The Grounded Naming Game

Luc Steels and Martin Loetzsch

This paper is the authors' draft and has now been officially published as:
Luc Steels and Martin Loetzsch (2012). The Grounded Naming Game. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*, 41–59. Amsterdam: John Benjamins.

Abstract

This chapter shows a concrete example of a language game experiment for studying the cultural evolution of one of the most basic functions of language, namely to draw attention to an object in the context by naming a characteristic feature of the object. If the object is a specific recognizable individual, then the name is called a proper name, and this is the case that is studied in this chapter. We investigate a concrete operational language strategy, with a conceptual as well a linguistic component, and show that a population of agents endowed with this strategy is able to self-organize a vocabulary of grounded proper names from scratch. The example provides a clear example of the role of alignment in stimulating self-organization and how expressive adequacy, cognitive effort, learnability, and social conformity act as selectionist forces, driving the population towards an effective language system.

1. Introduction

The Naming Game is the simplest possible kind of language game imaginable. It is a game of reference in which the speaker tries to draw the attention of the hearer to an object in the context by naming a characteristic feature of the object. If this feature is a unique individual identity, then the name is called a proper name. The speaker says for example ”(Where is) Pluto?” and the hearer points to the family dog walking by.

Proper names are only one of the ways in which the speaker can draw attention to an object. The speaker can also name a color, or a shape or some other perceptually grounded category or relation. As a matter of fact, referring expressions can become arbitrary complex, as in: “Where is this hairy monster that tried to bite in my leg yesterday evening?”. However, by focusing on proper names, many of the fundamental issues that come up in the emergence of a communication system, particularly the issue how a norm can become shared, can already be studied without raising complex questions about compositional semantics or grammar. It is therefore not surprising that the Naming Game, which was first introduced in (Steels, 1995), has become a model system for many studies in semiotic dynamics so that it has acquired a similar status as the Prisoner Dilemma game in agent-based socio-economic studies (Baronchelli et al., 2006).

Here is the scenario of the Naming Game: Assume a population `P` of agents, and a world `W` consisting of a set of individual objects. Two members are randomly selected from the population and take on the roles of speaker and hearer respectively. The context `C` contains a subset of the world `W`.

1. The speaker selects one object out of the context, further called the topic `T`.
2. The speaker categorizes which individual is involved and then names this individual.
3. The hearer looks up which individual is associated with this name in his memory and examines the context to find out whether there is an object which has the distinctive characteristics of this individual.
4. The hearer then signals to the speaker which object was named according to him, for example by pointing.
5. The speaker checks whether the hearer selected the same object as the one he had originally chosen. If they are the same, the game is a success, otherwise a failure.
6. The speaker signals the outcome of the game to the hearer.

It is a straightforward exercise to program a population of agents to play this game succesfully when we, as designers, supply each agent with the same vocabulary and when we postpone the difficulty of recognizing individual objects. Each agent needs to be able to store a bi-directional associative memory between individual objects and their names. The speaker can then look up the name given the individual and the hearer can look up the individual given the name. If we supply the complete vocabulary to all agents, then there is instant and persistent communicative success, at least as long as no new individual objects are introduced in the environment.

Things become interesting only when we don't do this, in other words, when the agents start *without* any prior names and so they have to invent and coordinate which names they are going to use. And the challenge is even greater when the agents start *without* any prior knowledge of which individual objects are present in the environment and what characteristic features distinguish one individual from another.

Identifying individual objects turns out to be a very non-trivial affair because individuals can take on many different appearances and so it is not possible to rely on visual appearance only (Kripke, 1980). Even for the geometric objects used in the experiments reported here, an object can look very differently, depending on the perspective from which it is observed, the position of the agent, which objects are next to or in front of it, and the light that is falling on the object. It is even not always clear whether two segments should be seen as belonging to one or more than one object.

The present chapter argues that a lot of the puzzles arising in naming can be solved by taking a *whole systems approach* which means that all subfunctions of a complete semiotic cycle (perception, conceptualization and language) are integrated tightly so that one component can make up for weakness in another one. It also shows the principles of *linguistic selection and self-organization* at work (Steels, 2012).

The first section focuses on the problem how a vocabulary may self-organize, assuming that agents can already recognize individuals. This is the Non-Grounded Naming Game. The second section turns to the Grounded Naming Game, in which agents do not know which individuals are present in their environment nor what the distinctive features are to recognize them. The concluding section discusses in which way these experiments instantiates the theory of language evolution by linguistic selection explored in this book.

2. The Non-Grounded Naming Game

A *strategy* for playing a particular language game implies fundamental decisions on what kind of conceptual structures and linguistic representations agents need, a set of *diagnostics* and *repairs* for inventing and adopting paradigmatic choices on the meaning or form side, and *alignment routines* for coordinating the conceptual and linguistic repertoire of the agent with that of others. What strategy might be appropriate for the Naming Game?

Because we are dealing with a distributed system without a central coordinator, there is a risk that one agent may invent a new name, not knowing that there is already a name floating around in the population for the same individual, and so *synonyms* (more than one form for the same meaning) unavoidably arise. There is also a risk, although it is small, that *homonyms* arise (more than one meaning for the same form), because one agent may purely accidentally invent a name which is already used by another agent for referring to another object. Because we do not assume that one agent can inspect the internal state of another, there is a risk that the hearer induces another meaning for an unknown word or grammatical construction than the meaning actually intended by the speaker, particulary for grounded or more complex language games. And this then introduces the risk of *ambiguity* or meaning uncertainty (more than one meaning for the same form).

It follows that for the linguistic component of a Naming Game strategy, agents should use a bi-directional associative memory as shown in Figure 1. Where a particular individual may be associated with more than one name (synonymy) or the same name with more than one individual (due to homonymy or meaning uncertainty).

Each agent has his own bi-directional memory. Figure 1 shows an example for the memories of Agent-1 and Agent-2 at some point in the evolution of their vocabularies. Individual-10 is named “sofido” or “boremi” by Agent-1, and “mikart” refers in his memory either to Individual-13 or Individual-9. “sofido” can also refer to Individual-2 for Agent-1. Agent-2 has associated the names “sofido” and “boremi” with Individual-10 but has “katlad” as name for Individual-9. Agent-2 has no name for Individual-2 or Individual-13 yet.

2.1. Diagnostics and Repairs for the Non-Grounded Naming Game

A possible Naming Game strategy for a *Non-Grounded Game* uses the following diagnostics and repairs:

1. *Speaker has no name*
    - Diagnostic: When the speaker does not have a name yet for the topic `T` (step 2 fails).
    - Repair: The speaker invents a new name `n` by creating a random combination of syllables and associates `T` with `n` in his memory.

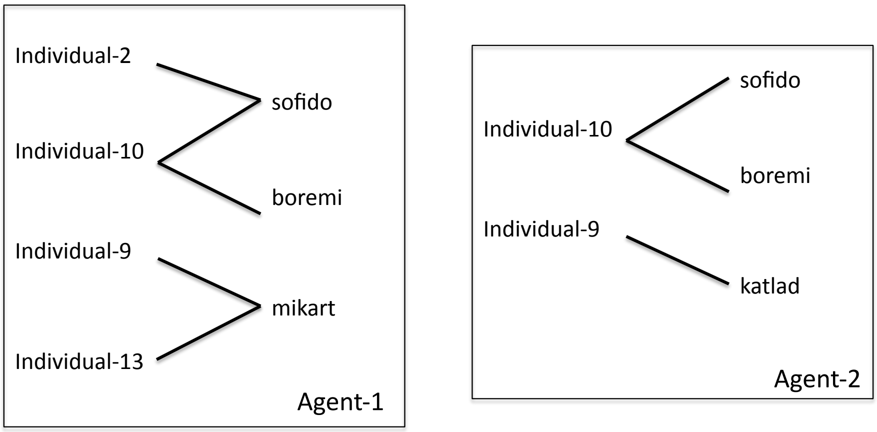

Figure 1. *To play the Naming Game, agents need to store bi-directional associations between meanings (in this case individual objects) and forms (in this case names). When the speaker (Agent-1) has to name an individual object (for example Individual-10) he selects randomly one of the names, for example "sofido", and the hearer (Agent-2) looks up this name in his memory to find back the individual.*

2. *Hearer does not know the name*
    - Diagnostic: When the hearer does not know a name `n` (step 3 fails).
    - Repair: The hearer signals failure and the speaker points to the topic `T`. The hearer then infers which individual was named and associates `T` with the unknown name `n` in his memory.
3. *Hearer uses name differently*
    - Diagnostic: The topic `T` chosen by the speaker and named `n` is not the object pointed at by the hearer (step 5 fails).
    - Repair: The speaker signals failure and points to `T`. The hearer infers which individual was intended and associates `T` with `n` in his memory.

The intended effect of these learning operators is that new names are invented and spread in the population. But is this actually the case? This is where computational simulations can be conclusive. Figure 2 shows the outcome of such a simulation by displaying average communicative success and average vocabulary size in a series of language games played by a population of 10 agents naming 5 different individual objects. All agents use the Naming Game strategy described above.

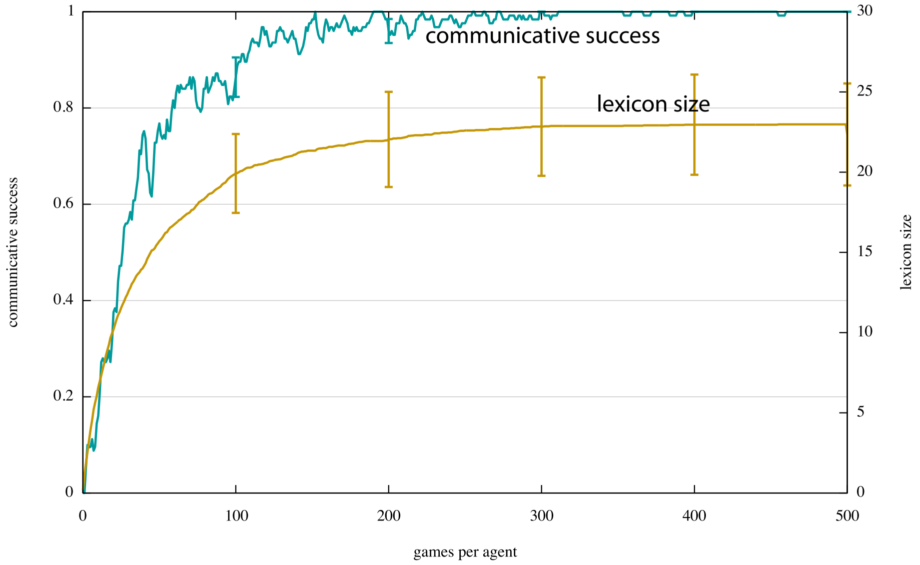

Figure 2. *Example of a graph plotting the outcome of 5 series of Non-Grounded Naming Games with a population of 10 agents and involving 5 distinct objects. The x-axis represents the number of games per agent and the y-axis the running average of communicative success (left axis) as well as average agent vocabulary size (right axis). Each agent plays about 500 language games either as speaker or as hearer.*

We see that the adopted strategy works in the sense that average communicative success rises steadily from zero to 100 % success. This is because agents invent in a first phase new names until all agents know at least one name, and then they keep acquiring names from others in the population until each agent knows all names that were ever invented. The fact that there is total communicative success implies that the system invented by the agents has the required expressive adequacy: There is at least one shared name for every individual that they may want to refer to. On the other hand, the vocabulary that emerged is far from optimal. There is a lot of variation, with about 5 synonyms for every object. This has a negative effect on cognitive effort and on learnability. Agents need to store many more names in memory than are necessary and subsequently the look-up time is greater than need be. Given that there are so many names, it takes much longer for a population to reach success, particularly if the population is very big. New agents entering the population have to learn a lot more names than if there would be only a single one for each individual object.

2.2. Alignment for the Non-Grounded Naming Game

An efficient shared language system only comes about when agents use an enhanded language strategy that includes *alignment*. There are many ways to operationalize alignment for the Naming Game. Here we will use a *lateral inhibition strategy*, which works as follows.

All bi-directional associations stored in the memory of the agents get a score. A new association gets an initial score σ_{init}. The speaker chooses the association with the highest score for naming an object (for the vocabulary in Figure 3 this is "sofido" for naming Individual-10) and the hearer does the same for interpreting a name.

- After a successful game, both speaker and hearer increase the score of the used association with $\delta_{success}$ and diminish competing associations with $\delta_{inhibit}$. Competing associations are assocations where the same individual is associated with a different name or associations where the same name is associated with a different individual. Inhibiting competing names is necessary to dampen synonyms and inhibiting competing meanings is necessary to dampen homonyms or ambiguity.

- After a failed game, both speaker and hearer decrease the score of the used association with δ_{fail}. When the score of an association becomes zero, it is still known but not counted as an active variant. When it is encountered again its score increases with $\delta_{success}$.

This lateral inhibition strategy can be implemented using a wide variety of neural network models of bi-directional associative memories (Kosko, 1988; Kohonen, 1982).

A computational simulation in which a population of 10 agents uses this enhanced strategy to name 5 objects is shown in Figure 4. $\sigma_{init} = 0.5$, $\delta_{success} = 0.1$, $\delta_{inhibit} = 0.2$, and $\delta_{fail} = 0.1$. We see again a phase where invention is dominant with a peak of about 12 names around 50 games per agent, until every agent has at least one name, but the phenomenology changes after that. Rather than a continued increase in the average number of names, there is a decrease. Agents are aligning

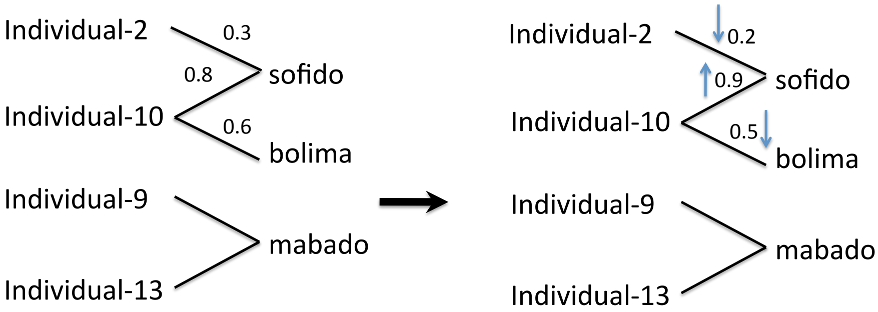

Figure 3. *Alignment is implemented by scoring all associations in memory. After a game, speaker and hearer update the scores of the relevant associations. This figure shows alignment after a successful game when agent-1 used the name "sofido" to name Individual-10.*

their vocabularies, even though there might still be some name adoptions going on. The average vocabulary size of the agents is decreasing until it reaches an optimal size (5 names for 5 objects). In other words, the agents have a fully coordinated efficient vocabulary. This vocabulary is the emergent structure that we were looking for. Alignment based on the outcome of communication triggered self-organization.

3. The Grounded Naming Game

So far we simply assumed that agents knew in advance which individual objects are present in their environment. We now take away this scaffold, and move from computer simulations to a Grounded Naming Game experiment with physically embodied agents (Steels et al., 2012). Experiments to teach robots names of objects have been carried out before (Roy & Pentland, 2002; Dominey, 2000; Roy, 2005; Plebe, 2007), although in these cases the human experimenters carefully prepare the learning data. They decide on identity and individual uniqueness and supply the names that they themselves, or other members of their language community, have decided upon. What we want to understand and model here is how identity, individual uniqueness, and new proper names emerge without a human experimenter already setting them up. We are therefore not trying to model how children acquire

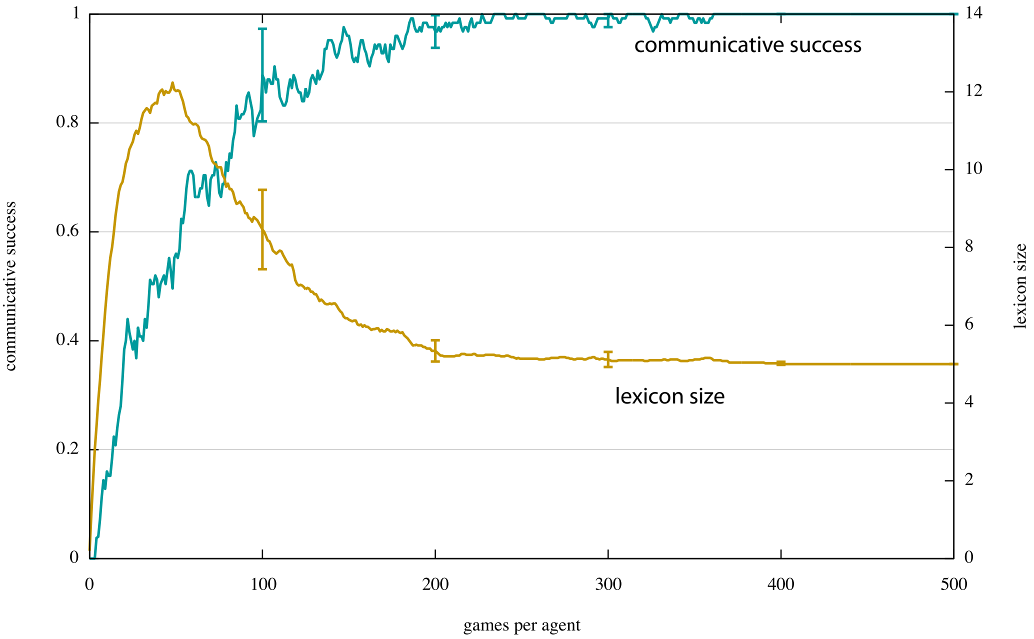

Figure 4. *Semiotic Dynamics of the Non-Grounded Naming Game with 10 agents using a lateral inhibition alignment strategy. The agents still reach 100 % communicative success but now variation gets damped and an optimal vocabulary of 5 names, one for each individual object emerges.*

pre-established names but how new names may come into existence and spread in a population.

The experiment uses embodied agents in the form of small humanoid robots (Fujita et al., 2003). See Figure 5 for an example. They come equiped with the necessary hardware (sensors, actuators, batteries, processing power, memory), but all the software needed for playing language games, including the vision system, the motor system, the conceptualization system, the language system, and the interaction script manager all had to be implemented (Loetzsch et al., 2012; Spranger et al., 2012). The experimental environment consists of an office environment with geometric blocks of various sizes and shapes. The objects are visually distinctive, otherwise there is not enough individual identity to make a strategy based on proper names viable and other language strategies, for example based on spatial language, would have to be used (as in the examples shown in a later chapter by Spranger, 2012).

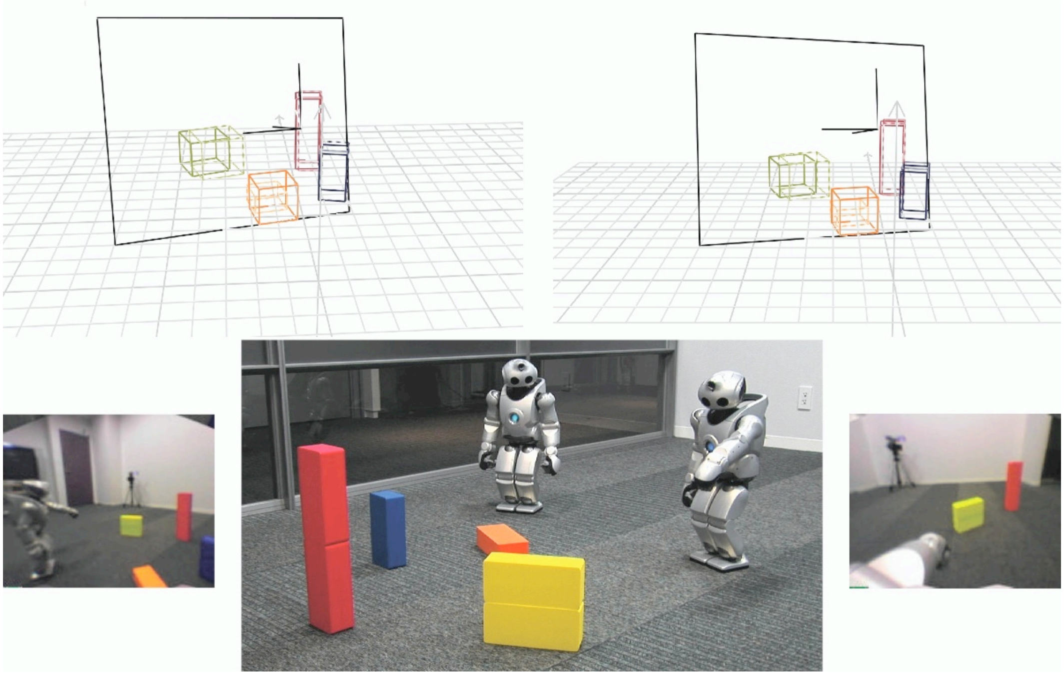

Figure 5. *Experimental setup for the Grounded Naming Game. The bottom middle picture shows two robots, with the one acting as hearer pointing to the object that it interpreted to be the referent. The left and right bottom picture shows one of the camera images of the left and right robot. The top left and right pictures shows the world models built up by each robot respectively. The square show what is in the robot's view and the corner shows the camera position.*

Robots walk around freely in this environment and thus see the situation from many different viewpoints and hence under different light conditions. This raises already the problem that the same object will inevitably have a different visual appearance for the speaker and the hearer. The robots have to build up and maintain in real time a world model which segments and detects the possible features of each object in the context. They also need to detect the position of each object with respect to themselves and the other robot so that pointing gestures can be interpreted. New objects can be added or removed at any time by the experimenter and objects move and appear in different positions. The agents do not know in advance how many objects there are nor their distinctive features. Dealing with unknown, open-ended environments is an important requirement of language game experiments because we need to show that the emerging language systems are adaptive to changing environments and changing ecological needs, indeed this is what makes human language fundamentally different from closed artificial languages such as logical calculi or programming languages.

3.1. Perception and Conceptualization

Let us first examine the way the robotic agents recognize and learn the visual characteristics of individual objects. More detail is given in (Loetzsch et al., 2012). Each agent builds and uses a *semiotic network* which links specific sensory experiences of the objects in the context to object models of these objects tracked over time. The semiotic network is entirely local to the agent and cannot be inspected or manipulated by another agent. Visual stimuli arrive as single *frames* through the camera at the rate of 15 per second. Each frame contains information about each of the picture elements in the scene as recorded at a particular instant of time. Each frame is *segmented* by grouping picture elements together based on color, motion, and texture, using standard vision techniques. The identified regions are compared to *object models* $\{O_1,...,O_n\}$ in short-term memory and an anchoring relation is established between the best fitting object models and each region in the image frame. An object model contains values for luminance (lightness), the hue opponent channels (yellow/blue and red/green), x and y-position, and height and width of an object. It therefore represents a point in a feature space, also known as a conceptual space (Gärdenfors, 2000). When no existing object model can be found for a particular region, a new object model based on the data extracted from the region is stored in short-term memory (see Figure 6 from (Steels et al., 2012).

The object models are then matched against the *prototypical views* $\{V_1,...,V_m\}$ stored in long-term memory. A prototypical view has the same dimensions as an object model. It is defined in terms of a typical value for each dimension and a minimum and maximum allowed deviation from the typical value. The distance in sensory space between the object model and the prototype is calculated by taking the Euclidian distance between the observed value and the typical value, weighted with the allowed deviation.

A prototypical view is associated in long-term memory with one or more *individual objects* $\{I_1,...,I_k\}$ which are then linked to possible names, thus implementing the bi-directional associative memory that was already used in the Non-

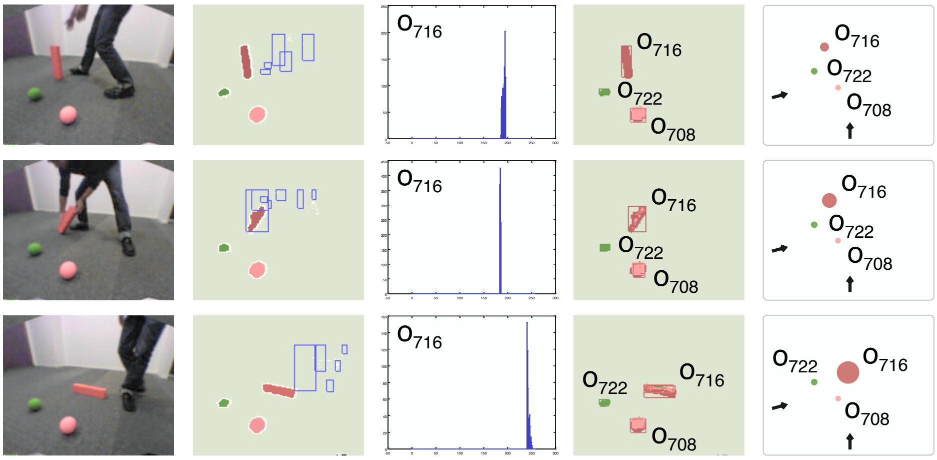

Figure 6. *Snapshot of the perceptual processing for three consecutive moments in time in which the experimenter changes the scene. This figure shows from left to right the visual experiences, segments, features (only the value for the luminance dimension is shown), object models, and finally the world model. Even if an object moves, the agent tracks that it is the same object by using a combination of visual appearance and movement. So even though the width of Object-716 changes as it moves from a vertical to horizontal position, it is still recognized as involving the same object.*

Grounded Naming Game discussed in the previous section. The same individual object may be associated with different prototypical views because an object may have different visual appearances, for example, a block can be upright or lying on its side, shaded or non-shaded due to the source of light or the position with respect to the object, etc.

The speaker typically establishes a chain from left to right, starting from sensory experiences and ending up with a name, whereas the hearer uses the network from right to left, starting from a name and trying to link that to the object-models found in the scene. All chains in the network are scored based on the score of each link and the chain with the highest score is used in the game.

Where do the associations between prototypes and individuals and individuals and names come from? Similarly to the Non-Grounded Naming Game, the links and their scores are updated through diagnostics, repairs, and alignment routines.

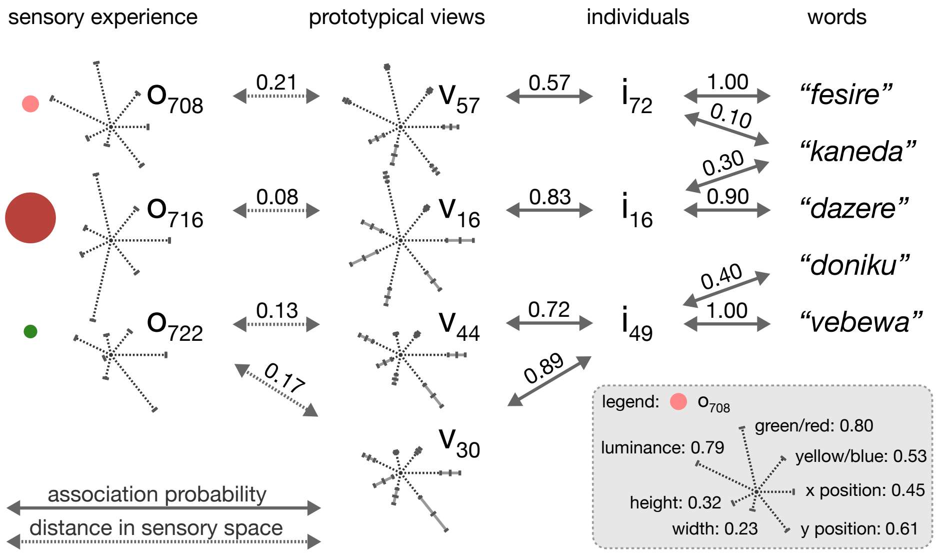

Figure 7. *A semiotic network relates object models of sensory experiences to prototypical views of individuals and names for these individuals. This picture shows an example network starting from the sensory experiences of three different regions perceived in a single image. Their object models are mapped to prototypical views, then to individuals and from there to possible names. Scores reflect either distance in sensory space (between a sensory experience, an object model, and a prototypical view) or the strength of association based on prior experience (between a prototypical view, an individual object, and a name).*

The same lateral inhibition strategy is used as in the Non-Grounded Naming Game for building and adapting the vocabulary (see 2.1). For acquiring individuals and their prototypical views the following diagnostics and repairs are added:

1. *Speaker introduces new prototypical view (a)*
   - Diagnostic: The object model does not match with any prototypical view already stored in memory and hence it cannot be linked to any individual.
   - Repair: The speaker introduces a new individual and a new prototypical view, using the object model for defining the prototype. The values for

all visual dimensions in the object model are taken as the prototypical values for the visual dimensions of the prototypes. The prototype is linked to the newly introduced individual in the semiotic network of the speaker with an initial score π_{init}.

2. *Speaker introduces new prototypical view (b)*
   - Diagnostic: More than one object model in the scene is matching with the same prototypical view. This means that the same individual object would appear twice, which is against one of the fundamental constraints on the real world (Spelke, 1990).
   - Repair: The speaker creates a new prototype for the object model that is less similar and associates it with a node for a new individual object.
3. *Hearer introduces new prototypical view*
   - Diagnostic: The hearer already knew the name used by the speaker but no object model in the scene matches with any of the prototypical views in his semiotic network.
   - Repair: The hearer signals failure and obtains information which object and hence object model constitutes the topic. The hearer uses this object model then as the seed for a new prototypical view for the individual already associated with the name and adds this to his semiotic network. Notice that this effectively introduces a top-down influence from language to concept formation.
4. *Hearer stores additional prototypical view*
   - Diagnostic: The object model matches with a prototypical view, but this prototype was associated with another individual than the topic `T` named by the speaker.
   - Repair: The hearer then links this prototypical view to `T`, even if this increases the uncertainty in the agent's conceptual system.

Besides diagnostics and repairs, agents also need routines for alignment. For the vocabulary part, they use the same one as discussed for the Non-Grounded Naming Game in the previous subsection (2.2). The following routines take care of conceptual alignment:

- Prototypes can be improved to reflect reality in a better way. After a successful game, speaker and hearer bring the prototypical view of the individual that was named closer to the data from the object model by moving the value of each dimension closer to its corresponding value in the object model. If the value is deviating further from the prototype then the minimum and maximum value can also be adjusted.

- The relation between a prototype and an individual is strengthened after a successful game and weakened when the game is not successful. A lateral inhibition dynamics is not useful here because the same individual may have different visual appearances. To weed out erroneous connections, a forgetting process is introduced. Agents diminish slightly the score of all associations and only the links that were necessary and successful in enough games survive.

3.2. Experimental results

Figure 8 (from Steels et al., 2012) shows the overall dynamics that is generated by the conceptual component of this Grounded Naming Game strategy. It shows on the left the prototypical values of the different visual dimensions that define a particular prototypical view for one agent. The values keep getting adjusted as new examples are seen. On the right, we see how the strength between prototypical views and individuals changes. Some prototypical views evolve towards a zero score. They were created but turned out to be irrelevant or ineffective and the forgetting process steadily decreased their strength.

Figure 9 (from (Steels et al., 2012)) shows the outcome of the complete Grounded Naming Game strategy. The left graph shows that agents consistently reach high success after about 1000 language games per agent. We notice however that they assume there are 15 different individuals instead of 10. Moreover, the number of prototypes is the same as the number of individuals which means that they are naming prototypes and fail to distinguish properly individual identities.

3.3. Adding Additional Heuristics

The beauty of the experimental approach is that different strategies can be compared, even for the same grounded data. Let us do this now by endowing the agents with an enhanced conceptualization strategy with two additional repairs:

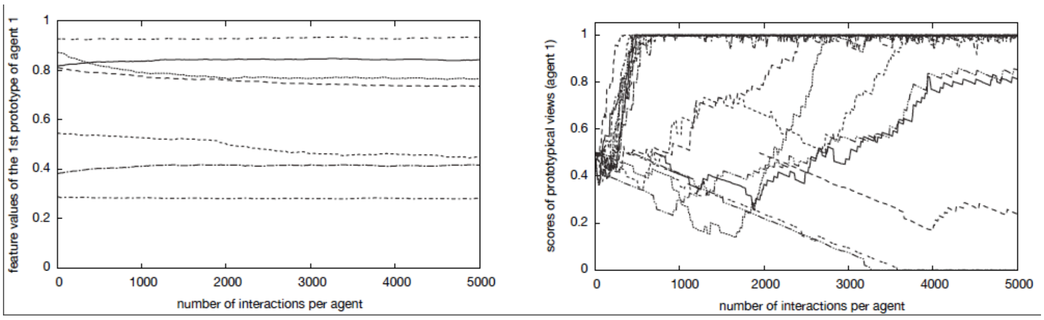

Figure 8. *Left: Change in the prototypical values of all features for one particular prototypical view of an individual for a single agent. Right: Change in the strength of the association between prototypical views and individuals in the semiotic network of a single agent. Graphs are not scaled for population size.*

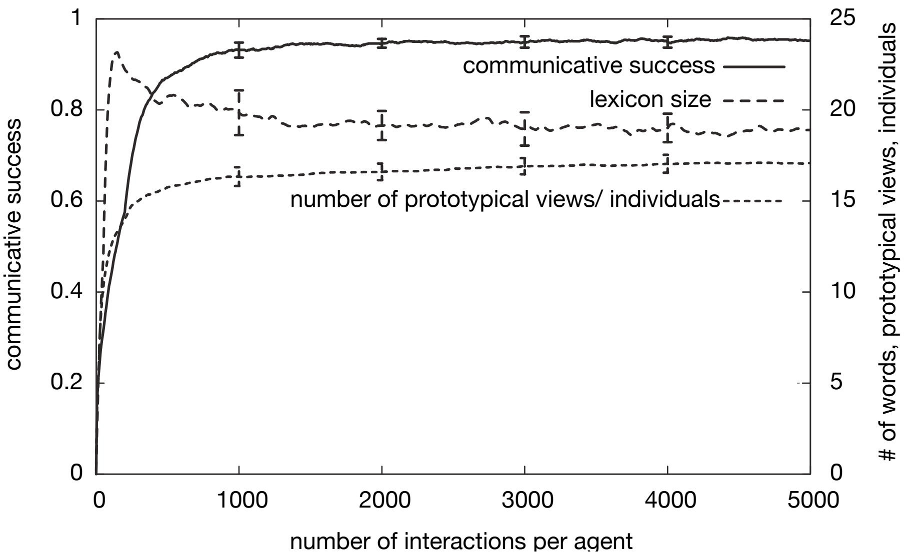

Figure 9. *Grounded Naming Game experiment with 10 distinctive individuals in the environment. Agents are able to bootstrap a system from scratch and reach above 90 % communicative success. However, they are naming prototypical views without linking different views to the same individual*

1. Agents track objects in real time, even if they move, based on the principle of object constancy (Spelke, 1990). If the objects are stationary, the object-model stays the same. But if the object moves (as in Figure 6 from top to bottom), it is possible that the object model changes quite drastically, and hence that the associated prototypical view changes as well. This is an opportunity for agents to find out that the same individual has more than one prototypical view. They should therefore create links from each of the prototypes to the same individual object, if those links did not exist yet.
2. The emergent language can itself be used as a heuristic to bootstrap the conceptual system. If a name associated with one individual in the semiotic network of the hearer is used by the speaker to refer to another individual, then this suggests to the hearer that the prototypical view involved might actually belong to this individual, at least for the speaker, and the hearer should therefore expand his network to take this into account.

When these repair actions are added, we see improved results as in Figure 10. Thanks to the first additional repair, the number of individuals is now only 11, almost equal to the number of individuals that were supplied, and thanks to the second additional repair, fewer words were created. The agents end up with a smaller and hence more efficient vocabulary and have a better notion of the different prototypical views that are associated with the same individual. Thanks to the additional repairs, agents increase the expressive adequacy, cognitive efficiency and learnability of their language system.

Figure 11 shows some visual snapshots of the internal semiotic networks of a single agent as it is playing grounded naming games. The network relates sensory experiences (shown as squares) with prototypical views (labeled v_1, v_2, etc.) shown as spiderweb diagrams linked to nodes for individuals (labeled as i_1, i_2,...) which are linked to words (as discussed in 7). The thickness of the line reflects the strength of the connection. We see that the network becomes denser as more sensory experiences of objects are encountered, synonyms get damped so that multiple associations between individuals and words get resolved. Each agent in the population maintains such a network and over time the networks of the different agents become more similar.

4. Conclusions

This chapter studied how a population of autonomous agents could self-organize a vocabulary of proper names for referring to individual objects in their environ-

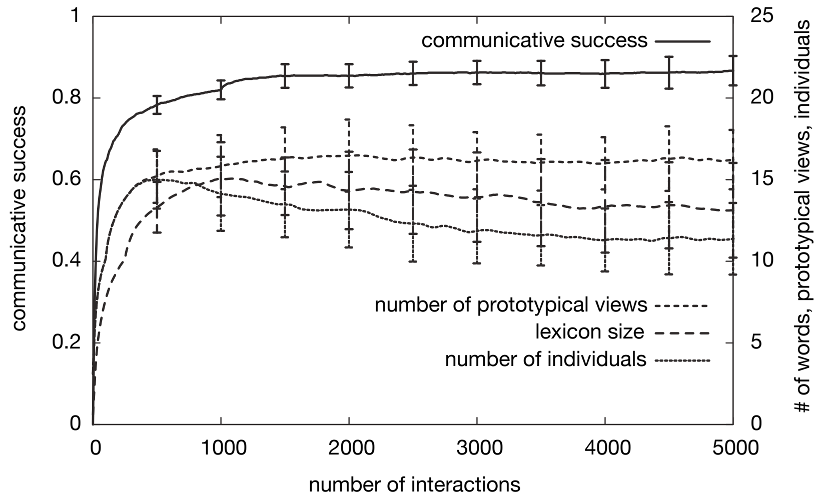

Figure 10. *An experiment for the same data as in figure 9 but with a more sophisticated strategy that makes use of additional heuristics. A more efficient vocabulary and a set of individuals more in tune with reality is obtained.*

ment. It proposed a language strategy, consisting of a conceptual component for the acquisition of grounded concepts of individuals and a linguistic component for the invention, adoption and alignment of names for these individuals. The consequences of this strategy were explored using computer simulations and experiments with humanoid robots.

It is important to see the general principles beyond the experimental specifics of this particular case study. The overall framework is that of linguistic selection (Steels, 2012). The strategies of the agents generate possible variants, both for prototypical views and proper names. Which variants survive in the population should reflect expressive adequacy, minimal cognitive effort, learnability and social conformity. Expressive adequacy means here that agents share at least one name for each individual in the environment and that these names name individuals rather than visual appearances of individuals. Cognitive effort is reduced when all agents use the minimal set of names and the minimal set of prototypical views for each individual. Learnability means that names invented by one agent can be acquired easily by others, after a failed dialog has been repaired by additional feedback from

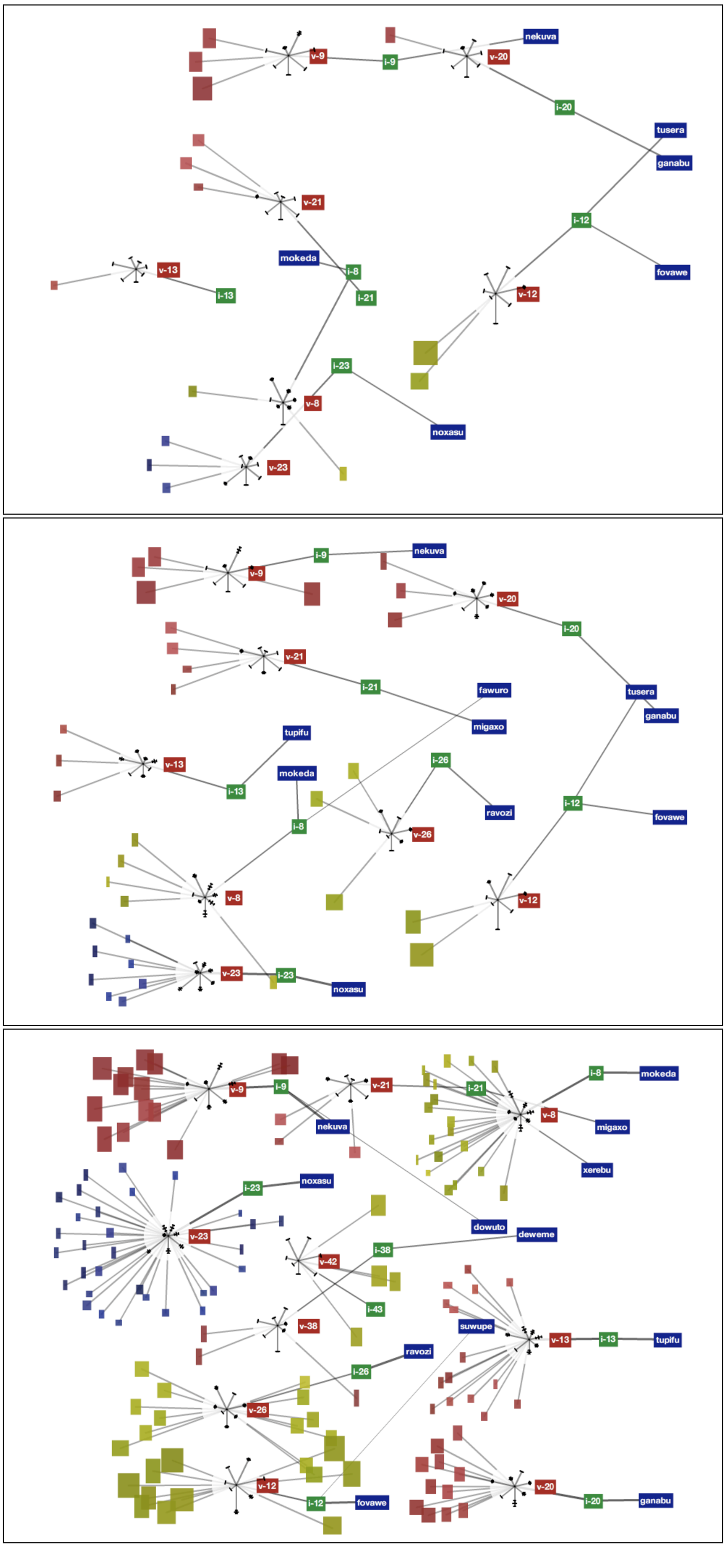

Figure 11. *Series of snapshots in time showing the internal semiotic networks of a single agent as it progressively develops an inventory of words for naming individuals.*

the speaker, and that agents do not need to learn more names than necessary. Social conformity means that agents start to use similar names than their peers.

These various factors are not explicitly engineered by the agent (indeed this is impossible because no agent has a global view or any insight into the internals of another agent) but they are a side effect of the strategy that they have adopted and the role of communicative success in choosing repairs. We see three elements at work:

- Cognition: When speakers fail to conceptualize or name an individual, they extend their language system with new prototypical views and new names. When hearers are confronted with an unknown name or with communicative failure due to a mismatch between speaker and hearer, then they extend their language system based on additional feedback. Speakers and hearers are constantly on the look out (for example using the two additional repairs introduced in section 3.3.) to improve their knowledge of individual objects, in particular which prototypical views are associated with the same individual.

- Self-organization: Alignment introduces a self-enforcing positive feedback loop between a particular paradigmatic choice (i.c. which name to use) and the communicative success with this choice. It was clearly demonstrated (compare in particular Figure 2 and 4) that the addition of alignment induces self-organization in the Naming Game, leading to a minimal language system and hence reduced cognitive effort, increased learnability, and higher social conformity.

- Selection: Even after finding the best possible strategy, agents will still be confronted with a great deal of uncertainty, particularly because we must take into account that agents only progressively learn about the environment, the environment can change, and new agents can enter any time. This is where selection is relevant, based on the outcome of a game. When a game succeeds, all language system components that were used are re-enforced so that the likelihood to reuse them in similar circumstances in the future increases. When a game fails, the opposite takes place.

The solutions proposed here are entirely general and scalable. Any kind of feature can be used to generate

Acknowledgements

This research was carried out at the AI Lab of the University of Brussels (VUB) and the Sony Computer Science Laboratory in Paris. We thank Michael Spranger for his help with the robotic setup and many discussions regarding the experiment. We also thank Tao Gong and anonymous reviewers for their helpful comments on the present paper.

References

Baronchelli, Andrea, Maddalena Felici, Vittorio Loreto, Emanuele Caglioti, Luc Steels (2006). Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics*, P06014.

Dominey, Peter (2000). Conceptual grounding in simulation studies of language acquisition. *Evolution of Communication*, 4(1), 57–85.

Fujita, Masahiro, Yoshihiro Kuroki, Tatsuzo Ishida, Toshi Doi (2003). Autonomous behavior control architecture of entertainment humanoid robot SDR-4X. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '03)*, vol. 1, 960–967. Las Vegas, Nevada.

Gärdenfors, Peter (2000). *Conceptual Spaces: The Geometry of Thought*. A Bradford Book. MIT Press.

Kohonen, Teuvo (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1), 59–69.

Kosko, Bart (1988). Bidirectional associative memories. *IEEE Transactions on Systems, Man and Cybernetics*, 18(1), 49–60.

Kripke, Saul (1980). *Naming and Necessity*. Oxford: Basil Blackwell.

Loetzsch, Martin, Michael Spranger, Luc Steels (2012). Object Naming and Identity. *submitted*.

Plebe, Alessio (2007). A neural model of object naming. *Proceedings of the World Academy of Science, Engineering and Technology*, 20, 232–237.

Roy, Deb (2005). Grounding words in perception and action: Computational insights. *Trends in Cognitive Sciences*, 9(8), 389–396.

Roy, Deb, Alex Pentland (2002). Learning words from sights and sounds: a computational model. *Cognitive Science*, 26(1), 113–146.

Spelke, Elizabeth (1990). Principles of object perception. *Cognitive Science*, 14(1), 29–56.

Spranger, Michael (2012). The Co-evolution of Basic Spatial Terms and Categories. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins Co.

Spranger, Michael, Martin Loetzsch, Luc Steels (2012). A Perceptual System for Language Game Experiments. In Luc Steels, Manfred Hild (Eds.), *Language Grounding in Robots*. New York: Springer.

Steels, Luc (1995). A self-organizing spatial vocabulary. *Artificial Life*, 2(3), 319–332.

Steels, Luc (2012). Self-organization and linguistic selection in language evolution. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins Co.

Steels, Luc, Martin Loetzsch, Michael Spranger (2012). A boy named sue. the semiotic dynamics of naming and identity. *submitted*.