Collective Learning of an Emergent Vocabulary: Naming Game with Reinforcement Learning

İsmet Adnan Öztürel
adnan@ii.metu.edu.tr

Orçun Orkan Özcan
orkan@ii.metu.edu.tr

Middle East Technical University
Cognitive Science Department

Central European Conference in Linguistics for Graduate Students
31 August 2011

Outline

Problem Definition

Can we capture emergent properties of language in a computational model of a community by presuming that language is a social tool?

Late Wittgensteinian language games conceptualize linguistic communication as a game like activity with two assumptions:

- Form of life hypothesis
- Rule following activity

i.e. PI2 provides a detailed description of an idealized minimal language game, namely the builders' language game. [*]

[*] Wittgenstein, L. (1953)

Complex Adaptive Systems (CAS)

Dynamic multi-agent systems can be classified as CAS if large number of local interactions create an adaptive behavior. *[*]*

i.e. Immune system, connectionist network models, etc.

Common properties of CAS:

- Evolution
- Aggregate Behavior
- Anticipation
- Simultaneous Interaction
- Self-Similarity

[*] Holland, J. (1992)

Emergence

A system exhibits emergence when there are coherent emergents at the macro-level that dynamically arise from the interactions between the parts at the micro-level. Such emergents are novel with respect to the individual parts of the system. *[*]*

Essential conditions for emergent behavior to arise:

- Micro-macro Effect
- Radical Novelty
- Coherence
- Interacting Parts
- Dynamicity
- Decentralized Control
- Two-Way Link
- Robustness and Flexibility

[*] De Wolf, T., & Holvoet, T. (2005)

Self-Organization

Self-organization is a dynamical and adaptive process where systems acquire and maintain structure themselves, without external control. [*]

Essential conditions for self-organization to arise:

- Increase in Order
- Autonomy
- Robustness
- Dynamical

[*] De Wolf, T., & Holvoet, T. (2005)

Emergence vs. Self-Organization

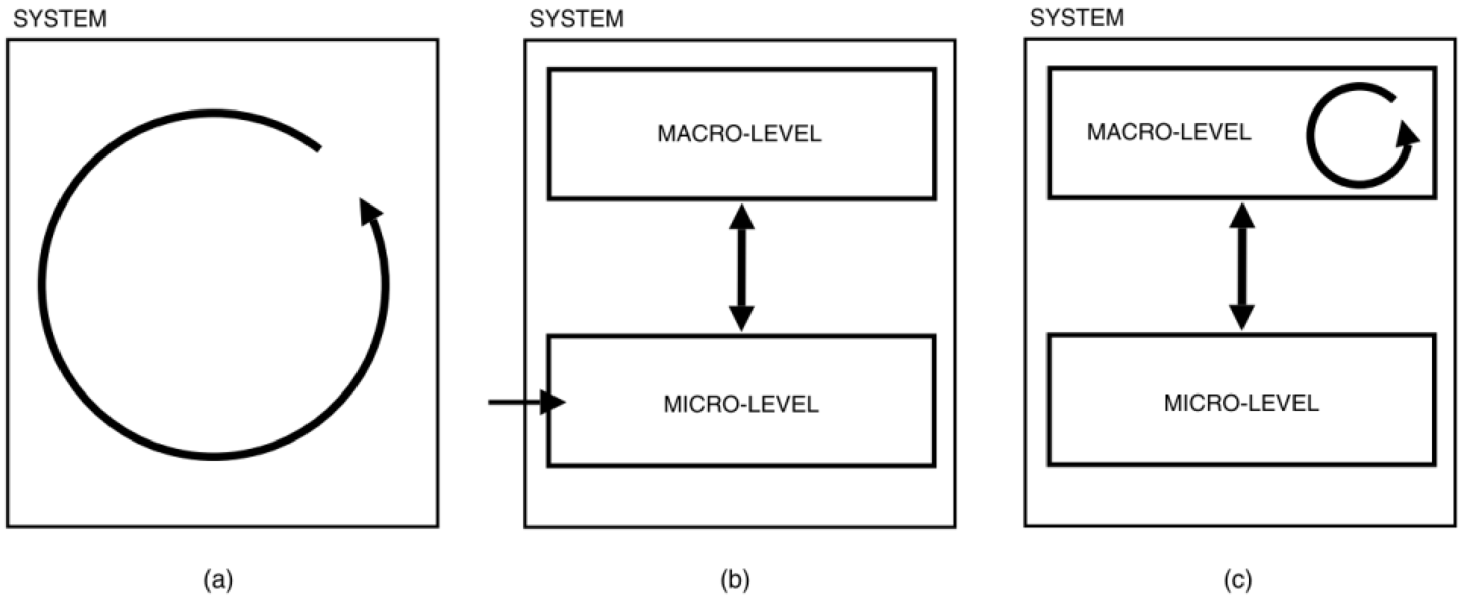

Fig.1. Comparison of Emergent and Self-Organizing Behaviors

(a) Solely Self-Organizing
(b) Solely Emergent
(c) Self-Organization and Emergence all together

Figure taken from De Wolf, T., & Holvoet, T. (2005)

Computational Modeling

Fundamental Methodology: **Classifier Systems** [*]

- Data driven and deterministic models are not suitable
- Computational exploratory research that facilitates computer based programs or simulations
- Models should include descriptions for agents and their interactions

Components of a Classifier System:

- Set of Classifiers
  - Non-static, can be expanded with rule discovery
  - Classifiers compete with each other
- Set of Reservoirs
- Effectors
- Detectors
- Set of Signals

[*] Holland, J. (2006)

Language as a Complex Adaptive System

Computational models of language games consist:

- Agents
  - Psychological Competence
  - Cognitive Competence
  - Enacting Script
  - Discrete Memory
- Environment
  - Non-restrictive
  - Communicative signals are token type
- Interactions are regulated with game rules

Language as a Complex Adaptive System

Broad generalization of language game rules:

1. A pair of speaker and hearer is selected randomly among the population.
2. Both parties attend to a topic, which could be anything related with the environment (e.g. an object). Topic must be perceivable by both parties.
3. Speaker transmits a feature about the topic to the hearer.
4. Hearer assesses whether it has that feature associated with the topic in its memory.
5. Both parties update their memory in accordance with the result of this assessment.

Discrimination Game

Rules of the discrimination game [*] :

1. Speaker attends to a portion of the perceptual space and shares this selection with the hearer as the *context*.
2. Within the *context* an object or a perceptual entity is chosen as the *topic* by the speaker.
3. Speaker transmits a word that represents a feature about the topic from its memory to the hearer.
4. Hearer tries to identify the topic in the context by using the word provided by the speaker.
5. Both parties are informed about the success of this identification and they perform memory updates accordingly.

[*] Steels, L. (2003).

Naming Game

Rules of the naming game [*]:

1. One speaker and one hearer is chosen randomly from the population.
2. Speaker transmits a name to the hearer
   (a) if its inventory is empty, a new word is invented.
   (b) if there is more than one name in the inventory for that object, then it is transmitted.
   (c) if there is more than one name, randomly transmits one of them.
3. Hearer process the uttered name.
   (a) if the uttered name is in hearers inventory, then the game is a *success*.
   (b) otherwise, game is a *failure*.
4. Final modifications on participants inventories.
   (a) if success both parties delete all the words from their inventories expect the one, which is agreed on.
   (b) if failure only hearer updates its inventory by adding the new word to its inventory.

[*] Baronchelli, A., Felici, M., Loreto, V., Caglioti, E., & Steels, L. (2006)

Model – Background Assumptions

- **No topology:**
  - Agents are equally distant to each other and they are equally probable to participate in an interaction.
- **Featureless object identification:**
  - Objects have no features and agents are capable of identifying objects in the same way.
- **Population opacity:**
  - Agents’ lexicons and their knowledge about the environment are inaccessible to their peers. Agents can not have a representation of the overall system behavior.
- **No parallel interactions:**
  - Interactions are concurrent because of the computational restrictions.

Model – Agents and Interactions

Population $A = \{a_1, \dots, a_n\}$ with size $N_A = n$
Environment $O = \{o\}$ with size $N_O = 1$
Inventory $L = \{\sigma_1, \dots, \sigma_m\}$ at a given time t

At the outset L = Ø

Generic Minimal Naming Game:

1. Speaker a_x and hearer a_y is chosen from A
2. a_x utters a word w from L with selection function α
   $$\alpha(L_x) = w = \text{random}(L)$$
3. a_y processes the word with assessment function β
   $$\beta(w) = \begin{cases} \text{success} & \text{if } w \in L_y \\ \text{failure} & \text{if not } w \in L_y \end{cases}$$
4. if success L_x and $L_y = \{w\}$
   if failure $L_y' = L_y + w$

Model – Agents and Interactions

Naming Game with Pair Selection Strategy:

Agents have a probability table *T* to remember how successful it communicated with each and every agent of the population.

Probability Table $T = \{P_1, \ldots, P_n\}$ with size $N_T = N_A = n$

1. Speaker a_x chosen randomly from *A*
2. A_x chooses a hearer counter part
   $a_y = max(T)$
   a_y = random(*A*) for once in each δ interaction
3. a_x utters a word *w* from *L* with selection function α
   $\alpha(L_x) = w = \text{random}(L)$
4. a_y processes the word with assessment function β
   $$\beta(w) = \begin{cases} \text{success} & \text{if } w \in L_y \\ \text{failure} & \text{if not } w \in L_y \end{cases}$$
5. Lexicon updates are the same with generic game
   T is updated with eGreedy algorithm

Model – Agents and Interactions

Naming Game with Word Selection Strategy:

Lexicon is static $L = \{\sigma_1, \ldots, \sigma_m\}$ throughout the game.
Agents have a probability table T to remember how successful it communicated with each word.

Probability Table $T = \{P_1, \ldots, P_n\}$ with size $N_T = N_L$

1. Speaker a_x and hearer a_y is chosen from A
2. A_x utters a word w from L
   $w = max(T)$
   w = random(L) for once in each δ interaction
3. a_y processes the word with assessment function β

$$\beta(w) = \begin{cases} \text{success} & \text{if } w \in L_y \\ \text{failure} & \text{if not } w \in L_y \end{cases}$$

4. Lexicon updates are the same with generic game
   In T_x value of w is updated with eGreedy

Model – Agents and Interactions

e-Greedy Learning Scheme:

The agent selects at each time step a random action with a fixed probability, $0 \leq \varepsilon \leq 1$, instead of selecting greedily one of the learned optimal actions:

$$\pi(s) = \begin{cases} \text{random action from } A(s) & \text{if } \xi < \varepsilon \\ \text{argmax}_{a \in A(s)} Q(s, a) & \text{otherwise} \end{cases}$$

where $0 \leq \xi \leq 1$ is a uniform random number drawn at each time step.

Model – Experiments

- **Test for absolute convergence on a shared vocabulary.**
  Generic, pair and word selection algorithms are compared according to the N_w , N_d and S, for 20 agents with an eGreedy exploration rate of 0.2, which is averaged over 200 runs.

- **Test for effects of exploration rate on learning.**
  The effects of rate of exploration examined on convergence trends of pair and word selection models, for 20 agents averaged over 200 runs.

- **Test for effects of learning rate on convergence.**
  The effects of reward/punishment rate on convergence trends of pair and word selection models, for 20 agents averaged over 200 runs.

Empirical Results – Generic Algorithm

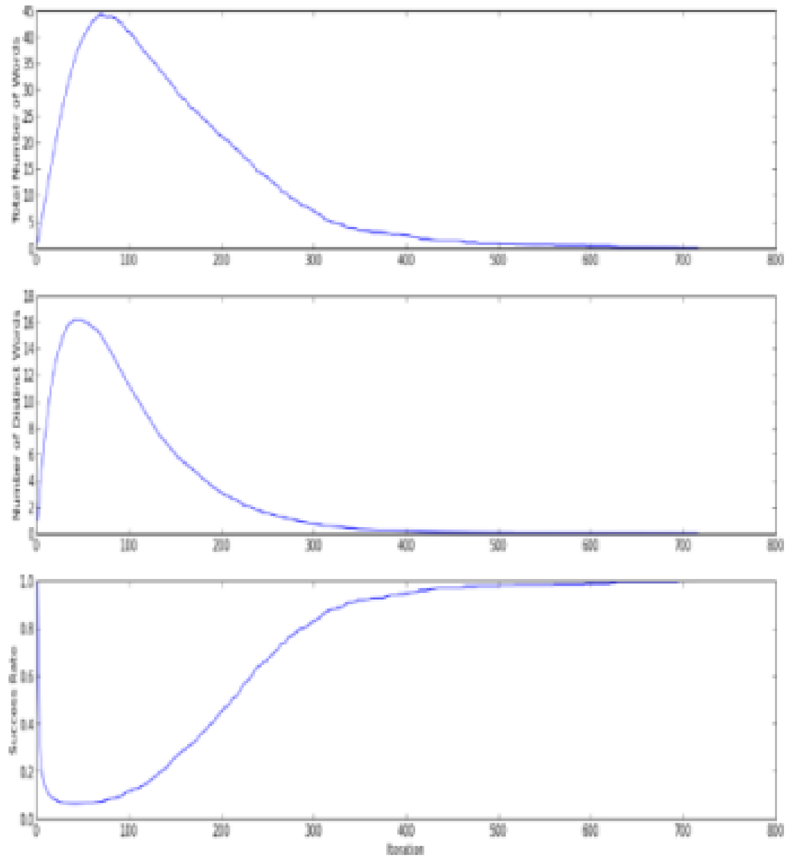

Fig. 2. N_w, N_d and S for generic naming game algorithm with exploration rate 0.2 (results are averaged over 200 runs for 20 agents).

Empirical Results – Modified Algorithms

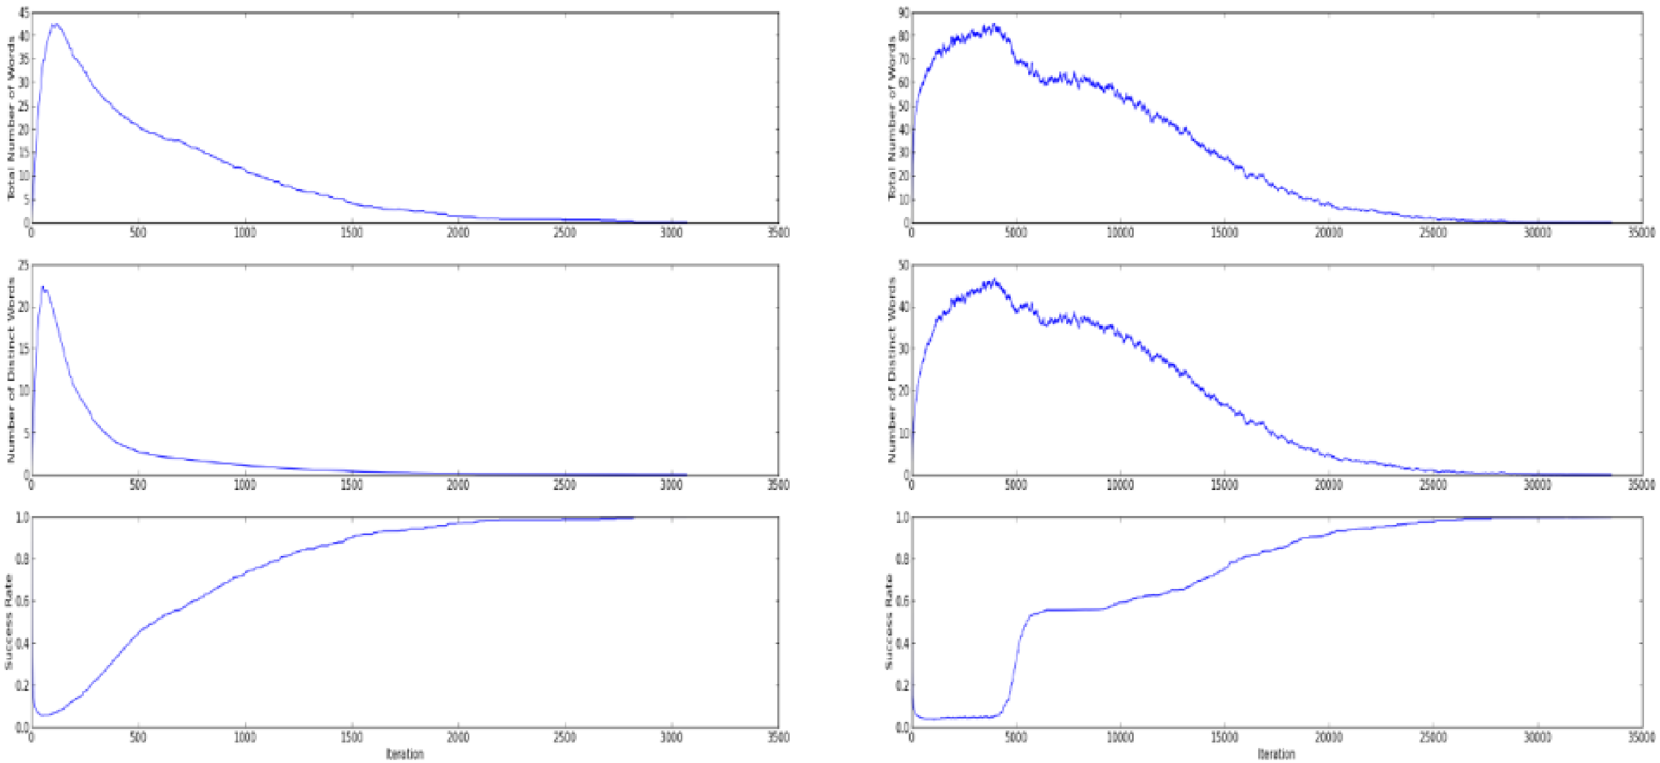

Fig. 3. N_w, N_d and S for pair (left) and word (right) selecting algorithms with exploration rate 0.2 (results are averaged over 200 runs for 20 agents).

Empirical Results – Exploration Rate

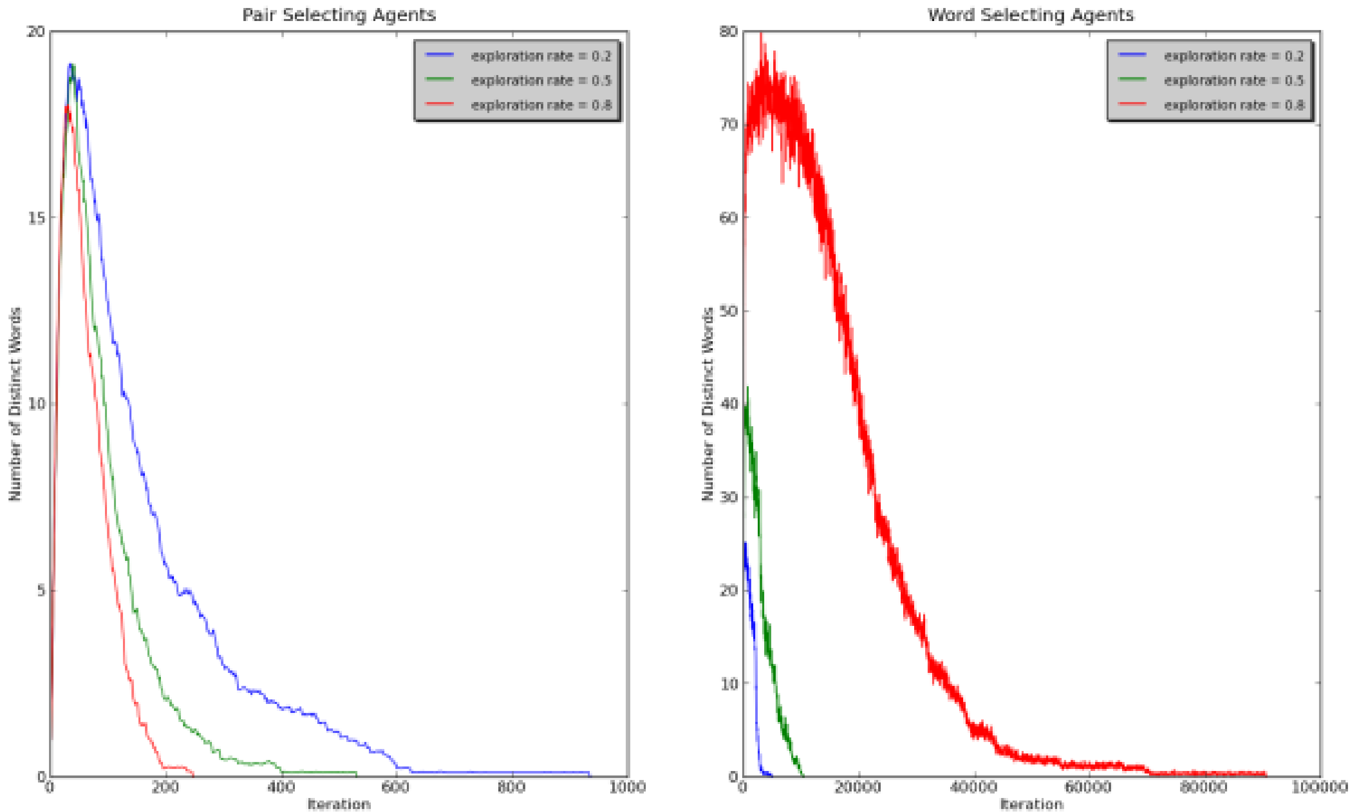

Fig. 4. N_d vs. *number of iterations* for pair (left) and word (right) selecting algorithms. Exploration rates are 0.2, 0.5 and 0.8 (results are averaged over 100 runs for 10 agents).

Empirical Results – Reward/Punishment Rate

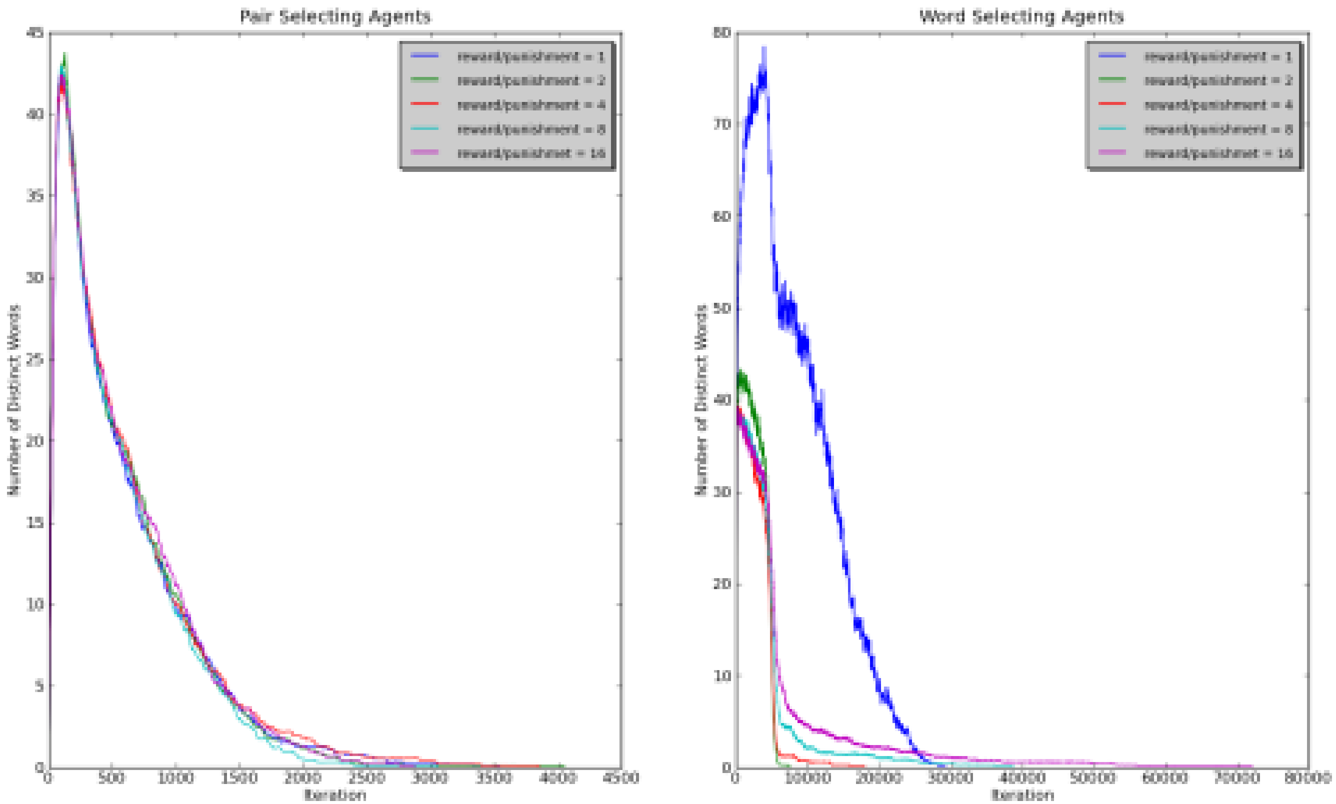

Fig. 5. N_d vs. *number of iterations* for pair (left) and word (right) selecting algorithms. Reward/Punishment rates are 1,2,4,8,16(results are averaged over 200 runs for 20 agents with exploration rate 0.2).

Conclusion

- An agreement on a shared vocabulary can be attained even if agents adopt restrictive strategies for agent and word selection throughout the naming game.
- Learning through exploration delays convergence, therefore reinforcement learning methods result in an overall deficiency in performance.
- Both pair selection and word selection algorithms result in a similar convergence trend, when compared with generic naming game. Yet, they are not preferable due to lack of performance.

Future Research:
A topology induced analysis could be carried out to see whether population groups under sub-communities. i.e. emergence of dialects?

References

- A. Baronchelli, L. Dall'Asta, A. Barrat, and V. Loreto: Strategies for fast convergence in semiotic dynamics. Arxiv preprint physics/0511201. (2005)
- A. Baronchelli, V. Loreto, L. DallAsta, and A. Barrat: Bootstrapping communication in language games: Strategy, topology and all that. Proceedings of the 6th International Conference on the Evolution of Language , pp. 1118. (2006)
- A. Baronchelli, M. Felici, V. Loreto, E. Caglioti, and L. Steels: Sharp transition towards shared vocabularies in multi-agent systems. Journal of Statistical Mechanics: Theory and Experiment, vol. 2006, p. P06014. (2006)
- A. Baronchelli, L. Dall'Asta, A. Barrat, and V. Loreto: The role of topology on the dynamics of the Naming Game. The European Physical Journal-Special Topics, vol.143, pp. 233235. (2007)
- A. Barrat, A. Baronchelli, L. DallAsta, and V. Loreto: Agreement dynamics on interaction networks with diverse topologies. Chaos: An Interdisciplinary Journal of Nonlinear Science, vol. 17, p. 026111. (2007)
- T. De Wolf and T. Holvoet: Emergence versus self-organisation: Different concepts but promising when combined. In Engineering self-organising systems: Methodologies and applications (Vol. 3464, p. 1-15). Springer-Verlag Berlin. (2005)
- J. Holland: Studying complex adaptive systems. Journal of Systems Science and Complexity, 19(1), 1-8. (2006).
- T. Gong, J. Ke, J.W. Minett, and W.S. Wang: A computational framework to simulate the coevolution of language and social structure. Artificial Life IX: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems, pp. 15864. (2004)
- T. Lenaerts, B. Jansen, K. Tuyls, and B. De Vylder: The evolutionary language game: An orthogonal approach. Journal of Theoretical Biology, vol. 235, pp. 566582. (2005)
- M.A. Nowak, J.B. Plotkin, and D.C. Krakauer: The evolutionary language game. Journal of Theoretical Biology, vol. 200, pp. 147162. (1999)
- X. Pan, J. Yang, and F. Jian-yong: Research on the Evolution of Corporate Culture Based on Naming Game. Computing, Communication, Control, and Management, 2008. CCCM'08. ISECS International Colloquium. (2008)
- L. Steels: Perceptually grounded meaning creation. Proceedings of the International Conference on Multi-Agent Systems, pp. 338344. (1996)
- R.S. Sutton and A.G. Barto: Reinforcement Learning: an introduction. MIT Press Cambridge, MA. (1998)
- P. Vogt: The emergence of compositional structures in perceptually grounded language games. Artificial Intelligence, vol. 167, pp. 206242. (2005)
- L. Wittgenstein: Philosophical investigations. Oxford: Blackwell. (1953)