin, Mark, 15
Bjørndalen, John Markus, 10
Blundell, Ben, 16
Bogacz, Rafal, 67
Bongard, Josh, 16
Bornhofen, Stefan, 16
Branciamore, Sergio, 16, 23
Brizzolara, Davide, 10
Bryden, John, 17, 18
Buckley, Christopher, 18
Buhmann, Thomas, 37
Bullinaria, John, 61
Bullock, Seth, 18, 28, 34, 37, 41
Burtsev, Mikhail, 18, 19
- Cañamero, Lola, 35, 51, 54
Cangelosi, Angelo, 65
Cannings, Chris, 36
Cariani, Peter, 19
Cartlidge, John, 20
- Cases, Blanca, 20
Caves, Leo, 23
Centler, Florian, 25
Chadwick, Matthew, 55
Chaggar, Satpal Singh, 21
Chalmers, Devin, 53
Channon, Alastair, 21
Chaumont, Nicolas, 8
Christensen, Anders, 67
Clarke, Donna, 49
Clark, James, 21
Cliff, Dave, 21, 34
Cohen, Michael, 22
Cohen, Netta, 75
Conibear, Tim, 53
Connor, Dustin, 60
Correia, Luís, 32
Costa, Ernesto, 30
Cowan, Lara, 23
Crailsheim, Karl, 59, 65
Crowder, Richard, 26
Cunningham, Alan, 51
Cussat-Blanc, Sylvain, 23
- Dale, Kyran, 23
d'Anjou, Alicia, 20
Danks, Gemma, 23
Dautenhahn, Kerstin, 46
Davey, Neil, 51
Davies, Mark S., 62
de Back, Walter, 16, 23, 39
Decraene, James, 24
Der, Ralf, 24
Di Paolo, Ezequiel, 12, 25, 57, 70
Dittrich, Peter, 25, 42
Dorigo, Marco, 67
Dorin, Alan, 25, 52
Doursat, Rene, 26
Dürr, Peter, 61
Duthen, Yves, 23
Dyke, James, 26
- Edlund, Jeffrey, 8
Endler, Lukas, 68
English, Stephen, 26
Etxeberria, Arantza, 27
Ewert-Sarmiento, Marcela, 45
- Fernandes, Carlos M., 47
Fernando, Chrisantha T., 37
Flamm, Christoph, 68
Floreano, Dario, 61
Francis, Charles, 53

Franks, Daniel W., 28
 Free, Andrew, 28

Gallori, Enzo, 16
 Gasser, Michael, 74
 Geard, Nicholas, 28
 Gerding, Enrico, 29
 Gerhardt, Holger, 14
 Gerlee, Philip, 29
 Ghio, Alessandro, 10
 Goldstein, Richard, 30
 Gómez, Jonatan, 45
 Gonçalves, Ângela, 30
 Gough, Jeffrey, 26
 Gowans, David, 30
 Graña, Manuel, 20
 Green, David, 52
 Greenwood, Dan, 31
 Griffith, Virgil, 76
 Grilo, Carlos, 32
 Grimm, Volker, 25
 Groom, Graeme, 32
 Gulyas, Laszlo, 39
 Güttler, Frank, 24

Hanczyc, Martin M., 33
 Hara, Masahiko, 11
 Haruna, Taichi, 33
 Harvey, Inman, 9, 33, 34, 43, 63, 70, 71
 Hawick, Ken, 59
 Head, Bryan, 53
 Hebborn, Tom, 34
 Hemelrijk, Charlotte, 51
 Heppenstall, Alison, 15
 Hernandez, Carmen, 20
 Hilbers, Peter, 42, 63
 Hintze, Arend, 8
 Hinze, Thomas, 42
 Hiolle, Antoine, 35
 Hofe, Robin, 35
 Hogeweg, Paulien, 64, 70
 Hogg, David, 18
 Holland, Owen, 43
 Houston, Alasdair, 67
 Hoverd, Tim, 54
 Howley, Enda, 35
 Husbands, Phil, 70

Iizuka, Hiroyuki, 36
 Ikegami, Takashi, 33, 36, 47, 60, 61, 64
 Irons, David, 36
 Izquierdo, Eduardo J., 37

Jacob, Christian, 71
 Jacyno, Mariusz, 37
 James, Emyr, 37

James, Richard, 28
 Jeanson, Francis, 38
 Jenkins, Dafyd, 38
 Jin, Yaochu, 39
 Joachimczak, Michal, 38
 Johnson, Alexis, 26
 Jones, Ben, 39

Kaleta, Christoph, 25
 Kampis, George, 23, 39
 Keijzer, Fred, 40
 Kelly, Ciaran, 40
 Knabe, Johannes F., 40
 Knight, Rob, 43
 Kok, Kasper, 69
 Kondo, Toshiyuki, 50
 Korb, Kevin, 25, 76
 Kraaijeveld, Alex, 53
 Krause, Jens, 66

Ladley, Daniel, 41
 Laredo, Juan Lu s Jim nez, 47
 Lattaud, Claude, 16
 Lehman, Joel, 41
 Lein, Adam, 42
 Lenser, Thorsten, 42
 Lenton, Tim, 74
 Liekens, Anthony, 42, 63
 Lipson, Hod, 56
 Litus, Yaroslav, 42
 Lizier, Joseph, 43
 Luga, Herve, 23
 Lundh, Torbjorn, 29

Macinnes, Ian, 9
 Malcolm, Grant, 73
 Manicka, Santosh, 43
 Marques, Hugo, 43
 Marques-Pita, Manuel, 44
 Marrow, Paul, 30
 Marshall, James, 67
 Matsumaru, Naoki, 25, 42
 Matthews, Jamie, 44
 Mattiussi, Claudio, 61
 McBurney, Peter, 29
 McGann, Marek, 44
 McMullin, Barry, 24, 40
 McNamara, John, 67
 McNeil, Amanda, 28
 Mendl, Mike, 67
 Mendoza, Ivan Alfredo, 45
 Merelo, Juan Julian, 47
 Metta, Giorgio, 65
 Miconi, Thomas, 45
 Mills, Rob, 32, 46, 49, 72, 76
 Miodownik, Mark, 22

Mirza, Naeem Assif, 46
 Mitchell, George G., 24
 Möller, Marco, 47
 Moore, Roger K., 35
 Mora, Antonio M., 47
 Moreno, Alvaro, 58
 Mozley, Olivia, 28
 Muethe, Patricia, 45

 Nagl, Sylvia, 48
 Nakajima, Kohei, 47
 Nehaniv, Chrystopher, 11
 Nehaniv, Chrystopher L., 40, 46
 Nene, Nuno, 48
 Newcombe, Richard, 43
 Nicholson, Ann, 76
 Nicholson, Daniel, 48
 Nishimura, Shin, 49
 Niu, Jinzhong, 29
 Noble, Jason, 21, 28, 37, 49
 Nolfi, Stefano, 66
 Nozawa, Takayuki, 50

 O'Brien, Darragh, 40
 Ochoa, Gabriela, 70
 Ogai, Yuta, 61
 Okanoya, Kazuo, 77
 Oosten, Johanneke, 51
 O'Riordan, Colm, 35, 51
 Oros, Nicolas, 51

 Palacios, Adrian Galo, 77
 Paperin, Gregory, 52
 Parisi, Domenico, 57
 Parodi, Giancarlo, 10
 Parsons, Simon, 29
 Paul, Liz, 67
 Peña, Jorge, 52
 Penn, Alexandra, 52, 53, 55, 72
 Pepper, Noah, 53
 Pestelacci, Enea, 54
 Philippides, Andrew, 15
 Pichler, Peter-Paul, 54
 Polack, Fiona, 54, 73
 Polani, Daniel, 11, 47, 55
 Popat, Sita, 18
 Powers, Simon, 53, 55, 72
 Prokopenko, Mikhail, 43, 55
 Psujek, Sean, 56

 Ramos, Vitorino, 47
 Rieffel, John, 56
 Rocha, Luis, 8, 44
 Rohde, Marieke, 57
 Rosa, Agostinho Claudio, 47
 Ruini, Fabio, 57

 Ruiz-del-Solar, Javier, 77
 Ruiz-Mirazo, Kepa, 58
 Ruxton, Graeme, 28, 66

 Sadedin, Suzanne, 52
 Saffre, Fabrice, 36, 58
 Sampson, Adam, 10, 54
 Sasai, Masaki, 49
 Sayama, Hiroki, 58
 Schermerhorn, Paul, 59
 Scheutz, Matthias, 59
 Schilstra, Maria J., 40
 Schmickl, Thomas, 59, 65
 Scogings, Chris, 59
 Sendhoff, Bernhard, 39
 Seth, Anil, 15, 60
 Shackleton, Mark, 58
 Shanahan, Murray, 60
 Silverman, Eric, 60, 61
 Smith, V. Anne, 61
 Soltoggio, Andrea, 61
 Sonneveld, Marthijn, 69
 Sorensen, Humphrey, 51
 Soyer, Orkun, 30
 Spanton, Robert, 26
 Speroni di Fenizio, Pietro, 25
 Sporns, Olaf, 76
 Spranger, Michael, 62
 Stafford, Richard, 62
 Stanley, Kenneth, 41
 Steels, Luc, 62
 Steijaert, Marvin, 42, 63
 Stekel, Dov, 38
 Stepney, Susan, 10, 23, 54, 73
 Steuber, Volker, 51
 Straatman, Bas, 63
 Studer, Gregory, 63
 Sun, Joanna, 26
 Suzuki, Keisuke, 64
 Suzuki, Reiji, 64

 Takeuchi, Nobuto, 64
 Tanev, Ivan, 43
 Tani, Jun, 65
 Tateson, Richard, 30
 te Boekhorst, René, 46
 ten Eikelder, Huub, 42, 63
 Thenius, Ronald, 65
 Tikhanoff, Vadim, 65
 Timmis, Jon, 10, 54
 Tomassini, Marco, 54, 70
 Tosh, Colin, 66
 Trianni, Vito, 66, 67
 Triesch, Jochen, 74
 Trimmer, Barry, 56
 Trimmer, Pete, 67

Tsuda, Soichiro, 12
Tuci, Elio, 67
Tufte, Gunnar, 67
Türkmen, Ulas, 67

Ullrich, Alexander, 68
Umerez, Jon, 58
Ursem, Bas, 70

van Apeldoorn, Dirk, 69
Vanderelst, Dieter, 69
van der Post, Daniel, 70
Vargas, Patricia A., 70
Vaughan, Richard, 42, 72
Veldkamp, Tom, 69
Verel, Sebastien, 70
Virgo, Nathaniel, 70, 71
Vladimirescu, Alexandru, 71
von Mammen, Sebastian, 71

Walker, Ian, 23
Wallis, Mick, 18
Warren, Douglas, 10
Watson, Richard, 32, 37, 46, 53, 55, 72
Wawerla, Jens, 72
Webb, Jeremy, 53
Webster, Matt, 73
Weeks, Andrew, 73
Welch, Peter, 10
Wheeler, Michael, 73
White, Roger, 63
Wiles, Janet, 74
Willadsen, Kai, 74
Williams, Gray A., 62
Williams, Hywel, 21, 74
Williams, Paul, 74
Winks, Oliver, 75
Womble, Stephen, 75
Woodberry, Owen, 76
Worgan, Simon, 76
Wróbel, Borys, 38

Yaeger, Larry, 76
Yamauchi, Hajime, 77
Yao, Xin, 39

Zagal, Juan Cristobal, 77
Zauner, Klaus-Peter, 12, 26
Zhang, Bing, 29
Zomaya, Albert, 43
Zugic, Radomir, 67