


LNAI 4211

Paul Vogt
Yuuga Sugita
Elio Tuci
Chrystopher Nehaniv (Eds.)

Symbol Grounding and Beyond

Third International Workshop on the Emergence
and Evolution of Linguistic Communication, EELC 2006
Rome, Italy, September/October 2006, Proceedings



Springer

74923-5
E53
2006
Paul Vogt Yuuga Sugita Elio Tuci
Chrystopher Nehaniv (Eds.)

Symbol Grounding and Beyond

Third International Workshop on the Emergence
and Evolution of Linguistic Communication, EELC 2006
Rome, Italy, September 30 – October 1, 2006
Proceedings



Springer



E200604081

Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

Volume Editors

Paul Vogt
Tilburg University, ILK/Language and Information Science
P.O. Box 90153, 5000 LE Tilburg, The Netherlands
E-mail: p.a.vogt@uvt.nl

Yuuga Sugita
RIKEN Brain Science Institute, Behavior and Dynamic Cognition Lab.
2-1 Hirosawa, Wako-shi, Saitama 351-0198, Japan
E-mail: sugita@bdc.brain.riken.go.jp

Elio Tuci
Université Libre de Bruxelles, IRIDIA
Avenue F. Roosevelt 50, CP 194/6, 1050 Brussels, Belgium
E-mail: etuci@ulb.ac.be

Chrystopher Nehaniv
University of Hertfordshire, School of Computer Science
Adaptive Systems & Algorithms Research Groups
College Lane, Hatfield, Hertfordshire, AL10 9AB, UK
E-mail: C.L.Nehaniv@herts.ac.uk

Library of Congress Control Number: 2006932800

CR Subject Classification (1998): I.2.11, I.2.6-7, I.6, F.1.1-2, H.3.1, K.4, H.5, J.4

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743
ISBN-10 3-540-45769-0 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-45769-5 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11880172 06/3142 5 4 3 2 1 0

Lecture Notes in Artificial Intelligence 4211

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 4213: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Eds.), Knowledge Discovery in Databases: PKDD 2006. XXII, 660 pages. 2006.
- Vol. 4212: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Eds.), Machine Learning: ECML 2006. XXIII, 851 pages. 2006.
- Vol. 4211: P. Vogt, Y. Sugita, E. Tuci, C. Nehaniv (Eds.), Symbol Grounding and Beyond. VIII, 237 pages. 2006.
- Vol. 4203: F. Esposito, Z.W. Ras, D. Malerba, G. Semeraro (Eds.), Foundations of Intelligent Systems. XVIII, 767 pages. 2006.
- Vol. 4201: Y. Sakakibara, S. Kobayashi, K. Sato, T. Nishino, E. Tomita (Eds.), Grammatical Inference: Algorithms and Applications. XII, 359 pages. 2006.
- Vol. 4196: K. Fischer, E. André, I.J. Timm, N. Zhong (Eds.), Multiagent System Technologies. X, 185 pages. 2006.
- Vol. 4188: P. Sojka, I. Kopeček, K. Pala (Eds.), Text, Speech and Dialogue. XIV, 721 pages. 2006.
- Vol. 4183: J. Euzenat, J. Domingue (Eds.), Artificial Intelligence: Methodology, Systems, and Applications. XIII, 291 pages. 2006.
- Vol. 4180: M. Kohlhase, OMDoc – An Open Markup Format for Mathematical Documents [version 1.2]. XIX, 428 pages. 2006.
- Vol. 4160: M. Fisher, W.v.d. Hoek, B. Konev, A. Lisitsa (Eds.), Logics in Artificial Intelligence. XII, 516 pages. 2006.
- Vol. 4155: O. Stock, M. Schaerf (Eds.), Reasoning, Action and Interaction in AI Theories and Systems. XVIII, 343 pages. 2006.
- Vol. 4149: M. Klusch, M. Rovatsos, T.R. Payne (Eds.), Cooperative Information Agents X. XII, 477 pages. 2006.
- Vol. 4139: T. Salakoski, F. Ginter, S. Pyysalo, T. Pahikkala, Advances in Natural Language Processing. XVI, 771 pages. 2006.
- Vol. 4133: J. Gratch, M. Young, R. Aylett, D. Ballin, P. Olivier (Eds.), Intelligent Virtual Agents. XIV, 472 pages. 2006.
- Vol. 4130: U. Furbach, N. Shankar (Eds.), Automated Reasoning. XV, 680 pages. 2006.
- Vol. 4120: J. Calmet, T. Ida, D. Wang (Eds.), Artificial Intelligence and Symbolic Computation. XIII, 269 pages. 2006.
- Vol. 4114: D.-S. Huang, K. Li, G.W. Irwin (Eds.), Computational Intelligence, Part II. XXVII, 1337 pages. 2006.
- Vol. 4108: J.M. Borwein, W.M. Farmer (Eds.), Mathematical Knowledge Management. VIII, 295 pages. 2006.
- Vol. 4106: T.R. Roth-Berghofer, M.H. Göker, H. A. Güvenir (Eds.), Advances in Case-Based Reasoning. XIV, 566 pages. 2006.
- Vol. 4099: Q. Yang, G. Webb (Eds.), PRICAI 2006: Trends in Artificial Intelligence. XXVIII, 1263 pages. 2006.
- Vol. 4095: S. Nolfi, G. Baldassarre, R. Calabretta, J.C. T. Hallam, D. Marocco, J.-A. Meyer, O. Miglino, D. Parisi (Eds.), From Animals to Animats 9. XV, 869 pages. 2006.
- Vol. 4093: X. Li, O.R. Zaïane, Z. Li (Eds.), Advanced Data Mining and Applications. XXI, 1110 pages. 2006.
- Vol. 4092: J. Lang, F. Lin, J. Wang (Eds.), Knowledge Science, Engineering and Management. XV, 664 pages. 2006.
- Vol. 4088: Z.-Z. Shi, R. Sadananda (Eds.), Agent Computing and Multi-Agent Systems. XVII, 827 pages. 2006.
- Vol. 4087: F. Schwenker, S. Marinai (Eds.), Artificial Neural Networks in Pattern Recognition. IX, 299 pages. 2006.
- Vol. 4068: H. Schärfe, P. Hitzler, P. Øhrstrøm (Eds.), Conceptual Structures: Inspiration and Application. XI, 455 pages. 2006.
- Vol. 4065: P. Perner (Ed.), Advances in Data Mining. XI, 592 pages. 2006.
- Vol. 4062: G. Wang, J.F. Peters, A. Skowron, Y. Yao (Eds.), Rough Sets and Knowledge Technology. XX, 810 pages. 2006.
- Vol. 4049: S. Parsons, N. Maudet, P. Moraitis, I. Rahwan (Eds.), Argumentation in Multi-Agent Systems. XIV, 313 pages. 2006.
- Vol. 4048: L. Goble, J.-J.C. Meyer (Eds.), Deontic Logic and Artificial Normative Systems. X, 273 pages. 2006.
- Vol. 4045: D. Barker-Plummer, R. Cox, N. Swoboda (Eds.), Diagrammatic Representation and Inference. XII, 301 pages. 2006.
- Vol. 4031: M. Ali, R. Dapoigny (Eds.), Advances in Applied Artificial Intelligence. XXIII, 1353 pages. 2006.
- Vol. 4029: L. Rutkowski, R. Tadeusiewicz, L.A. Zadeh, J.M. Zurada (Eds.), Artificial Intelligence and Soft Computing – ICAISC 2006. XXI, 1235 pages. 2006.
- Vol. 4027: H.L. Larsen, G. Pasi, D. Ortiz-Arroyo, T. Andreassen, H. Christiansen (Eds.), Flexible Query Answering Systems. XVIII, 714 pages. 2006.
- Vol. 4021: E. André, L. Dybkjær, W. Minker, H. Neumann, M. Weber (Eds.), Perception and Interactive Technologies. XI, 217 pages. 2006.

- Vol. 4020: A. Bredendfeld, A. Jacoff, I. Noda, Y. Takahashi (Eds.), *RoboCup 2005: Robot Soccer World Cup IX*. XVII, 727 pages. 2006.
- Vol. 4013: L. Lamontagne, M. Marchand (Eds.), *Advances in Artificial Intelligence*. XIII, 564 pages. 2006.
- Vol. 4012: T. Washio, A. Sakurai, K. Nakajima, H. Takeda, S. Tojo, M. Yokoo (Eds.), *New Frontiers in Artificial Intelligence*. XIII, 484 pages. 2006.
- Vol. 4008: J.C. Augusto, C.D. Nugent (Eds.), *Designing Smart Homes*. XI, 183 pages. 2006.
- Vol. 4005: G. Lugosi, H.U. Simon (Eds.), *Learning Theory*. XI, 656 pages. 2006.
- Vol. 3978: B. Hnich, M. Carlsson, F. Fages, F. Rossi (Eds.), *Recent Advances in Constraints*. VIII, 179 pages. 2006.
- Vol. 3963: O. Dikenelli, M.-P. Gleizes, A. Ricci (Eds.), *Engineering Societies in the Agents World VI*. XII, 303 pages. 2006.
- Vol. 3960: R. Vieira, P. Quaresma, M.d.G.V. Nunes, N.J. Mamede, C. Oliveira, M.C. Dias (Eds.), *Computational Processing of the Portuguese Language*. XII, 274 pages. 2006.
- Vol. 3955: G. Antoniou, G. Potamias, C. Spyropoulos, D. Plexousakis (Eds.), *Advances in Artificial Intelligence*. XVII, 611 pages. 2006.
- Vol. 3949: F. A. Savacı (Ed.), *Artificial Intelligence and Neural Networks*. IX, 227 pages. 2006.
- Vol. 3946: T.R. Roth-Berghofer, S. Schulz, D.B. Leake (Eds.), *Modeling and Retrieval of Context*. XI, 149 pages. 2006.
- Vol. 3944: J. Quiñonero-Candela, I. Dagan, B. Magnini, F. d'Alché-Buc (Eds.), *Machine Learning Challenges*. XIII, 462 pages. 2006.
- Vol. 3930: D.S. Yeung, Z.-Q. Liu, X.-Z. Wang, H. Yan (Eds.), *Advances in Machine Learning and Cybernetics*. XXI, 1110 pages. 2006.
- Vol. 3918: W.K. Ng, M. Kitsuregawa, J. Li, K. Chang (Eds.), *Advances in Knowledge Discovery and Data Mining*. XXIV, 879 pages. 2006.
- Vol. 3913: O. Boissier, J. Padget, V. Dignum, G. Lindemann, E. Matson, S. Ossowski, J.S. Sichman, J. Vázquez-Salceda (Eds.), *Coordination, Organizations, Institutions, and Norms in Multi-Agent Systems*. XII, 259 pages. 2006.
- Vol. 3910: S.A. Brueckner, G.D.M. Serugendo, D. Hales, F. Zambonelli (Eds.), *Engineering Self-Organising Systems*. XII, 245 pages. 2006.
- Vol. 3904: M. Baldoni, U. Endriss, A. Omicini, P. Torroni (Eds.), *Declarative Agent Languages and Technologies III*. XII, 245 pages. 2006.
- Vol. 3900: F. Toni, P. Torroni (Eds.), *Computational Logic in Multi-Agent Systems*. XVII, 427 pages. 2006.
- Vol. 3899: S. Frintrop, *VOCUS: A Visual Attention System for Object Detection and Goal-Directed Search*. XIV, 216 pages. 2006.
- Vol. 3898: K. Tuyls, P.J. 't Hoen, K. Verbeeck, S. Sen (Eds.), *Learning and Adaption in Multi-Agent Systems*. X, 217 pages. 2006.
- Vol. 3891: J.S. Sichman, L. Antunes (Eds.), *Multi-Agent-Based Simulation VI*. X, 191 pages. 2006.
- Vol. 3890: S.G. Thompson, R. Ghanea-Hercock (Eds.), *Defence Applications of Multi-Agent Systems*. XII, 141 pages. 2006.
- Vol. 3885: V. Torra, Y. Narukawa, A. Valls, J. Domingo-Ferrer (Eds.), *Modeling Decisions for Artificial Intelligence*. XII, 374 pages. 2006.
- Vol. 3881: S. Gibet, N. Courty, J.-F. Kamp (Eds.), *Gesture in Human-Computer Interaction and Simulation*. XIII, 344 pages. 2006.
- Vol. 3874: R. Missaoui, J. Schmidt (Eds.), *Formal Concept Analysis*. X, 309 pages. 2006.
- Vol. 3873: L. Maicher, J. Park (Eds.), *Charting the Topic Maps Research and Applications Landscape*. VIII, 281 pages. 2006.
- Vol. 3864: Y. Cai, J. Abascal (Eds.), *Ambient Intelligence in Everyday Life*. XII, 323 pages. 2006.
- Vol. 3863: M. Kohlhase (Ed.), *Mathematical Knowledge Management*. XI, 405 pages. 2006.
- Vol. 3862: R.H. Bordini, M. Dastani, J. Dix, A.E.F. Seghrouchni (Eds.), *Programming Multi-Agent Systems*. XIV, 267 pages. 2006.
- Vol. 3849: I. Bloch, A. Petrosino, A.G.B. Tettamanzi (Eds.), *Fuzzy Logic and Applications*. XIV, 438 pages. 2006.
- Vol. 3848: J.-F. Boulicaut, L. De Raedt, H. Mannila (Eds.), *Constraint-Based Mining and Inductive Databases*. X, 401 pages. 2006.
- Vol. 3847: K.P. Jantke, A. Lunzer, N. Spyratos, Y. Tanaka (Eds.), *Federation over the Web*. X, 215 pages. 2006.
- Vol. 3835: G. Sutcliffe, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning*. XIV, 744 pages. 2005.
- Vol. 3830: D. Weyns, H. V.D. Parunak, F. Michel (Eds.), *Environments for Multi-Agent Systems II*. VIII, 291 pages. 2006.
- Vol. 3817: M. Faundez-Zanuy, L. Janer, A. Esposito, A. Satue-Villar, J. Roure, V. Espinosa-Duro (Eds.), *Nonlinear Analyses and Algorithms for Speech Processing*. XII, 380 pages. 2006.
- Vol. 3814: M. Maybury, O. Stock, W. Wahlster (Eds.), *Intelligent Technologies for Interactive Entertainment*. XV, 342 pages. 2005.
- Vol. 3809: S. Zhang, R. Jarvis (Eds.), *AI 2005: Advances in Artificial Intelligence*. XXVII, 1344 pages. 2005.
- Vol. 3808: C. Bento, A. Cardoso, G. Dias (Eds.), *Progress in Artificial Intelligence*. XVIII, 704 pages. 2005.
- Vol. 3802: Y. Hao, J. Liu, Y.-P. Wang, Y.-m. Cheung, H. Yin, L. Jiao, J. Ma, Y.-C. Jiao (Eds.), *Computational Intelligence and Security, Part II*. XLII, 1166 pages. 2005.
- Vol. 3801: Y. Hao, J. Liu, Y.-P. Wang, Y.-m. Cheung, H. Yin, L. Jiao, J. Ma, Y.-C. Jiao (Eds.), *Computational Intelligence and Security, Part I*. XLI, 1122 pages. 2005.
- Vol. 3789: A. Gelbukh, Á. de Albornoz, H. Terashima-Marín (Eds.), *MICA 2005: Advances in Artificial Intelligence*. XXVI, 1198 pages. 2005.

¥359.00元

Preface

This volume is the collection of papers and abstracts of the Third Annual International Workshop on the Emergence and Evolution of Linguistic Communication (EELCIII), held in Rome from September 30 to October 1, 2006. This workshop was the third in line after previous editions held in Kanazawa (Japan) in 2004 and in Hatfield (UK) in 2005. Although the previous events were published as post-proceedings, this event was the first to have its proceedings published at the workshop. Three types of papers were elicited: full papers, invited full papers and invited abstracts. All full papers were peer-reviewed by the International Programme Committee.

The workshop's focus was on the evolution and emergence of language. This is a fast-growing interdisciplinary research area with researchers coming from disciplines such as anthropology, linguistics, psychology, primatology, neuroscience, cognitive science and computer science. Although most papers focus on evolution, a number of papers focus more on language acquisition. This was highly welcomed, since research on language acquisition (both from psychology and artificial intelligence) is extremely important in gaining insights regarding language evolution and, not least, regarding the theme of this workshop 'Symbol Grounding and Beyond.'

Despite the interdisciplinarity of the field and – in principle – of the EELC series, most contributions stem from computer science (mainly artificial intelligence and artificial life). This is not surprising, because this was also the case in previous workshops and because this workshop was part of the 'Simulation of Adaptive Behavior' conference (SAB 2006), a.k.a. 'From Animals to Animats.'

We would like to thank those involved in the organisation of SAB 2006, especially Stefano Nolfi, for their assistance in organising the workshop, the members of the Programme Committee for their assistance in reviewing the papers, the invited speakers (Peter Gärdenfors, Naoto Iwahashi, Elena Lieven, Deb Roy and Luc Steels) and, of course, all authors of the contributions in this collection.

June 2006

Paul Vogt
Yuuya Sugita
Elio Tuci
Chrystopher Nehaniv

Organisation

EELCIII was organised by the Organising Committee as part of the Simulation of Adaptive Behavior conference.

Organising Committee

Paul Vogt (Tilburg University, The Netherlands)	Chair
Yuuya Sugita (RIKEN BSI, Japan)	Co-chair
Elio Tuci (Université Libre de Bruxelles, Belgium)	Co-chair
Chrystopher Nehaniv (University of Hertfordshire, UK)	Co-chair

Programme Committee

Takaya Arita (University of Nagoya, Japan)
Tony Belpaeme (University of Plymouth, UK)
Bart de Boer (University of Groningen, The Netherlands)
Angelo Cangelosi (University of Plymouth, UK)
Tecumseh Fitch (University of St. Andrews, UK)
Takashi Hashimoto (JAIST, Japan)
Jim Hurford (University of Edinburgh, UK)
Takashi Ikegami (University of Tokyo, Japan)
Simon Kirby (University of Edinburgh, UK)
Caroline Lyon (University of Hertfordshire, UK)
Davide Marocco (ISTC, National Research Council, Italy)
Chrystopher Nehaniv (University of Hertfordshire, UK)
Stefano Nolfi (ISTC, National Research Council, Italy)
Kazuo Okanoya (RIKEN BSI, Japan)
Tetsuo Ono (Future University Hakodate, Japan)
Domenico Parisi (ISTC, National Research Council, Italy)
Akito Sakurai (Keio University, Japan)
Andrew Smith (University of Edinburgh, UK)
Kenny Smith (University of Edinburgh, UK)
Luc Steels (Vrije Universiteit Brussel, Belgium)
Yuuya Sugita (RIKEN BSI, Japan)
Jun Tani (RIKEN BSI, Japan)
Satoshi Tojo (JAIST, Japan)
Elio Tuci (Université Libre de Bruxelles, Belgium)
Paul Vogt (Tilburg University, The Netherlands)

Table of Contents

A Hybrid Model for Learning Word-Meaning Mappings	1
<i>Federico Divina, Paul Vogt</i>	
Cooperation, Conceptual Spaces and the Evolution of Semantics	16
<i>Peter Gärdenfors, Massimo Warglien</i>	
Cross-Situational Learning: A Mathematical Approach	31
<i>Kenny Smith, Andrew D.M. Smith, Richard A. Blythe, Paul Vogt</i>	
Dialog Strategy Acquisition and Its Evaluation for Efficient Learning of Word Meanings by Agents	45
<i>Ryo Taguchi, Kouichi Katsurada, Tsuneo Nitta</i>	
Evolving Distributed Representations for Language with Self-Organizing Maps	57
<i>Simon D. Levy, Simon Kirby</i>	
How Do Children Develop Syntactic Representations from What They Hear?	72
<i>Elena Lieven</i>	
How Grammar Emerges to Dampen Combinatorial Search in Parsing	76
<i>Luc Steels, Pieter Wellens</i>	
Implementation of Biases Observed in Children's Language Development into Agents	89
<i>Ryo Taguchi, Masashi Kimura, Shuji Shinohara, Kouichi Katsurada, Tsuneo Nitta</i>	
Lexicon Convergence in a Population With and Without Metacommunication	100
<i>Zoran Macura, Jonathan Ginzburg</i>	
Operational Aspects of the Evolved Signalling Behaviour in a Group of Cooperating and Communicating Robots	113
<i>Elio Tuci, Christos Ampatzis, Federico Vicentini, Marco Dorigo</i>	
Propositional Logic Syntax Acquisition	128
<i>Josefina Sierra-Santibáñez</i>	

Robots That Learn Language: Developmental Approach to
Human-Machine Conversations 143
 Naoto Iwahashi

Simulating Meaning Negotiation Using Observational Language
Games 168
 Tiina Lindh-Knuutila, Timo Honkela, Krista Lagus

Symbol Grounding Through Cumulative Learning 180
 Samarth Swarup, Kiran Lakkaraju, Sylvian R. Ray, Les Gasser

The Human Speechome Project 192
 Deb Roy, Rupal Patel, Philip DeCamp, Rony Kubat,
 Michael Fleischman, Brandon Roy, Nikolaos Mavridis,
 Stefanie Teller, Alexia Salata, Jethran Guinness, Michael Levit,
 Peter Gorniak

Unify and Merge in Fluid Construction Grammar 197
 Luc Steels, Joachim De Beule

Utility for Communicability by Profit and Cost of Agreement 224
 Ryuichi Matoba, Makoto Nakamura, Satoshi Tojo

Author Index 237

A Hybrid Model for Learning Word-Meaning Mappings^{*}

Federico Divina¹ and Paul Vogt^{1,2}

¹ Induction of Linguistic Knowledge / Language and Information Science
Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands

² Language Evolution and Computation Research Unit, School of Philosophy,
Psychology and Language Sciences, University of Edinburgh, UK
{f.divina@uvt.nl, paulv}@ling.ed.ac.uk

Abstract. In this paper we introduce a model for the simulation of language evolution, which is incorporated in the New Ties project. The New Ties project aims at evolving a cultural society by integrating evolutionary, individual and social learning in large scale multi-agent simulations. The model presented here introduces a novel implementation of language games, which allows agents to communicate in a more natural way than with most other existing implementations of language games. In particular, we propose a hybrid mechanism that combines cross-situational learning techniques with more informed feedback mechanisms. In our study we focus our attention on dealing with referential indeterminacy after joint attention has been established and on whether the current model can deal with larger populations than previous studies involving cross-situational learning. Simulations show that the proposed model can indeed lead to coherent languages in a quasi realistic world environment with larger populations.

1 Introduction

For language to evolve, the language has to be transmitted reliably among the population, which is only possible if the individual agents can learn the language. In human societies, children have to learn for instance the sounds, words and grammar of the target language. In the current paper, we focus solely on the evolution and acquisition of word-meaning mappings. The way children acquire the meanings of words still remains an open question. Associating the correct meaning to a word is extremely complicated, as a word may potentially have an infinite number of meanings [1].

Different mechanisms that children may adopt when acquiring the meanings of words have been suggested, see, e.g., [2] for an overview. For example, Tomasello has proposed that *joint attention* is a primary mechanism [3]. According to this

^{*} This research and the New Ties project is supported by an EC FET grant under contract 003752. We thank all members the New Ties project for their invaluable contributions. Opinions and errors in this manuscript are the authors' responsibility, they do not necessarily reflect those of the EC or other New Ties members.

mechanism, children are able to share their attention with adults on objects, e.g., through gaze following or pointing. Moreover, children can learn that adults have control over their perceptions and that they can choose to attend to particular objects or aspects of a given situation. This allows children to focus their attention on the same situation experienced by adults, thus reducing the number of possible meanings of a word.

This mechanism, however, is not sufficient, because it is still uncertain whether a word relates to the whole situation, to parts of the situation or even to a completely different situation. This is known as the *referential indeterminacy* problem illustrated by Quine [1] with the following example: Imagine an anthropologist studying a native speaker of an unfamiliar language. As a rabbit crosses their visual field, the native speaker says “gavagai” and the anthropologist infers that “gavagai” means *rabbit*. However, the anthropologist cannot be completely sure of his inference. In fact, the word “gavagai” can have an infinite number of possible meanings, including *undetached rabbit parts*, *large ears*, *it’s running*, *good food* or even *it’s going to rain*.

To overcome this problem, additional mechanisms have been proposed to reduce the referential indeterminacy. Among these is a representational bias known as the *whole object bias* [4], according to which children tend to map novel words to whole objects, rather than to parts of objects. Another mechanism that children appear to use is the *principle of contrast* [5], which is based on the assumption that if a meaning is already associated with a word, it is unlikely that it can be associated with another word.

There is also evidence that children can acquire the meanings of words more directly by reducing the number of potential meanings of words across different situations [6,7]. This *cross-situational learning* can work statistically by maintaining the co-occurrence frequencies of words with their possible meanings [8,9] or simply by maintaining the intersection of all situations in which a word is used [10,11]. Crucially, cross-situational learning depends on observing a sufficient degree of one-to-one mappings between words and meanings. Although theoretically, the level of uncertainty (i.e. the number of confounding – or background – meanings) in situations may be quite large, this may have a large impact on the time required to learn a language [11].

Cross-situational learning yields poor results when the input language is less consistent regarding the one-to-one mapping. This has been found in simulation studies of language evolution with increased population sizes [9]. In such simulations, different agents create many different words expressing the same meaning when they have not yet communicated with each other. So, the more agents there are, the more words can enter a language community during the early stages of evolution. In models that use explicit meaning transfer, there are positive feedback loops that reduce the number of words sufficiently over time, allowing the language to converge properly [12]. However, when there is no positive feedback loop, as is the case with cross-situational learning, there appears to be no efficient mechanism for reducing the number of words in the language. A possible solution

to this problem could be to include an additional mechanism that imposes a bias toward one-to-one mappings between words and meanings [13].

In this paper we propose a hybrid model for the evolution of language that combines joint attention, cross-situational learning and the principle of contrast as mechanisms for reducing the referential indeterminacy. In addition, a feedback mechanism and related adaptations are used as a synonymy damping mechanism. This model is used to investigate the effect that context size has on the development of language, but more importantly it is used to investigate how this model can deal with large populations. The model is embedded in the New Ties project¹, which aims at developing a benchmark platform for studying the evolution and development of cultural societies in very large multi-agent systems [14].

The paper is organised as follows: in the next section, we provide a brief description of the proposed model (for details, consult [14,15]). In Section 3 we present some experiments, whose aims are to show that the proposed hybrid model can lead to the evolution of a coherent lexicon in large population sizes and with varying context sizes. The results are discussed in Section 4. Finally, Section 5 concludes.

2 The Model

2.1 New Ties Agent Architecture

The New Ties project aims at developing a platform for studying the evolution and development of cultural societies in a very large multi-agent system. In this system, agents are inserted in an environment consisting of a grid world in which each point is a location. The world, which is inspired by Epstein & Axtell's [16] sugar scape world, is set up with tokens, edible plants, building bricks, agents, different terrains of varying roughness, etc. The aim for the agents is to evolve and learn behavioural skills in order for the society to survive over extended periods of time. As part of these skills, language and culture are to develop.

At each time step each agent receives as input a set of perceptual features and messages, which constitute the context of an agent, and outputs an action (see Fig. 1 for the basic agent architecture). These actions are collected by the environment manager, and when all agents have been processed, the collected actions are executed and the environment is updated.

The perceptual features an agent receives represent both objects and actions that occur in its visual field. These features are processed with a categorisation mechanism based on the discrimination game [17] (a detailed description of this mechanism is given in [14,18]). Basically, each object is mapped onto a set of categories, where each category corresponds to a feature. So, if an object is described by n features, it will be categorised into n categories. Messages are processed with a language interpretation module, described in Section 2.2, and also yield a set of categories. All these categories are stored in the short-term memory (STM), which can be accessed by the control module, as well as all other modules.

¹ New Ties stands for New Emerging World models Through Individual, Evolutionary and Social learning. See <http://www.new-ties.org>

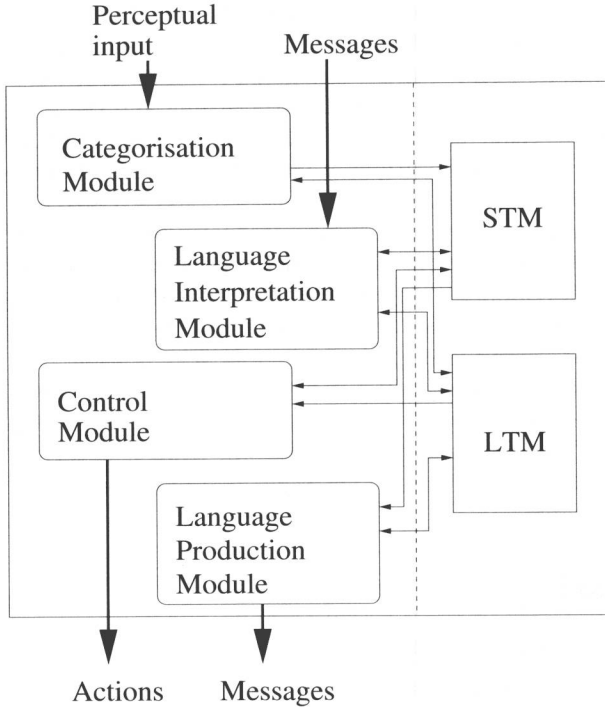


Fig. 1. The basic architecture of a New Ties agent. Perceptual features of objects and actions are processed by the categorisation module, while messages are interpreted with the language interpretation module. The control module outputs actions and the language production module produces outgoing messages. Various sources of knowledge are stored in the short- and long-term memories.

Once the perceptual features and messages have been processed, the controller is used to determine the action to perform. This controller is represented by a *decision Q-tree* (DQT), which is a decision tree that can change during an agent's lifetime using reinforcement learning [14]. The possible actions include, among others, *move*, *turn left*, *turn right*, *mate*, *talk*, *shout*, ... In case the output of the DQT is either the *talk* or *shout* action, the agent must produce a message, which is done by the language production module, described below. Each action performed costs a certain amount of energy, and when an agent's energy level decreases to zero or below, it dies. Energy levels can be increased by eating plants. Agents also die when they reach a predefined age.

Agents start their life with a small initial DQT, which, as mentioned above, can be changed by reinforcement learning. This initial DQT is the result of evolution. When two agents reproduce, they produce an offspring who inherits its genome from its parents, subject to cross-over and mutations. This genome carries the code for producing the initial DQT and other biases, which regulate, for instance, the 'socialness' of the agent. This socialness gene is a bias for

	$m_1 \dots m_N$		$m_1 \dots m_N$
w_1	$\sigma_{11} \dots \sigma_{1N}$	w_1	$P_{11} \dots P_{1N}$
\vdots	$\vdots \quad \vdots \quad \vdots$	\vdots	$\vdots \quad \vdots \quad \vdots$
w_M	$\sigma_{M1} \dots \sigma_{MN}$	w_M	$P_{M1} \dots P_{MN}$

Fig. 2. A simplified illustration of the lexicon. The lexicon consists of two matrices that associate meanings m_j with words w_i . The left matrix stores association scores σ_{ij} and the right matrix stores co-occurrence probabilities P_{ij} .

an agent to be social; the more social an agent is, the more frequently it will communicate and the more likely it is to provide more information regarding the meaning of a message. Unlike standard evolutionary algorithms, reproduction is not processed cyclical, but acyclical, i.e., two agents can reproduce when they decide to, but only if they are of different sex and in nearby locations.

2.2 Communication and Learning Word-Meaning Mappings

The language evolves in the society by agents' interacting through language games. While doing so, each individual constructs its own lexicon, which is represented in the long-term memory (LTM) by two association matrices (Fig. 2). Each matrix associates words w_i with meanings m_j . The first matrix stores association scores σ_{ij} , while the second stores co-occurrence probabilities P_{ij} . The former is updated based on feedback the agents may receive regarding the effectiveness (or success) of their interaction. However, as this feedback is not always available, the agents also maintain the co-occurrence frequencies of words and the potential meanings as they co-occur in a given situation (or context). The two matrices are coupled via the *association strength*, $strL_{ij}$, which is calculated as:

$$strL_{ij} = \sigma_{ij} + (1 - \sigma_{ij})P_{ij}. \quad (1)$$

This coupling allows the agents to infer the right word-meaning mappings across different situations using the co-occurrence probabilities when there has been little feedback. However, when there has been sufficient feedback on the language use of the agents, the association score σ_{ij} may become high enough to overrule the co-occurrence probabilities.

Both matrices are updated after each language game. If a language game is considered successful based on the feedback mechanism, the association score σ_{ij} of the used association is increased by

$$\sigma_{ij} = \eta \cdot \sigma_{ij} + 1 - \eta, \quad (2)$$

where $\eta = 0.9$ is a constant learning parameter. In addition, the scores of competing associations are laterally inhibited by

$$\sigma_{ij} = \eta \cdot \sigma_{ij}. \quad (3)$$

An association α_{nm} is competing if either the word is the same ($n = i$) or the meaning ($m = j$), but not both. If the game has failed according to the feedback mechanism, σ_{ij} is also decreased this way. The association score is unchanged if no feedback is processed.

In each game, irrespective of its outcome, the co-occurrence frequencies f_{ij} of words with potential meanings in that situation are increased, thus affecting the co-occurrence probabilities:

$$P_{ij} = \frac{f_{ij}}{\sum_i f_{ij}}. \quad (4)$$

The reason for adopting this dual representation is that earlier studies have indicated that using the mechanism for updating the association scores (Eqs. 2 and 3) work much better than for updating the co-occurrence probabilities (Eq. 4) if there is feedback, while the opposite is true for cross-situational learning [19].

Unlike standard implementations, such as [17,18], a language game is initiated by an agent when its controller decides to talk or shout², or otherwise with a certain probability proportional to socialness gene. This agent (the speaker) then selects an arbitrary object from its context as a *target object*³ and decides on how many words it will use to describe the object. This number, expressed in the *task complexity* T_c , is determined by generating a random number between 1 and 5 following a Gaussian distribution with the average age of the target audience in tens of ‘New Ties years’ (NTYrs)⁴ as its mean and a standard deviation of 0.75. This way, the agent will tend to produce shorter messages when addressing a young audience and longer messages when addressing an older audience.

Depending on this task complexity, the agent selects arbitrarily T_c different categories that represent the object. Recall that each category relates to one perceptual feature of an object, such as the object’s colour, shape, distance or weight. For each category, the speaker then searches its lexicon for associations that have the highest strength $strL_{ij}$. If no such association is found, a new word is invented as an arbitrary string and added to the lexicon. Each word thus found is then appended to the message which is distributed to the agent(s) in the speaker’s vicinity.

On certain occasions, for instance, when the hearer had signalled that it did not understand the speaker, the speaker may accompany the message with a pointing gesture to draw the attention to the target (such a gesture is only produced with a probability proportional to the socialness gene mentioned earlier). This way, the agents establish joint attention, but still the hearer does not necessarily know exactly what feature of the object is signalled (cf. Quine’s problem).

² The ‘talk’ action is directed to only one visible agent, while ‘shout’ is directed to all agents in the audible vicinity of the initiator.

³ In later studies we intend to make this selection depending on the decision making mechanism determined by the DQT, so the communication will be more functional with respect to the agent’s behaviour.

⁴ In the current paper, a year in ‘New Ties time’ equals to an unrealistic 365 time steps.

When an agent receives a message, its language interpretation module tries to interpret each word in the message by searching its lexicon for associations with the highest strength $strL_{ij}$. If the association score σ_{ij} of this element exceeds a certain threshold (i.e., $\sigma_{ij} > \Theta$, where $\Theta = 0.8$), then the hearer assumes the interpretation to be correct. If not, the hearer may – with a certain probability proportional to the socialness gene – consider the interpretation to be incorrect and signal a ‘did not understand’ message, thus soliciting a pointing gesture; otherwise, the hearer will assume the interpretation was correct.

In case the interpretation was correct, the hearer may – again with a probability proportional to its socialness gene – signal the speaker that it understood the message, thus providing feedback so that both agents increase the association score of used lexical entries and inhibit competing elements as explained above. In all cases, the co-occurrence probability P_{ij} is increased for all categories in the context that have an association with the expressed words. In case the speaker had pointed to the object, this context is reduced to the perceptual features of this object. Otherwise, the context contains all categories of all visible objects, which may differ from those the speaker sees – including the target object. All interpretations are added to the STM, which the controller uses to decide on the agent’s next action.

When no interpretation could be found in the lexicon, the agent adds the novel word to its lexicon in association with all categories valid in the current context (i.e., either all objects and events perceived or the object that was pointed to). The frequency counters of these associations are set to 1 and the association scores σ_{Nj} are initialised with:

$$\sigma_{Nj} = (1 - \max_i(\sigma_{ij}))\sigma_0, \quad (5)$$

where $\max_i(\sigma_{ij})$ is the maximum association score that meaning m_j has with other words w_i , $\sigma_0 = 0.1$ is a constant, and $i \neq N$. This way, if the agent has already associated the meaning (or category) m_j with another word w_i , the agent is biased to prefer another meaning with this novel word. Hence, this implements a notion of the principle of contrast [5]. Note again that the hearer may not have seen the target object and thus may fail to acquire the proper meaning.

3 Experiments

In the experiments we test the effectiveness of the model described in the previous section. In particular, we are interested to see whether reasonable levels of communicative accuracy can be reached with relatively large populations. In addition, we investigate the influence of considering a different number of perceptual features that agents have at their disposal for inferring word-meaning mappings. In order to focus on these questions, the evolutionary and reinforcement learning mechanisms were switched off. So, although agents could reproduce, each agent has exactly the same hand-crafted controller that did not change during their lifetimes. As a result, in the simulations reported here, agents only