AISB Journal

The Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour

Volume 1 – Number 4 – December 2003

The Journal of the Society for the Study of Artificial Intelligence and the Simulation of Behaviour http://www.aisb.org.uk

Published by The Society for the Study of Artificial Intelligence and Simulation of Behaviour

http://www.aisb.org.uk/

ISSN 1476-3036 © December 2003

Contents

Imitation in Animals and Artifacts303Kerstin Dautenhahn & Chrystopher Nehaniv
A longitudinal study of immediate, deferred, and synchronic imitation through the second year
Mark Nielsen & Cheryl Dissanayake
Plausible Roles for Social and Individual Learning in the Speciation and Evolution of Cichlid Fish
Enhancing Autonomous Agents Evolution with Learning by Imitation
Creativity in Arts and Science
Optimality Principles for Conceptual Blending: A First Computational Approach 351 Francisco C. Pereira & Amílcar Cardoso
Scientific Methods for the Analysis of Agent-Environment Interaction
The Behaviour of a Mobile Robot Is Chaotic
On the Role of Robot Simulations in Embodied Cognitive Science

AISB Journal

Imitation in Animals and Artifacts

Kerstin Dautenhahn and Chrystopher Nehaniv

Department of Computer Science, University of Hertfordshire College Lane, Hatfield, Hertfordshire AL10 9AB, UK K.Dautenhahn@herts.ac.uk ; C.L.Nehaniv@herts.ac.uk

Editors' Introduction

This special issue of the AISB Journal includes a section comprising fully refereed, extended journal versions of the three most outstanding contributed papers selected from amongst those that were presented at the AISB'03 Second International Symposium on Imitation in Animals and Artifacts but which were not co-authored by symposium programme committee members. The symposium was organized by Prof. Kerstin Dautenhahn and Prof. Chrystopher L. Nehaniv (Adaptive Systems Research Group, University of Hertfordshire, U.K.), who are pleased to acknowledge the U.K. Engineering and Physical Sciences Research Council for support of the interdisciplinary meeting under EPSRC Grant GR/S57907/01 entitled Social Learning and Artificial Intelligence.

Mark Nielsen and Cheryl Dissanayake's article investigates three different types of imitation that can be observed in children: immediate imitation, deferred imitation and synchronic imitation. In a longitudinal study infants are studied repeatedly during their second year of life. While immediate and deferred imitation could be observed from 12 months onwards along a similar developmental trajectory, synchronic imitation only emerged after 18 months. The authors suggest that synchronic imitation is primarily a communicative behaviour. This contribution stresses the importance of a developmental perspective in the study of human behaviour and intelligence, as well as the need for carefully controlled experiments that can aid the understanding of the complexity and richness of animal and human behaviour.

The contribution by Ross Clement addresses the possible roles of individual versus social learning in the speciation and evolution of cichlid fish from the African Great Lakes where different species co-exist, occupying very similar ecological niches. Agent-based computer simulations study mechanisms of sympatric speciation, varying different environmental and behavioural parameters. Results point out the importance of learned behaviour in creating barriers between different populations. This paper gives an example of possible contributions of agent-based simulations on social and individual learning towards a better understanding of specific processes in biological evolution.

Elhanan Borenstein and Eytan Ruppin investigate interactions between evolution and imitative behaviour where within a generation agents can learn from other members in the population, 'teachers' (via horizontal transmission only). The Imitation Enhanced Evolution (IEE) framework is applied to three tasks, namely the parity problem, the triangle classification problem, and a foraging task. Computer simulations demonstrate that lifelong adaptive learning via learning by imitation can significantly enhance the performance of evolving agent populations. This work exemplifies the evolutionary potential of learning by imitation in Artificial Intelligence and autonomous agents research. The full proceedings of the symposium contain many other excellent contributions and are published in K. Dautenhahn and C. L. Nehaniv (Eds.), Proc. Second International Symposium on Imitation in Animals and Artifacts (7-11 April 2003; Aberystwyth, Wales; AISB'03), The Society for the Study of Artificial Intelligence and Simulation of Behaviour (ISBN 1-902956-30-7). Also, an edited book based on the topics of the AISB symposium and tentatively entitled Models and Mechanisms of Imitation and Social Learning: Behavioural, Social and Communicative Dimensions is in preparation for publication by a well-known scientific publisher.

A longitudinal study of immediate, deferred, and synchronic imitation through the second year

Mark Nielsen Early Cognitive Development Unit, School of Psychology, University of Queensland and Cheryl Dissanayake School of Psychological Science, La Trobe University *nielsen@psy.uq.edu.au; c.dissanayake@latrobe.edu.au*

Abstract

The aim in this study was to investigate the emergence of immediate imitation, deferred imitation and synchronic imitation in human infants through the second year. A longitudinal study was conducted of eighty-six infants at three monthly intervals from 12-through to 24-months-of-age. At each session the infants were assessed for immediate, deferred and synchronic imitation. Immediate and deferred imitation was evident in the infants from 12-months onwards. In contrast, it was not until 18-months that synchronic imitation began to emerge. Moreover, the infants' engagement in both immediate imitation and deferred imitation showed a common developmental trajectory that was distinct from the trajectory for synchronic imitation. It is argued that synchronic imitation, unlike the other forms of imitation assessed here, is primarily a communicative act.

1 Introduction

The vast corpus of experimental research conducted on the imitative abilities of infants and young children over the course of the last century has established that observing adults is an especially powerful way of learning in human children. The majority of this research has focused on the ability of infants and children to recreate the actions of others, whether the actions comprise facial and manual gestures or actions on objects.

Though somewhat controversial, there is evidence that immediate imitation of facial and manual gestures is present from birth (Meltzoff & Moore, 1977, 1983, 1989, 1994)¹. Immediate imitation and deferred imitation of actions on objects appear concomitantly towards the end of the first year (Meltzoff, 1985, 1988b). Deferred imitation refers to the actions of a model that are reproduced after a delay without the benefit of immediate practice.

In addition to facilitating the acquisition of new skills and behaviours, imitation also provides infants with a means by which they can acquire an awareness of sharing, not only specific actions, but also the intentions of others (Chapman, 1991). Imitation can therefore function to increase interaction between the imitator

¹ Though see Anisfield (1991)

and the imitated and, in so doing, provides an environment in which infants can engage in sustained communicative sequences.

Towards 18 months of age, certain behavioural and cognitive skills (Mitchell, 2002; Suddendorf & Whiten, 2001) develop that permit infants to coordinate their own actions with the thematic specifics of a social partner's play, and this in turn helps generate and sustain valued forms of cooperative action (Eckerman, Davis, & Didow, 1989; Eckerman & Didow, 1989; Nadel & Baudonnière, 1980, 1982; Nadel, Baudonnière, & Fontaine, 1983; Nadel & Fontaine, 1989). A common expression of this achievement in nonverbal communication is the interest infants show in sustaining dyadic play through what has been labelled synchronic imitation.

At approximately 18 months, infants begin to show a preference for engaging with objects that are similar to ones chosen by their play partner, and they use the common object in a similar postural, motoric, and symbolic way. Such play is characterised by its inherent reciprocity whereby the partners do not solely adopt one role but consistently alternate between model and imitator (Nadel et al., 1983). A feature of synchronic imitation is that infants not only reproduce the behaviour of their play partner but continue to do so in concert with him or her, all the while watching the other closely and showing clear signs of enjoyment.

The ability of infants to engage in synchronic imitation has been identified as a foundation for the onset of verbal language. That is, Nadel and her colleagues (2002; Nadel, Guérini, Pezé, & Rivet, 1999) argue that synchronic imitation is a form of pre-linguistic communication. As Chapman (1991) has noted, "infants perform particular actions as a means of signalling to their partners that they wish them to act in a similar manner, and their partners' actions are understood in the same way" (p. 214). When infants engage in synchronic imitation they are therefore able to build long-lasting interactions about an object through the alternation of imitating and being imitated. More specifically, the long-lasting interactions that often characterise synchronic imitation episodes in dyadic play lead to the construction of shared topics with a co-referent and, thereby, present a means by which infants can learn how to pre-verbally invite participation in play. Turn taking in synchronic imitation thus provides an important tool for social exchange, and hence a means for developing both referential and inferential communication.

Despite the important role that has been attributed to synchronic imitation, no study to date has investigated the developmental changes in synchronic imitation through the second year, the period when this ability becomes a feature of typical infant behaviour. Moreover, the development of synchronic imitation has yet to be appraised with reference to the development of infants' maturing capacities for reproducing actions on objects. The aim in the research reported here was to address this gap in the literature.

2 Method

2.1 Participants

Eighty-six infants infants (47 male, 39 female) were seen five times at intervals of three months from 12 through to 24 months of age. The infants were recruited from Maternal and Child Health Centres in suburbs surrounding La Trobe University, Melbourne, Australia. The infants were predominantly Caucasian, of middle socio-economic status, and participated in this study as part of a larger

http://www.aisb.org.uk

Nielsen and Dissanayake

longitudinal investigation (Nielsen & Dissanayake, submitted; Nielsen, Dissanayake, & Kashima, 2003). All testing was conducted within one month of the relevant target age.

2.2 General Procedure

The sessions were conducted in two playrooms at the Child Development Unit at La Trobe University. Playroom 1 was 5.35×4.60 metres and contained a video camera (that could be operated remotely) positioned above a 1.54×1.60 metre play mat placed on the floor. A basket containing a range of toys (see Appendix A) was placed on the edge of the mat. Playroom 2 measured 2.60×2.45 metres and contained a table (120 X 60cm) that was placed between two cameras. These cameras were used to videotape the infant and experimenter during administration of the tasks. All infants were tested individually with their primary carer present.

Upon arrival at the Child Development Unit, the infant and carer were escorted into Playroom 1. The carer and the experimenter sat on the play mat together with the infant who was allowed to explore the room. This warm-up stage was terminated once the infant appeared comfortable with both the playroom environment and the experimenter (usually less than 5 minutes). The free-play episode was then conducted followed by the synchronic imitation task. The carer and infant were then escorted into Playroom 2 and the deferred imitation task was administered.

2.3 Measures

2.3.1 Immediate Imitation

A Free-Play episode was conducted on the playmat in Playroom 1 at each session. The carer was asked to play with his or her infant, using the toys from the basket, exactly as if they were at home. They were also requested to encourage the infant to return to the playmat if he or she ventured off it. The infant and carer were left to play alone with the toys. The ensuing fifteen minutes were videotaped and the exhibition of immediate imitation was later coded continuously from the videotapes of these episodes. This provided a naturalistic measure of the infant's engagement in imitation.

The measure of immediate imitation was coded at each of the six sessions and was derived from Hart and Fegley (1994). Following Hart and Fegley, to be coded as an exemplar of immediate imitation the following criteria had to be met:

- 1. The infant had to observe the modelled behaviour from beginning to end.
- 2. The carer had to pause at the completion of the behaviour as an indication that the infant was expected to perform the action, or alternatively had to communicate the expectation verbally (e.g., "you do it").
- 3. Repeated attempts by the carer to model the behaviour within one minute were not coded.
- 4. The infant had to perform the action within one minute of observation of the carer, and repeated attempts by the infant at imitating the same action were not coded.

Scoring. All exemplars of immediate imitation were coded for the frequency of carer initiation and frequency of infant response. As the number of items imitated by the infant is dependent on the number of actions modelled by the

carer, the dependent variable of immediate imitation was the percentage of modelled actions adopted by the infant. Where the response of the infant was considered to be ambiguous, no imitation was coded.

2.3.2 Deferred Imitation

The deferred imitation task was adapted from Meltzoff (1985; Meltzoff, 1988a, 1988b). All testing took place with the experimenter sitting across from the infant at the table. The procedure was the same for each session although the objects and actions used were changed (a complete list of the objects and modelled actions is included in Appendix B). The procedure consisted of four distinct stages.

Baseline stage. A baseline stage was introduced to ensure that no infant spontaneously produced the target actions. The experimenter placed the first object on the table in front of the infant. If necessary the experimenter attracted the infant's attention to the object by making neutral comments that did not relate to the task. The infant was given 30 seconds to explore the object. This procedure was repeated until the infant had been exposed to three objects. The order in which the objects were presented was counterbalanced across infants. None of the target actions were spontaneously produced during this phase.

Modelling stage. Immediately following the baseline stage, the experimenter modelled the target actions associated with each object in the same order that the objects had been presented. Each action was modelled four times within a 30 second period. The experimenter then removed the object from the table thus ensuring that the infant was given no opportunity to engage with it until the response stage.

Intervention stage. The last object was placed away and the infant was engaged in other unrelated tasks for eight minutes.

Response stage. The infant was presented with the objects under the same conditions as the baseline stage. He or she was given 30 seconds with each object to produce the target behaviour.

Scoring. All coding was conducted from videotape. The coder was required to judge whether the target behaviour was produced in the response stage and, if so, the infant was awarded one point for each action. Therefore, for each session the infant could receive a deferred imitation score ranging from zero to three.

2.3.3 Synchronic Imitation

The procedure for synchronic imitation was adapted from Asendorpf, Warkentin, and Baudonnière (1996). The experimenter sat on the play mat opposite the infant. The infant's carer sat behind the infant and was instructed to avoid engaging with him or her. The experimenter placed a small tray of four duplicate toys/objects behind him. The episode commenced when the experimenter took the first object, offered the duplicate object to the infant, and began the first activity with that object. The experimenter continuously modelled an action for 15 seconds and then performed a second action with the same object for a further 15 seconds. Following completion of the second modelling period, the experimenter placed the object and its duplicate in the tray behind him and repeated the procedure until the infant had been exposed to all four objects and their accompanying actions. The order of administration was counterbalanced across subjects. If the infant stopped paying attention to the experimenter, the sequence was paused and recommenced

Nielsen and Dissanayake

when the infant reoriented to the experimenter. A complete list of the objects and modelled actions is included in Appendix C.

Scoring. The combined amount of time each infant spent engaged in synchronic imitation during the four trials of each session was coded from the videotapes. Following Asendorpf et al. (1996) a synchronic imitation sequence was determined to have commenced when the infant took the duplicate object, looked at the experimenter within 3 seconds, and imitated the action of the experimenter for at least 3 seconds. The duration of the sequence was coded for as long as the infant maintained imitation of the modelled action and continued to look at the experimenter at least once every ten seconds. A synchronic imitation sequence was considered to have terminated when either the experimenter stopped the activity or when the infant stopped the activity for more than 3 seconds. For each session the infant could receive a synchronic imitation score ranging from 0 to 120 seconds.

2.3.4 Reliability

The first author was the primary coder of all measures. A second trained coder, who was blind to the specific hypotheses of the study, independently observed and coded the videotapes of 12 randomly selected infants (2 from each session). Intraclass Correlation Coefficients (Shrout & Fleiss, 1979) were calculated between the scores of the two coders for each measure. The inter-rater reliability on all measures was good. The individual coefficients were as follows: Immediate Imitation 0.75; Deferred Imitation 0.96; and Synchronic Imitation 0.98.

3 Results

3.1 Immediate Imitation

The mean immediate imitation scores (and 95% confidence intervals) for each session are presented graphically in Figure 1. At 12-months, the infants adopted approximately two-fifths of the imitative bids offered by their carers and this number increased to slightly less than two thirds in the 15-month session. The proportion of bids adopted then stabilised at this level for the remainder of the study (though a slight but non-significant rise was observed at 21-months).



Figure 1. Mean percent (and 95% confidence intervals) of immediate imitation bids adopted by infants from 12- to 24-months-of-age.

http://www.aisb.org.uk

3.2 Deferred Imitation

The means (and 95% confidence intervals) for the number of items imitated (out of a total of 3) in each session are presented in Figure 2. Similar to the data for immediate imitation, the number of items imitated increased from the 12- to the 15-month session, after which performance reached a plateau.



Figure 2. Mean number (and 95% confidence intervals) of deferred imitation (DI) items exhibited by infants from 12- to 24-months-of-age.

3.3 Synchronic Imitation

The mean duration (and 95% confidence intervals) of synchronic imitation exhibited by the infants at each session are presented graphically in Figure 3. Unlike immediate imitation and deferred imitation, the infants exhibited little synchronic imitation prior to the 18-month session. However, a marked increase in the duration of synchronic imitation was observed from the 18-month session onwards. By the 24-month session, infants were spending approximately one third of the 120-second episode engaging in sustained imitative sequences.



Figure 3. Mean duration (seconds) (and 95% confidence intervals) of synchronic imitation exhibited by the infants from 12- to 24-months-of-age.

3.4 Summary

Infants demonstrated immediate and deferred imitation from the 12-month session onward. The infants' exhibition of these two skills increased to the 15-month session and then remained relatively constant thereafter. In contrast to these measures of imitation, the infants exhibited little synchronic imitation prior to the 18-month session. However, from 18-months onwards a marked increase in synchronic imitation was observed.

http://www.aisb.org.uk

Nielsen and Dissanayake

In order to enhance comparability across measures, the cumulative percentage of infants demonstrating at least one example of each measure was calculated for each session (see Figure 4). Immediate and deferred imitation emerged in close correspondence with each other. Synchronic imitation emerged later, and also showed a developmental trajectory that was distinct from the remaining two measures of imitation.



Figure 4. The cumulative percentage of infants showing evidence of immediate imitation (Immd), deferred imitation (Defd) and synchronic imitation (Synch) at each session.

4 Discussion

The majority of infants demonstrated a capacity for immediate imitation and deferred imitation from the 12-month session onward². Performance on both measures of imitation increased to the 15-month session, and then remained relatively constant thereafter. In addition, the percentage of infants demonstrating an ability to engage in immediate imitation and deferred imitation from 12- through to the 24-month session was almost identical. These findings are consistent with Meltzoff's (1988b) report that the capacity for deferred imitation is a concomitant of the capacity for immediate imitation. The concordance between the current findings and those of Meltzoff is notable given that Meltzoff evaluated immediate imitation through an experimenter-elicited task while 'spontaneous' caregiver elicited imitation was measured in the current study.

The pattern of results reported here indicates that there is a developmental disjunction in the emergence of synchronic imitation and both immediate and deferred imitation. This disjunction indicates that the exhibition of synchronic imitation is not solely an outcome of the ability of infants to reproduce the actions of a modelling adult. In this context, it is worthwhile noting that during administration of the synchronic imitation task in the 12- and 15-month sessions it was not uncommon for the infants to reproduce the target actions of the experimenter, but not to do so continuously. That is, they imitated but did not synchronically imitate.

² As noted previously a capacity for immediate imitation and deferred imitation is likely to exist earlier.

Immediate, deferred, and synchronic imitation

A potential explanation for the developmental disjunction between synchronic imitation and both immediate and deferred imitation is that, unlike the latter two skills, synchronic imitation relies on the infant's ability to attribute 'intentionality' to others (Asendorpf, 2002; Asendorpf et al., 1996; Suddendorf & Whiten, 2001). Asendorpf (2002) argues that in order to engage in synchronic imitation infants must be able to take the perspective of their interaction partner regarding the on-going use of objects. That is, the infant must appreciate the experimenter's intention for the infant to continue to copy his behaviour and to do so for a sustained period.

It is therefore argued that the exhibition of synchronic imitation is underpinned by an ability to attribute intentions to others. This argument would be severely compromised if the attribution of intentionality was first evident in children either well before or well after the middle of the second year. Fortunately this is not the case. A growing body of literature has identified the middle of the second year as the period when children develop a capacity for reading intentionality into the behaviour of others (Baldwin & Baird, 2001; Bellagamba & Tomasello, 1999; Carpenter, Akhtar, & Tomasello, 1998; Meltzoff, 1995; Moore & Corkum, 1998; Repacholi & Gopnik, 1997). For example, 18- but not 12-month-olds will reproduce an action an experimenter meant to do (but never actually did) as often as children who saw a successful demonstration of the complete target action (Bellagamba & Tomasello, 1999; Meltzoff, 1995). Meltzoff (1995) argues that by the middle of the second year infants do not interpret the behaviour of others solely in terms of physical movements or motions but rather represent the behaviour of others within a psychological framework that involves goals and intended acts. The data presented here on synchronic imitation supports this claim.

It is also noteworthy that while infants can imitate the movements of inanimate objects they do not treat them with a sense of intentionality in the same way that they treat human models (McConnell & Slaughter, 2001, July; Meltzoff, 1995). In this context, assessing the tendency of infants to synchronically imitate a humanoid robot will provide an interesting test of the argument that reading intentionality into the behaviour of others underpins synchronic imitation.

There has been considerable debate throughout the last decade over the capacity of children with autism and nonhuman primates to engage in imitation. However, mention of synchronic imitation in this debate has been largely absent. It is thus of interest to speculate on the possible nature of this novel form of imitation in children with autism and in nonhuman primates.

Children with autism show impairments in their ability to imitate the behaviour of others (Rogers, 1999; Williams, Whiten, Suddendorf, & Perrett, 2001). However, the results from a growing number of studies suggest that the imitation impairments in autism are reduced when these children are provided with a structured experimenter-elicited setting (Charman & Baron-Cohen, 1994; Charman et al., 1997). Responding to these discrepant findings, Nadel et al. (1999) argue that it is the socially laden aspect of imitation that is impaired in children with autism. According to this view, children with autism are capable of reproducing the actions of a model, but do not enter into the social aspects that characterise a typical imitative exchange. This view concurs with the position that the primary deficits associated with autism stem from socio-affective and/or socio-communicative dysfunction (Dawson & Lewy, 1989; Hobson, 1990, 1993; Moore, Hobson, & Lee, 1997).

Nielsen and Dissanayake

Given the role of social communication in synchronic imitation, the sociocommunicative dysfunction characteristic of children with autism should inhibit their tendency to engage in synchronic imitation. Consonant with this speculation, Nadel and her colleagues (Nadel & Pezé, 1993; Nadel & Revel, 2003, April) have reported subtle differences in the ways that children with autism engage in synchronic imitation when compared with normally developing children, For example, when placed in a synchronic imitation situation, children with autism reproduced the actions of their play partners but failed to develop the role switching behaviour (i.e., alternating between imitatee and imitator) characteristic of normally developing children. These children were also deficient in their attempts at inviting their partner to imitate. Nonetheless, there remains a paucity of research investigating synchronic imitation in children with autism, and this is clearly an area in need of continued investigation.

With regard to nonhuman primates, there remains considerable debate in the literature about the precise nature of social learning in these animals. First, no nonhuman primate species outside of the great apes (bonobos, common chimpanzees, gorillas, orangutans) has demonstrated a clear capacity for reproducing the behaviours of others via observation (Visalberghi & Fragaszy, 2002)³. Given this finding it is highly unlikely that synchronic imitation will be evident in monkeys or other primate species outside of the great apes.

To date, no study has been published that has systematically evaluated synchronic imitation in primates. The only potential source of information regarding this ability comes from Chantek, a sign-language trained orangutan, who in the context of a complex on-going interaction requested his carer to "do what I do" (Miles, Mitchell, & Harper, 1996). Nonetheless, debate remains as to whether the great apes should be classified as 'true imitators' (whereby they reproduce both the goal of a particular behaviour they have seen performed by another as well as the specific actions that brought about that goal) or 'emulators' (whereby they learn about the features and affordances of an object or situation by observing the behaviour of another but devise their own behavioural strategies to bring about the observed goal) (Byrne & Russon, 1998; Tomasello, 1996; Tomasello & Call, 1997; Whiten, 1998). However, if the great apes are 'emulators', rather than 'true imitators', and do not pay attention to the specific actions that bring about a goal, synchronic imitation is unlikely to be evident. That is, long-lasting interactions about an object through the alternation of imitating and being imitated are unlikely to be built if the focus of social learning exchanges is on the acquisition of knowledge about the features and affordances of an object or situation.

The line between the imitative abilities of human infants and the imitative abilities of our closest evolutionary relatives remains somewhat blurred. The use of imitation as a form of communication may yet prove to be a distinguishing feature between humans and our closest living relatives. This hypothesis remains to be tested.

The longitudinal data reported here provide further empirical support for the contention that that the exhibition of synchronic imitation does not emerge until around 18 months of age. Thus, it may not be until the middle of the second year that infants begin to use imitation not only as a means of acquiring new behaviours

³ Although see Custance, Whiten and Fredman, 1999; Voelkl and Huber, 2000.

and/or skills but as a form of communication. The long-term importance of synchronic imitation to normally developing children has yet to be determined. Moreover, the exhibition of this skill in children with autism and in nonhuman primates has not been systematically charted. In this context, further investigation of this novel form of imitation is clearly warranted.

Acknowledgements

The work described in this study was funded by a La Trobe University Postgraduate Research Scholarship and University of Queensland Postdoctoral Fellowship to the first author and formed part of his doctoral thesis under the supervision of the second author. The authors would like to thank the carers and children who gave their time to participate in this study. The authors also thank Virginia Slaughter, Thomas Suddendorf and Gabrielle Simcock for their comments and suggestions on this paper. A version of this paper first appeared in K. Dautenhan and C. L. Nehaniv (Eds.) *Proceedings of the AISB'03 Second International Symposium on Imitation in Animals and Artifacts* (pp. 131-137). Wales, UK: The Society of Artificial Intelligence and the Simulation of Behaviour.

References

Anisfield, M. (1991). Neonatal imitation. Developmental Review, 11, 60-97.

- Asendorpf, J. B. (2002). Self awareness, other awareness, and secondary representation. In A. Meltzoff & W. Prinz (Eds.), *The imitative mind: Development, evolution, and brain bases* (pp. 63-73). Cambridge, UK: Cambridge University Press.
- Asendorpf, J. B., Warkentin, V., & Baudonnière, P.-M. (1996). Self-awareness and other-awareness II: Mirror self-recognition, social contingency awareness, and synchronic imitation. *Developmental Psychology*, 32, 313-321.
- Baldwin, D. A., & Baird, J. A. (2001). Discerning intentions in dynamic human action. *Trends in Cognitive Sciences*, 5, 171-178.
- Bellagamba, F., & Tomasello, M. (1999). Re-enacting intended acts: Comparing 12and 18-month-olds. *Infant Behaviour and Development*, 22, 277-282.
- Byrne, R. W., & Russon, A. E. (1998). Learning by imitation: A hierarchical approach. *Behavioural and Brain Sciences*, 21, 667-721.
- Carpenter, M., Akhtar, N., & Tomasello, M. (1998). Fourteen- through eighteenmonth-old infants differentially imitate intentional and accidental actions. *Infant Behaviour and Development*, 21, 315-330.

Nielsen and Dissanayake

- Chapman, M. (1991). The epistemic triangle: Operative and communicative components of cognitive development. In M. Chandler & M. Chapman (Eds.), *Criteria for competence: Controversies in the conceptualization and* assessment of children's abilities (pp. 209-228). Hillsdale, NJ: Erlbaum.
- Charman, T., & Baron-Cohen, S. (1994). Another look at imitation in autism. Development and Psychopathology, 6, 404-413.
- Charman, T., Swettenham, J., Baron-Cohen, S., Cox, A., Baird, G., & Drew, A. (1997). Infants with autism: An investigation of empathy, pretend play, joint attention and imitation. *Developmental Psychology*, 33(781-789).
- Custance, D., Whiten, A., & Fredman, T. (1999). Social learning of an artificial fruit task in capuchin monkeys (Cebus apella). *Journal of Comparative Psychology*, *113*, 13-23.
- Dawson, G., & Lewy, A. (1989). Arousal, attention, and the socio-emotional impairments of individuals with autism. In G. Dawson (Ed.), Autism: Nature, diagnosis, and treatment (pp. 49-74). New York: Guilford Press.
- Eckerman, C. O., Davis, C. C., & Didow, S. M. (1989). Toddlers' emerging ways of achieving social coordinations with a peer. *Child Development*, 60, 440-453.
- Eckerman, C. O., & Didow, S. M. (1989). Toddlers' social coordinations: Changing responses to another's invitation to play. *Developmental Psychology*, 25, 794-804.
- Hart, D., & Fegley, S. (1994). Social imitation and the emergence of a mental model self. In S. T. Parker & R. W. Mitchell & M. L. Boccia (Eds.), Selfawareness in animals and humans: Developmental perspectives (pp. 149-165). New York: Cambridge University Press.
- Hobson, R. P. (1990). On the origins of self and the case of autism. *Development* and *Psychopathology*, 2, 163-181.
- Hobson, R. P. (1993). Autism and the development of mind. Hillsdale, NJ: Erlbaum.
- McConnell, D., & Slaughter, V. (2001, July). Does imitation reflect recognition of the intentions of others? Evidence from 12- and 18-month-old infants. Paper presented at the 12th Biannual Australasian Human Development Conference, Brisbane, Australia.
- Meltzoff, A. N. (1985). Immediate and deferred imitation in fourteen and twentyfour-month-old infants. *Child Development*, 56, 62-72.

http://www.aisb.org.uk

- Meltzoff, A. N. (1988a). Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology*, 24, 470-476.
- Meltzoff, A. N. (1988b). Infant imitation and memory: Nine-month-olds in immediate and deferred tests. *Child Development*, 59, 217-225.
- Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31, 838-850.
- Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, 198, 75-78.
- Meltzoff, A. N., & Moore, M. K. (1983). Newborn infants imitate adult facial gestures. *Child Development*, 54, 702-709.
- Meltzoff, A. N., & Moore, M. K. (1989). Imitation in newborn infants: Exploring the range of gestures imitated and the underlying mechanisms. *Developmental Psychology*, 25, 954-962.
- Meltzoff, A. N., & Moore, M. K. (1994). Imitation, memory, and the representation of persons. *Infant Behaviour and Development*, *17*, 83-99.
- Miles, H. L., Mitchell, R. W., & Harper, S. (1996). Simon says: The development of imitation in an enculturated orangutan. In A. Russon & K. Bard & S. T. Parker (Eds.), *Reaching into thought: The minds of the great apes* (pp. 278-299). Cambridge: Cambridge University Press.
- Mitchell, R. W. (2002). Imitation as a perceptual process. In K. Dautenhahn & C. L. Nehaniv (Eds.), *Imitation in animals and artifacts* (pp. 441-469). Cambridge, MA: MIT Press.
- Moore, C., & Corkum, V. (1998). Infant gaze following based on eye direction. British Journal of Developmental Psychology, 16, 495-503.
- Moore, D. G., Hobson, R. P., & Lee, A. (1997). Components of person perception: An investigation with autistic, non-autistic retarded and typically developing children and adolescents. *British Journal of Developmental Psychology*, 15, 401-423.
- Nadel, J. (2002). Imitation and imitation recognition: functional use in preverbal infants and nonverbal children with autism. In A. Meltzoff & W. Prinz (Eds.), *The imitative mind: Development, evolution, and brain bases* (pp. 63-73). Cambridge: Cambridge University Press.

Nielsen and Dissanayake

- Nadel, J., & Baudonnière, P.-M. (1980). L'imitation comme mode préponderant d'échange entre pairs au cours de la troisième année [Imitation as a predominant mode for exchange among three-year-old peers]. *Enfance*, 1-2, 77-90.
- Nadel, J., & Baudonnière, P.-M. (1982). The social function of reciprocal imitation in 2-year-old peers. *International Journal of Behavioural Development*, 5, 95-109.
- Nadel, J., Baudonnière, P.-M., & Fontaine, A. M. (1983). Les comportements sociaux imitatifs [Imitative social behaviours]. *Recherches de Psychologie Sociale*, 5, 15-29.
- Nadel, J., & Fontaine, A. M. (1989). Communicating by imitation: A developmental and comparative approach to transitory social competence. In B. H. Schneider & G. Attili & J. Nadel & R. P. Weissberg (Eds.), Social competence in developmental perspective (pp. 131-144). Dordrecht, The Netherlands: Kluwer.
- Nadel, J., Guérini, C., Pezé, A., & Rivet, C. (1999). The evolving nature of imitation as a format for communication. In J. Nadel & G. Butterworth (Eds.), *Imitation in infancy* (pp. 209-234). Cambridge: Cambridge University Press.
- Nadel, J., & Pezé, A. (1993). What makes immediate imitation communicative in toddlers and autistic children? In J. Nadel & L. Camaioni (Eds.), *New perspectives in early communicative development* (pp. 139-156). London, NY: Routledge.
- Nadel, J., & Revel, A. (2003, April). *How to build an imitator?* Paper presented at the Second International Symposium on Imitation in Animals and Artifacts, Aberystwyth, Wales.
- Nielsen, M., & Dissanayake, C. (submitted). Imitation, pretend play and mirror selfrecognition: A longitudinal investigation through the second year.
- Nielsen, M., Dissanayake, C., & Kashima, Y. (2003). A longitudinal investigation of self-other discrimination and the emergence of mirror self-recognition. *Infant Behavior and Development*, 26, 213-226.
- Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: Evidence from 14- and 18-month-olds. *Developmental Psychology*, 33, 12-21.
- Rogers, S. J. (1999). An examination of the imitation deficit in autism. In J. Nadel & G. Butterworth (Eds.), *Imitation in infancy* (pp. 253-283). Cambridge: Cambridge University Press.

http://www.aisb.org.uk

- Suddendorf, T., & Whiten, A. (2001). Mental evolution and development: Evidence for secondary representation in children, great apes and other animals. *Psychological Bulletin, 127*, 629-650.
- Tomasello, M. (1996). Do apes ape? In C. M. Heyes & B. G. Galef, Jr. (Eds.), *Social learning in animals: The roots of culture* (pp. 319-346). San Diego, CA: Academic Press.
- Tomasello, M., & Call, J. (1997). *Primate cognition*. New York: Oxford University Press.
- Visalberghi, E., & Fragaszy, D. M. (2002). "Do monkeys ape?" Ten years after. In K. Dautenhahn & C. L. Nehaniv (Eds.), *Imitation in animals and artifacts* (pp. 471-499). Cambridge, MA: The MIT Press.
- Voelkl, B., & Huber, L. (2000). True imitation in marmosets. *Animal Behaviour, 60*, 195-202.
- Whiten, A. (1998). Imitation of the sequential structure of actions by chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, 112, 270-281.
- Williams, J. H. G., Whiten, A., Suddendorf, T., & Perrett, D. I. (2001). Imitation, mirror neurons and autism. *Neuroscience and Biobehavioral Reviews*, 25, 287-295.

Plausible Roles for Social and Individual Learning in the Speciation and Evolution of Cichlid Fish

Ross Clement

Harrow School of Computer Science, University of Westminster Northwick Park, Harrow, HA1 3TP, UK *clemenr@wmin.ac.uk*

Abstract

Two agent-based simulation systems were created to investigate open problems in the speciation of cichlid fish from the African Great Lakes. A behaviour-free system has demonstrated that incompletely isolated subpopulations are still likely to speciate. A system modelling both social and individual learning was built to investigate the role that social and individual learning may play in creating partially isolated subpopulations. These populations have limited genetic exchange, a prerequisite for speciation in typical models of evolution. Individual learning does produce non-overlapping subpopulations, but social learning does not. Social learning can lead to populations of (virtual fish) failing to fully exploit all the resources in an environment, particularly when it is combined with individual learning. Incomplete exploitation of resources is a plausible explanation for the much larger than expected numbers of cichlid species inhabiting rocky reefs in lakes.

1 Introduction

There are still many mysteries remaining concerning exactly how evolution occurs, and particularly on how new species arise. That is, the process of 'speciation' where a single ancestor species gives rise to a number of new, 'daughter' species. Theories of evolution and the origin of species were inspired by natural systems, such as Darwin's observations of finches in the Galapogos Islands (Darwin, 1859). In current research on evolution, the natural system that poses the greatest number of questions about the process of speciation are the cichlid fish (Barlow, 2000) of the Great Lakes of Africa (Victoria, Tanganyika, and Malawi). The first 'mystery' concerning these fish is the rapidity at which new species arise. Lake Victoria completely dried up approximately 14,000 years ago (Johnson et al, 1996). Hundreds of species of cichlid fish are found in Lake Victoria that are found nowhere else. Some genetic and comparative evidence supports the view that these fish evolved in the lake, from one or two ancestors, in this time, though there are currently arguments that cichlids from Lake Kivu, rather than river cichlids, may have seeded the Lake Victoria cichlid flock. (See (Kocher, 2003) for a brief discussion of this argument). Cichlid fish also appear to invalidate the competitive exclusion principle. Different cichlid fish coexist in the same location, despite no, or extremely small, differences between the ecological niches occupied by these species (Konings, 2001). Finally, cichlid fish are one of the most common examples used to support the concept of sympatric speciation. The 'standard' model by which organisms speciate is 'allopatric' speciation. In short, a population is separated by a physical barrier (e.g. a mountain range), and continue breeding independently. Genetic drift, and/or selection pressures, cause the two populations to diverge, and by the time the barrier disappears, the two populations are unable to interbreed, and are now separate species. In sympatric speciation (Maynard Smith, 1966) (Via, 2001), an ancestor divides into two (or more) daughter species in the absence of physical barriers. The cichlids of Lake Barombi-Mbo (Schiewen et al, 1994) appear to have undergone sympatric speciation, as a single ancestor species has divided into a number of daughter species, in this small $(0.6km^2)$, smooth-walled crater lake. However, sympatric speciation has been viewed very narrowly in the past, as solely speciation with a lack of physical distance between the two populations. For example, in (Via, 2001), a parasitic insect is given as an example of sympatric speciation. The insect lives its entire life cycle, including meeting mates and mating, on a single species of host plant. When the insect jumped to another host found in the same area, two populations were formed which did not interbreed, and hence can be considered different species. However, while there is a lack of distance separating the two populations, the host-plant specificity is a strong physical barrier. Possibly as strong a barrier as a mountain range separating the two populations.

A great deal of research is continuing to investigate these open questions concerning cichlid biology and speciation/evolution of cichlid fish. Among the approaches being used in efforts to understand both speciation itself, and cichlid speciation in particular, computer simulation and modelling has been used several times. The need for computer simulation and mathematical modelling is clear, as speciation itself is probably too slow a process to be observed directly, or generated by experiment. Computer simulation has been used to investigate sympatric speciation (and is cited as strong support for it in (Via, 2001)). (Kondrashov & Kondrashov, 1999) and (Dieckmann & Doebeli, 1999) show that in the presence of appropriate environmental conditions (e.g. two different food sources requiring incompatible adaptation), a covariance can emerge between food source adaptation (trophic adaptation), and a species signalling/selection device (male colour and female colour preference). A population that divides into two in this manner is very consistent with speciation among real species.

(Turner & Burrows, 1995) modelled a population of fish where simultaneous mutations in both colour (from light to dark) and female mate colour preference resulted in speciation. In (Lande *et al*, 2001), sex-reversing mutations lead to a division of one species into two.

Computer simulations specifically aimed at modelling African cichlids have tended to ignore environmental effects (such as food sources), and concentrate on pure genetic models. The motivation for this is clear, as closely related (and presumably recently speciated) species tend to be nearly identical in trophic adaptation. This suggests that environmental factors may have little effect on speciation. In (Sturmbauer 1998), mention is made of 'intrinsic' and 'extrinsic' properties of organisms that can be involved in speciation. Intrinsic properties are properties of the organism itself that leads to speciation, while extrinsic properties are the properties of the environment that promote speciation. The genetic based models of cichlid speciation have investigated intrinsic properties alone can explain cichlid speciation. The cichlids that initially seeded the lakes were river fish, and no mass speciation has occurred in the rivers. However, exactly the same fish undergo mass speciation in lakes. This suggests that some difference in these two environments enables or triggers mass speciation.

Clement

The abstract simulations used to investigate sympatric speciation mix both intrinsic (a breeding model where females have a colour preference) and extrinsic (multiple food sources in the environment which require differential adaptation) properties. These three factors (colour, colour preference, and feeding adaptation) plus gender are the only factors that can vary during evolution in the genetic model. However, there are many other factors that combine to make up the full spectrum of cichlid biology. Notably absent from any of the previous models of speciation (cichlid or not) is behaviour. Cichlids are advanced organisms, and behaviourally complex. However, there seems to have been little investigation of any roles behaviour may have played in speciation. Most research on evolution and behaviour (Boyd & Richerson, 1985). In (Barlow, 2000) it is mentioned that behavioural changes will precede genetic adaptation, as there is no advantage in adaptation to a food source unless that food source is already being consumed. However, Barlow discusses only the effect of behaviour on differential adaptation after a species of fish has been divided into separate populations.

In this paper (an expanded version of (Clement, 2003a)), we use agent-based modelling techniques to investigate the conjecture that behaviour, in particular social and individual learning, can lead to population segmentation that can enable speciation without physical distance between subpopulations. This is, we believe, a few form of sympatric speciation, where a behavioural, rather than geographical barrier arises between subpopulation. Furthermore, we look at the possibility that social learning can result in a species not harvesting all resources present in an environment, leaving open the possibility of other species exploiting these 'overlooked' resources.

2 Methods

Agent-based modelling has been used to investigate the properties of cichlid populations with varying amounts of social (by observing other fish) and individual (from a fish's own experience) learning. There are of course many types of social learning (?). The style of social learning modelled in this paper most closely matches 'social facilitation'.

At present there are two agent-based modelling systems. A more 'traditional' model has been built (Clement, 2003c). This system can be considered an extension of the systems of (Kondrashov & Kondrashov, 1999) and (Dieckmann & Doebeli, 1999), intended to situate the genetic models proposed in those systems in a reasonably detailed simulation of rock living cichlids and their lake environment. This system will be referred to as the 'genetic' system, as speciation proceeds solely according to population genetics. In addition, a second system, the 'behavioural' system, has been implemented which does not allow variation in 'genetic' factors, but only behavioural factors. The 'genetic' system is the target of long-term development, and the 'behavioural' system is being used to experiment with behavioural models before their eventual incorporation into the genetic system.

Agent-based modelling was chosen because of the potential for flexibility in building, and configuring models. It is the central thesis of this paper that low-level day to day effects (e.g. a fish observing the actions of another fish) can have large cumulative effects on processes occurring at quite different time-scales (e.g. evolution and speciation). Small changes to the day to day characteristics of the model may have large (and unpredictable) effects on the longer term characteristics. By using an agent based model, it is possible to perform experiments over a wide range of systems by configuring and choosing agents. This is particularly the case in the genetic system, which allows the choice of types of agents, over and above setting simulation (and agent) parameters.

2.1 Population Segmentation

Experiments with the genetic system showed some interesting properties which inspired the search for alternative, more subtle methods of dividing populations than solid physical barriers. Rockfish in the African Great Lakes (hereafter this natural system will be referred to as 'the wild') live on rocky reefs, which can be quite small, and separated by sandy regions that the fish are very reluctant to cross. This leads researchers to propose that these reefs could enable a 'micro-allopatric' model of speciation, where very small isolated populations exist on reefs. Micro-allopatric situation is discussed and criticised in (Knight, 1999). While rare, fish do cross between reefs, as observed in (Konings, 1990). The genetic system allows the creation of a number of (virtual) reefs, and the specification of probabilities for fish crossing (migrating) between reefs. Initially, a single reef was created, with two food sources. These food sources were different, but not different enough for sympatric speciation to occur (i.e. the system stabilised with a single, generalist, species harvesting both food sources). After confirmation (by simulation) that this was the case, a new experiment was created with two reefs, where each food source was placed on a different reef. With no migration, each population stabilised as a specialist consumer of the food source present on their reef. After sufficient time (roughly equivalent to a virtual 7700 years), genetic drift gave two populations that differed not only in trophic adaptation, but also in colour and female mate preference, i.e. different species. Since the initial fish populating both reefs were identical, the experiment demonstrated allopatric speciation. i.e. we have a situation where speciation will occur with division of the population into independent subpopulations, but will not occur if the population remains as a single group.

Further experiments were performed with increasing probabilities of migration. After each experiment, the final population was investigated to see whether the result was a single generalist species, or two specialist species. This was achieved by writing a diagram in the native format of the xfig drawing/diagramming package (http://www.xfig.org), mapping the numerical phenotypes to (x,y) coordinates on an A4 page. The distribution of phenotypes was then investigated by eye. For high probabilities of migration, the two populations had acted as a single population. For lower, but not zero, probabilities of migration, the populations had acted as independent populations. Migration probabilities were made equal for *reef1* to *reef2* and migrations in the opposite direction. Probabilities represent the probability that an individual fish will migrate during each time period (roughly one week). The results are graphed in Figure 1.

The x-axis of this graph shows the probability of one fish migrating to the other population during one time period (one week). For each migration probability, 20 experiments were performed. A probability of speciation was calculated by counting the number of times speciation occurred, and dividing by 20. This is the p(speciation) plotted on the y-axis.

The point at which the populations ceased behaving as independent populations was approximately at a migration probability of 0.02 (2 percent) While this may appear a low number, the average lifespan of a fish was set at 100 steps (slightly under two years), meaning that the average fish living to maturity will switch populations twice. Or, that since the population sizes typically number about 100 fish per population, that four fish will swap populations every week. That speciation still occurs in this situation shows that division of populations need not be complete for speciation to be possible.

Clement



Figure 1: Speciation of Partially Divided Populations

2.2 Behavioural System

In the wild, different species of cichlids are typically found at different water depths. Seehausen (1996) describes the cichlid species found at a single reef, showing how the large number of species present and the characteristic depths where they live. As well as dividing themselves into overlapping, but partially distinct populations, different food sources (and methods of harvesting them) are found at the different depths. For example, algae grows profusely in shallow water, but is far less frequent in the relatively dim depths. Hence, if a single species of fish invades a new reef, and somehow divides itself into populations that are sufficiently (if not completely) divided, then all of the factors necessary for speciation (genetic separation and divergent survival pressures) are present, and speciation should follow.

Before experimentation, both social and individual learning appeared to be plausible mechanisms of achieving this division. For individual learning, a fish may associate a given depth and mode of harvesting a certain type of food with success in obtaining food. This may encourage the fish to remain at that depth. With social learning, a fish may observe fish in its own vicinity tending to remain at that depth and harvesting food in a particular manner. This may encourage fish to (again) remain at that depth following a particular lifestyle.

Social learning in fish has been observed by several researchers. For example (Dugatkin, 1996) observed that social learning affected mate choice in guppies. In (Brown & Laland, 2002), the effects of social learning on the foraging behaviour of Atlantic salmon was described. The models to be presented in this paper most resemble simple forms of social learning, similar to that described as 'social facilitation' in (Zentall, 1996). This is a simple application of Occam's razor. If division of populations can be demonstrated in the simplest forms of social learning, those most supported by hard evidence, then there is no need for the adoption of more sophisticated models such as 'true imitative learning' (Thorpe, 1963). The models of the effects of copying mate choice described in (Boyd & Richerson, 1985) were also a major inspiration for the work described in this paper. Boyd and Richerson describe optimal ratios for social and individual learning in organisms given different properties of their environment. This is particularly relevant to our work, as we suppose that cichlids colonising a lake will have a mix more suited to river environments than lakes.

2.3 Behaviour Model

In order to implement agents (representing fish) that learn, a model is needed to represent behaviour that fish can perform, and allow individual fish to vary that behaviour in response to individual experience, and by learning from observation of other fish. Rather than try to develop a model of flexible behaviour that would allow fish to create new behaviours not encoded in the original model, a simple model was created. In this model, based on that in (Nehaniv & Dautenhahn, 2002), all possible behaviours are encoded in a tree. Each node in the tree is considered to be a state, and each move from one node to a daughter node is an action. Therefore, a path in the tree to some other node in the tree becomes a sequence of actions, which in Nehaniv and Dautenhahn's terminology characterises a 'behaviour'. Given that the root of the tree is considered to be a resting state, a tree can be built representing all sequences of actions that a fish might undertake before returning to the resting state. Note that behaviours can end at any node marked as a possible end state (bold rectangles in Figure 2).

A simple set of possible behaviours was created for experiments. The 'reef' water space is divided into two spatial locations, *shallow* and *deep*. Within these locations, there are two food sources: a, and b, perhaps roughly equivalent to algae and invertebrates. A fish can either crudely scrape the rock, harvesting (slightly inefficiently) both a and b, or it can adopt one of two specialist modifications of that behaviour, harvesting specifically a or b alone, with higher efficiency. a1 refers to food source a in shallow water, while a2 is the same food present in deep water. Similar b1 and b2 refer to food type b in shallow and deep water. The numbers in Figure 1 (e.g. $a1 \ 0.03$, $b1 \ 0.0$) refers to a specialist from food source a1, but nothing from food source b1. By omission, this behaviour also harvests nothing from food sources a2 and b2.

This model is common to all fish, and describes the universe of possible behaviours available to fish. Individual fish store their own behavioural characteristics as a set of probabilities. Using Figure 2 as an example, a fish in state x1 has three possible choices. It can cease its behaviour in that state $(stop_{x1})$, it can continue the behaviour by performing action s3 leading to state x3, or it can perform action s4, leading to state x4. Hence, the set of possible decisions for state x0 is:

$$D_{x1} = \{p(s3), p(s4), p(stop_{x1})\}$$
(1)

As these are an exhaustive list of decisions for state x1, we have:

$$p(s3) + p(s4) + p(stop_{x1}) = 1.0$$
(2)

All behaviours begin at state x0. Since this is not a possible final node the probabilities stored for this state are:

$$D_{x0} = \{p(s1), p(s2)\}$$
(3)

Clement



Final Node	Reward	
x1	0.0075 a1	0.0075 b1
x2	0.03 a1	0.0 b1
x3	0.0 a1	0.03 b1
x4	0.0075 a2	0.0075 b2
x5	0.03 a2	0.0 b2
x6	0.0 a2	0.03 b2

Figure 2: Behaviour Model and Rewards

with p(s1) + p(s2) = 1.0. Conversely state x3 has no actions leading from it, and hence

$$D_{x3} = \{p(stop_{x3})\}$$
(4)

with $p(stop_{x3}) = 1.0$.

At the beginning of a behaviour, a fish is in state x0. It then makes a choice of action according to the probabilities in D_{x0} , moving to another state (called x_N for convenience.) It then examines the options for actions at that state, and continues until the *stop* action is chosen. Since the universe of possible behaviours is a tree, the fish must eventually end the behaviour. In the universe of behaviours of Figure 2, the maximum number of actions in a behaviour is 2.

At the end of a behaviour, the fish will *bite* at resources, and will harvest resources in amounts controlled by both the maximum reward for that behaviour (Figure 2), and also the amount of the resources still left in the environment.

2.4 Learning

Given the behaviour model described, learning is achieved by modifying the weights. Sources of learning include: recent individual experience, social learning (observation of others, i.e. horizontal transmission of behaviour), and parental learning (vertical transmission).

Individual learning is the simplest, as it only requires access to an individual's past experience in order to update the weights. Behaviours are performed in batches (typically of 100 behaviours), and learning is performed at the end of each of these batches. First, the success of all actions performed is calculated. This is performed by finding the rewards received for every action that passed through, or stopped at, each state in the tree. For example, a fish may have performed 58 behaviours that included state x1, resulting in a total energy gain of 0.7. The same fish will have performed 42 behaviours including state x2, resulting in a total energy gain of 0.5. To calculate the weights of experience, these weights are raised to the power of a learning parameter (called *power*) to allow a controlled degree of non-linearity in learning. Hence, if $reward(state_n)$ is the

actual amount of reward received from behaviours including $state_n$, in the past batch of behaviours, then the weight for $state_n$ is:

$$weight(state_n) = reward(state_n)^{power}$$
⁽⁵⁾

A new set of probabilities can be calculated from these weights. We call the set of probabilities according solely to recent experience p'. Assuming that *action* is a state change from $state_n$ to $state_m$, and $p'(stop_n)$ is the probability of stopping in $state_n$ (which is 0 for nodes that are not potential final nodes), and A_n is the complete set of actions available from $state_n$, then:

$$p'(state_n \to state_j) = \frac{weight(state_j)}{weight(state_n)}$$
(6)

and:

$$p'(stop_n) = \frac{weight(state_n) - \left(\sum_{(state_n \to state_j) \in A_n} weight(state_j)\right)}{weight(state_n)}$$
(7)

Finally, the *strength* of learning controls how quickly a fish learns. That is, if p(x) is a probability from a fish's previous behaviour, and p'(x) is the probability solely due to recent experience, then the updated behaviour (probability for any action) after one step of learning is:

$$p^*(x) = (1.0 - strength) \times p(x) + strength \times p'(x)$$
(8)

strength can range between 0.0 and 1.0. A *strength* of 1.0 would create a fish that only acted according to its recent experience, while a *strength* of 0.0 results in a fish that never learns. Though, note that the previous expression is only correct if the fish is only performing individual learning, and the weighted expression is more complex when more than one type of learning is being performed at once.

Social learning (either vertical or horizontal transmission) proceeds similarly, except that the weights assigned to states in the tree are calculated by different methods. For parental transmission, if $count(parent, state_n)$ is the number of times that a parent (either *father* or *mother*) had performed a behaviour including $state_n$ in the last batch of behaviours, $weight(state_n)$ is calculated as.

$$weight(state_n) = count(father, state_n)^{power} + count(mother, state_n)^{power}$$
(9)

Once these weights are calculated, then probabilities p'(x) can be calculated according to (5) and (6).

For horizontal transmission, the situation is more complicated again. A fish learns from its peers by adjusting its weights according to the behaviours that it sees other fish performing. In our model, we assume that the environment is divided into two locations, *shallow* and *deep* water. A fish only observes behaviour in the location that it is in at any particular time. However, fish do not necessarily spend all their time for a whole batch of behaviours in the same location. The first step of any behaviour is effectively a choice to be in *shallow* or *deep* water, and hence a fish can (if it so chooses) move back and forth between the available locations. Fish choose behaviours (and hence locations) statistically independently from each other. Therefore, if $L = \{shallow, deep\}$ is the set of locations, and time(f, l) is the proportion of time that a fish f spends in location l, then the expected degree of spatial overlap between any two fish f_1 and f_2 is:

$$overlap(f_1, f_2) = \sum_{l \in L} time(f_1, l) \times time(f_2, l)$$
(10)

Clement

And, if F is the full population, $count(f, state_n)$ is the number of behaviours including $state_n$ performed by fish f in the last batch, then the weight assigned to $state_n$ (from the viewpoint of the learning fish $f_{learning}$) is:

$$weight(state_n) = \left(\sum_{f \in F} count(f, state_n) \times overlap(f, f_{learning})\right)^{power}$$
(11)

Note that these weights must be calculated anew for every fish in the population, as each fish has its own independent view of the rest of the population.

There are three different types of learning that can be performed by fish, and these are performed simultaneously. Assuming that $p'_i(x)$, and $p'_s(x)$ are the two sets of new probabilities calculated for recent individual and social learning respectively, and, p(x) and $p^*(x)$ are the previous probability and the new (updated) probability as before, then:

$$p^{*}(x) = strength_{i} \times p_{i}'(x) + strength_{s} \times p_{s}'(x) + (1.0 - strength_{i} - strength_{s}) \times p(x)$$
(12)

Given this equation, a framework for individual action, and learning has been created. Fish have the opportunity to compete for resources, and adjust their behaviours according to social and/or individual learning.

2.5 Other Parameters

The behavioural system starts a simulation by reading a configuration file, and creating a population of fish with random behavioural probabilities. Food sources are also created, matching those in the reward table of the behavioural tree. At each step in the simulation, all fish perform a fixed number of 'behaviours'. Each behaviour results in a *bite* at food sources, which transfers energy from the food source to the biting fish. If during this process, any food sources 'run out' (i.e. available energy falls to zero), then any further bites at that food source are ineffective. This repeats until all fish (performing behaviours in round robin style) have performed the fixed (currently 100) number of behaviours in a batch. Next, all fish lose a fixed amount of energy, representing energy expended on metabolism, and behaviour. Any fish whose energy reserves fall below zero dies of starvation. Fish then increment their 'age' parameter by one step. Any fish that has reached the end of its natural lifespan then dies of old age. Food sources then regenerate by gaining a fixed amount of energy, set in the configuration file. If any food source has exceeded its set maximum of available energy (simulating properties such as algae covering all available rock surfaces and being unable to spread further), then the energy is reduced to the maximum. Finally, fish are allowed to mate. If a female fish has reached maturity, and is ready to breed, then it does so. If the female is un-mated (or has a mate that has died), then it chooses a mate from among the un-mated males. This is done solely according to the amount of time that the female and the males have spent together over the last time period. If U is the set of un-mated males, then the chance of the female f choosing a male m is:

$$p(f,m) = \frac{overlap(f,m)}{\sum_{m' \in U} overlap(f,m')}$$
(13)

Once a mate is chosen, this pair of fish remains mated for life. This pair will then breed frequently (every few weeks), laying a small number of eggs each time, with one juvenile surviving for each clutch of eggs. This is a slightly optimistic (no divorces occur) and over-regular model of the 'trickle breeding' Tanganyikan rock-living *Julidochromis* species. In the wild pairs of fish will defend a single cave, effectively 'setting down' in a fixed location. In our model, we do not limit the movement of the fish in a pair, or deliberately make the pair 'stay together' in a single location. This is so that co-located pairs can be shown to be (or not be) an emergent property of the behaviour system being modelled.

Note that breeding produces more fish (parameters are set so that the population will grow), but the fixed input of energy per time period effectively defines a carrying capacity (the maximum number of fish the environment will support). In the absence of factors preventing full exploitation of the resources, the population will grow to the carrying capacity, and oscillate around that figure.

Every step, a large number of properties of the population and its actions are written to a trace file. As well as sum totals indicating the number of times each action was performed, two summary variables are output: the total population size, and the average overlap between mated pairs. This property has been used as the most reliable indication of population segmentation. For average overlap to approach 1.0, both of the pair must spend all or nearly all of their time in one environment, and all or nearly all fish must have chosen mates from the same environment.

One major omission in the current model is predation and predators. As well as cichlids that feed on algae and invertebrates, all the African lakes contain a wide variety of cichlids that feed on cichlids. For the rocky reefs of Lake Malawi, the aptly named *Tyrannochromis* predates on *mbuna* (the typical algae grazing rock cichlids). *Nimbochromis venestus* has an advanced predation strategy, pretending to be a dead, rotting, fish laying on the lake bed, and eating small fish that attempt to eat this 'carrion'. Even in the relatively young Lake Victoria, cichlid predators were common before their decimation after the introduction of the Nile Perch (Goldschmidt, 1996).

Importantly for simulation studies, it is theorised that predators can stabilise populations of competing species (Kooi & Kooijman). While Kooi and Kooijman did not even mention speciation, or the particular situation of cichlids in African lakes, it seems intuitive that their models may have something to say about African cichlids. For example, they may provide an alternate contributory or complete explanation for the unexpectedly large number of apparantly competing species found in African lake cichlid species. It is planned that future experiments will investigate the effects of recursive (fish speciating into predators consuming fish from a common ancestor), and non-recursive (specialist 'predation' agents without the capability of evolution) predation. It is felt that trying to add all possible factors affecting evolution into a single model at once is a bad research strategy. A solid understanding of the properties of biological factors studied independently (at least very difficult using real organisms) is necessary. Without this it is likely to be very difficult to quantify the contribution of these properties to patterns observed in any 'kitchen sink' simulation approach.

2.6 Experiments

Experiments are performed by creating a configuration file. There are huge numbers of variables which can be adjusted, and it is impossible to test all possible combinations of values. Hence, a single behavioural model (Figure 2) has been applied, and 'reasonable' values have been chosen for all parameters apart from the *strength* and *power* of each type of learning. These are: initial population of 50 individuals, maximum age of 1000 steps, and a probability of a female breeding per step of 0.01. All trials were run for 6000 steps (roughly 115 years).

Given the above fixed (across all experiments) parameters, an individual experiment

Clement

is characterised solely by the *strength* and *power* of the type of learning investigated. Note that (11) and (12) show how *power* and *strength* are used in the model. Experiments performed were:

- 1. No learning.
- 2. Weak social learning (strength=0.01, power=1.0)
- 3. Social learning (strength=0.1, power=2.0)
- 4. Strong social learning (strength=0.1, power=2.0)
- 5. Very strong social learning (strength=0.1, power=4.0)
- 6. Individual learning (strength=0.01, power=2.0)
- 7. Strong individual learning (strength=0.1, power=2.0)
- 8. Both social (strength=0.1, power=2.0) and individual (strength=0.1, power=2.0) learning.

Each experiment was run for 2000 time steps. The simulation system wrote measures of population size, population segmentation, and summaries of behaviour to a plain text trace file. This file was parsed by a program that converted the plain trace into comma separated value (csv) files. These were then loaded into a spreadsheet application (OpenOffice Calc), and graphs created for visual inspection.

3 Results

Population sizes are graphed in Figure 3. The x-axis represents time (each step being approximately one week), and the y-axis graphs the total number of individuals alive at each time period. Without learning, the population rises to, and stabilises at, 200 individuals. Weak social learning, and moderate individual learning give similar results, showing that the available resources are being fully exploited by populations implementing these learning strategies.

Strong social learning, very strong social learning, and both types of learning led to a stable population size about half the size of the full carrying capacity of the environment, indicating that at least some of the available resources were being ignored. Finally, strong individual learning leads to an unstable situation, with fluctuating population sizes. Investigation of simulation traces show that the population is gravitating to one food source (whichever is more plentiful at the time), and becoming fixated on that source. Due to excessive competition, the population crashes, and allows the resource to regenerate.

Population segmentation results are presented in Figure 4. Due to the method of calculation, 0.5 is the minimum amount of segmentation that can be expected (indicating random movement of fish between the two depths). 1.0 is the maximum possible, if all fish spend all their time at one depth, and only mate with other fish spending all their time at the same depth. Pure social learning (strong, average, or weak) does not lead to noticeable segmentation of the population, with the exception of very strong social learning, which leads to cyclic behaviour where the populations split apart, then merge back together again. Individual learning shows progressively more segmentation as with increasing amounts of learning. Strong individual learning segments the population to a high degree (>90%). Combining social and individual learning appears to give perfect



Figure 3: Population Sizes

segmentation of the population. However, close examination of the final trace shows that nearly all fish are spending all their time in *deep* water, leading to a high degree of mate similarity, even though there is no real 'pre-speciation'. High degrees of social learning result in population segmentation, even though resources from both locations are being utilised to their full, indicating a true splitting of the population.

In interpreting these results it is important to remember that we are considering two different methods by which learned behaviour may affect speciation in African cichlids. A population may split into two parts with limited genetic exchange between the two. Individual learning appears to provide these conditions with both high segmentation of the population, and full exploitation of the available resources as evidenced by the large population size. (Individual traces were examined to verify that this was indeed the case). Another method by which learned behaviour may help explain African cichlids is if learned behaviour leads to a population failing to exploit all the resources in an environment. This would leave those resources available for other species, without requiring the other species to have a significantly different ecological niche. This pattern (again, verified by examining traces) is present in social (and combined) learning experiments, shown by a small population, and a high measured population segmentation, showing that the species is choosing only one of the available locations.

4 Conclusions and Future Work

Segmentation of the population due to flexible, and learned, behaviour has been demonstrated. It was originally expected that social learning would be the main influence behind any segmentation of the population, but this has been contradicted by the results of the

Clement



Figure 4: Population Segmentation

experiments. Social learning has a tendency to focus the population on a single resource, without dividing the population.

Sympatric speciation is frequently criticised for the lack of mechanisms that could prevent gene flow between diverging populations. The results presented in this paper suggest that learned behaviour is a plausible factor which could provide a barrier to gene flow. While the barrier created is not absolute (as an impassible mountain range would be), an absolute barrier is not necessary for speciation. This suggests that non-geographical barriers, such as behavioural barriers, are plausible explanations for apparent examples of sympatric speciation.

Agent-based simulations often have the problem that the model is described as a large number of explicit low-level parameters, while the characteristics of interest (such as population segmentation) are implicit. For properties such as population size, it is trivial to calculate the property. But, tracking the rise (and possible fall) of species is a more difficult task. Fish agents are not labelled with a crisp species label, and to do so would be to pre-judge the results of the simulation. In the work reported in this paper, measurable factors (similar location preferences between mates) assumed to be synonymous with the implicit factor being searched for (separation of subpopulations) were used in place of an exact measure of the desired property. Currently most development effort in this research is being devoted to finding model-independent methods of tracking and visualising speciation, with current research described in (Clement, 2003b).

While the current results appear informative, great care is needed when evaluating the relevance of any results from simulation to real natural systems. As discussed in (Clement, 2003c), there is a major problem in designing simulations of natural systems, that of knowing exactly where to set the values of crucial parameters. For example, in

this paper we describe certain parameter settings as being 'strong individual learning', or 'very strong social learning'. While these parameters are felt to be 'intuitively correct', further research is definitely required to put these figures on a more concrete footing. An example is the research reported in papers such as (Kondrashov & Kondrashov, 1999). Simulation models show that given a set of entirely reasonable assumptions, sympatric speciation neatly follows. These are: multiple food sources requiring different adaptations, female selection of males according to some characteristic, and plasticity in both the female's preference, and the characteristic displayed by males. While there is significant argument about whether sympatric speciation really occurs in nature, it is at best a rare event. Clearly at least one of the assumptions leading to sympatric speciation in these simulation models must be incorrect. One possibility is that the differences in adaptation required by the multiple food sources are greater than those found in real environments.

At present biological simulation allows us to undertake experiments impossible in real life. However, it must be remembered that real life, in the form of biological field research, remains the ultimate model upon which results should be evaluated.

Acknowledgements

This research has benefited greatly from discussions with a number of colleagues, in particular George Turner, Mark Gardner, and the comments of anonymous referees.

References

- Barlow, G. W. (2000). *The Cichlid Fishes: Nature's Grand Experiment in Evolution*. Perseus, M.A.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press, Chicago.
- Brown, C. and Laland, K. (2002) Social learning and social inhibition in a novel foraging task in Atlantic salmon (Salmo salar). *Journal of Fish Biology* 61(4): 987-998.
- Clement, R. P. (2003a). Plausible Roles for Social Learning in the Speciation and Evolution of Cichlid Fish. Proceedings of the Second International Symposium on Imitation in Animals and Artifacts, K. Dautenhahn and C. L. Nehaniv (Eds.), Aberystwyth. 138-144.
- Clement, R. P. (2003b). Visualising Speciation in Models of Cichlid Fish. Proceedings of the 17th European Simulation Multiconference, Nottingham, UK. 344-348.
- Clement, R. P. (2003c). ESS2003, 15th European Simulation Symposium and Exhibition. Deft.
- Darwin, C. (1859). On the Origin of Species by Means of Natural Selection.... Project Gutenberg. http://www.ibiblio.org/gutenberg/cgi-bin/sdb/t9.cgi/t9.cgi?entry=1228.
- Dautenhahn, K., & Nehaniv, C. L. (Eds) (2002). *Imitation in Animals and Artifacts*. MIT Press.
- Dieckmann U, & Doebeli M. (1999). On the origin of species by sympatric speciation. Nature 400 (6742): 354-357.

Clement

Dugatkin, L. A. (1996). Copying and mate choice. In (Heyes & Galef, 1996), 85-105.

Goldschmidt, T. (1996). Darwin's Dreampond: Drama in Lake Victoria. MIT Press.

- Heyes, C. (2002). Transformational and Associative Theories of Imitation. In (Dautenhahn & Nehaniv, 2002), 501-523.
- Heyes, C. & Galef, B. G. (Eds) (1996). Social Learning in Animals: The Roots of Culture. Academic Press.
- Johnson T. C., Scholz C. A., Talbot M. R., Kelts K., Ricketts R. D., Ngobi G., Beuning K., Ssemmanda I. I., & McGill J. W. (1996). Late Pleistocene Desiccation of Lake Victoria and Rapid Evolution of Cichlid Fishes. *Science* 273(5278) 1091-1093.
- Knight, M. E. (1999). *Sexual Selection and Speciation in Lake Malawi Cichlids*. Ph.D. thesis. University of Southampton.
- Kocher, T. D. (2003). Fractious phylogenies. Nature 423: 489-491.
- Kondrashov, A. S., & Kondrashov, F. Y. (1999). Interactions among qualitative traits in the course of sympatric speciation. *Nature* 400: 351-354.
- Konings, A. (2001). Malawi Cichlids in their Natural Habitat 3rd Edition. Cichlid Press.
- Konings, A. (1991). Koning's Book of Cichlids and All Other Fishes of Lake Malawi. TFH Publications.
- Kooi, B. W., & Kooijman (2000). Invading Species can Stabilize Simple Trophic Systems. *Ecological Modelling* 133 57-72.
- Laland, K., Richerson, P., & Boyd, R. (1996). Developing a theory of animal social learning. In (Heyes & Galef, 1996).
- Lande, R., Seehausen, O., & van Alphen, J. J. (2001). Mechanisms of rapid sympatric speciation by sex reversal and sexual selection in cichlid fish. *Genetica* 112-113: 435-443.
- Mamedov, A. & Udalov, S. (2002). A computer tool to develop individual-based models for simulation of population interactions. *Ecological Modelling* 145, 23-68.
- Maynard Smith, J. (1966). Sympatric speciation. Am. Nat. 100 637-650.
- Nehaniv, C. L. & Dautenhahn, K. (2002). The Correspondence Problem. In (Dautenhahn & Nehaniv, 2002), 41-62.
- Schiewen, U. K., Tautz, D., & Pääbo, S. (1994) Sympatric speciation suggested by monophyly of crater lake cichlids. *Nature* 368 (6472), 629-632.
- Seehausen, O. (1996). Lake Victoria Rock Cichlids. Verduijn Cichlids, Germany.
- Sturmbauer, C. (1998). Explosive speciation in cichlid fishes of the African Great Lakes: a dynamic model of adaptive radiation. *Journal of Fish Biology* 53 (Supplement A), 18-36.
- Thorpe, W. H. (1963). Learning and Instinct in Animals. Methuen.

- Turelli, M., Barton, N., & Coyne, J. (2001). Theory and speciation. *Trends in Ecology and Evolution.* 16: 330-343.
- Turner, G. F. & Burrows, M. T. (1995). A model of sympatric speciation by sexual selection. *Proc. Royal Soc. Series B* 260 (1359), 287-289.
- Via, S. (2001). Sympatric speciation in animals: the ugly duckling grows up. *Trends in Ecology and Evolution* 16: 381-390.
- Zentall, T. R. (1996). An analysis of imitative learning in animals. In (Heyes & Galef, 1996), 221-243.
- Zentall, T. R., & Akins, C. (2001). Imitation in Animals: Evidence, Function, and Mechanisms. In *Avian Visual Cognition*, Cook R. G. (Ed). http://ww.pigeon.psy.tufts.edu/avc/toc.htm.
Enhancing Autonomous Agents Evolution with Learning by Imitation

Elhanan Borenstein* and Eytan Ruppin[†]

* School of Computer Science, Tel Aviv University, Tel-Aviv 69978, Israel, borens@post.tau.ac.il

[†] School of Computer Science and School of Medicine, Tel Aviv University, Tel-Aviv 69978, Israel, *ruppin@post.tau.ac.il*

Abstract

This paper presents a new mechanism to enhance the evolutionary process of autonomous agents through lifetime adaptation by imitation. Imitation is an effective method for learning new traits and is naturally applicable within the evolutionary paradigm. We describe a set of simulations where a population of agents evolve to solve a certain task. In each generation, individuals can select other agents from the population as models (teachers) and imitate their behavior. In contradistinction to previous studies, we focus on the interaction between imitation and evolution when imitation takes place only across members of the same generation, and does not percolate across generations via vertical (cultural) transmission. We show how this mechanism can be applied to successfully enhance the evolution of autonomous agents, when other forms of learning are not applicable.

1 Introduction

A large body of work in recent years has studied *the interaction between lifetime learning and genetic evolution* when lifetime adaptations, acquired by learning, are not inherited. Hinton and Nowlan (1987) introduced a simple model that demonstrates how learning can guide and accelerate evolution. Nolfi et al. (1994) presented experimental results supporting this view, even when the learning task differs from the evolutionary task. Other researchers (Nolfi and Parisi, 1997; Floreano and Mondada, 1996) studied the interaction between learning and evolution in robots and artificial agents systems. These studies employed various sources of training data such as external oracles, regularities in the environment or "self-generated" teaching data. There is, however, an additional source of training data; one which is naturally available within the evolutionary paradigm - the knowledge possessed by other members of the population. This knowledge can be harnessed to improve the evolutionary process in the form of *learning by imitation*.

The motivation for using learning by imitation to enhance evolution is twofold. First, imitation is an effective and robust way to learn new traits by utilizing the knowledge already possessed by others. The existence of true imitative behavior in the animal kingdom is still in debate, however, social learning can be found in a variety of species providing clear benefits over other forms of learning (Kawamura, 1963; Whiten and Ham, 1992; Zentall, 2001). Second, while oracles or other forms of supervised training data are scarce

in agent environments, learning by imitation is still a valid option, using other members of the population as teachers.

Extending these studies further, our goal is to put forward a novel framework for merging these two approaches and study learning by imitation within the scope of the interaction between learning and evolution. We wish to explore learning by imitation as an alternative to conventional supervised learning and to apply it as a tool to enhance genetic evolution. We will label this framework as imitation enhanced evolution (IEE).

Learning by imitation has already been applied by researchers in the fields of artificial intelligence and robotics in various experiments. Hayes and Demiris (1994) presented a model of imitative learning to develop a robot controller. Billard and Dautenhahn (1999) studied the benefits of social interactions and imitative behavior for grounding and use of communication in autonomous robotic agents. For an up-to-date introduction to work on imitation in both animals and artifacts see the cross-disciplinary collection (Dautenhahn and Nehaniv, 2002). Furthermore, various frameworks that study the interaction between cultural transmission and evolution have already been well established (e.g. Boyd and Richerson, 1985; Cavalli-Sforza and Feldman, 1981; Laland, 1992). Gene-culture co-evolution accounts for many adaptive traits (Feldman and Laland, 1996). Studies and simulations of the evolution of language (Ackley and Littman, 1994; Kirby and Hurford, 1997; Arbib, 2002) assume, by definition, some sort of cultural transmission.

It is important to realize though, that in contradistinction to these studies, our framework does not employ cultural evolution. In fact, we preclude culture from evolving in the first place. Following in the footsteps of the studies of the interaction between learning and evolution cited above, we thus avoid any direct form of acquired-knowledge transfer between generations either genetically or culturally. We work in a strict Darwinian framework, where lifetime adaptations are not directly inherited (although, as demonstrated in some of the studies cited above, they may be genetically assimilated through the Baldwin effect, 1896) and may affect the evolutionary process only by changing the individual's fitness, and thus the number of its offsprings. In terms of cultural transmission (see Boyd and Richerson, 1985, for a detailed definition), we allow horizontal transmission alone (where individuals of the same generation imitate each other) and exclude any form of vertical transmission (where members of the current generation transmit their knowledge to members of the next generation). Numerous field studies suggest that at least in nonhuman societies, horizontal transmission is far more common than vertical transmission (Laland, 1992). Furthermore, to prevent any form of cultural evolution from taking place, within each generation, only innate behaviors are imitated; that is, we prevent behaviors acquired by imitation from being imitated again by another member.

A simple model that fits this framework has been studied before by Best (1999). He demonstrated an extension of the computational model presented in Hinton and Nowlan (1987), introducing social learning (namely *imitation*) as an additional adaptive mechanism. The reported results exemplify how horizontal cultural transmission can guide and accelerate the evolutionary process in this simplified model. Best has also demonstrated how social learning may be superior to conventional learning and yield faster convergence of the evolutionary process. However, Best's model has several limitations. The evolutionary fitness function (which is the one used in Hinton and Nowlan, 1987) represents a worst-case scenario where only the exact solution has a positive fitness value. There is no probable path that a pure evolutionary search can take to discover this solution. Additionally, there is no distinction between genotypes and phenotypes and thus no real *phenotypic* adaptation process. Imitation is carried out simply by copying certain *genes* from the teacher's genome to the student.

We wish to generalize this framework and study the effects of learning by imitation

Borenstein and Ruppin

in a more realistic scenario of autonomous agents evolution (see Ruppin, 2002, for a general review). The definition of imitation in the literature varies considerably (Billard and Dautenhahn, 1999), but for the purpose of this paper we use imitation (or learning by imitation) in the sense of having an individual (student) being able to match its behavior to that of a demonstrator (teacher). In particular, using autonomous agents to model the population members, this form of imitation is implemented by using the teacher's output for each sensory input, as the target output in a back-propagation training algorithm. We focus on the effects that imitation may have on the genetic evolutionary process, starting with the most basic question: can imitation enhance the evolution of autonomous agents (in the absence of vertical transmission), in an analogous manner to the results previously shown for supervised learning, and how? Although it was shown that learning can guide the evolutionary process (e.g., via the Baldwin effect), the contribution of *imitation* to evolution is not obvious; while in late stages of the evolutionary process the best agents may already possess sufficient knowledge to approximate a successful teacher, in early stages of the process it may be the case of "the blind leading the blind", resulting in a decrease of the population's average fitness.

This paper presents a set of simulations, where lifetime learning by imitation was used to adapt individuals that go through an evolutionary process. The results are compared with those of a simple evolutionary process, where no lifetime learning is employed, and with those of an evolutionary process that employs conventional supervised learning.

The remainder of this paper is organized as follows. We begin in Section 2 with a brief overview of the effect of lifetime adaptation on the evolutionary process. In Section 3 we present the *IEE* model in details. To validate the effectiveness of our model we introduce in Section 4 a set of tasks which were used to test our model and the experimental results in Section 5. The paper concludes with a discussion of future work and a short summary.

2 The Effects of Lifetime Adaptation on Genetic Evolution

Studies of the interaction between lifetime learning and evolution (Hinton and Nowlan, 1987; Nolfi et al., 1994; Nolfi and Parisi, 1997; Floreano and Mondada, 1996) have shown that learning can accelerate and guide the genetic evolutionary process. These studies demonstrated (through both theoretical analysis and simulations) how the dynamics of the lifetime adaptation process can account for this positive effect. The phenotypic modifications that take place in an individual subject to lifetime adaptation (e.g. learning), significantly depend upon its innate configuration. Individuals which initially have a low fitness value, may attain higher fitness through learning. The expected fitness gain though, will be higher for individuals which are initially closer to the optimum configuration. As illustrated in Figure 1, learning can thus help to reveal the innate potential of each individual in the population. One may consider lifetime adaptation as a local search process that can enhance the global search (evolution) by determining which configurations lie in the vicinity of the global optimum solution and are thus worthwhile retaining in the population (as they have a better chance to produce successful offsprings). From a mathematical standpoint, lifetime adaptation can be conceived as a functional that can potentially transform an initially ragged fitness function into a smoother function, making the evolutionary process more effective.

Our hypothesis is that learning by imitation, that is, using the best individuals in the population as teachers, may be sufficient to reveal the innate *potential* of the population

members. The results reported in the following sections clearly validate this assumption.

In this study we focus on the simple case where the learning (imitation) task is similar to the evolutionary task. This case most probably does not closely represent the imitation processes found in nature. Lifetime adaptation in humans and other cultural organisms operates on high-level traits which are not coded directly in their genome. However, we believe that this simple scenario can provide valuable insights into the roots of imitative behavior. We further discuss this topic in Section 6.



Figure 1: An illustration of the effect that lifetime adaptation may have on the genetic evolutionary process. Both agents start with the same innate fitness value (indicated by the black dots). Applying lifetime adaptation (illustrated as a simple hill climbing process) will result in the selection of agent A which is closer to the optimal solution. Inspired by Nolfi and Floreano (1999)

3 The Model

A haploid population of agents evolve to solve various tasks. Each agent's neurocontrollers is a simple feed-forward (FF) neural network (Hertz et al., 1991). The initial weights of the network synapses are coded directly into the agent's genome (the network topology is static throughout the process). The initial population is composed of 100 individuals, each assigned randomly selected connection weights from the interval [-1,1]. The *innate fitness* of each individual is determined by its ability to solve the specific task upon birth. Within the pure evolutionary process, the innate fitness will determine the reproductive probability of this individual. Each new generation is created by randomly selecting the best agents from the previous generation according to their innate fitness, and allowing them to reproduce (Mitchell, 1996). During reproduction, 10% of the weights are mutated by adding a randomly selected value from the interval [-0.35,0.35]. The genomes of the best 20 individuals are copied to the next generation without mutation.

When conventional supervised learning is applicable (i.e., an explicit oracle can be found) we also examined the effect of supervised learning on the evolutionary process. Each individual in the population goes through a lifetime learning phase where the agent employs a back-propagation algorithm (Hertz et al., 1991), using the explicit oracle as a teacher. Its fitness is then reevaluated to determine its *acquired fitness* (i.e., its fitness level after learning takes place). In order to simulate the delay in fitness acquisition associated with acquired knowledge, we use the average of the innate and acquired fitness values as

Borenstein and Ruppin

the agent's *final fitness* value. This fitness value is then used to select the agents that will produce the next generation.

In the IEE paradigm, agents do not use conventional supervised learning, but rather employ learning by imitation. In every new generation of agents, created by the evolutionary process, each agent in the population selects one of the other members of the population as an imitation model (teacher). Teachers are selected stochastically, where the probability of selecting a certain agent as a teacher is proportional to its *innate* fitness value (i.e., its initial fitness levels before learning takes place). The agent employs a back-propagation algorithm, using the teacher's output for each input pattern as the target output, mimicking a supervised learning mode. The imitation phase in each generation can be conceived as happening simultaneously for all agents, preventing behaviors acquired by imitation from being imitated. Only the *innate* behavior of the teacher is imitated by the student. The *acquired fitness* and *final fitness* are evaluated in the same method that was described in the case of conventional learning.

As stated above, acquired knowledge does not percolate across generations. Each time a new generation is produced, all lifetime adaptations possessed by the members of the previous generation are lost. Newborn agents inherit only the genome of their parents which does not encode the acquired network adaptations that took place during the parent's lifetime. Successful individuals that were copied from the previous generation also go through a new genotype-to-phenotype ontogenetic development process and thus lose all adaptations acquired during the previous generation.

To summarize, learning by imitation in a population of evolving agents (IEE) works as follows:

- 1. Create the initial population. Assign the network weights of each individual with randomly selected values.
- 2. Repeat:
 - (a) For each individual in the population:
 - i. Evaluate the innate fitness F_i .
 - (b) For each individual S in the population:
 - i. Set S to be the student.
 - ii. Select a teacher T from the population. The probability of selecting a certain individual as a teacher is proportional to its innate fitness value F_i .
 - iii. Train S with back-propagation algorithm. Use the output of T as the desired output (when computing the output of T, use the innate configuration of T).
 - iv. Evaluate the acquired fitness F_a of S.
 - (c) For each individual in the population:

i. Evaluate the final fitness $F_f = \frac{F_i + F_a}{2}$.

(d) Create the next generation by selecting the best individuals according to F_f and allow them to reproduce as described above.

4 The Tasks

The model described in the previous section was tested on three different tasks. The first two are standard classification benchmark problems. The third is an agent-related task

used in previous studies of the interaction between learning and evolution.

4.1 The Parity Problem

The agents evolved to solve the five bit parity problem. A network topology of 5-6-2-1 was used (i.e., 5 input neurons, two hidden layers, the first with 6 neurons and the second with 2, and 1 output neuron), with an additional threshold unit in each layer. All 32 possible input patterns were used both for evaluating the network performance and for training.

4.2 The Triangle Classification Problem

A simple two-dimensional geometrical classification problem was used in this task. The network receives as input a point from the unit square and should determine whether it falls within the boundaries of a predefined triangle. A network topology of 2-5-1 was used (with an additional threshold unit in each layer). The test set and training set consisted of 100 points randomly selected from the unit square.

4.3 Foraging

The task in this simulation is similar to the one described by Nolfi et al. (1994). An agent is placed on a two-dimensional grid-world (Figure 2). A number of food objects are randomly distributed in the environment. As its sensory input the agent receives the angle (relative to its current orientation) and distance to the nearest food object. The agent's output determines one of four possible actions: turn 90 degrees left, turn 90 degrees right, move forward one cell, or do nothing (stay). If the agent encounters a food object while navigating the environment, it consumes the food object. The agent's fitness is the number of food objects that were consumed during its lifetime. Each agent lives for 100 time steps



Figure 2: The foraging task: The agent (triangle) navigates in a 2D grid-world. Food objects (stars) are randomly distributed in the world. The agent can turn 90 degrees left, turn 90 degrees right, move one cell forward, or stay. Each time the agent encounters a food object, it consumes the food object and gains one fitness unit. Inspired by Nolfi and Floreano (1999)

Borenstein and Ruppin

in a 30x30 cells world which initially contains 30 food objects. A network topology of 2-6-2 was used (with an additional threshold unit in each layer).

In this task, unlike the previous ones, there is no explicit oracle we can use to train the agent. Nolfi et al. (1994) used available data to train the agent on the task of predicting the next sensory input, which differs, but is in some sense still "correlated" with that of finding food (the evolutionary task). In our model, we can still use the same mechanism of learning by imitation to train the agent on the original evolutionary task, using the best individuals in the population as teachers.

There are several strategies we can apply to determine which sensory input patterns should be used for training. Randomly selecting arbitrary input patterns, as we did in previous tasks, is not a suitable strategy here as the real input distribution that an agent encounters while navigating the environment may differ considerably from a uniform distribution. However, two behaviorally motivated strategies may be considered: a *query* model and an *observational* model. In the query model, the student agent navigates in the environment and for each sensory input pattern it encounters, the student queries the teacher to obtain the teacher's output for this pattern. The teacher's output is than used as the target output in back-propagation training of that pattern. In the observational model, the student "observes" the teacher agent as the teacher navigates in the environment and uses the sensory input patterns encountered by the teacher as training patterns (again, using the teacher's output for the back-propagation algorithm). Using this model we can further limit the observed patterns to those which occur during time steps that precede the event of finding food. This constraint will allow the student to imitate only useful behavioral patterns. We will label this strategy as *reinforced agent imitation (RAIL)*.

5 Results

We first studied IEE in the two classification tasks described in Sections 4.1 and 4.2, where conventional supervised learning can still be applied. In these tasks we were able to compare the effects that both lifetime adaptation mechanisms (i.e., learning and imitation) have on the evolutionary process. The results clearly validate that the IEE model



Figure 3: The triangle classification task: the innate fitness of the best individual in the population as a function of generation.

consistently yields an improved evolutionary process. The *innate fitness* of the best individuals in populations generated by applying learning by imitation is significantly higher than that produced by standard evolution.

Figure 3 illustrates the *innate* performances of the best agent as a function of generation, in populations evolved to solve the triangle classification problem (Section 4.2). To evaluate the agent's classification accuracy we use the Mean-Square Error (MSE) measure to calculate the distance between the network predicted classification and the true classification, averaged over all the patterns in the test set. Fitness is defined as (1-Error). The results of a simple evolutionary process (dashed line) and of an evolutionary process that employs conventional supervised learning (dotted line) are compared with those of an evolutionary process that employs learning by imitation (solid line). Each curve represents the average result of 4 different simulation runs with different, randomly assigned, initial connection weights. The results presented in Figure 3 demonstrate how applying either of the learning paradigms yields better performing agents than those generated by a simple evolutionary process. In fact, applying learning by imitation produces practically the same improvement throughout the process as does conventional supervised learning.

When facing the 5-bit parity task, the effect of applying lifetime adaptation is even more surprising. Figure 4 illustrates the *innate* performances of the best agent as a function of generation, in populations evolved to solve the 5-bit parity problem. Each curve represents the average result of 10 different simulation runs with different, randomly assigned, initial connection weights. While simulations applying the IEE model still outperform the simple evolutionary process, using conventional supervised learning actually results with a significant decrease in performances. The problematic nature of this specific task may account for these poor results. The parity problem, although often used as a benchmark, is considered to be a difficult and atypical classification problem (Fahlman, 1989). Learning algorithms facing this task tend to get trapped in local minima. However, learning from an imperfect teacher, as is the case in learning by imitation, induces a certain level of noise into the learning process and may thus help to prevent the process from getting stuck.

Evidently, learning by imitation has a similar (if not superior) effect on the evolution-



Figure 4: The 5-bit parity task: the innate fitness of the best individual in the population as a function of generation.

Borenstein and Ruppin

ary process to the one that was previously shown for conventional supervised learning. The knowledge possessed by the best members of the population can be used as alternative training data for other members, even in the early stages of the evolutionary process. We then turned to use IEE to enhance evolution where explicit training data is not available. This is the case in the foraging task described in Section 4.3.



Figure 5: The foraging task: the average *innate fitness* of the population as a function of generation. The results of a simple evolutionary process are compared with those of simulations that employed lifetime imitation with two distinct adaptation forces (2 and 8 learning iterations).

Figure 5 illustrates the results of the simulations in which the agents faced the foraging task. The average *innate fitness* of the population in a simple evolutionary process is compared with the average *innate fitness* of populations that applied learning by imitation. The agents in this simulation employed the *RAIL* strategy of imitation. Fitness is measured as the number of food objects an agent consumes during its lifetime. Each curve represents the average result of 10 different simulation runs with different, randomly assigned, initial connection weights. As can be seen in Figure 5, autonomous agents produced by our model demonstrate better performances than those generated by the simple evolutionary process; that is, their *innate* capacity to find food in the environment is superior.

We also examined the effect of employing different *adaptation forces*. In our experiments, the adaptation force is implemented simply as the number of learning (back-propagation) iterations we apply in each lifetime adaptation phase. The results illustrated in Figure 5 also demonstrate that a higher adaptation force (i.e., a higher number of iterations in each imitation phase) further improves the performance of the resulting agents. This effect coincides with an analogous effect reported by Best (1999) where higher transmission force resulted with faster convergence of the evolutionary process.

To further explore the effects of lifetime imitation on evolution, we examined the improvement in fitness during lifetime as a function of generation. The improvement can be evaluated by calculating the difference between the *acquired fitness* and the *innate fitness* (i.e., $F_a - F_i$) in every generation. The results illustrated in Figure 6 clearly demonstrate that in very early stages of the evolutionary process, the best agents in the population already possess enough knowledge to improve the fitness of agents that imitate them. In fact, the contribution of imitative learning decreases as the evolutionary process

Enhancing Autonomous Agents Evolution with Learning by Imitation



Figure 6: The foraging task: the improvement of the population average fitness gained by lifetime imitation as a function of generation.

proceeds, probably due to population convergence to high performance solutions.

An additional observation on the interaction between lifetime adaptation and evolution can be obtained from examining the diversity of the population throughout the evolutionary process. The average genome variance of the population, i.e., the variance among the population members, in the value of each gene (encoding a certain network weight) averaged over all genes, can serve as a measure of the population's diversity. As demonstrated in Figure 7, during the first few generations, the population's initial diversity decreases rapidly due to the selection pressure of the evolutionary process. However, throughout most of the following generations, the diversity found in populations subject to lifetime adaptation by imitation is higher than the diversity of populations undergoing a



Figure 7: The foraging task: the average genome variance as a function of generation with and without imitation. Populations that employ lifetime adaptation, maintain a higher diversity throughout the evolutionary process.

Borenstein and Ruppin

simple evolutionary process. Allowing members of the population to improve their fitness through lifetime adaptation before natural selection takes place facilitates the survival of suboptimal individuals and helps to maintain a diversified population. This feature can partly account for the benefit gained by applying lifetime adaptation to agents evolution.

6 Discussion

This paper demonstrates how learning by imitation can be applied to an evolutionary process of a population of agents, utilizing the knowledge possessed by members of the population. Our IEE model proves to be a powerful tool that can successfully enhance evolutionary computation simulations in agents.

In our model, the agents' ability and incentive to imitate is assumed to be instinctive. Quoting Billard and Dautenhahn (1999), "our experiments address learning by imitation instead of learning to imitate". The imitation paradigm presented in this paper additionally assumes that the agents can estimate the fitness of their peers (i.e., more successful agents are larger and look healthier, etc.). More specifically, the RAIL strategy, where agents imitate only successful behavior, assumes that agents can detect significant changes in the fitness of their peers during their lifetime or identify specific activities that may contribute to their fitness. The model presented in Section 3 can provide a framework to explore ways in which these assumptions can be relaxed. Coding the imitative behavior patterns themselves into the genome might result in the spontaneous emergence of imitative behavior in a population of agents. Behavior patterns that can be coded may include attributes such as the imitation model selection scheme, imitation strategy, imitation period, etc. Our model can also be extended to study the incentive that should be provided to an agent to make it assume the role of a teacher. Teaching, or even allowing someone else to imitate one's actions is, by definition, an altruistic behavior, and might have various costs associated with it. We wish to explore the conditions which may lead to the emergence of active teaching even in the presence of a fitness penalty for such a behavior. Such favorable teaching conditions may arise when the fitness associated with various actions is correlated with the frequency of these actions in the population (see also Boyd and Richerson, 1985, for a discussion of frequency-dependent bias). A good example of this case can be found in the emergence of normative behaviors (Axelrod, 1986; Flentge et al., 2001). Since the IEE model presented here entails a relatively simple form of cultural transmission, confined to horizontal transmission of innate behaviors, it can serve as a solid testbed for future studies of the emergence, evolution and prevalence of imitation.

7 Summary

Our study focuses on the effects of imitation on the evolution of agents in the absence of cultural evolution. We show that introducing the adaptive mechanism of lifetime learning by imitation can significantly enhance the evolutionary processes, resulting in better performing agents. This paradigm is particulary useful in evolutionary simulations of autonomous agents, when conventional supervised learning is not possible. Our model can serve as a theoretical and experimental framework to further explore central issues concerning the interaction between imitation, learning and evolution.

Acknowledgements

We wish to thanks the anonymous reviewers for their helpful comments. We are grateful to Daniel Polani for his valuable insights.

References

- Ackley, D. and Littman, M. (1994). Altruism in the evolution of communication. In Brooks, R. and Maes, P., editors, *Artificial Life IV: Proceedings of the International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press.
- Arbib, M. (2002). The mirror system, imitation, and the evolution of language. In Dautenhahn, K. and Nehaniv, C., editors, *Imitation in Animals and Artifacts*. The MIT Press.
- Axelrod, R. (1986). An evolutionary approach to norms. *American Political Science Review*, 80(4):1095–1111.
- Baldwin, J. (1896). A new factor in evolution. American Naturalist, (30):441-451.
- Best, M. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 7(3/4):289–306.
- Billard, A. and Dautenhahn, K. (1999). Experiments in learning by imitation: grounding and use of communication in robotic agents. *Adaptive Behavior*, 7(3/4):411–434.
- Boyd, R. and Richerson, P. (1985). *Culture and the evolutionary process*. The University of Chicago Press, Chicago.
- Cavalli-Sforza, L. and Feldman, M. (1981). *Cultural transmission and evolution: a quantitative approach*. Princeton University Press.
- Dautenhahn, K. and Nehaniv, C. (2002). *Imitation in Animals and Artifacts*. MIT Press, Cambridge, Mass., USA.
- Fahlman, S. E. (1989). Faster-learning variations on back-propagation: An empirical study. In *Proceedings of the 1988 Connectionist Models Summer School*, Los Altos, CA. Morgan-Kaufmann.
- Feldman, M. and Laland, K. (1996). Gene-culture coevolutionary theory. Trends in Ecology and Evolution, 11(11):453–457.
- Flentge, F., Polani, D., and Uthmann, T. (2001). Modelling the emergence of possession norms using memes. *Journal of Artificial Societies and Social Simulation*, 4(4).
- Floreano, D. and Mondada, F. (1996). Evolution of plastic neurocontrollers for situated agents. In Maes, P., Mataric, M., Mayer, J., Pollack, J., and Wilson, S., editors, *From Animals to Animates*, volume IV, Cambridge, MA. MIT Press.
- Hayes, G. and Demiris, J. (1994). A robot controller using learning by imitation. In *Proceedings of the 2nd International Symposium on Intelligent Robotic Systems*.
- Hertz, J., Krogh, A., and Palmer, R. (1991). *Introduction to the theory of neural computation*. Santa Fe Institute.

Borenstein and Ruppin

- Hinton, G. and Nowlan, S. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.
- Kawamura, S. (1963). The process of sub-culture propagation among japanese macaques. In Southwick, C. and Nostrand, V., editors, *Primates Social Behaviour*, pages 82–90, New York.
- Kirby, S. and Hurford, J. (1997). Learning, culture and evolution in the origin of linguistic constraints. In Husbands, P. and Harvey, I., editors, *4th European Conference on Artificial Life*, volume IV, pages 493–502, Cambridge, MA. MIT Press.
- Laland, K. (1992). A theoretical investigation of the role of social transmission in evolution. *Ethology and Sociobiology*, 13(2):87–113.
- Mitchell, M. (1996). An introduction to genetic algorithms. MIT Press.
- Nolfi, S., Elman, J., and Parisi, D. (1994). Learning and evolution in neural networks. *Adaptive Behavior*, 1(3):495–502.
- Nolfi, S. and Floreano, D. (1999). Learning and evolution. *Autonomous Robots*, 7(1):89–113.
- Nolfi, S. and Parisi, D. (1997). Learning to adapt to changing environment in evolving neural networks. *Adaptive Behavior*, 1:99–105.
- Ruppin, E. (2002). Evolutionary autonomous agents: A neuroscience perspective. *Nature Reviews Neuroscience*, 3(2):132–141.
- Whiten, A. and Ham, R. (1992). On the nature and evolution of imitation in the animal kingdom: Reappraisal of a century of research. In Slater, P., Rosenblatt, J., Beer, C., and Milinski, M., editors, *Advances in the Study of Behavior*, pages 239–283, San Diego, CA. Academic Press.
- Zentall, T. (2001). Imitation in animals: evidence, function, and mechanisms. *Cybernetics and Systems*, 32(1-2):53–96.

AISB Journal

Creativity in Arts and Science

Simon Colton* and Pablo Gervás[†]

- * Department of Computing, Imperial College, London, 180 Queen's Gate, South Kensington, London SW7 2AZ, UK, *sgc@doc.ic.ac.uk*
- [†] Departamento de Sistemas Informáticos y Programación, Facultad de Informática, Universidad Complutense de Madrid, 28040 Madrid, Spain, *pgervas@sip.ucm.es*

Editors' Introduction

The fifth AISB symposium on creativity in AI and Cognitive Science was intended to bring together researchers in all aspects of creativity research within AI and cognitive science, an area that is flourishing and progressively attaining respectability. Along with philosophical discussions about the nature of creativity in both arts and sciences, it included the presentation of research discussing implemented systems and frameworks within which implementations can be undertaken, assessed and applied. The research presented covered philosophical approaches to machine creativity, including a study of epistemic and ethical mediators and the functioning of tropic communication. Frameworks for machine creativity were also covered with papers on optimality principles for conceptual blending, and the use of concept maps for natural language generation. As usual, computer music was well represented at the symposium, with papers discussing macrostructure discovery for musical signals, driving chord progressions using tension, fugal exposition composition, improving the prediction of PPM variants, music generation from statistical models and an approach to improvisation using multi-swarms. Finally, computer art motivated by human perception was represented with a paper on symbolic pointillism."

AISB Journal

Optimality Principles for Conceptual Blending: A First Computational Approach

Francisco C. Pereira and Amílcar Cardoso

Centro de Informática e Sistemas da Universidade de Coimbra (CISUC) Dep. Engenharia Informática, Pólo II, Pinhal de Marrocos, 3030 Coimbra *camara@dei.uc.pt*; *amilcar@dei.uc.pt*

Abstract

We propose an implementation of the eight Optimality Principles from the framework of Conceptual Blending, as presented by (Fauconnier and Turner, 1998). Conceptual Blending explains several cognitive phenomena in the light of the integration of knowledge from different *mental spaces* onto a single *mental space*: the Blend. The Optimality Principles express general pressures that compete in the generation of the Blend.

The work we present now corresponds to the Constraints module of our computational model of Conceptual Blending, also described in other papers.

1 Introduction

One big challenge for Computational Creativity is the generation of new concepts and a very interesting source of inspiration for approaching this issue comes from the framework of Conceptual Blending, from (Fauconnier and Turner, 1998). The Conceptual Blending (CB) framework suggests processes and principles that may be at work within the cognitive integrations of knowledge needed in reasoning. We are developing a model of computational creativity that takes an approach to CB as a fundamental pilar. This system has been described in earlier papers ((Pereira and Cardoso, 2001), (Pereira and Cardoso, 2003)) as well as its first experiments ((Pereira and Cardoso, 2002)), where the need for developing the Optimality Principles of Conceptual Blending was evident as an urgent improvement. These are the general guidelines that drive the process of blending and allow the differentiation between a "good" and a "bad" blend. Here, we propose a formal realization for each of the eight principles. Experimentations around these principles have already been published in (Pereira and Cardoso, 2003). Although in the present paper we provide some results of these experiments, we strongly advise the interested reader to explore this previous publication.

In the first sections, we give an overview of Conceptual Blending, so as to provide the reader with the motivation and background for our approaches to the Optimality Principles. We made some exploratory experiments which are informally reported at the end of the discussion of each principle. We finalize this paper with reflections regarding the relevant aspects of this work for Computational Creativity.

2 Conceptual Blending

Conceptual Blending (CB) was initially proposed by (Fauconnier and Turner, 1998) as part of a major framework concerning cognition and language and had the role of explaining the integration of knowledge coming from distinct sources onto a single, independent and coherent unit, the Blend. A blend is a concept or web of concepts whose existence and identity, although attached to the pieces of knowledge that participated in its generation (the inputs), conquers gradual independence through time and use.

We find examples of blends in many sorts of situations. A blend can be an effective way to get attention and curiosity towards advertising a product (e.g. Sony's AIBO robot uses all sorts of Sony products behaving as if it were a real human) or spreading a message (e.g. the Marlboro cowboy with impotence problems). People have been making blends with creatures from the times of Greek mythology (e.g. *pegasus*) till today (e.g. *the pokemons*), natural language discourse (e.g. "John digested the book", "Sue sneezed the napkin off the table"), poetry (see (Freeman, 1999)). Many more examples and situations could be listed and studied in detail, demonstrating the ubiquity of CB as observed by these researchers.

It is noticeable that Conceptual Blending and its research community are growing, and possibly still in its early stages. It is an elegant proposal for modelling a creative process and its relationships with language and cognition, but it carries formal vagueness across its several aspects, making it difficult even to be considered as a theory, in a *Popperian* sense. Indeed, if the big hole between the general model and the specific examples does not show any incoherence, it also does not allow falsifiability, leaving very much undefined the boundary between what is and what is not a blend. These criticisms intend to motivate work that, from our point of view, is fundamental. They also support our own motivation for the present project, which is that of contributing with a formal model and implementation based on Conceptual Blending. We hope it can be useful in shedding some light on a number of issues.

In order to understand the CB framework, we must introduce a fundamental concept: the *mental space*. According to (Fauconnier and Turner, 1998), mental spaces are partial structures that proliferate when we think and talk, allowing a fine-grained partitioning of our discourse and knowledge structures. For simplifying, let us consider a mental space as a partial selection of knowledge from a domain, a memory of a situation, an imagined scenario or entity, essentially a knowledge structure of inter-related concepts that is explicitly or implicitly necessary for a reasoning. As AI researchers, we see a mental space representable as a *semantic network*, a *frame*, a *case* or any other symbolic knowledge structure that gathers a set of inter-related concepts towards a specific situation. In the *four space model*, Conceptual Blending is described as involving two input mental spaces that, recurring to a cross-space mapping between them and a generic space that has general knowledge relevant for both the input domains, will generate a fourth one, called Blend. This new domain will maintain partial structure from the input domains and have emergent structure of its own.

The process of generation of a blend can be summarized according to three general steps ((Fauconnier, 1997)): Composition, where new relations become available that did not exist in the separate inputs; Completion, when generic knowledge is projected into the blend, to "complete" the emergent structure; and Elaboration, in which cognitive work is performed in the blend, according to its own emergent logic. The order of these steps may be changed and several iterations of the process may be necessary. There is a set of *governing principles*, the *Optimality Pressures*, that should drive the process of generating a "good blend" ((Fauconnier and Turner, 2002)):



Figure 1: The Conceptual Blending Model

- Integration The blend must constitute a tightly integrated scene that can be manipulated as a unit. More generally, every space in the blend structure should have integration.
- Pattern Completion Other things being equal, complete elements in the blend by using existing integrated patterns as additional inputs. Other things being equal, use a completing frame that has relations that can be the compressed versions of the important outer-space vital relations between the inputs.
- Topology For any input space and any element in that space projected into the blend, it is optimal for the relations of the element in the blend to match the relations of its counterpart.
- Maximization of Vital Relations Other things being equal, maximize the vital relations in the network. In particular, maximize the vital relations in the blended space and reflect them in outer-space vital relations. Turner and Fauconnier identify 15 different vital relations: change, identity, time, space, cause-effect, part-whole, representation, role, analogy, disanalogy, property, similarity, category, intentionality and uniqueness.
- Intensification of Vital Relations Other things being equal, intensify vital relations.
- Web Manipulating the blend as a unit must maintain the web of appropriate connections to the input spaces easily and without additional surveillance or computation.
- Unpacking The blend alone must enable the understander to unpack the blend to reconstruct the inputs, the cross-space mapping, the generic space, and the network of connections between all these spaces
- Relevance Other things being equal, an element in the blend should have relevance, including relevance for establishing links to other spaces and for running the blend. Conversely, an outer-space relation between the inputs that is important for the purpose of the network should have a corresponding compression in the blend.

As far as we know, there is no work yet towards an objective study of the optimality pressures, measuring examples of blends or specifying these principles in detail. This, we believe, disturbs considerably the appreciation and application of Conceptual Blending in scientific research, and therefore testing and specifying a formal proposal for the optimality pressures is also a particular motivation for this work.

3 Basic Notions from our model

We present now some basic notions that are necessary to understand this document. We use a specific type of mental space, that is static, modeless and non attached to discourse, which is closer to the notion of *domain knowledge* in AI. We call it a *domain*, comprising a *theory* and a set of *instances*. In this paper, we consider only the theory, which is represented by a *concept map* that explains the structural and causal organization of the domain. A concept map is a semantic network, with binary directed relations between the concepts¹. In CB, there is a set of vital relations that take a special role in the Blending process. We also consider these relations, but allow the choice of a different set of relations. In principle, these vital relations can be the source for establishing *mappings* between the input spaces, fundamental for the projection operation (two objects mapped to each other can be projected to the same concept, e.g. horse and bird projected to pegasus). Each mapping projection consists of a ternary relation m/3. For example, a mapping algorithm based on "property" could try to map pairs of concepts that make a valid pair of property/object (e.g. "m(property, dark, bird)"), a different one based on analogy could link pairs of analogical counterparts (e.g. "m(analogy, leg, wing)"). As far as our research goes, we are applying mapping algorithms that find 1-to-1 structures mapping between the concept maps of the two domains based on analogy, identity, space and time². Knowing this is a strong limitation, we hope to address other mapping algorithms in future developments. Currently, for the sake of validation and experimentation of the system, we allow user-defined mappings, so as to allow conclusions independent from the mapping choice.

Finally, another important notion is that of *frames*. A frame comprises a set of concepts and relations that should be tightly integrated according to a situation, structure, cause-effect or any other relation that ties a set of concepts onto one, more abstract or broad, composite concept. We envisage different kinds of frames, in terms of level of abstraction. The more specific ones correspond to concepts easily identifiable as familiar objects or situation (which we can see as a kind of *prototypes*). For example, the (very much simplified) frame of "transport means" corresponds to a set of concepts and relations that, when connected together, represent something that has a container and a subpart (e.g. an engine) that serves for locomotion.

 $\begin{array}{rcl} frame(transport_means(X)): & & & \\ have(X, container) \wedge \\ carrier(X, people) & \longleftarrow & have(X, Y) \wedge \\ & & & purpose(Y, locomotion) \wedge \\ & & & \\ drive(_, X) \end{array}$

¹First order logic predicates with arity 2, e.g. isa(bird, aves); purpose(wing, fly))

²Analogy: the algorithm searches for a structure alignment between the two domains; Identity: equal or synonymous elements from the two domains get mapped together; Space: Spatial data (e.g. location) at the same level of detail are mapped together; Time: Temporal data (e.g. a date) at the same level are mapped together.

Pereira and Cardoso

The more abstract frames can consist of top-level decisions or directives that decide the underlying philosophy of the blend construction. For example, if a blend satisfies the "aframe" frame, it means it maintains the structure of the input domain 1. If it satisfies the "bprojection" frame, then the concepts of input space 2 all get projected unchanged to the blend (e.g. "bird" is projected to "bird", "wing" is projected to "wing"). If a blend satisfies "aframe" and "bprojection" simultaneously, then it should³ have the concepts from input space 2 organized according to the structure of input space 1. It is more complex to design these more abstract frames, thus we allow the use of programming (in *prolog* language, inside curly brackets) within a frame specification. Below we can see the programming of the "aframe".

 $\begin{array}{rl} frame(aframe(A)):\\ aframe(A,Blend) &\longleftarrow \\ \{stats(domain1,A), current_blend(Bl)\}, \land \\ \{findall(R/X/Y, (rel(A,XA,R,YA), pro-jection(Bl,A,XA,X), projection(Bl,A,YA,Y)), L)\} \land op(exists(L)) \end{array}$

Basically, "aframe" searches for all relations of domain 1 (A) and obtains a list (L) containing their projections to the blend. The op(exists(L)) condition is an *operator* that, when, interpreted by the frame processor, expands L into a set of concept map relations (like those in the "transport_means" frame). For "aframe" to be totally satisfied, it is necessary that all relations from input space 1 also exist in the blend.

There are also frames of intermediate level of abstraction, which aren't either as specific as "transport_means" or as abstract as "aframe". For example, "new creature" is concerned with finding a "creature" (thus having its specific properties - e.g. being a "living being") that didn't exist before in either domains or existed but not as a "creature" (e.g. a "flying snout").

When we say that a frame f is satisfied in the blend b, we mean all its premises are true in the domain b. We see frames as *information molds* and building a blend for a given situation should depend much on the choice of these structures, either being structures towards which the blend self-organizes or as pragmatic goals or query specifiers that the blend is expected to accomplish.

4 Optimality Principles

Following the F&T notion of Optimality Principles, the pressures that should lead towards stable, integrated new blends, we propose now a set of measures that should reflect as much as possible the rationale behind each principle. In order to give a clearer idea of its individual effect in our blending system, we present a brief report of experiments we made. These experiments consisted of running a parallel search method, a genetic algorithm, to retrieve blends from the search space. The input domains were the domains of *horse* and *bird* (see tables 1 and 2 - in the end of this paper), meaning that the expected results range from the unchanged copy of one (or both) of the concepts to a horse-bird (or bird-horse) which is a combination of selected features from the input domains. The construction of these domains was subject to the following constraints: they should be

³Depending on the mappings, the concepts considered in aframe and bprojection may become separate in the inputs (and so there wouldn't be any systematic relation between concepts from input spaces 1 and 2), yet this would receive little value in the measures presented in this paper.

concept networks in which nodes are concepts and arcs are relations; the concepts should be connected to the ontology in the generic space through an "isa" relation; the relations used should be present in the Generalized Upper Model (GUM) hierarchies ((Bateman et al., 1995)) or be subtypes of them. GUM is a general top-level ontology that has two hierarchies (elements and relations) that comprise abstract relations, properties, spatial relationships, among others. Although allowing a normalization of the concept maps, the constraints in the construction of the domains do not avoid, *per se*, the biasing or ingenious tailoring. For this reason, in this paper and in the exploratory experiments we show, we do not give special attention to a qualitative *reading* of the results or use them to demonstrate its validity, instead we are interested in reporting the effects that each measure has on the results and on the search landscape.

The generic domain (in tables 3 and 4, in the end of this paper) consists of a simple general ontology, a set of frames and integrity constraints. We applied 3 different mappings (figure 2), all generated automatically (via a structure alignment algorithm - read (Pereira and Cardoso, 2001) to know more). These mappings range from very small (only four mapping correspondence) to large (21 mapping correspondences), from non-surprising associations (e.g. "animal" and "animal") to nonsense (e.g. snout and lung). For each mapping, we tested the seven optimality pressures, each of these comprising 30 runs⁴.

			vegetable_food	\leftrightarrow	vegetable
			food	\leftrightarrow	food
ear	\leftrightarrow	wing	horse	\leftrightarrow	bird
snout	\leftrightarrow	bird	equidean	\leftrightarrow	aves
eve	\leftrightarrow	lung	animal	\leftrightarrow	animal
mouth	\leftrightarrow	feathers	human setting	\leftrightarrow	house
1110uun ງ		2	wildernoss		wildornoss
1	$\overline{\nabla}$	2	wildefiless	$\overline{\nabla}$	· ·
hear	\leftrightarrow	ffy	ruminant	\leftrightarrow	oviparous
			run	\leftrightarrow	fly
	1		cargo	\leftrightarrow	pet
			neigh	\leftrightarrow	chirp
			snout	\leftrightarrow	lung
			mane	\leftrightarrow	feathers
			tail	\leftrightarrow	beak
mouth	\leftrightarrow	beak	leg	\leftrightarrow	eye
snout	\leftrightarrow	bird	hoof	\leftrightarrow	wing
eye	\leftrightarrow	lung	4	\leftrightarrow	2
ear	\leftrightarrow	feathers	eye	\leftrightarrow	leg
eat	\leftrightarrow	eat	ear	\leftrightarrow	claw
			hear	\leftrightarrow	catch
	2		grass	\leftrightarrow	grass
				3	

Figure 2: The three mappings used in the experiments

 $^{^{4}\}mathrm{A}$ run is an entire evolutive cycle, from the initial population to the population in which the algorithm stopped

4.1 Integration

Frames have a natural integration role because they gather knowledge around abstractions, tightening the links between concepts. Assuming the set F of frames that are satisfied in a blend, we define the *frame coverage* of a domain to be the set of relations from its concept map that belong to the set of conditions of one or more frames in F. The larger the frame coverage of the blend, the more it is integrated. Yet, a blend that is covered by many frames should be less integrated than a frame with the same coverage, but with less frames. In other words, if a single frame covers all the relations of a blend, it should be valued with the maximal integration, whereas if it has different frames being satisfied and covering different sets of relations, it should be considered less integrated. The intuition behind this is that the unity around an integrating concept (the frame) reflects the unity of the domain. The Integration measure we propose varies according to this idea. It also takes integrity constraints into account so that, when a frame violates such a constraint, it is subject to penalty.

Definition 4.1 For a frame f with a set C of conditions $B_i \bigcup \neg B_i$, a blend b, with a concept map CM_b , a blendoid⁵ CM_{B^+} , the concept map of the blendoid and VI, the set of integrity constraints⁶ that are violated in the frame, the integration value, I_f is defined by:

$$I_f = \left(\frac{\#C}{\#CM_b} \times (1-\iota)^{\#VI}\right) \times \left(1 + \frac{\#CM_b}{\#CM_{B^+}}\right)/2$$

where ι is a penalty factor between 0 and 1, a value that penalizes a frame for each violation to integrity constraints. An integrity constraint is violated if its premises are true. In the context of the integration measure of frame f above, f violates integrity *ic* if the conditions Cic of *ic* are verified and $Cic \cap C \neq \emptyset$. In other words, f needs to violate *ic* in order to be integrated.

We would like to clarify the above formula: the first factor represents the ratio of coverage of b w.r.t. f; the second factor means that each integrity constraint violation implies an exponential discount; the third factor serves the purpose of maximizing the size of the blend (if two frames have the same ratio of coverage, the one that contains more relations should have higher integration); the division by 2 aims to normalize the result between 0 and 1.

While the value for a single frame integration is described above, the integration measure of a domain w.r.t. a set of frames is not necessarily straightforward. At a first sight, it is appealing to just sum the values of integration of all frames, or of the union of them, or even their intersection. But this would lead to wrong results, because a set of frames could not be reduced to a single frame from the point of view of integration. In this measure, we want to stimulate unity, coverage and take into account the strength of each frame individually. In terms of unity, we argue that the set of relations that make the "core" of all the frames that are satisfied, i.e. the intersection of the sets C of conditions of all frames, should be highly valued. On the other side, the coverage of this "core" will be smaller than the overall coverage (or equal, if the frames have equivalent C sets), which leads us to take into account the disjoint sets of relations of the frames. Finally, the integration of each individual frame (as defined above) should also be present in the overall measure. These last two issues (the overall coverage and the integration of individual frames) are

⁵The concept map that contains all the possible relations that the blend may have

⁶Rules with *false* conclusion as in table 3



Figure 3: The role of frame coverage in Integration value

subject to a disintegration factor because they reflect the existence of different, not totally intersected, frames. We propose this factor, α , to be a configurable value from the interval [1, 0]. It is now time to present our proposal for the *IntegrationMeasure* of a blend:

Definition 4.2 For a set of frames $f_i \in F_b$, where F_b is the set of the frames that have their conditions (C_i) satisfied in the blend b

$$Integration = I_{\bigcap_{0}^{i}C_{i}} + \alpha \times Uncoverage \times \sum_{0}^{i}I_{f_{i}}$$

The *Uncoverage* value consists of the ratio of relations that do not belong to the intersection of all frames w.r.t. the total number of relations considered in the frames:

$$Uncoverage = \frac{\# \bigcup_{0}^{i} C_{i} - \# \bigcap_{0}^{i} C_{i}}{\# \bigcup_{0}^{i} C_{i}}$$

We think the integration measure is a fundamental brick of the blending process. It leads the choice of the blend to something *recognizable* as a whole, fitting patterns that help to determine and understand what a *new* concept is.

In order to illustrate this reasoning, in figure 3, we show 4 blends and the respective frame coverage. Blend A gets clearly the highest Integration value (all the relations are covered by a single frame); B is also totally covered, but by two different frames; Blend C should get lower Integration value than B because it does not cover every relation (Uncoverage is bigger than 0); finally, blend D would possibly get the lowest value (depending on the value of α) because, although covering every relation, there is a high dispersion of frames.

Experiments: An immediate conclusion about the effect of Integration is that frames behave as *attractor* points in the search space. Moreover, the frames with a larger coverage tend to be preferred, although when too large (like *aprojection* or *aframe*) they are dropped away. The evolution is directed to a compromise of coverage and satisfiability. More specifically, when it satisfies a frame like *pw_based_explanation*, the resulting Integration value is a local maximum or a point in its neighborhood (because sometimes other, related, frames were also found) and "jumping" to another area of the search space becomes difficult.

Another conclusion to draw concerns to the observation that the complexity of the search space landscape grows with mapping size. In fact, when we have a mapping of size 2, the algorithm only finds two different solutions and the better rated (possibly a global maximum) is achieved in 77% of the runs, but with a mapping of size 5, it returns six different blends, with the best choice retrieved only 43% of the times. To confirm this conclusion, the mapping of size 21 led the algorithm to 16 different maxima, with

Pereira and Cardoso

the best one found only 7% of the times. A good compensation for this apparent loss of control is that the returned values are clearly higher (0.68, for the best) than in the small mappings (0.22), meaning that with big mappings there are many more possibilities to find integrated blends.

The resulting concept maps consist of exactly the relations that are covered by the satisfied frame or combination of them, more specifically there were two frames that were very persistent: *pw_based_explanation* and *purposeful_subpart*. In some blends, the former was present multiple times (e.g. a part-whole explanation of a horse, with part-whole explanations of its subparts) and both were plenty of times combined.

4.2 Topology

The Topology optimality pressure brings *inertia* to the blending process. It is the constraint that drives against change in the concepts because, in order to maintain the same topological configuration as in the inputs, the blend should maintain exactly the same neighborhood relationships between every concept, ending up being a projected copy of the inputs. In *real blends*, this pressure is normally one that is disrespected without big loss in the value of the blend. This is due to the *imagination* context that normally involves blends, i.e., novel associations are more tolerable.

In our Topology measure, we follow the principle that, if a pair of concepts, x and y, is associated in the blend by a relation r, then the same relation must exist in the inputs between the elements from which x and y were projected. We say that such relation, r(x, y), is *topologically correct*. Thus, the value of Topology corresponds to the ratio of topologically correct relations in the concept map of the blend.

Definition 4.3 For a set $TC \subseteq CM_b$ of topologically correct relations, defined as

$$TC = \{r(x, y) : r(x, y) \in CM_1 \cup CM_2)\}\$$

where CM_1 and CM_2 correspond to the concept maps of inputs 1 and 2, respectively⁷. The topology measure is calculated by the ratio:

$$Topology = \frac{\#TC}{\#CM_b}$$

Intuitively, this measure represents the amount of relations from the inputs that got projected onto the topologically equivalent position in the blend. At the moment, the only way to violate topology is by having a pair of concepts projected to the same one (e.g. "horse" and "bird" projected to "horse"), bringing a new relation that was exclusive to one of the domains (e.g. ability(bird, fly) projects to ability(horse, fly)). Topology thus decreases as fusion or transfer projections are made.

Experiments: In all the experiments with Topology, the final results were valued 100%, meaning that this constraint is easily fully accomplished, independently of the mapping. An interesting fact is that there is a multitude of solutions in the *search land-scape* of Topology, showed by the amount of different final results in each mapping. Intuitively, and observing the short duration of each run, this means that, wherever the search starts, there is always a Topology optimal point in the neighborhood.

Topology is more an inertia than a transformation force because it values knowledge that remains the same. In our horse-birds, this pressure projects wings, beaks and claws

⁷In other words, TC is the intersection of the concepts maps of the blend and the input spaces

(i.e. concepts from the bird domain) to the blend but isolates them unless there is strong evidence to connect to horses, legs and snouts (i.e. concepts from the horse domain).

4.3 Pattern Completion

The Pattern Completion pressure brings the influence of patterns, being them present in the *inputs* or coming from the *generic* space. Sometimes a concept (or a set of concepts) may seem incomplete but making sense when "matched against" a pattern.

At present, in the context of this work, a pattern is described by a frame, i.e. we do not distinguish these two concepts, and therefore pattern completion is basically frame completion. Here, as in the definition of this principle, the completing knowledge becomes available from "outside", not as a result of projection. This means that the act of completing a frame consists of asserting the truth of the ungrounded premises, a process that happens only after a sufficient number of premises is true. We call this the *evidence threshold*. The evidence threshold e of a frame f_i with regard to a blend b is calculated according to the following.

$$e(f_i, b) = \frac{\#Sat_i}{\#C_i} \times (1 - \iota)^{\#VI}$$

where Sat_i contains the conditions of each f_i that are satisfied in b, ι is the integrity constraint violation factor and VI the set of violated integrity constraints.

As in the integration pressure, we have the problem of taking into account multiple frames. This time, given that we are evaluating possible completion of subsets of relations, instead of sets of relations that are actually verified in the domain, it is difficult to find such a linear rationale (e.g. would two patterns each with individual completion x value higher than three having each slightly less than x?). As a result, we propose to find the union of the patterns and then estimate its own evidence threshold:

Definition 4.4 *The Pattern Completion measure of a blend b with regard to a set F with n frames is calculated by*

$$PatternCompletion = e(\cup_{0}^{n} f_{i}, b)$$

This measure has a very important role in increasing the potential of the blend, for it brings the "seeds" that may be used in the Completion and Elaboration phases. In figure 4, we illustrate Pattern Completion with two examples. Assuming a frame with three conditions, on the left it has an evidence threshold of 66,6% (two relations are already accomplished), whereas on the right the evidence threshold is only 33,3%. For both, since we consider only one frame (i.e. one *pattern*), the value of Pattern Completion is the same as of the evidence threshold.

Experiments: The first conclusion to take from the experiments with Pattern Completion is that the size of the resulting concept maps tend to grow as the evolution progresses, although there is no linear correlation found. This has a simple interpretation, given that this measure stimulates the appearance of patterns (frames) that are only partially completed. In doing so, it drives the blend to partially complete (i.e., instantiate partially its conditions) the highest possible number of frames, leading, in each case, to several sets of relations that fit into those frames without satisfying them (e.g. wings are projected, they serve to fly, but they are not attached to anything).

In which respects to the *search landscape*, it seems to be very rich in local maxima. This fact was not unexpected, considering the discussion of the previous paragraphs, the



Figure 4: Pattern Completion examples

number of different frames available (9) and all their different combinations. The most constant results came from mapping 2, with 13% of the best result obtained and 20% of the second best. An interesting remark is that the resulting local maxima always fall within a very strict range of values (of maximum amplitude 0.11).

4.4 Maximization of Vital Relations

For the maximization of vital relations, we estimate the impact of the vital relations to the blend calculated by the ratio of vital relations w.r.t. the whole set of possible vital relations, the *blendoid*. The blendoid is the largest possible blend that can be obtained from a given mapping and is calculated by projecting every concept to it (i.e., there is no selective projection) regardless of integrity constraint violation or any other constraint. Since it has the largest set of potential relations, it also has the maximum possible of vital relations.

Definition 4.5 Let Υ be a set of vital relations. From the concept map of the blend b, we may obtain the set of vital relations in b, B_{VR} :

$$B_{VR} = \{r(x, y) : r(x, y) \in CM_b \land r \in \Upsilon\}$$

From the blendoid (the largest possible blend), B^+ , we have B_{VB}^+ :

$$B_{VR}^+ = \{ r(x, y) : r(x, y) \in CM_B^+ \land r \in \Upsilon \}$$

Finally, the Maximization of Vital Relations measure is calculated by the ratio

$$Maximization_VR = \frac{\#B_{VR}}{\#B_{VR}^+}$$

Experiments: The influence of Maximization of Vital Relations in the results is straightforward, given that its highest value (1) reflects the presence, in the blend, of all the vital relations that exist in the inputs, independently of the projections of the concepts or non-vital relations (which become *noise* in the sense that these appear randomly and making no difference to the value of the measure, yet confusing the "reading" of the concept map). As the evolution goes on in each run, the value grows until reaching the maximum reasonably early. For each set of 30 runs, it reached the value 1 a minimum of 93% of the times, and the remaining times achieved at least a value of 0.95. As in Topology, the search space of Maximization of Vital Relations is very *simple* since there is a global maximum in the neighborhood of (almost) every point. On the other hand, since there is no control on *noise*, the resulting concept maps show unity only in the subset of vital relations.

4.5 Intensification of Vital Relations

Intensification of Vital Relations is the principle that maximizes the concentration around a specific vital relation. I.e., while the Maximization of Vital Relations favors the creation in the blend of vital relations in general as opposed to "regular" relations, Intensification is based on focussing a specific vital relation. The former relates "new" vital relations with "new" relations in the blend; the latter relates vital relations with themselves. Thus, we need a notion of "intensity" of a vital relation. For such, we argue that a vital relation is considered more "intense" when there is more evidence of its strength. This evidence should be dependent on the kind of vital relation we are dealing with. For example, an "analogy" vital relation between two concepts is stronger when there is also a systematical association between the neighborhood concepts (the systematicity principle). In fact, systematicity is the only "intensity" heuristic we have now and its calculation ($Int_{analogy}$) is straightforward. For a mapping (of size n)

$$Int_{analogy} = \frac{\#analogical_transfers}{n}$$

where an *analogical transfer* consists of a projection of a concept x from input space 1 to y in the blend (where y is the analogical counterpart of x in the input space 2).

Considering several different "intensity" heuristics, the evaluation of this pressure takes the point of view that a blend that applies only one vital relation, with intensity x, should have higher measure than a blend with n vital relations, each with intensity x/n (the sum would thus be x). So we want to favor "concentration" of vital relations.

Definition 4.6 Let $v_i \in \Upsilon$ be a vital relation and the set VRB_{v_i} , of the instances of vital relation v_i in the blend, defined by $VRB_{v_i} = \{v_i(x, y) : v_i(x, y) \in CM_b\}$

Assuming a value Int_{v_i} of intensity of the vital relation v_i , the measure of Intensification of Vital Relations is calculated by:

$$Intensification_VR = \frac{\sum_{0}^{n} Int_{v_{i}}^{2}}{(\sum_{0}^{n} Int_{v_{i}})^{2}}, \ n = \#\Upsilon$$

Intensification is thus higher when there is more "concentration" (e.g. $Int_{V_1} = 2$, $Int_{V_2} = 2 \Rightarrow Intensification = 8/16$; $Int_{V_1} = 4 \Rightarrow Intensification = 1$). In figure 5, we intend to illustrate the reasoning behind this measure. On the left, we see that the only mapping function (vital relation) used in building the projections of A and B is *analogy*, whereas on the right there are two kinds of connections (*analogy* and *property*). Assuming that the number of mapped concepts of both examples remains the same, the one on the left gets a higher *Intensification_V R* value because every mapping contributes to the intensity of a single vital relation, whereas on the right, we see that intensity is spread over two different relations. As a result, the former is more "concentrated" than the latter, yielding a bigger value.

Experiments: The behavior of Intensification of Vital Relations is similar to that of Maximization in the sense that the search landscape is not very complex. From wherever the search starts, there is a high probability (of at least 70%, in the worst case) of finding the maximum value in the neighborhood. It is important to point, though, that we are only applying one heuristic (analogical transfer), and so the results could not be different. Given this fact, we must say that we cannot claim or discuss much about this Intensification of V.R. proposal unless we find other heuristics than based on analogy systematicity.



Figure 5: Intensification VR examples

4.6 Unpacking

Unpacking is the ability to reconstruct the whole process starting from the blend. In our view, such achievement underlies the ability to reconstruct the input spaces, specifically. The complete reconstruction of the input spaces from the blend would demand the assessment of the cross-space mappings, the generic space and other connections. However, Unpacking should take the point of view of the "blend reader", i.e., someone or something that is not aware of the process of generation, thus not having access to the actual projections. This "reader" should look for patterns pointing to the "original" concepts.

Thus, what we are proposing is that Unpacking can be reduced to the ability to reconstruct the inputs. This is so because there is no way to properly reconstruct the inputs without a reconstruction of the cross-space mappings, generic space and the connections between spaces.

Once again we use the idea of *frames*, more specifically the *defining frame* of a concept, which comprises its immediately surrounding concepts and relations. For example, if the concept "wing" were projected onto x in blend, the defining frame with regard to the "bird" domain would consist of purpose(x, fly), *conditional*(x, fly), *quantity*(x, 2) and pw(x, bird). The more of these relations are found in the blend, the more likely it is that the "reader" will find easy to understand the relationship between x and "wing".

Definition 4.7 Given a blend b and an input space d, the concept x (which is the projection of the element x^d of input space d to the b), has a defining frame $f_{x,d}$ in d consisting of

$$f_{x,d} = C_0, C_1 \dots C_n \longrightarrow true$$

where $C_i \in \{r(x, y) : r(x^d, y) \in CM_d\}.$

Assuming that k is the number of conditions (C_i) of $f_{x,d}$ that are satisfied in the blend, the unpacking value of x with regard to d (represented as $\xi(x,d)$) is

$$\xi(x,d) = \frac{k}{n}$$

We calculate the *total estimated unpacking value* of x corresponding to the average of the unpacking values with regard to the input spaces. Thus, having input spaces 1 and 2, we have

$$\xi(x) = \frac{\xi(x,1) + \xi(x,2)}{2}$$

Definition 4.8 Let X be the set of n concepts of the blend b, generated from input spaces 1 and 2. The Unpacking value of b is calculated by

$$Unpacking = \frac{\sum_{i=0}^{n} \xi(x_i)}{n}, x_i \in \mathcal{X}$$



Figure 6: Unpacking examples

In Figure 6, we present the defining frame for "horse", in the "Horse" domain. In Blend 1, the concept "horse-bird" (the projection of "horse") will have the highest Unpacking value because it fits exactly its defining frame. In Blend 2, the value is lower because there are two new relations (with "fly" and "wings"), meaning it is not the exact same concept. Blend 3 will get the lowest Unpacking value of all three because it also lacks some relations (e.g. with "run" and "grass").

Experiments: The results of the Unpacking measure show that it has a notorious side effect on the size of the blend, since it drives it to very small sets (between 0 and 5) of relations. The interpretation here is straightforward: the ratio of *unpackable* concepts is highly penalized in bigger sets because of the projected relations that come as side effect of the projection of (*unpackable* or not) concepts. These relations *confuse* the unpacking algorithm so that it leads the evolution to gradually select the smaller results.

The maxima points also correspond to the value 1, but it seems, from the experiments, that there is a very limited set of such individuals, achieved in the majority (at least 93% for each mapping) of the experiments.

4.7 Web

The Web principle concerns to being able to "run" the blend without cutting the connections to the inputs. In our opinion, this is not an independent principle, it is co-related to those of Topology and Unpacking because the former brings a straightforward way to "maintain the web of appropriate connections to the input spaces easily and without additional surveillance or computation" and the latter measures exactly the work needed to reconstruct the inputs from the blend. It is not to say that Web is the same as Topology or Unpacking, what we are arguing is that, on one side, Topology provides a pressure to maintain the most fundamental connection to the input: the same structure; on the other side, Unpacking evaluates the easiness of reestablishing the links to the inputs. These two values combined in a weighted sum yield, we propose, an estimation of the strength of the web of connections to the inputs:

$$Web = \alpha \times Topology + \beta \times Unpacking$$

with $\alpha + \beta = 1$.

Since this is not an independent variable, making independent experiments with the Web measure would not add any valuable conclusion. In a subsequent publication, we plan to focus on correlation of measures, where we may explore the behavior of this measure.

4.8 Relevance

The notion of "relevance" or "good reason" for a blend is tied to the context and goal of the blending generation. A blend, or a part of it, may be more or less relevant dependent of what it is for. Once again, frames take a fundamental role as "context" specifiers, (i.e., the set of constraints within a frame describe the context upon which the frame is fulfilled). Therefore, having a set of goal frames, which could be selected from any of the existent domains or specified externally, a blend gets the maximum Relevance value if it is able to satisfy all of them.

An aspect of the goal frames is that they allow the application of queries. For example, if we want to find a concept that "flies", we could build a goal frame with the relation ability(x, fly). The blends that satisfy this frame would have high relevance.

Definition 4.9 Assuming a set of goal frames, F_g , the set F_b of the satisfied frames of blend b and the value PCN_F for the pattern completion of a set of frames F in blend b, as described in section 4.3, we have

$$Relevance = \frac{\#(F_g \cap F_b) + \#F_u \times PCN_{F_u}}{\#F_g}$$

where F_u , the set of unsatisfied goal frames, consists of $F_u = F_g - F_b$.

From the point of view of creativity, we propose the use of Relevance as a "usefulness" measure, as we already made in some experiments (Pereira, 2003).

Experiments: The first part of the test on Relevance focussed on making a single relation query. In this case, we asked for "something that flies" (ability($_$, fly)). The results were straightforward in any mapping, accomplishing the maximum value (1) in 100% of the runs, although the resulting concept maps did not reveal necessarily any overall *quality* or unity. In other words, the evolution took only two steps: when no individual has a relation "ability($_$, fly)", therefore with value 0; when a relation "ability($_$, fly)" is found, yielding a value 1, independently of the rest of the concept map.

The second part of the test on Relevance, by adding a frame (*ability explanation*) to the query, revealed similar conclusions. There was no sufficient knowledge in any of the input domains to satisfy this new frame completely, so the algorithm searched for the maximum satisfaction and reached it 100% of times in every mapping. So the *landscape* seems to have one single global and no local maxima, reflecting the integration of the two parts of the query. If there were separate frames, local maxima might be expected. Intuitively, the *search landscapes* of Integration and Relevance seem to be similar.

5 Discussion and further work

A first conclusion we draw from this work is that the eight optimality principles can be reduced to seven (since Web is not independent). Even more, given the power of language that we use in frames, some of the principles can be *coded* within a frame, namely Topology and Unpacking, and accomplished via a query measured by Relevance. This reduces our number to five. Yet, we do not know whether this reduction reflects a fault in the CB framework or in our interpretation of it.

It is important to say that the system is being already tested and some the experiments were published (e.g. in (Pereira, 2003), we used the same set of nouns used in the C^{3} Conceptual Combination framework (Costello and Keane, 2000) to generate noun-noun combinations), the general conclusions confirming the same behavior we showed in this paper.

Regarding the creative aspects of this system, there are some issues we may discuss concerning the role of the optimality principles. Since different configurations may lead to considerably different results, an obvious question emerges: Is there any "specific" configuration (or configurations) that produces more creative results? If we assume a "creative result' as something that is novel (i.e., not existent in the list of previous results) and useful (according to a context), as a simple definition, we may at the least try to understand what configurations bring better results, and which optimality principles are more influential. From the experiments we have already taken, we can say that there are no precise patterns for "creative results". Although the individual behavior regarding each of these principles remained faithful to what we presented here, their relative weight in the the value of a blend may vary considerably. For example, in some experiments with classical examples of Blending literature, in which we sought for the best weight configurations that the system needed to achieve the "correct blend", we observed a variability in the importance of the principles (e.g. in the "Buddhist Monk" example, Topology and Relevance were central; in "Kant Debate", it was Relevance and Pattern Completion). However, it has been clear that usefulness is very much dependent on Relevance and, at a lesser degree, Integration. Novelty seems to be related to how strong an optimality principle is in a given context. For example, if Topology is important for the generation of a well known kind of blends (e.g. creatures), then diminishing its weight on the configuration will bring novel results more often. Naturally, these also tend to be less useful. This competition is a possible justification for the variability we have been observing in experiments so far, given that different kinds of blends will behave differently with the optimality principles.

We have already stated that the main motivation of our system is to generate new concepts out of previous knowledge. Two works under development include the application of the system to a game environment and the study of ways to increase the emergence of knowledge within the blend. The latter will certainly improve the creativity potential of the system since it consists of the generation of new relations and concepts (not present in the inputs) within the blend. Since this new knowledge will come from outside (from the generic space and other domains in the knowledge base), we expect it will open the search space for the transfer of knowledge from distant sources.

Acknowledgements

We would like to thank the reviewers for their constructive comments and people with whom we have been discussing some of the ideas presented in this paper, namely Pablo Gervás, Margaret Freeman, Tony Veale and Penousal Machado.

Pereira and Cardoso

Appendix

isa(horse,equinae) isa(equinae,mammal) existence(horse, farm) existence(horse, wilderness) pw(snout, horse) pw(mane, horse) pw(tail, horse) quantity(hoof, 4) pw(eye, snout) pw(ear, snout) pw(mouth,snout) motion_process(horse,walk)	pw(leg, horse) purpose(leg, stand) pw(hoof, leg) purpose(horse, traction) eat(horse, grass) ability(horse, run) carrier(horse, human) quantity(leg, 4) quantity(eye, 2) quantity(ear, 2) purpose(eye, see) ride(human, horse)	purpose(horse, food) sound(horse, neigh) purpose(mouth, eat) purpose(ear, hear) color(mane, dark) size(mane, long) material(mane, hair)
motion_process(horse,walk)	ride(human, horse)	
taxonomicq(horse, ruminant)	purpose(horse, cargo)	

Table 1: The domain theory of *horse*

isa(bird,aves)	existence(bird, house)	isa(parrot, bird)
isa(aves,oviparous)	purpose(bird, pet)	isa(nest, container)
lay(oviparous, egg)	existence(bird, wilderness)	role_playing(bird, freedom)
purpose(bird, food)	purpose(eye, see)	ability(parrot, speak)
smaller_than(bird, human)	purpose(beak, chirp)	purpose(claw, catch)
pw(lung, bird)	motion_process(bird, fly)	purpose(wing, fly)
purpose(lung, breathe)	quantity(eye, 2)	pw(claw, leg)
isa(paradise_bird, bird)	quantity(wing, 2)	pw(beak, bird)
isa(owl, bird)	quantity(claw, 2)	pw(eye, bird)
ability(bird, fly)	pw(wing, bird)	quantity(leg, 2)
pw(feathers, bird)	conditional(wing, fly)	pw(leg, bird)
purpose(beak, eat)	sound(bird, chirp)	purpose(leg, stand)
pw(straw, nest)	_	

Table 2: The domain theory of *bird*

isa(mammal, animal)
isa(animal, living_entity)
isa(oviparous, animal)
isa(behavior, property)
isa(human_setting, setting)
isa(bird, existence, wilderness)
isa(wilderness, setting)
isa(farm, human_setting)
isa(equinae, mammal)

isa(house, human setting) isa(setting,space location) isa(space location,spatial entity) isa(physical entity,spatial entity) isa(physical object,physical entity) isa(property measure,information entity) isa(property,information entity) isa(living entity,physical entity)

shape(X, Y),shape(X,Z), $Y \neq Z \rightarrow$ false quantity(X, Y), quantity(X,Z), $Y \neq Z \rightarrow$ false behavior(X, friendly), behavior(X, dangerous) \rightarrow false actor(X,_), not isaN(X, action) \rightarrow false pw(A,A) \rightarrow false

Table 3: The generic space concept map and integrity constraints

Frame name	Conditions
aframe	The blend contains identical structure from input 1
aprojection	The blend contains the same concepts of input 1
bframe	The blend contains identical structure from input 2
bprojection	The blend contains the same concepts of input 2
pw_based_explanation	The blend contains a concept that has associated a set of
	part-whole relations (i.e. it is explained by a set of these
	relations)
transport_means	The blend contains a concept that has associated the set of
	features of a generic transport means
purposeful_subpart	The blend contains a a concept that has a subpart that has
	associated a set of relations that justify its existence
	(e.g. purpose, cause-effect)
new_ability	A concept has an ability relation not existent in any of the
	inputs
new_creature	A concept is a living thing that did not exist (or wasn't such)
	in any of the inputs
new_feature	A concept has a feature relation not existent in any of the
	inputs
	Table 4: Frames of the generic space

Pereira and Cardoso

References

- Bateman, J., Magnini, B., and Fabris, G. (1995). The generalized upper model knowledge base: Organization and use. *Towards very large knowledge bases: knowledge building and knowledge sharing*.
- Costello, F. J. and Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, 24(2):299–349.

Fauconnier, G. (1997). Mappings in Thought and Language. Cambridge University Press.

- Fauconnier, G. and Turner, M. (1998). Conceptual integration networks. *Cognitive Science*, 22(2):133–187.
- Fauconnier, G. and Turner, M. (2002). The Way We Think. Basic Books.
- Freeman, M. (1999). The role of blending in an empirical study of literary analysis. In *Proceedings of the 6th International Cognitive Linguistics Conference*.
- Pereira, F. C. (2003). Experiments with free concept generation in divago. In Cardoso, A., Bento, C., and Gero, J., editors, *Proceedings of the Third Workshop on Creative Systems*. IJCAI-03.
- Pereira, F. C. and Cardoso, A. (2001). Knowledge integration with conceptual blending. In *Proceedings of AICS 2001*.
- Pereira, F. C. and Cardoso, A. (2002). The boat-house visual blending experience. In *Proceedings of the Symposium for Creativity in Arts and Science of AISB 2002*. AISB.
- Pereira, F. C. and Cardoso, A. (2003). The horse-bird creature generation experiment. *AISB Journal*, 1(3).

AISB Journal
Scientific Methods for the Analysis of Agent-Environment Interaction

Ulrich Nehmzow* and Gregor Schöner[†]

- * Department of Computer Science, University of Essex, Wivenhoe Park, Colchester Essex CO4 3SQ, UK, *udfn@essex.ac.uk*
- [†] Ruhr University Bochum, Institut für Neuroinformatik, Ruhr-Universität-Bochum, Universitätsstr. 150, 44801 Bochum, Germany, *Gregor.Schoener@neuroinformatik.ruhr-uni-bochum.de*

Editors' Introduction

The motivation behind the symposium on "Scientific Methods for the Analysis of Agent-Environment Interaction" was the desire to finally start with some of the "mopping up" activities that characterise "normal science" (Kuhn, *The Structure of Scientific Revolutions*), the attempt "to force nature into the pre-formed and relatively inflexible box that a paradigm supplies." The current paradigms underlying autonomous agents research still seem too weak to support "normal science" *sensu* Kuhn. When we first identified this as a problem within the discipline of autonomous agents in general, and autonomous mobile robotics specifically, we hadn't realised that this concern was so widely shared within the community; but the call for papers for the AISB symposium elicited numerous encouraging comments and sufficient contributions for a viable meeting, so that the first symposium (ever?) on scientific methods in mobile robotics went ahead.

Arguably, research in agent-environment interaction hasn't even reached the stage of "normal science" yet, because paradigms forming the foundation of our research are still emerging and formulated in vague terms, theories of agent-environment interaction that would make testable predictions and thus structure the "mopping up" are still lacking.

We have few precise instruments (such as for instance quantitative descriptions of agent-environment interaction) at our disposal, and are consequently left to present qualitative descriptions of experiments and existence proofs. These are by no means futile activities, but perhaps the time has come to aim for more rigorous experimental methods that allow independent replication and verification of experimental results.

The fact that this symposium has happened is very encouraging. Slowly, but inexorably the emergence of a new branch of research in behaving agents, namely "scientific methods", is gathering momentum, and for the first time it was discussed within the research community what we actually mean by "scientific methods for the analysis of agent-environment interaction".

The papers presented at the AISB workshop on Scientific Methods for the Analysis of Agent-Environment Interaction¹ covered topics ranging from characterisation of behaviour through quantitative descriptions and computer modelling of agents to theoretical cognitive science, cognitive robotics and computational neuroscience. Two papers — on

¹Proc. AISB convention 2003, Aberystwyth, ISBN 1 902956 32 3.

Agent-Environment Interaction

neuro-robotic simulation of minimally cognitive behaviour to investigate cognitive scientific theories, and on quantitative descriptions of robot-environment interaction using dynamical systems theory — were selected by an independent jury chaired by the programme chairman for inclusion in this special issue of AISBJ. Paper presentations at the workshop were augmented by animated and fruitful discussions, resulting in a stimulating workshop that defined some important research topics for the future, and gave indications on how these might be addressed.

The Behaviour of a Mobile Robot Is Chaotic

Ulrich Nehmzow* and Keith Walker[†]

* Department of Computer Science, University of Essex, *Colchester CO4 3SQ*

[†] Department of Physics and Engineering, Point Loma Nazarene University, San Diego CA 92106

Abstract

We present a method to describe the behaviour of a mobile robot *quantitatively*, using methods from dynamical systems theory, time series analysis and deterministic chaos theory.

Experimental results obtained with a Pioneer II mobile robot demonstrate the use of the method, and show that robot behaviour exhibits deterministic chaos, and is substantially influenced by the control program executed by the robot, while changes to the environment have far less influence.

1 Background

1.1 Motivation

Research in mobile robotics to date has, with very few exceptions, been based on trialand-error experimentation and the presentation of existence proofs. Task-achieving robot control programs are obtained through a process of iterative refinement, typically involving the use of computer models of the robot, the robot itself, and program refinements based on observations made using the model and the robot. This process is iterated until the robot's behaviour resembles the desired behaviour to a sufficient degree of accuracy. Typically, the results of these iterative refinement processes are valid within a very narrow band of application scenarios, they constitute "existence proofs". As such, they demonstrate that a particular behaviour *can* be achieved, but not, how that particular behaviour can *in general* be achieved for *any* experimental scenario.

The purpose of this paper is to introduce quantitative measures of robot behaviour, as components of a theory of robot-environment interaction. Using dynamical systems theory and methods of analysis derived from chaos theory, we investigate quantitatively in what way the behaviour of a mobile robot changes if a) the robot's environment is modified, and b) the robot's control code is modified. Underlying this research, however, is the fundamental question: How can the interaction of a mobile robot with its environment be characterised quantitatively?

1.2 Robot-Environment Interaction

The behaviour of a mobile robot cannot be discussed in isolation: it is the result of properties of the robot itself (physical aspects — the "embodiment"), the environment ("situatedness"), and the control program (the "task") the robot is executing (see figure 1). This triangle of robot, task and environment constitutes a non-linear system, whose analysis is the purpose of any theory of robot-environment interaction.



Figure 1: The fundamental triangle of robot-environment interaction.

Rather than speaking solely of a robot's behaviour, it is therefore necessary to speak of *robot-environment interaction*, and the robot's behaviour resulting thereof.

1.3 The Role of Quantitative Performance Measures

1.3.1 Measurement: The Backbone of Science

Measurement is the backbone of science, and supports

- the precise documentation of experimental setups and experimental results,
- the principled modification of experimental parameters,
- independent verification of experimental results,
- theoretical design of artefacts without experimental development, and
- predictions about the behaviour of the system under investigation.

The experimental scenarios within mobile robotics, using *quantitative* performance measures, are manifold and expressed in figure 1. Basically, if any two of the three components shown in figure 1 remain unaltered, then the quantitative performance measure will characterise the third, modified component. This would allow the investigation of, for instance, i) the effect of modifications of the robot, ii) the influence of the robot control program on robot behaviour, or iii) the effect of modifications to the environment on the overall behaviour of the robot.

The purpose of the experiments reported in this paper was to demonstrate how robotenvironment interaction could be characterised quantitatively, and how such quantitative measures could be used to determine the influence of i) a change in the robot controller, and ii) a change of environment.

1.3.2 Three Theses

The experimental work reported in this paper is based on three theses:

Thesis 1 A mobile robot, interacting with its environment, is essentially an analog computer that "computes" *behaviour* (the output) from the three inputs *robot morphology, environmental characteristics* and *executed task* (see figure 2).



Figure 2: Robot-environment interaction as analog computation: Behaviour is the result of analog computation, taking robot morphology, task and environmental properties as "input".

Thesis 2 The behaviour exhibited by a robot, interacting with its environment while executing some particular task, is encapsulated, as a first approximation, in the robot's *trajectory*.

Thesis 3. One suitable method to describe a robot's trajectory quantitatively is to analyse the trajectory taken, using methods of time-series analysis and dynamical systems theory.

Based on these theses, we analysed the (x, y) components of logged robot trajectories, using time-series analysis methods from chaos theory.

1.4 Related Work

So far, quantitative measures of robot-environment interaction are not widely used in mobile robotics, and the amount of related work is consequently limited. The most important references are probably the work by Schöner et al. and Smithers.

(Schöner et al., 1995) used dynamical systems theory to investigate robot-environment interaction, and (Smithers, 1995) discussed the use of quantitative performance measures as a tool of scientific mobile robotics research.

Regarding the methods of characterisation used, the most important references are found in the area of time series analysis using chaos theory. Computing Lyapunov exponents, we used Arbarbanel's [(Abarbanel, 1996)] and Wolf's et al. [(Wolf et al., 1995)] methods. We computed correlation dimensions again using Arbarbanel's method, as well as the method discussed by Kaplan and Glass [(Kaplan and Glass, 1995)]. Background reading to the methods applied are (Peitgen et al., 1992) and (Kantz and Schreiber, 1997).

1.5 Experimental Setup

To investigate the two questions posed in section 1.3.1 — how does a change in robot control program influence robot behaviour, and what influence does a change of environmental conditions have? — we conducted experiments with a Pioneer II autonomous mobile robot (figure 3), executing a number of different control programs in a number of different environments.



Figure 3: The Pioneer II mobile robot used.

The robot's trajectory was logged, using an overhead camera. Every 250 ms the robot's x and y-coordinate were recorded. Figure 4 gives an example of the kind of trajectory obtained by this method, the dataset in that figure contains just over 26000 data points, recording 109 minutes of robot operation.

Figure 5 shows part of the x and y coordinate of that trajectory — trajectories like these were used for subsequent computation of Lyapunov exponent and correlation dimension (see section 2).

2 Quantitative Measures of Behaviour using Chaos Theory

2.1 Reconstructing the Attractor

The behaviour of dynamical systems, such as a mobile robot interacting with its environment, is characterised by velocity and position along each degree of freedom of the system — the phase space. Because a direct representation of the phase space of our robot-environment system is not available to us, the phase space has to be reconstructed from an observed time series such as the one shown in figure 5.



Figure 4: Robotic billiard ball (obstacle avoidance) behaviour — (entire trajectory (top) and 150 data points (bottom)).



Figure 5: x (top) and y coordinate (bottom) of the data shown in figure 4 (part).

To achieve this, we used the time-lag embedding method (discussed, for instance, in (Peitgen et al., 1992) or (Kaplan and Glass, 1995, p. 309)), embedding the attractor in *p*-dimensional space, using the coordinates given in equation 1.

$$\mathbf{D}_{t} = [D_{t}, D_{t-h}, D_{t-2d}, \dots, D_{t-(p-1)h}],$$
(1)

where D_t is the *p*-dimensional embedding of the attractor at time *t*, *p* the embedding dimension, and *h* the embedding lag.

As an example, figure 6 shows a three-dimensional phase space reconstruction of the attractor underlying the robotic billiard ball behaviour shown in figure 4.

2.1.1 Determining Embedding Lag and Embedding Dimension

In order to reconstruct the attractor, using the embedding method given by equation 1, embedding dimension p and embedding time lag h have to be determined first. To determine a suitable embedding lag, we used two methods:

- 1. The autocorrelation method presented in (Kaplan and Glass, 1995, p. 353): according to Kaplan and Glass "a good choice for the embedding lag is [that value] at which the autocorrelation function falls to $e^{-1} \approx 0.37$ ". Figure 7 shows the autocorrelation function for the x-coordinate of the data shown in figure 5, it can be seen that for t = 31 the autocorrelation has fallen below 0.37. Because t = 1represents an embedding lag of h = 0, we select an embedding lag of h = 30 to reconstruct the attractor.
- The Mutual Information Method: This procedure is associated with non-linear statistics while the autocorrelation method is based on linear statistics. In essence, mutual information is analogously a nonlinear autocorrelation function. Mutual

3D reconstruction of phase space - lag:30



Figure 6: Phase space reconstruction of robotic billiard ball experiment, using the x coordinate, embedding dimension 3 and an embedding lag of 30.

Information has its origin from the treatment of chaos as a source of information [(Gallager, 1968)]. The mutual information essentially measures the degree to which knowledge of one measurement determines another measurement h seconds later. (Fraser and Swinney, 1986) argue that one should take the first minimum that occurs in the mutual information function as the ideal time lag for constructing an embedded attractor. This minimum in the mutual information function corresponds to the time separation between two measured values so that there is a minimum of redundancy of information connecting the two measurements in the time series. Figure 8 gives the Mutual Information of the x-coordinate data shown in figure 4. The first minimum of this curve is obviously at a time lag of 30, thereby giving the same result as the autocorrelation function.

It is known [(Takens, 1981)] that the dynamics of a reconstructed attractor are geometrically identical to the original dynamics (for both continuous and discrete systems), if the relationship $p = 2\nu + 1$ holds (ν is the dimension of the original attractor). In order to establish a suitable embedding dimension p to reconstruct the attractor, it is therfore necessary to determine the dimension ν of the original attractor. One common measure of the dimension of an attractor is the correlation dimension ν .

In order to determine the correlation dimension, one uses the *correlation integral* C(r), defined in equation 2.

$$C(r) = \frac{\theta}{N(N-1)},\tag{2}$$

where θ is the number of times that $|\mathbf{D}_i - \mathbf{D}_j| < r$, *i* and *j* are two different times at which the embedding \mathbf{D} is taken, and *r* is an arbitrary distance measure, the "correlation distance". N(N-1) the maximum number of cases where $|\mathbf{D}_i - \mathbf{D}_j| < r$ is theoretically



Figure 7: Autocorrelation of the x-coordinate of the data shown in figure 5. For a lag of h = 30 the autocorrelation has dropped below e^{-1} .



Figure 8: Mutual information of the x-coordinate of the data shown in figure 5. The first minimum occurs for a lag of h = 30.



Figure 9: Correlation dimension versus correlation distance for the data shown in figure 5, for various embedding dimensions and various correlation distances. h = 30.

possible (the trivial case i = j is excluded).

The important measure is not the correlation integral C(r) itself, but how this integral changes with respect to r, i.e. the slope of the curve C(r) versus r. This slope is the correlation dimension ν .

The mechanism to determine a suitable embedding dimension to reconstruct the attractor is based on these considerations. Obviously, the computation of the correlation dimension is dependent upon the chosen embedding dimension p and the correlation distance r. To compute both p and ν from the same process is an ill-defined problem, and the goal therefore is to find a range of parameters p and r for which ν is computed identically (a so-called "scaling region"). In other words, one aims to find a region where the computation of the correlation dimension ν is not critically dependent upon the choice of embedding dimension p and correlation distance r.

To find a scaling region, we plot the correlation dimension ν as a function of correlation distance r for all embedding dimensions p between 1 and 10 (see (Kaplan and Glass, 1995, page 323)). Figure 9 shows ν versus p and r for the x-coordinate shown in figure 5.

Choosing a large correlation distance r is somewhat like looking at an object from a great distance: it will have dimension zero (Kaplan and Glass, 1995, p. 323). On the other extreme, choosing too small an r will result in signal noise being amplified, which is equally undesirable. Figure 9 reveals that for $r \approx 110$ the computed correlation dimension is similar for a large number of embedding dimensions, in other words, increasing the embedding dimension no longer changes the computed correlation dimension. The correlation dimension for r = 110 is around 2, figure 10 shows that in fact $\nu \approx 1.9$.

This means that an embedding dimension of 5 is a good choice (according to Takens' theorem). To confirm this finding, we used the 'false nearest neighbours' method as well (Abarbanel, 1996), the results are shown in figure 11 and confirm that indeed p = 5 is a



Figure 10: Ave Correlation dimension versus embedding dimension for dataset 240601, for scaling region between r = 103 - 127. Embedding lag h = 30. The computed correlation dimension is $\nu \approx 1.9$.

suitable choice.



Figure 11: Percentage of False Neighbours for the x-coordinate of dataset 240601. Zero false neighbours occurs at an embedding dimension of 5.

2.2 Characterising the Attractor

Systems that exhibit deterministic chaos are characterised by four properties:

- 1. They are stationary,
- 2. they are deterministic,
- 3. they are aperiodic, and
- 4. they show sensitivity to initial conditions.

Stationarity and determinism are easily established, the first by ascertaining that the means and standard deviations of different portions of the signal do not differ significantly (Kaplan and Glass, 1995, p.314), the latter by demonstrating that the first half of the data can be used as a predictor for the second half (Kaplan and Glass, 1995, p.324ff). In purely stochastic signals the best prediction possible is simply the mean of the signal, whereas in deterministic signals better predictions can be obtained by using parts of the data to predict other parts. All our data was confirmed to be stationary and deterministic.

Aperiodicity and sensitivity to initial conditions, however, are the truly interesting determinants for the purpose of this paper, as they supply a *quantitative* description of the attractor.

2.2.1 Aperiodicity

One main characteristic of a dynamical system exhibiting chaos is that the state variables never return to their exact previous values: The trajectory in phase space lies on a strange attractor. There is, however, variation from system to system in how close state variables return to previous values, and it is therefore desirable to quantify this degree of "proximity".

The measure to quantify the degree of aperiodicity is the correlation dimension ν , which was discussed in detail above. Besides being a quantitative description of the attractor, the correlation dimension indicates whether data is aperiodic or not, and to what degree: Periodic data has an integer correlation dimension, while chaotic attractors have a fractal (non-integer) correlation dimension.

2.2.2 Sensitivity to Initial Conditions

Another distinctive characteristic of a chaotic system is its sensitivity to a variation in the system's variables: Two trajectories in phase space that started close each other will diverge from one another as time progresses, the more chaotic the system, the greater the divergence.

Consider some state S_o of a deterministic dynamical system and its corresponding location in phase space. As time progresses the state of the system follows a deterministic trajectory in phase space. Let another state S_1 of the system lie arbitrarily close to S_o , and follow a different trajectory, again fully deterministic. If d_o is the initial separation of these two states in phase space at time t = 0, then their separation d_t after t seconds can be expressed as $d = d_o e^{\lambda t}$.



Figure 12: Computation of the Lyapunov exponent of the x-coordinate of the data shown in figure 5, using the method outlined by (Wolf et al., 1995).

 λ is known as the Lyapunov exponent. For a m-dimensional phase space, λ will have m-components. If any one or more of those components are positive, then the trajectories of nearby states diverge exponentially from each other in phase space and the system is deemed chaotic. Since any system's variables of state are subject to uncertainty, a knowledge of what state the system is in can quickly become unknown if chaos is present. The larger the positive Lyapunov exponent, the quicker knowledge about the system is lost. One only knows that the state of the system lies somewhere on one of the trajectories traced out in phase space, i.e., somewhere on the strange attractor.

The Lyapunov exponent is one of the most useful quantitative measures of chaos, since it will reflect directly whether the system is indeed chaotic, and will quantify the degree of that chaos. Also, knowledge of the Lyapunov exponents becomes imperative for any analysis on prediction of future states.

Determination of the Lyapunov exponents from a time series such as the robot trajectories used in this research or, more importantly, the existence and value of any positive exponents, has been discussed extensively in the literature [(Wolf et al., 1995; Peitgen et al., 1992; Kantz and Schreiber, 1997; Abarbanel, 1996; Kaplan and Glass, 1995; Abarbanel et al., 1993)]. Several academic and commercial software packages have been made available to compute Lyapunov exponents from time series, for example (Wolf et al., 1995; Applied Nonlinear Sciences LLC and Randle Inc, 2003; Kantz and Schreiber, 2003). In our analyses we have used the method proposed by (Wolf et al., 1995), as well as Abarbanel's commercially available software package (Applied Nonlinear Sciences LLC and Randle Inc, 2003) to determine the largest Lyapunov exponent.

Figure 12 shows the result of the computation of the Lyapunov exponent for the data shown in figure 5, using the algorithm described in (Wolf et al., 1995).

It can be seen that the algorithm converges to λ just above 0.3, which indicates that this data exhibits deterministic chaos.



Figure 13: Wall following behaviour of the Pioneer Robot — entire trajectory (top) and 100 consecutive positions (bottom).

3 Experiments: Analysing Robot-Environment Interaction

We stated above that robot-environment interaction is governed by the triangle of robot, task and environment (figure 1). If, therefore, two of the three components remain constant, Lyapunov exponent and correlation dimension will characterise the third component and its influence on the overall behaviour of the robot.

Of the many experiments we conducted, we would like to discuss two experiments in this paper:

- 1. In experiment 1 we kept the environment constant, and altered the robot control code.
- 2. In experiment 2, we kept the control code constant, but altered the environment.

In both experiments the same robot was used.

3.1 Changing the Robot Task

In the first experiment, the Pioneer II operated in an empty environment, surrounded by plasterboard walls. The robot executed two different behaviours: wall following (13000 data points, 54 minutes of operation, see figure 13) and "quasi-billiard-ball" obstacle avoidance (26000 data points, 109 minutes of operation, see figure 4).

For the wall following behaviour, we computed a correlation dimension of $\nu \approx 1.4$ and a Lyapunov exponent of $\lambda \approx 0.1$. For the quasi-billiard ball behaviour noticeably different values were computed, i.e. $\nu \approx 1.9$ and $\lambda \approx 0.3$. In other words: The quasibilliard ball behaviour exhibited a considerably larger degree of deterministic chaos than the wall following behaviour.



Figure 14: Entire trajectory of quasi-billiard ball behaviour in an environment with a central obstruction (top), and 200 consecutive positions (bottom).

3.2 Changing the Environment

In the second environment, we kept both robot and task constant (using the quasi-billiard ball behaviour, but changed the environment by introducing a central obstruction to the arena). The trajectory obtained for that case is shown in figure 14.

Here, both the correlation dimension ($\nu \approx 1.7$) and the Lyapunov exponent ($\lambda \approx 0.3$) remained essentially the same to the values obtained when no central obstruction was present. In other words: The same degree of deterministic chaos was present, irrespective of whether the central obstruction was in place or not.

4 Conclusion

4.1 Towards a Theory of Robot-Environment Interaction

Mobile robotics research to date is still largely dependent upon trial-and-error development, and results often are existence proofs, rather than generalisable, fundamental results. Furthermore, independent replication and verification of experimental results are largely unknown in mobile robotics research practice.

One reason for this is the dearth of quantitative measures of behaviour. We have therefore investigated the application of dynamical systems theory — specifically deterministic chaos theory — to the analysis of robot environment interaction.

Using Lyapunov exponent and correlation dimension of the (strange) attractor underlying our Pioneer robot's behaviour, we found that the robot's behaviour indeed exhibited deterministic chaos. Moreover, we found that the degree of deterministic chaos increased noticeably when the robot's control program was changed from wall following to quasi-billiard ball obstacle avoidance, while it remained essentially constant when the environment was changed, but the control program remained the same.

However, we see the main contribution of this paper not in the specific analysis of our Pioneer II mobile robot, but in presenting a method of analysing robot-environment



Figure 15: Possible alternative environment to investigate wall following behaviour.

interaction quantitatively. Quantitative descriptions form the backbone of any scientific theory, and are therefore the first step towards a theory of robot-environment interaction, that allows the formation of hypotheses, precise reporting and independent verification of results, and the theoretical design of robot systems.

4.2 Interpretation of Results

The specific interpretation of our experimental results is that in the experimental environment we have used, the robot control code has far more influence on the behaviour of the robot than the environment itself. The wall following behaviour exhibits a far lower degree of deterministic chaos than the quasi-billiard ball behaviour, while the quasi-billiard ball behaviour shows a similar degree of chaos in two different environments.

4.3 Future Work

The obvious extension to the experiments reported here is to conduct the wall following behaviour in an environment that is different to the one that is shown in figure 13 — for example an environment similar to the one shown in figure 15.

Our hypothesis is that in such an environment we would obtain a correlation dimension of $\nu \approx 1.5$ and a Lyapunov exponent of $\lambda \approx 0.1$.

Acknowledgements

This project was supported by a grant from the Research Promotion Fund of the University of Essex. Keith Walker was supported by a grant from the Point Loma Nazarene University Alumni Association. We acknowledge this support gratefully.

References

- Abarbanel, H. (1996). Analysis of Observed Chaotic Data. New York: Springer-Verlag.
- Abarbanel, H., Brown, R., Sidorowich, J., and Tsimring, L. (1993). The analysis of observed chaotic data in physical systems. *Rev. Mod. Phys.*, 65:1331–1392.
- Applied Nonlinear Sciences LLC and Randle Inc (2003). CSP time series analysis software package. http://www.zweb.com/apnonlin/csp.html.
- Fraser, A. and Swinney, H. (1986). Independent coordinates for strange attractors from mutual information. *Phys. Rev. A.*, 33:1134–1140.
- Gallager, R. (1968). *Information Theory and Reliable Communication*. New York: John Wiley and Sons.
- Kantz, H. and Schreiber, T. (1997). *Nonlinear Time Series Analysis*. Cambridge University Press.
- Kantz, H. and Schreiber, T. (2003). *TISEAN Nonlinear time series analysis*. http://ls11www.cs.uni-dortmund.de/people/hermes/NLDdocs/docs/.
- Kaplan, D. and Glass, L. (1995). *Understanding Nonlinear Dynamics*. Springer Verlag London.
- Peitgen, H.-O., Jürgens, H., and Saupe, D. (1992). Chaos and Fractals. Springer Verlag.
- Schöner, G., Dose, M., and Engels, C. (1995). Dynamics of behavior: theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, 16.
- Smithers, T. (1995). On quantitative performance measures of robot behaviour. *Robotics* and Autonomous Systems, 15:107–133.
- Takens, F. (1981). Detecting Strange Attractors in Turbulence, volume 898 of *Lecture Notes in Mathematics*. Berlin: Springer-Verlag.
- Wolf, A., Swift, J., Swinney, H., and Vastano, J. (1995). Determining lyapunov exponents from a time series. *Physica 16D*.

On the Role of Robot Simulations in Embodied Cognitive Science

Tom Ziemke

Department of Computer Science, University of Skövde PO Box 408, 54128 Skövde, Sweden tom@ida.his.se

Abstract

Research in embodied cognitive science emphasizes that a close interaction of brain, body and environment is central to the emergence of cognitive processes. Much work on embodied artificial intelligence has therefore shifted focus from purely computational modeling to autonomous mobile robotics. Many researchers emphasize the importance of working with real robots rather than simulations which usually cannot fully capture the complexities of the physical world. However, from a cognitive science point of view, robot simulations nevertheless have an important, complementary role to play, due to the fact that in many cases they allow for more extensive, systematic experimentation as well as for experiments, e.g. with evolving robot morphologies, that can only be carried out in very limited form on real robots. Furthermore, it will be argued in this paper, robot simulations are very useful tools in experimentation with active adaptation of non-trivial environments, an aspect that is still largely ignored in much embodied artificial intelligence research.

1 Introduction

As several authors have pointed out, mobile robotics, or autonomous agents research in general, can be viewed from at least two different, though intertwined perspectives: that of engineering, mostly concerned with the design of artefacts, and that of science, mostly concerned with the understanding of natural systems. Furthermore, the latter can of course be broken down according to the different scientific fields that use robots and/or other autonomous agents as modeling tools, e.g. cognitive science (Pfeifer, 1995; Pfeifer and Scheier, 1999), neuroscience (Voegtlin and Verschure, 1999; Ruppin, 2002), or the study of animal behavior (Webb, 2000; Webb, 2001).

While these distinctions appear fairly obvious, they seem to receive surprisingly little attention in discussions of methodology in the field(s) of autonomous robotics and embodied/behavior-oriented AI, where overly general statements such as "simulations are useless" or "Khepera robots are not real robots" or "existence proofs just don't do it" often can be heard. While from an engineering point of view these statements might very well be more or less correct, they do not necessarily apply equally generally to the scientific use of autonomous agents as models of natural organisms. Steels, for example, explains the skepticism towards simulations as follows: 'The goal is to build artifacts that are "really" intelligent, that is, intelligent in the physical world, not just intelligent in a virtual world. This makes unavoidable the construction of robotic agents that must sense the environment and can physically act upon the environment, particularly if sensorimotor competences are studied. This is why behavior-oriented AI researchers insist so strongly on the construction of physical agents ... Performing simulations of agents ... is, of course, an extremely valuable aid in exploring and testing out certain mechanisms, the way simulation is heavily used in the design of airplanes. But a simulation of an airplane should not be confused with the airplane itself. ' (Steels, 1994)

Obviously, Steels has a point here, and nobody would seriously question the view that simulations, however good they are, cannot fully capture the complexities of the physical world. Hence, simulations might have limited value in robot engineering, and might only be second choice in the modeling of animal behavior in cases where a real robot can made to interact with (roughly) the same physical environment as the modeled animal (Lambrinos et al., 1997; Webb, 2000). However, from a cognitive science point of view, robot simulations nevertheless have an important, complementary role to play, due to the fact that in many cases they allow for more extensive, systematic experimentation, which simply takes less time in simulation, as well as for experiments, e.g. with evolving robot morphologies, that can only be carried out in very limited form on real robots.

This is particularly important from the perspective of embodied cognitive science which emphasizes that a close interaction of brain (or more generally, nervous system), body and environment is central to the emergence of cognitive processes (Chiel and Beer, 1997; Clancey, 1997; Clark, 1997; Clark, 1999; Lakoff and Johnson, 1999; Pfeifer and Scheier, 1999; Varela et al., 1991). Hence, instead of focusing on one experiment or a few experiments with a real robot, many questions are more suitably addressed through a large number of experiments allowing for more variations of agent morphologies, control architectures and/or environments. Finally, it should be added that in many studies simulated and physical robot experiments, as also mentioned by Steels above, can play complementary roles, as for example in the field of evolutionary robotics (Nolfi and Floreano, 1999), where it is not uncommon to run a large number of experiments in simulation first, and then transfer the most interesting experiments or their results (e.g. evolved controllers) to physical robots.

Concerning the role of existence proofs, again we have to distinguish between engineering and scientific robotics. While from an engineering point of view existence proofs certainly are of limited value (for example, who would want to fly in an air plane that has been tested successfully exactly once?), from a cognitive science point of view they can be very valuable in the development of theories. Much connectionist cognitive modeling research, for example, has been concerned with providing concrete examples of neural networks exhibiting properties such as systematicity (e.g. Boden and Niklasson, 2000), which on purely theoretical grounds they had been argued not to be able to exhibit (Fodor and Pylyshyn, 1988). This is just one example, where existence proofs constrain and thus aid the development of cognitive-scientific theories. For this type of research robot simulations are very useful tools in agent-based modeling (Schlesinger and Parisi, 2001), paying more attention to the interaction of agents and environments than traditional computational cognitive modeling of mostly internal processes.

Much of the cognitive science and AI research in our group (Ziemke and Niklasson, 2003) is concerned with theories and models of situated and embodied cognition, with a focus on the interaction and co-adaptation of 'internal' mechanisms (such as thought,

Ziemke

memory, internal "representations", etc.), sensorimotor/bodily mechanisms, and external "scaffolds", such as artifacts, other agents, or sociocultural structures. Theoretical work on the mechanisms underlying human situated and social cognition (Susi and Ziemke, 2001; Susi et al., 2003; Lindblom and Ziemke, 2003) is complemented by experimental work such as neuro-robotic studies of what Beer referred to as "minimally cognitive behavior" (Beer, 1996), i.e. "the simplest behavior that raises cognitive interesting issues", and the detailed analysis of the mechanisms underlying such behavior (Biro and Ziemke, 1998; Ziemke, 1999; Ziemke, 2000; Ziemke and Thieme, 2002).

Using two example neuro-robotic simulation studies, of delayed response tasks and predator-prey co-evolution respectively, the following sections exemplify how we combine qualitative and quantitative analyses, based on a detailed correlation of simulated data (e.g. sensory inputs, neural activation values, morphological parameters, behavior, etc.) at each point in time. This allows us to analyze in detail, at different time scales, the interaction between often minimal internal mechanisms and the environmental or behavioral dynamics they exploit.

2 Example 1: Delayed Response Tasks

Delayed response tasks are a standard paradigm in experimental psychology studies of short-term memory, and they have also been used in a number of autonomous agents experiments (Ulbricht, 1996; Jakobi, 1997; Rylatt et al., 1998; Rylatt and Czarnecki, 2000; Linaker and Jacobsson, 2002). Using a simulated version (Michel, 1996) of the Khepera robot (Mondada et al., 1994), we carried out a large number of experiments with four different neural control architectures, trained by evolutionary algorithms, in six different types of delayed response tasks (Thieme and Ziemke, 2002; Ziemke and Thieme, 2002). Figure 1 illustrates the simulated robot and one of the tasks, in which the agent has to navigate a maze in which light sources (or, to be exact, the sides on which they appear, i.e. to the left or the right of the corridor) indicate the correct turning directions in the junctions that it later encounters on its way to the goal. In the particular multiple delayed response task illustrated here, the most complex in our series of experiments, the first light source indicates the correct turning direction in the first junction, whereas the second one indicates the wrong direction in the second junction. That means, the robot has to deal with two overlapping delays of different lengths and stimuli with varying meaning.

What we found in our experiments was that, irrespective of which neural control architecture they were using, the robots tended to evolve to maximize their use of the environment as an external "scaffold" to guide their behavior reliably (despite the usual random variations of exact starting positions, orientations, light source locations, etc.). In fact, even purely reactive feed-forward networks without internal short-term memory turned out to solve the simpler delayed response tasks (not the one illustrated here though), e.g. by following particular walls, thus using their own position in the environment as some form of "external memory" (Thieme and Ziemke, 2002).

Figure 2 illustrates how a robot, using a particular recurrent neural control architecture (the only one solving all task reliably) that realizes short-term memory by dynamic modulation of otherwise reactive sensorimotor mappings (Ziemke and Thieme, 2002) solves a particular instantiation of the task. It reactively follows the first corridor up to the first junction, using weight matrix (a), and then switches to another reactive behavior, lefthand wall following, realized by weight matrix (b), that takes the robot all the way to the goal. The (type of) data illustrated in figure 2 allows us to analyze in detail how in this case, and many others, the robots solve their tasks through maximal exploitation of



Figure 1: Top: Example of a multiple delayed response task. The agent has to find its way from the starting position (indicated by the arrow) to the goal (white circle) using the information provided by the light sources. In this particular case the first light source (the side on which it appears) indicates the correct turning direction in the first junction, the second light source indicates the wrong direction in the second junction. Bottom: The simulated Khepera robot used in the delayed response experiments, equipped with eight distance/proximity sensors (D), two ambient light sensors (L), and two motors (M).

Ziemke

environmental regularities and the dynamics of their own behavior (Ziemke and Thieme, 2002).



Figure 2: Top: Sensory/neural activation values at each point in time during a successful robot trajectory. For abbreviations see Figure 1. Bottom: Sensorimotor connection weight matrices, mapping sensory inputs to motor outputs (white: positive weights; black: negative weights; size reflects weight magnitude). Matrix (a), used up to the first junction, generates simple corridor-following behavior, whereas matrix (b), to which the agent switches in the first junction, generates left-hand wall-following behavior. Adapted from Ziemke and Thieme (2002), for details see the original paper.

As in the vast majority of similar experiments in the field, the robots here had to adapt themselves (in this case their neural control mechanisms) to their environments. Although humans, as well as other animals, instead of adapting themselves in many cases can choose to adapt their environments to their own needs (Kirsh, 1995), that alternative has been studied relatively little in robot experiments, except for robotic studies of stigmergy (Beckers et al., 1994; Bonabeau, 1999), the simplest form of environment adaptation as used, for example, by social insects. In future work we therefore aim to address more complex forms of active adaptation of the environment, both by individual agents, e.g. by moving objects, and in the course of 'cultural evolution' of, e.g., artefacts and en-

vironmental structures. For the reasons discussed in the introductory section, we believe that robot simulations can be a powerful tool in such studies, and a more natural starting point than physical robot experiments.

3 Example 2: Predator-Prey Co-Evoluation

In a series of twenty-one experiments integrating competitive co-evolution (CCE) of neural robot controllers with 'co-evolution' of robot morphologies and control systems we aimed to extend previous work on CCE (Nolfi and Floreano, 1999) by systematically investigating the tradeoffs and interdependencies between morphological parameters and behavioral strategies through a series of predator-prey experiments in which increasingly many aspects are subject to self-organization (Buason and Ziemke, 2003b; Buason and Ziemke, 2003a; Buason and Ziemke, 2003c).

Again, these experiments were carried out using a modified Khepera simulator (Carlsson and Ziemke, 2001). Figure 3 (top) illustrates the basic setup with one predator and one prey robot, both equipped with short-range infrared sensors and simple long-range cameras. Neural control systems (standard recurrent networks) were evolved together with a number of parameters of the robots' respective visual morphologies (view range, view angle, camera direction). In these experiments, the predator's task was simply to catch (make contact with) the prey as quickly as possible, whereas the prey obviously had to avoid being caught as long as possible.

Figure 3 also illustrates the results of one particular experiment in which task-dependent visual morphologies were evolved in the competing robot species (with the maximum speed being constrained by the view angle, i.e. the larger the visual field, the slower the robots) (Buason and Ziemke, 2003a). As the graphs illustrate, the predators tended to have a relatively long view range and narrow view angle, whereas prey (right) tended to minimize their view angle, irrespective of view range, i.e. they 'preferred' speed over vision.

This experiment illustrates one particular case of a close coupling between morphology, sensorimotor processes and behavioral strategies (the latter is not illustrated here) and their mutual adaptation, i.e. the type of interdependencies that are considered crucial by embodied cognitive theories and, for obvious reasons, are difficult to investigate in physical robot experiments.

Compared to the previous experiment, this one is more interesting in the sense that at least part of the environment, i.e. the competitor, is no longer static and passive, but actually co-adapting in a competitive fashion. As before, however, the rest of the environment is completely static and of even more limited complexity. In current and future work we are therefore investigating the role that a more complex and dynamic environment can play in such co-adaptive scenarios, (a) as a constraint for both species and (b) as a resource that the robots/species might be able to adapt to their own needs.

4 Summary and Conclusion

We have in this paper argued that robot simulations, despite their limitations from an engineering point of view, have an important role to play in the scientific modeling of adaptive behavior and embodied cognitive processes. This has been illustrated with two examples of robot simulation studies of 'minimally cognitive behavior' that allowed for



Figure 3: Top: Environment (left) and basic morphology of simulated predator and prey robots (right). Middle: Average view ranges (bold lines) and view angles (thin lines) for predator (left) and prey (right) over 250 generations. Bottom: An illustration of the species' 'preferences' and their respective fitness (diamond grey level). Predators (left) tend to have a long view range with a relatively narrow view angle, whereas prey (right) tend to minimize their view angle, irrespective of view angle, i.e. they 'prefer' speed over vision. Adapted from Buason and Ziemke (2003a), for details see the original paper.

extensive experimentation with systematic variations of task complexities, neural control systems, morphological parameters, environments and other constraints.

Although some of these experiments, in particular in the delayed response task study, certainly also could have been carried out with physical robots under more realistic circumstances, this would have been much more time-consuming and more difficult to analyze. In the second case, of predator-prey co-evolution, obviously large parts of the experiments could hardly be carried on physical robots with current technology. Furthermore, we have argued that future experiments should pay more attention to the role of the environment, in particular to their (more or less) active adaptation through individual agents or 'cultural evolution'. Again, robot simulations seem a natural starting point for this.

In sum, it is certainly true that robot simulations have a number of more or less obvious limitations. For example, it is not at all clear to what degree the results of the discussed experiments could be transferred or generalized to physical robots. It is, however, also true that simulations of the type(s) discussed here can be powerful tools for the realization and analysis for experiments in the study of agent-environment interaction, and thus they can aid the further development of theories in embodied cognitive science.

Acknowledgements

The experiments discussed in this paper were performed by Mikael Thiemke (Thieme and Ziemke, 2002; Ziemke and Thieme, 2002) and Gunnar Buason (Buason and Ziemke, 2003b; Buason and Ziemke, 2003a; Buason and Ziemke, 2003c), respectively. Furthermore, the author would like to thank the organizers and participants of the AISB '03 symposium on *Scientific Methods for the Analysis of Agent-Environment Interaction* for providing useful feedback on an earlier version of this paper (Ziemke, 2003).

References

- Beckers, R., Holland, O., and Deneuborg, J. (1994). From local actions to global tasks: Stigmergy and collective robotics. In *Artificial Life IV*, pages 181–189, Cambridge, MA. MIT Press.
- Beer, R. A. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J. B., and Wilson, S. W., editors, *From animals to animats 4 Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 421–429, Cambridge, MA. MIT Press.
- Biro, Z. and Ziemke, T. (1998). Evolution of visually-guided approach behaviour in recurrent artificial neural network robot controllers. In Rolf Pfeifer et al., editor, *From animals to animats 5 - Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior*, pages 73–76, Cambridge, MA. MIT Press.
- Boden, M. and Niklasson, L. (2001). Semantic systematicity and context in connectionist networks. *Connection Science*, 12(2):1–31.

Bonabeau, E. (1999). Special issue: Stigmergy. Artificial Life, 5(2).

Ziemke

- Buason, G. and Ziemke, T. (2003a). Co-evolving task-dependent visual morphologies in predator-prey experiments. In Cantu-Paz et al., editor, *Genetic and Evolutionary Computation Conference*, Berlin. Springer Verlag. To appear.
- Buason, G. and Ziemke, T. (2003b). Competitive co-evolution of predator and prey sensory-motor systems. In Raidl et al., editor, *Applications of Evolutionary Computing - EvoWorkshops 2003*, pages 605–615, Berlin. Springer Verlag.
- Buason, G. and Ziemke, T. (2003c). Explorations of task-dependent visual morphologies in competitive co-evolutionary experiments. In *Proceedings of the European Conference on Artificial Life*, Berlin. Springer Verlag. To appear.
- Carlsson, J. and Ziemke, T. (2001). Yaks Yet Another Khepera Simulator. In Rückert, U., Sitte, J., and Witkowski, M., editors, Autonomous Minirobots for Research and Edutainment - AMiRE2001 - Proceedings of the 5th International Heinz Nixdorf Symposium, pages 235–241, Paderborn, Germany. Heinz Nixdorf Institut.
- Chiel, H. J. and Beer, R. A. (1997). The brain has a body: Adaptive behavior emerges from interactions of nervous system, body, and environment. *Trends in Neurosciences*, 20:553–557.
- Clancey, W. J. (1997). Situated Cognition: On Human Knowledge and Computer Representatiosns. Cambridge University Press, NeW York.
- Clark, A. (1997). *Being There Putting Brain, Body and World Together Again.* MIT Press, Cambridge, MA.
- Clark, A. (1999). An embodied cognitive science? *Trends in Cognitive Science*, 3(9):345–351.
- Fodor, J. A. and Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28:3–71.
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6(2):325–368.
- Kirsh, D. (1995). Adapting the environment instead of oneself. *Adaptive Behavior*, 4(3-4):415–452.
- Lakoff, G. and Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to western thought*. Basic Books, New York.
- Lambrinos, D., Marinus, M., Kobayashi, H., Labhart, T., Pfeifer, R., and Wehner, R. (1997). An autonomous agent navigating with a polarized light compass. *Adaptive Behavior*, 6(1):131–161.
- Linaker, F. and Jacobsson, H. (2002). Learning delayed response tasks through unsupervised event extraction. *International Journal of Computational Intelligence and Applications*, 1(4):413–426.
- Lindblom, J. and Ziemke, T. (2003). Social situatedness of natural and artificial intelligence: Vygotsky and beyond. *Adaptive Behavior*, 11(2). To appear.
- Michel, O. (1996). Khepera simulator package version 2.0. Freeware mobile robot simulator, available from http://diwww.epfl.ch/lami/team/michel/khep-sim/.

- Mondada, F., Franzi, E., and Ienne, P. (1994). Mobile robot miniaturisation: A tool for investigating in control algorithms. In Yoshikawa, T. and Miyazaki, F., editors, *Third International Symposium on Experimental Robotics*, 1993, pages 501–513, Berlin. Springer Verlag.
- Nolfi, S. and Floreano, D. (1999). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA.
- Pfeifer, R. (1995). Cognition perspectives from autonomous agents. *Robotics and Autonomous Systems*, 15:47–70.
- Pfeifer, R. and Scheier, C. (1999). Understanding Intelligence. MIT Press, Cambridge, MA.
- Ruppin, E. (2002). Evolutionary autonomous agents: A neuroscience perspective. *Nature Reviews Neuroscience*, 3(2):132–142.
- Rylatt, R. M. and Czarnecki, C. (2000). Embedding connectionist autonomous agents in time: The 'road sign problem'. *Neural Processing Letters*, 12:145–158.
- Rylatt, R. M., Czarnecki, C. A., and Routen, T. W. (1998). Beyond physical grounding and naive time: Investigations into short-term memory. In *From animals to animats* 5 - Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior, pages 22–31, Cambridge, MA. MIT Press.
- Schlesinger, M. and Parisi, D. (2001). The agent-based approach: A new direction for computational models of development. *Developmental Review*, 21:121–146.
- Steels, L. (1994). The artificial life roots of artificial intelligence. Artificial Life, 1:75-110.
- Susi, T., Lindblom, J., and Ziemke, T. (2003). Beyond the bounds of cognition. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*. To appear.
- Susi, T. and Ziemke, T. (2001). Social cognition, artefacts, and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. *Cognitive Systems Research*, 2(4):273–290.
- Thieme, M. and Ziemke, T. (2002). The road sign problem revisited: Handling delayed response tasks with neural robot controllers. In B. Hallam et al., editor, *From Animals to Animats 7 - Proceedings of the Seventh International Conference on the Simulation of Adaptive Behavior*, pages 228–229, Cambridge, MA. MIT Press.
- Ulbricht, C. (1996). Handling time-warped sequences with neural networks. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J. B., and Wilson, S. W., editors, *From animals* to animats 4 - Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, pages 180–189, Cambridge, MA. MIT Press.
- Varela, F., Thompson, E., and Rosch, E. (1991). *The Embodied Mind Cognitive Science and Human Experience*. MIT Press, Cambridge, MA.
- Voegtlin, T. and Verschure, P. (1999). What can robots tell us about brains? a synthetic approach towards the study of learning and problem solving. *Reviews in Neuroscience*, 10(3-4):291–310.

Ziemke

- Webb, B. (2000). What does robotics offer animal behavior? *Animal Behaviour*, 60:545–558.
- Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24(6).
- Ziemke, T. (1999). Remembering how to behave: Recurrent neural networks for adaptive robot behavior. In Medsker, L. and Jain, L. C., editors, *Recurrent Neural Networks: Design and Applications*, pages 359–389. CRC Press, New York.
- Ziemke, T. (2000). On 'parts' and 'wholes' of adaptive behavior: Functional modularity and diachronic structure in recurrent neural robot controllers. In Jean-Arcady Meyer et al., editor, *From Animals to Animats 6 - Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior*, Cambridge, MA. MIT Press.
- Ziemke, T. (2003). On the analysis of robot behavior and its contribution to theoretical cognitive science. In *Proceedings of the AISB '03 Symposium on Scientific Methods for the Analysis of Agent-Environment Interaction*, pages 31–37, UK. Society for the Study of Artificial Intelligence and the Simulation of Behavior.
- Ziemke, T. and Niklasson, L. (2003). Cognitive science and AI research in Skövde, Sweden. In Schmalhofer, F., Young, R., and Katz, G., editors, *Proceedings of EuroCogSci 03 - The European Cognitive Science Conference 2003*, pages 471–472, Hillsdale, NJ. Lawrence Erlbaum. To appear.
- Ziemke, T. and Thieme, M. (2002). Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks. *Adaptive Behavior*, 10(3-4):185–199.

AISB Journal