The Naming Game

September 29, 2013

The concept of language games was introduced by the 20th century philosopher Ludwig Wittgenstein in his philosophical investigations (Wittgenstein, 1967). In this work he provides what must be the most famous of all language games, which is, perhaps not by accident, about building.

> Let us imagine a language ... The language is meant to serve for communication between a builder A and an assistant B. A is building with building-stones; there are blocks, pillars, slabs and beams. B has to pass the stones, and that in the order in which A needs them. For this purpose they use a language consisting of the words 'block', 'pillar', 'slab', 'beam'. A calls them out; –B brings the stone which he has learnt to bring at such-and-such a call. – Conceive this as a complete primitive language. (Wittgenstein, 1967, passage 2)

Wittgenstein did not believe this is the way language works. In fact he provided this example to criticise a certain conception of language (that of Augustine). He saw language games much more as a tool, used to shed light on some aspect of a linguistic interaction. Sometimes these aspects can be brought out through similarities to real world language use but often (and especially in the earlier language games) through their dissimilarities. The Naming Game, which is the language game discussed in this chapter, also bears little similarity to real world language use. Language games should not be taken as representative of even a part of language, or as Wittgenstein put it, as if they were first approximations, ignoring friction and air-resistance (Wittgenstein, 1967, passage 130).

From the mid nineties computational implementations of increasing complexity of language games have been proposed. They are known as *Naming Games* (Steels and McIntyre, 1999), *Guessing Games* (De Beule et al., 2006), *Description Games* (van Trijp, 2008) and so forth. The game discussed in this chapter, the Naming Game, very much resembles the game described by Wittgenstein in the quote above.

1 Definition and history of the Naming Game

Although the Naming Game is arguably the most restrictive language game, it was not the first language game to be operationalised. Trying to determine

the very first reported Naming Game implementation turned out to be quite a challenge. For example the very definition of a Naming Game has undergone quite some change over the past fifteen years.

In this paper the focus is on the *Non Grounded Naming Game*, all experiments are done in simulation on computers without the agents being embodied in robots. The Grounded Naming Game (Steels and Loetzsch, 2012) shares only limited features with its Non Grounded variant. In all the following whenever a reference to *the* Naming Game is made, it is meant to be the Non Grounded Naming Game. More often than not statements would be false if the Grounded version was to be included.

Instead of presenting a strict definition of the Naming Game I will start with presenting features found in many but not necessarily all Naming Games.

1. A Naming Game involves two participants, a listener and a speaker.

2. In a Naming Game a speaker utters a single name to refer to one object in a context.

3. Objects do not show internal complexity and names can only refer to objects as a whole and not to categories or features.

4. At the end of the game both participants know the intended referent and can thus unambiguously make a pairing between name and object. Uncertainty about the referent of a name does not or only marginally occurs.

5. Different names for the same referent (form competition) can be introduced in the population and increase with its size.

With the aid of this list it is possible to map out a short history of the Naming Game and define the game as it will be used in this chapter. The very first operationalized language games (Steels, 1995, 1996), although referred to as Naming Games in (Steels and Kaplan, 1998a), do not show many of the features above. In Steels (1996) the game is played by more than two participants, there are multiple referents (here called meanings instead of objects) per game and the referent of a single word is ambiguous. In Steels (1995) there is a single referent (which is an object in a context), but this object is described through spatial terms and thus shows internal complexity. Words do not refer to objects but to spatial terms. These games miss crucial constraints essential to still be categorised as a Naming Game today.

Only two years later in (Steels and Kaplan, 1998a,b; Steels, 1998; Steels and McIntyre, 1999) a game is introduced, explicitly called a Naming Game, which shares features 1,2,3 and 5 of the above list. The setup in these games does allow uncertainty at the end of the game through different variants of "noise" parameters or interpretation probabilities, as such invalidating the fourth constraint. In 2005 and 2006 the Naming Game got renewed attention from a far more theoretical perspective (Lenaerts et al., 2005; De Vylder and Tuyls, 2006;

Baronchelli et al., 2006b). It was only in these papers that feature 4 (no uncertainty) became standard part of the Naming Game. De Vylder and Tuyls (2006) and Baronchelli et al. (2006b) argued that when features 3 and 4 are present the script of the game can be simplified. There is no need to simulate multiple objects per game, one object suffices. They went even further and showed that, without loss of generality, there doesn't even have to be more than one object at all. Whether the agents have to agree on names for 1 or $n \geq 1$ objects has no impact on the dynamics when feature 3 and 4 are present. This more theoretical approach to the Naming Game has been investigated in depth over the past few years (Lu et al., 2008; Liu et al., 2009; Lipowski and Lipowska, 2008; Baronchelli et al., 2008).

The Naming game as discussed in this chapter follows the approach of De Vylder and Tuyls (2006) and Baronchelli et al. (2006b). All five features (or restrictions) are present. In this thesis the script of the game is kept closer to that of Steels and Kaplan (1998a) with multiple objects per context so that it is more similar to that of games introduced in later chapters. To keep simulations fast and since it doesn't impact the results in a significant way the number of objects is kept small at five.

Also note that in a Grounded Naming game only features 1 and 2 tend to be present, bringing with it all kinds of problems not present in the Naming Game as discussed in this chapter.

2 Script of a Naming Game

The central question in a Naming Game experiment is how a population of agents can bootstrap and maintain a shared set of names for a set of objects. In a Naming Game experiment the agents can only engage in Naming Games, which are local to the participating agents only. The agents cannot negotiate names in any other way. Let us start by taking a closer look at the script of a single Naming Game.

The Naming Game as discussed in this chapter requires objects $o_i \in O$ and agents $a_1 \cdots a_n \in P$ (population). For every game a context is generated and two agents a_s and a_l are chosen from the population. One agent is assigned the role of speaker, the other of listener[1]. Both agents are confronted with the same context C of n objects. When these conditions are met the game can start (also see Figure 1).

1. The speaker mentally picks one object from the context (called the *topic*) and utters a name for it.

2. The listener interprets the name and points to the object he believes the speaker intended.

3. The speaker either agrees or disagrees with the listener.

[1] Agents are chosen at random, there is no network topology (see (Baronchelli et al., 2007; Van Segbroeck et al., 2010) for more on the role of network topologies).

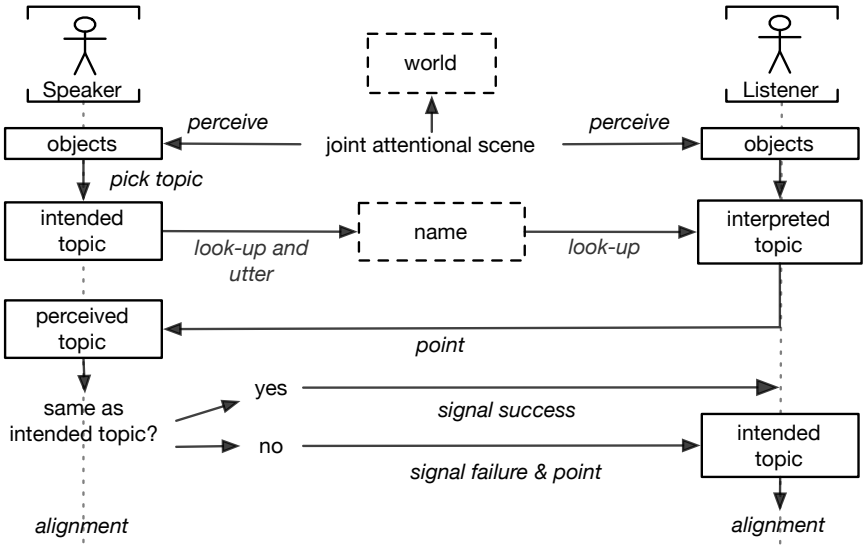

Figure 1: Script of a Naming Game. Squared boxes represent data entities. Italicised text refers to processes. The look-up processes can fail and therefore have diagnostics and repair strategies associated with them. A Naming Game experiment consists of thousands of these interactions.

4. Both agents have the opportunity to adapt their inventory of names.

The script can be made less verbose without any loss of generality. For example it is not required that the listener points to the object since per definition (feature 4) it is impossible that the listener would have associated a different meaning to the name.

On the *representational* level not many options are available in strategies for the Naming Game. The nature of the game requires meanings to be represented as non-compositional concepts. Likewise the form of the names is also assumed to be represented as non-modifiable strings[2]. Each agent a needs a bi-directional memory M_a pairing names n_j with objects o_i. One option left open by the specification of the Naming Game is whether agents can score each name-object pairing. The idea is that maintaining such a score helps the agent to better determine the communicative outcome when using that particular name. The strategies discussed in this chapter will differ in this respect.

What kind of language strategy do the agents need to share in order to arrive at the system of a shared set of names? A language strategy needs to supply

[2]In all the experiments in this thesis word forms are assumed to already be segmented and holistic in nature. These strings are transferred to the other agent without relying on speech. A long tradition of computational simulations does exist for the domain of speech (de Boer, 2000; Oudeyer, 2005) [see (de Boer and Fitch, 2010) for an overview].

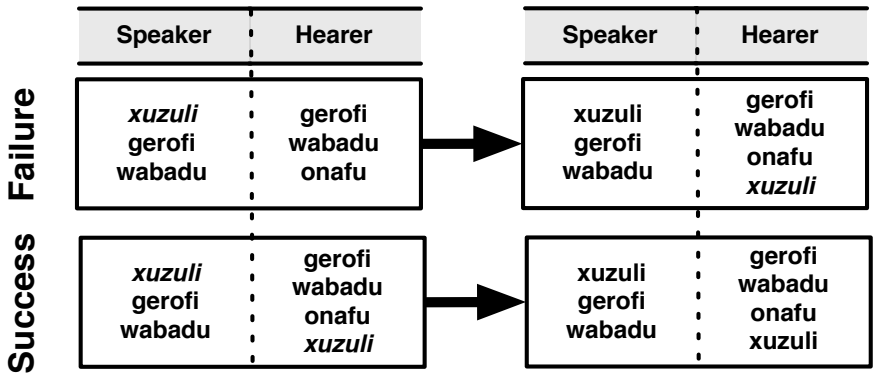

Figure 2: Update mechanics for the Baseline NG Strategy. Only the names for the topic object are shown. From a list of available names the speaker utters "xuzuli". In case of failure the listener does not know the name and adopts the new name. In case of success agents do not update their inventories.

both representational and processing aspects.

3 Minimal strategies for the Naming Game

In a Naming Game experiment a population of agents should develop a *shared* and *minimal* set of names for the set of objects. Minimal means there is only one name used for each object and all agents share this name.

In this section we investigate three *minimal strategies*, so called because they keep representational and processing capabilities as minimal as possible. Most prominently, none of the minimal strategies implements a scoring mechanism.

3.1 The Baseline NG Strategy

Since the Baseline NG Strategy is a minimal strategy it does not allow agents to score their name-object pairings. All that is assumed here is that an agent can store pairs $< n, o >$ in his memory M_a. Such a memory can be implemented in many different ways, ranging from a simple set of pairs to a hash-table or more structured inventories.

For baseline processing each agent a uses two look-up functions $f_{produce}(o, a)$ and $f_{interpret}(n, a)$ which access the memory M_a of agent a and look-up an object o or a name n respectively. Since in our flavor of the Naming Game uncertainty about meaning is absent $f_{interpret}(n, a)$ returns either one object in case the agent already encountered the word once or no object if it is the first exposure. The production function $f_{produce}(o, a)$ however returns a set of names, which can be either empty, have one element, or multiple elements depending

on whether several names are competing for the same object. In case multiple names are returned the agent will, before uttering, pick a random name from this set of competitors. The implementation of these two functions depends on the implementation of the memory architecture M_a and should not impact the dynamics of the strategy as long as it conforms to the input output behaviour just explained. The goal of the Naming Game is that agents agree on a *single* preferred name per object. This means all agents should always return the same name from $f_{produce}(o, a)$.

All agents start out with empty inventories and thus require a capacity to introduce a new name if needed. When the speaker does not yet know a name for his chosen topic t he *invents* a random name n[3], links it to topic t and stores the pairing $< n, t >$ in his memory M_a. As listener, an agent needs a similar capability to store or *adopt* new names. When $f_{interpret}(n, a)$ fails to return an object o the listener first waits for the speaker to point to the intended topic t and adds the association $< n, t >$ linking name and object. Production and interpretation are thus augmented with two learning functions $f_{invent}(t, a)$ and $f_{adopt}(n, t, a)$.

The two processing functions $f_{produce}$ and $f_{interpret}$ and the two learning functions f_{invent} and f_{adopt} make up the Baseline NG language strategy for the Naming Game. Figure 2 schematically shows the updating mechanics just explained.

3.2 Measures for naming strategies

In evaluating strategies for the Naming Game we use the following measures:

Communicative success: A Naming Game is communicatively successful when the listener points to the correct object. Due to the restrictions of this game, when a listener knows a name he will point correctly, thus communicative success can be simplified to whether the listener knew the spoken name (i.e. $f_{interpret}(n, a)$ returns an object). In the graphs a running average of 25 is used.

Name competition and alignment: With name competition the average number of competing names per object is measured. Name alignment measures the level of alignment towards one preferred name for each object. Given the set of all expressible objects O, the population P and the production function $f_{produce}(o, a)$ which returns the set of names that agent $a \in P$ prefers for object $o \in O$.

$$\text{Name competition} = \langle | \bigcup_{a_j \in P} f_{produce}(o_i, a_j) | \rangle_{o_i \in O} \tag{1}$$

$$\text{Name alignment} = \frac{1}{\text{Name competition}} \tag{2}$$

[3]The algorithm guarantees that the same name cannot be invented twice ruling out "accidental" homonymy or meaning competition.

In a fully aligned population, where all agents prefer the exact same name for each object, name alignment equals 1.

Alignment success: This measure combines communicative success and name alignment. It is, like communicative success, a Boolean value which is only true when the listener knows the name spoken by the speaker (i.e. communicative success) and it is also his own preferred name for that object (if he were speaker he would utter it as well). In other words $f_{produce}(t, a_{speaker}) = f_{produce}(t, a_{listener})$ and $|f_{produce}(t, a)| = 1$ for both agents.

All Boolean measures, like communicative and alignment success, are plotted as running averages over a window of 25 games. In what follows I will sometimes talk of *full* Communicative success or full Alignment success (or any other measure) to denote that the population has reached 100% for that measure.

All the above measures only depend on the output behaviour of $f_{interpret}$ and $f_{produce}$ and not on the internal representations of the agents, which makes them independent of a particular strategy.

3.3 Experimental results of the Baseline NG Strategy

We are now in a position to run the Baseline NG Strategy and see how it performs. Unless otherwise stated all following Naming Game experiments are run with a population of 50 agents and a total of 5 objects. As explained earlier this number does not impact in a significant way the dynamics of the strategies. The only difference is that reaching coherence takes slightly longer since instead of a single name the agents have to establish five names. Context size is irrelevant in the type of Naming Game that is discussed in this chapter. In most graphs (e.g. Figure 3) the x axis represents the number of games played either per agent or in total. With a straightforward calculation one can switch between total number of games or games per agent.

$$\text{games played per agent} = 2 \times \frac{\text{total number of games}}{\text{size of population}} \tag{3}$$

Results are averaged (mean) over 100 experimental runs with error bars showing the 5 and 95 percentile[4].

Figure 3a shows that the Baseline NG Strategy is sufficient for a population to reach full communicative success. Full communicative success entails that the agents know all names used within the population, which in turn means that invention and adoption must have come to a halt which is indeed the case as shown in Figure 3b. In the baseline strategy invention stops as soon as each agent has played one game about each object. Adoption takes much longer because every name that was ever invented has a random chance to be spoken

[4]The error is shown either by bars or by a transparent filled curve surrounding the graph, everything inside this filled curve is in the 5 to 95 percentile. An example of both is shown in Figure 3.

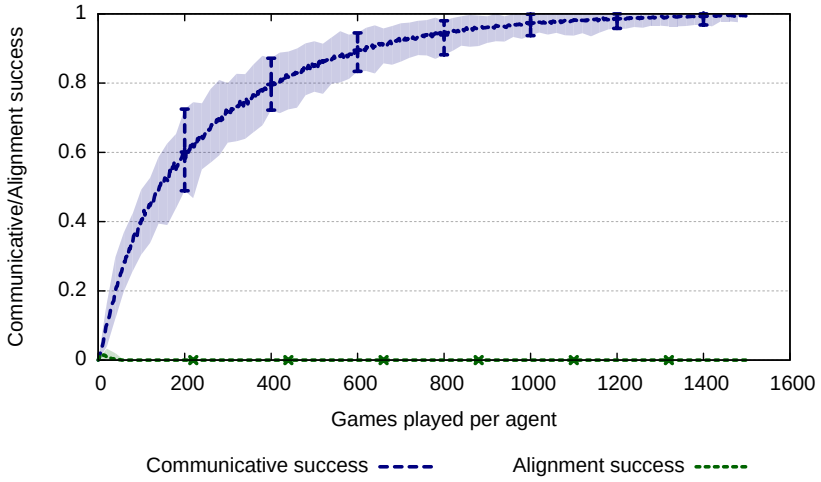

(a) Communicative and alignment success

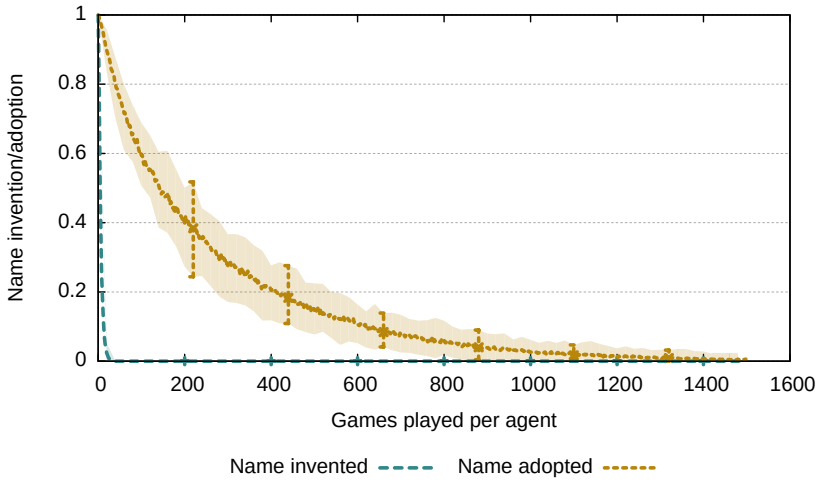

(b) Invention and adoption

Figure 3: Experimental results for the Baseline NG Strategy. (a) Agents reach full communicative success because they store and thus recognize all invented names. Alignment success, however, remains at zero signifying that there is no shared preferred name for each object. (b) Both invention and adoption come to a halt, with adoption taking significantly longer. Parameters of (a) and (b): population size: 50, total object size: 5, number of simulations: 100, error: 5 to 95 percentile.

for its associated object and thus in the long term will get adopted by all agents. Also note that for a population of n agents, a single invention requires $n-1$ adoptions.

Although the population reaches communicative success the emergent language system is wanting with respect to alignment success. The agents invent and need to maintain a large set of names per object. Each game is local to only the two participating agents leading to the introduction of different names for the same object, a problem we call *name* or *form competition*. For example assume there are 10 agents $a_1 \cdots a_{10}$ and in the first game speaker a_3 invents a name n_{first} for object o_5 and transmits it to listener a_7. None of the other agents is aware of the pairing $< n_{first}, o_5 >$. Only if an agent participates in a game with either a_3 or a_7 as speaker and topic o_5 can they be confronted with this name. In the meantime this agent might play games with o_5 as topic, prompting him to also invent or adopt new competing names for it.

As long as agents maintain multiple competing names for an object, without a single shared preference, alignment success is bound to remain low, which is indeed what Figure 3a shows. Only in the very beginning minimal alignment is measured but this is only due to small chances that some agents interact and they have only established one and the same name. The baseline strategy thus fails at establishing an aligned naming system in which all agents agree on a single shared name per object.

The problem of name competition increases with the population size because names spread through the population much slower prompting other agents to invent competing names in the process. In fact, the average number of names per object is more or less equal to half of the size of the population as shown in Figure 4a and b. Only in the special case when the population consists of two agents do the agents converge on the optimal set of names. Indeed, in this case both agents participate in all games and thus no names can be introduced outside of their knowledge.

Figure 4c illustrates this point most clearly. Over the course of 1600 Naming Games we tracked all names in use by the population for a single object. Per name n_i (each line in the graph) the percentage of agents for which $n_i \in f_{produce}(o, a)$ is measured. When a name n reaches 1 for the object o it means that for all agents $n \in f_{produce}(o, a)$. What we see is that for the Baseline NG Strategy all names are accepted by all agents. The strategy does not implement a winner takes all dynamics.

With a bi-directional memory, the processing functions $f_{produce}$ and $f_{interpret}$ and the learning functions f_{invent} and f_{adopt} as outlined for the Baseline NG Strategy a population is capable of reaching a shared lexicon but not a minimal one. We not turn to two minimal strategies that do reach a shared and minimal set of name-object pairings.

3.4 The Imitation Strategy

What would be the most minimal adaptation of the Baseline NG strategy that would allow the population of agents to reach communicative success and align

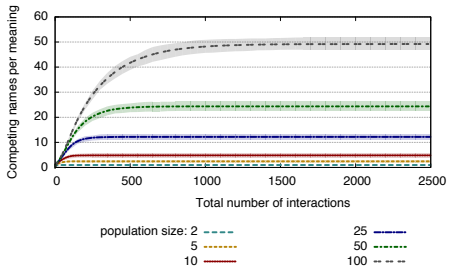

(a) Evolution of name competitors

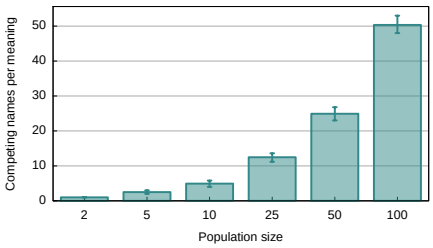

(b) Final number of name competitors

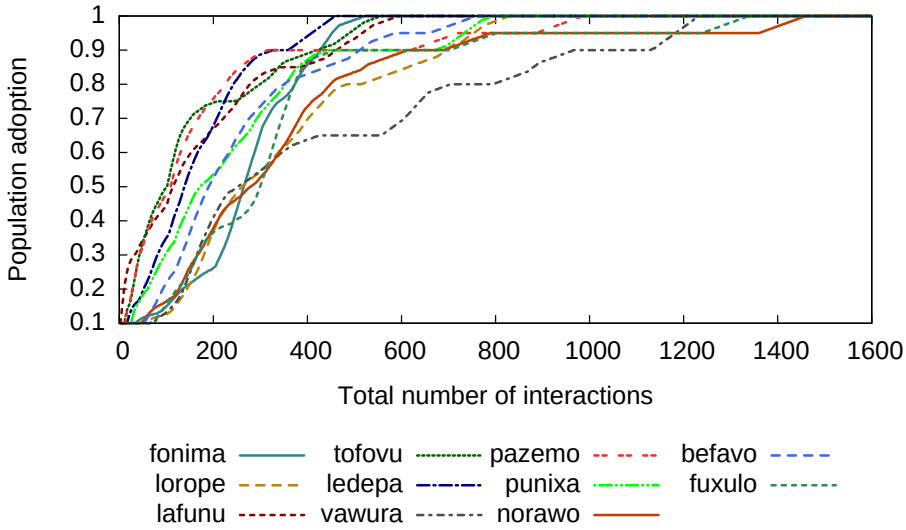

(c) Name adoption by population

Figure 4: Results for the Baseline NG Strategy without alignment. (a) and (b) The number of competing forms per meaning when the size of the population is increased from 2 to 100. (c) The names floating around for one object o. Each line corresponds to a name n_i and shows the percentage of agents for whom $n_i \in f_{produce}(o, A)$. For sake of clarity, population size was reduced to 20 agents.

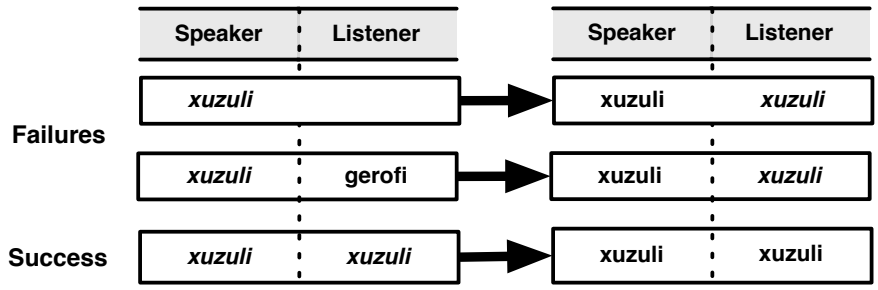

Figure 5: Update mechanics for the Imitation Strategy. The speaker utters the word "xuzuli" for the chosen topic. In case of failure the listener adopts the spoken name.

on a minimal set of names? In other words how can the population reach full alignment success? Interestingly one solution consists in further *simplifying* the Baseline NG Strategy by disallowing agents to remember different names per object. What if an agent could store only a single name per object?

With this in mind we propose a strategy where a listener, when adopting a new name overwrites any name already stored for the object. The adoption function $f_{adopt}(n, t, a)$ is thus simplified so that instead of extending the list of names for t the agent only remembers the pairing new $< n, t >$ for object t, removing the previous pairing if there was one. The listener will thus imitate the speaker when he has to talk about the same object, which is why we call this strategy the *Imitation Strategy*. Representationally the strategy can be implemented using an even more basic bi-directional memory than the Baseline NG Strategy since only one name per object needs to be remembered. With regard to invention, and look-up it remains the same as the baseline, except that $f_{produce}(o, a)$ can never return more than one name. This strategy corresponds to the Voter Strategy introduced very recently by Baronchelli et al. (2011). Its updating scheme is schematically shown in Figure 5.

Because of the restrictive nature of the Imitation Strategy the measure for communicative success and alignment success, exceptionally, returns identical results. Indeed when agents reach communicative success it means the listener knows the spoken name, but since agents can only store one name per object it is per definition also his preferred name and thus the agents also reach alignment success.

Experimental results (see Figure 6) show that the Imitation Strategy achieves both full communicative success and full alignment success. The reason for this success lies in the *winner takes all* dynamics. From the competing names for an object one emerges as the sole preferred name by the whole population as visualised in Figure 7. This figure clearly shows the self organizing dynamics of the Imitation Strategy, leading to population-wide preference of the same name per object.

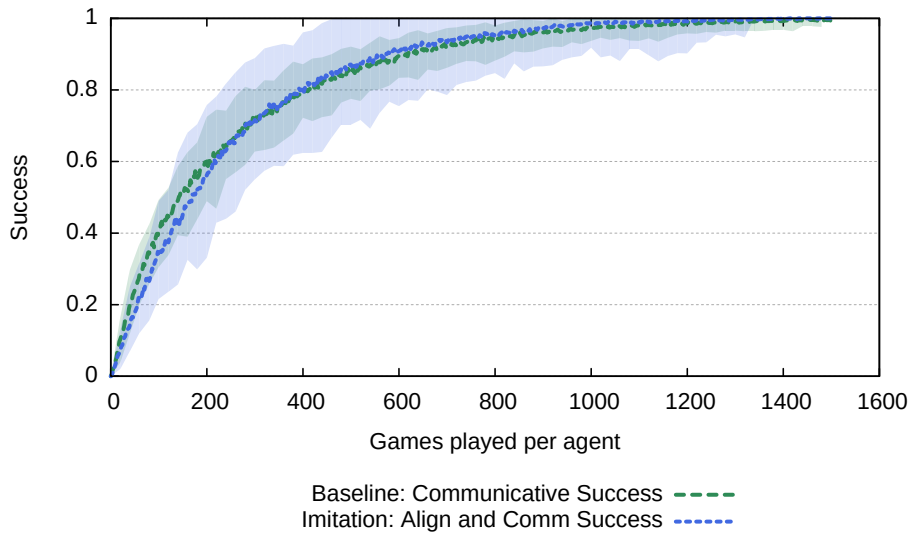

Figure 6: Communicative success and alignment success for the Imitation Strategy. The population reaches full communicative and alignment success. Parameters: population size: 50, total object size: 5, number of simulations: 100, error: 5 to 95 percentile.

A crucial feature of the Imitation Strategy is the *adaptivity* of the agents in the sense that, as listeners, agents change their own inventory when contradicted by a speaker. Imagine a similar but more "stubborn" variant of the strategy where $f_{adopt}(n, t, a)$ is modified so that in case the listener already knows another name, he does *not* adopt the new pairing $< n, t >$. Agents only adapt their lexicons when they do not yet know a name for the object (see Figure 8a). This minor modification to the strategy breaks the winner takes all dynamics and results in populations not capable of even reaching communicative success as shown in Figure 8b[5]. A non-adaptive strategy does not yield the desired result because agents are not willing to change their preference.

3.5 The Minimal NG Strategy

The last minimal strategy discussed is better known as the *Minimal Naming Game* first introduced by Baronchelli et al. (2006a,b), hence we refer to it as the *Minimal NG Strategy*. The Minimal NG Strategy has its roots in opinion dynamics (Castellano et al., 2009) and is related to the models of Fu and Wang (2008); Blythe (2009). For a detailed exposition of the most important aspects of the Minimal NG Strategy please see Baronchelli (2012).

[5]Only in one exceptional case can a population reach communicative success. For each object the first invented name should spread through the entire population before a second invention for that object takes place.

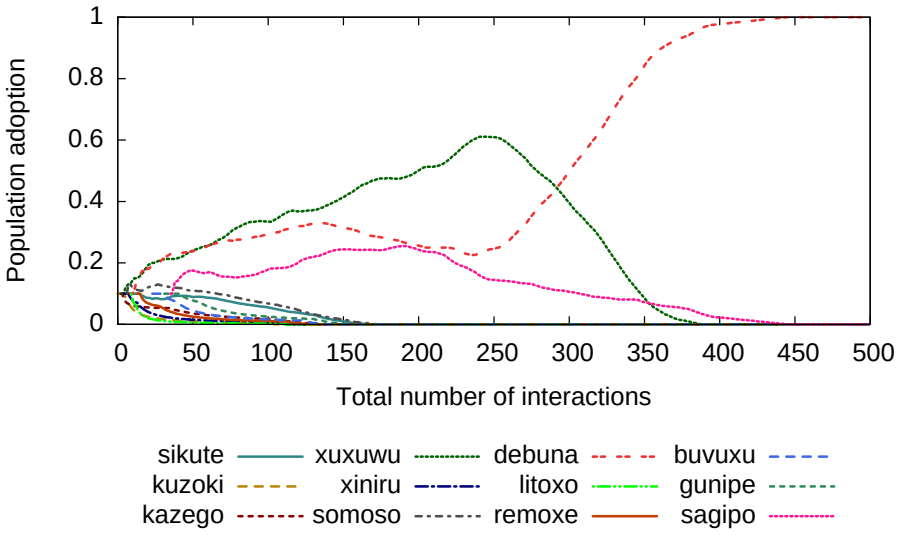

Figure 7: Name competition for the Imitation Strategy. Through self-organization the name "debuna" comes out as single preferred name by the full population. A total of 12 names have been invented for the object. Just as for Figure 4(c), a population of 20 agents was used. Comparison to this figure clearly illustrates the difference in dynamics.

The Minimal NG strategy brings together the Baseline NG Strategy and the Imitation Strategy. From the baseline strategy it keeps the ability to store multiple names per object with the difference that in case of a successful game (i.e. the listener knows the name), both agents remove all other competing names for that object as shown in Figure 9. Since this update cannot be seen as part of f_{invent} or f_{adopt} a fifth function $f_{align}(n, t, a)$ needs to be added to the strategy. Alignment takes place at the end of a language game, in this case, by both participating agents.

In the Minimal NG Strategy alignment depends on the communicative outcome of the game. At the end of a successful game $f_{align}(n, t, a)$ removes all pairings (n, o) for which $o \neq t$. In successful games the Minimal NG Strategy behaves identical to the Imitation Strategy whereas in failed games it behaves like the Baseline NG Strategy.

In the Minimal NG Strategy communicative success and alignment success no longer overlap. It is possible for agents to communicate successfully about an object and not yet have aligned on a preferred name for that object. In Figure 10a alignment success starts out slower than communicative success but finally catches up and both reach full success. Compared to the Imitation Strategy alignment converges drastically faster as shown in Figure 10b. The Minimal NG Strategy exhibits a fast transition as discussed by Baronchelli et al. (2006b) where the Imitation Strategy does not.

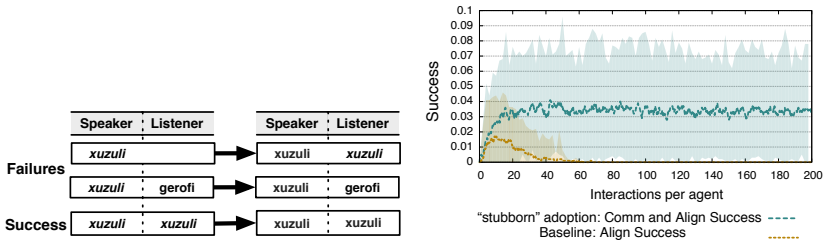

(a) Update mechanics (b) Experimental results

Figure 8: (a) Update mechanics for the "stubborn" variant of the Imitation Strategy. The speaker utters the word "xuzuli" for the chosen topic. In case of failure the listener only adopts when he did not know a name yet. (b) Communicative and alignment success for the Strict Adoption Strategy. The strategy fails at reaching acceptable levels of communicative success.

So far we have presented three minimal strategies to complement the Baseline NG Strategy. Two strategies tried to reach alignment by limiting the agents to store multiple names per object. From these two only the adaptive agents of the Imitation Strategy were able to reach alignment and communicative success. The Minimal NG strategy showed faster alignment than the Imitation Strategy by combining properties of the Baseline NG Strategy and the Imitation Strategy.

The reason why the Minimal NG Strategy converges significantly faster than the Imitation Strategy is because it implements a somewhat extreme version of *lateral inhibition*. Agents inhibit all competing names by removing them. We now turn our attention to strategies implementing a more gradual inhibition scheme.

4 Lateral inhibition strategies

The term "lateral inhibition" comes from neurobiology and refers to the capacity of an excited neuron to reduce the activity of its neighbours (Hartline et al., 1956). It is popular in connectionist modeling of pattern learning (Amari, 1977). As a mechanism in language games it was first proposed by Steels and Kaplan (1999) for the Talking Heads Experiment (Steels, 1999). In the context of the Naming Game it was introduced by Steels (2000). Lateral inhibition has since been widely used to dampen both form and meaning competition (Steels and Kaplan, 2002; Vogt, 2000; Vogt and Coumans, 2003)Before this, either ad hoc strategies or a form of frequency strategy (discussed in the next section) was used to tackle the problem of name competition. Keep in mind that the first minimal strategies were introduced only in 2005. Chronologically lateral inhibition strategies thus preceded the minimal strategies.

Lateral inhibition strategies aim to improve the previous strategies by extending the representational capabilities of the agents. Each pairing of a name

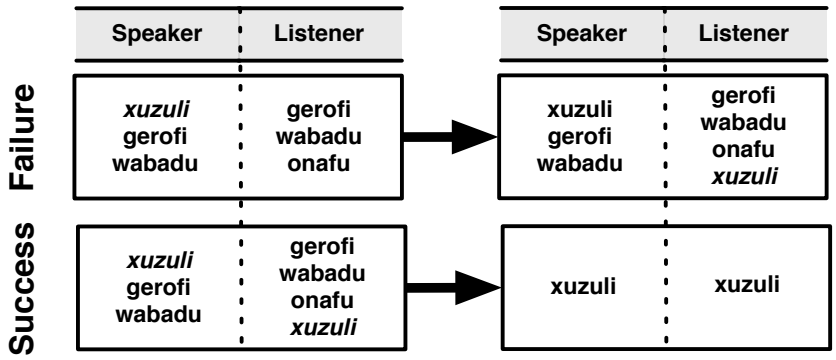

Figure 9: Update mechanics for the Minimal NG Strategy. The speaker utters the word "xuzuli" for the chosen topic. In case of failure the mechanics are the same as those form the Baseline NG Strategy (see Figure 2). In case of success they correspond to that of the Imitation Strategy (see Figure 5).

with an object is scored by a real value in the interval $]0, 1]$. When multiple names compete for the same object identifier the score keeps track of the most popular or conventional name. Names are removed from the lexicon when their scores become equal to or lower than zero. Lateral inhibition refers to the process of agents actively inhibiting competing names at the end of each game. This inhibition is achieved, not by removing the competitors as in the Minimal NG Strategy, but by lowering their score.

Production $f_{produce}$ takes the score into account by preferring the name with the highest score when multiple possibilities are available. When the highest score is shared among multiple names a random name form this list is chosen. Invention and adoption initialize the score to $s_{initial}$. Interpretation remains the same as in the Baseline NG Strategy.

We present two lateral inhibition strategies which share the above changes but differ in their alignment f_{align}. The first is based on the proposal spelled out by Steels and Belpaeme (2005). This strategy, which we call the *Basic Lateral Inhibition Strategy*, introduces three parameters, δ_{inc}, δ_{inh}, and δ_{dec}, all influencing the scores of names during alignment. At the end of a failed game the speaker lowers the score of the used name by δ_{dec}, the listener adopts with the initial score $s_{initial}$. In case of communicative success both agents increase the score of the used name by δ_{inc} and decrease the score of all competitors by δ_{inh}. By default $s_{initial}$ is 0.5, scores are capped at 1.0 and names with $s \leq 0$ are removed. Figure 11 shows the update mechanics visually. Note that with parameters $s_{initial} = \delta_{inh} = 1$ and $\delta_{inc} = \delta_{dec} = 0$ the Basic Lateral Inhibition Strategy is identical to the Minimal NG Strategy.

The second lateral inhibition strategy differs only in that instead of incrementing or decrementing the score of a name with a given δ, it interpolates

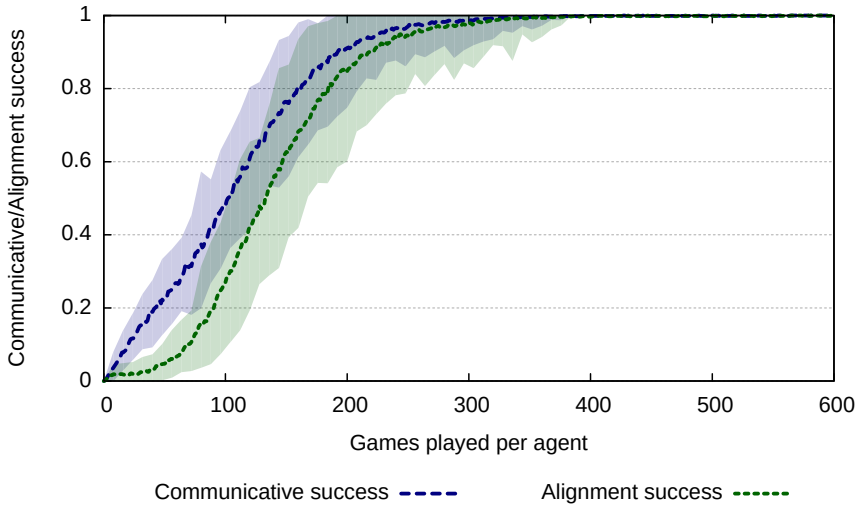

(a) Minimal NG Strategy

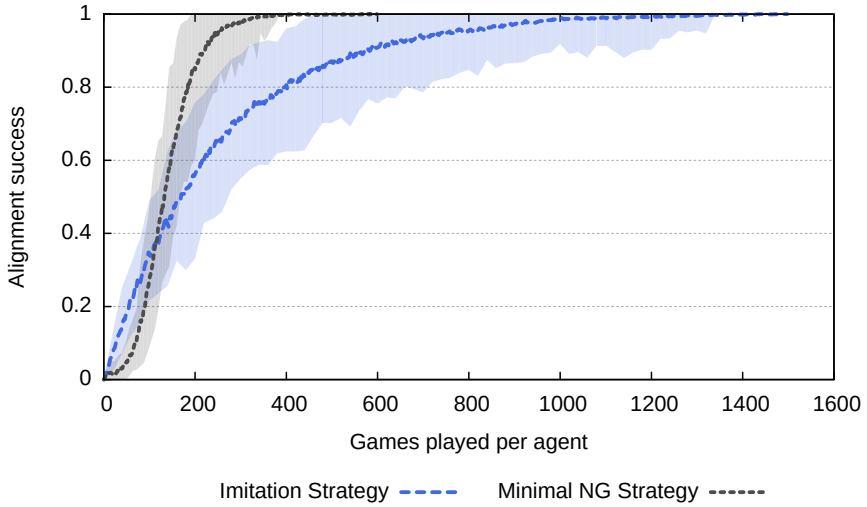

(b) Comparison with the Imitation Strategy

Figure 10: (a) Communicative success and alignment success for the Minimal NG Strategy. The agents reach full communicative success and alignment success. (b) Comparison of Imitation and Minimal NG Strategy in terms of alignment success. (a) and (b) population size: 50, object size: 5, number of simulations: 100, error: 5th to 95th percentile.

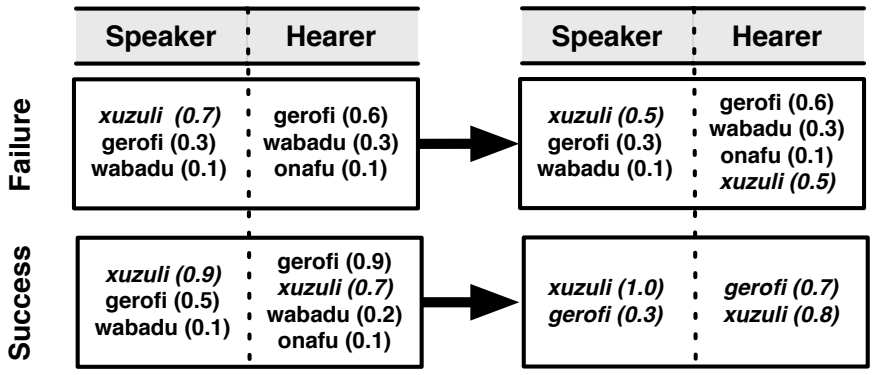

Figure 11: Update mechanics for the Basic Lateral Inhibition Strategy. The speaker utters the word "xuzuli" for the chosen topic. In case of failure the listener adopts the new name with $s_{initial}$ and the speaker decrements the spoken name with δ_{dec}. In case of success both agents increase the score of the spoken name with δ_{inc} and decrease (inhibit) the scores of the competitors with δ_{inh}. Here $s_{initial} = 0.5$, $\delta_{dec} = \delta_{inh} = 0.2$, and $\delta_{inc} = 0.1$.

towards either 1 or 0 as follows:

$$\text{reinforcement:} \qquad s \leftarrow s + \delta_{inc}(1 - s) \tag{4}$$

$$\text{inhibition:} \qquad s \leftarrow s - \delta_{inh} s \tag{5}$$

This update scheme for lateral inhibition was used for the Naming Game by Vogt and Coumans (2003) and later De Beule and De Vylder (2005) among others. It can be interpreted as the Rescorla-Wagner/Widrow-Hoff rule as described by Sutton and Barto (1981). A consequence of interpolations is that scores never reach 0 or 1 and agents thus never "forget" names, just like the Baseline NG Strategy. However, the agents develop preferences using the scores and should thus reach alignment success.

4.1 Experimental results for lateral inhibition strategies

Expectations are that the more subtle updating mechanics of the lateral inhibition strategies should result in faster alignment when compared to the Minimal NG Strategy. As explained the Minimal NG Strategy is a special case of the Basic LI Strategy which, compared to that strategy, limits the memory capacity of the agents by resorting to a complete removal of competitors instead of a gradual inhibition. The impact of lateral inhibition, however, turns out to be relatively small. Figure 12 shows alignment success for both lateral inhibition strategies and the Minimal NG Strategy (see caption for experimental parameters). Although convergence is faster for both in the beginning and only for the

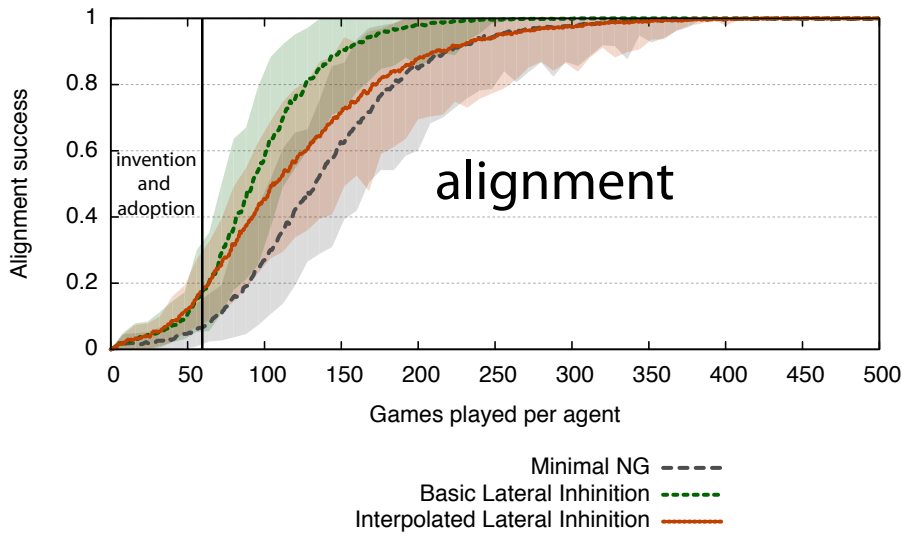

Figure 12: Alignment success for the lateral inhibition strategies. Parameters: population size: 50, object size: 5, total games: 10000, number of simulations: 100, error: 5 to 95 percentile. For both lateral inhibition strategies $\delta_{inc} = \delta_{dec} = 0.1$, $\delta_{inh} = 0.5$ and $s_{initial} = 0.5$.

Basic LI Strategy in the end, the increase remains limited. This is all the more surprising since alignment, which is what has changed between these strategies, is doing the bulk of the convergence and not invention or adoption. In the same figure the vertical dashed line shows the point at which for all three strategies invention and adoption have come to a complete stop.

In the Interpolated Lateral Inhibition Strategy alignment slows down considerably, even to such an extent that the Minimal NG Strategy takes over. This can be understood by looking more closely at the update rules for both lateral inhibition strategies. Interpolation implies that the further away from an extreme the more the score will jump toward that point if possible. For example with $\delta_{inh} = \delta_{inc} = 0.3$ and $s_{initial} = 0.5$ it takes 11 exposures (i.e. increments) to reach a score $s > 0.99$ while it takes only 2 inhibitions to drop from that score of 0.99 to below 0.5 which is the initial score. Agents thus become ever more adaptive while many studies have shown, on the contrary, that frequency effects directly impact entrenchment (Bod et al., 2000; Bybee, 1998).

The lateral inhibition strategies have parameters influencing their performance. The inhibition parameter δ_{inh} regulates the strength of the inhibition at the end of a successful game. The higher its value the more strongly competitors are inhibited. For example in the Basic Lateral Inhibition Strategy, with $\delta_{inh} = 0.5$ a name needs only be inhibited twice in a row to be certain of deletion and only once starting from $s_{initial} = 0.5$. On the other hand a value of 0.0 for δ_{inh} effectively disables inhibition and disallows the agents from

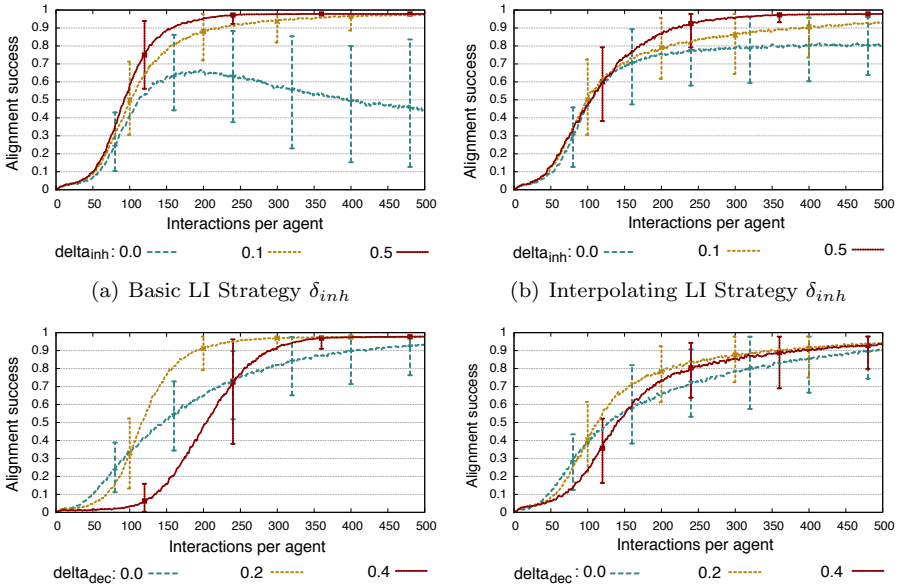

(a) Basic LI Strategy δ_{inh}

(b) Interpolating LI Strategy δ_{inh}

(c) Basic LI Strategy δ_{dec}

(d) Interpolating LI Strategy δ_{dec}

Figure 13: Impact of δ_{inh} and δ_{dec} on the Basic and Interpolating Lateral Inhibition Strategy. Alignment success for the Basic (a) and Interpolating (b) Lateral Inhibition Strategy for different values of δ_{inh}. In both (a) and (b) $\delta_{dec} = \delta_{inc} = 0.1$ and δ_{inh} takes values 0.0, 0.1 and 0.5. (c and d) Alignment success for the Basic (c) and Interpolating (d) Lateral Inhibition Strategy for different values of δ_{dec}. In both (c) and (d) $\delta_{inc} = \delta_{inh} = 0.1$ and δ_{dec} takes values 0.0, 0.2 and 0.4. (a-d) population size: 50, object size: 5, total games: 10000, number of simulations: 100, error: 5 to 95 percentile.

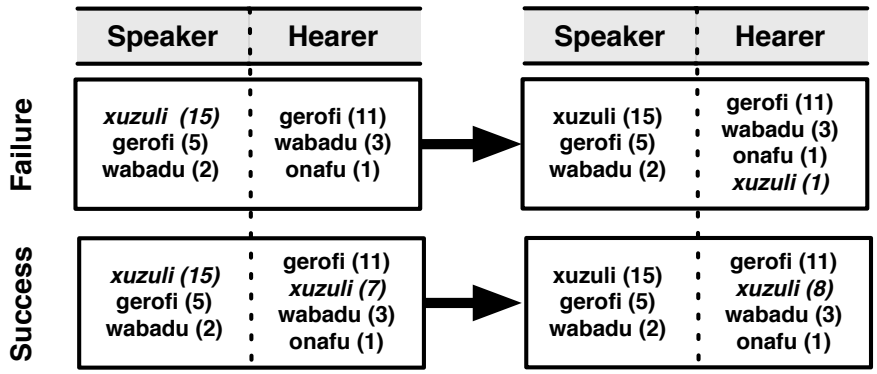

Figure 14: Update mechanics for the Frequency Strategy. Only the names for the topic object are shown. In case of failure the listener adopts the new name with frequency 1. In case of success the listener increments the frequency of the name with 1. Note that only the listener ever updates frequencies.

reaching alignment success for both strategies. Results for different values of the inhibition parameter are shown in Figure 13(a-b) for both the Basic and Interpolating LI Strategy. The tendency shows that for both strategies a higher value improves alignment, although from the results for the Minimal NG Strategy we have learned that extreme values (e.g. $\delta_{inh} = 1.0$) again negatively impact alignment.

The story is different for δ_{dec} as shown in Figure 13(c-d). Disabling this functionality by setting $\delta_{dec} = 0.0$ is detrimental for alignment in the long run. As opposed to δ_{inh} a higher value for $delta_{dec}$ does not result in faster alignment. For example with $delta_{dec} = 0.4$ alignment takes significantly longer to reach its transition point.

5 The Frequency Strategy

The Frequency Strategy was first proposed by Steels and McIntyre (1999). Instead of keeping a name-object pair $< n, o >$ agents following a frequency strategy keep a triplet $< n, o, f >$ with f an integer in the range $[1, \infty]$ representing the number of times that name was heard (not spoken).

Starting from the baseline strategy the production function $f_{produce}(o, a)$ is extended so that in case there are multiple names paired with o the one with the highest frequency is uttered. In the exceptional case that multiple names share the highest frequency a random one of those names is chosen. The learning functions f_{invent} and f_{adopt} differ only from the Baseline Strategy in that they initialize the frequency f to 1. Just like in the Minimal NG, the Frequency Strategy also needs an alignment function $f_{align}(n, t, a)$ which does nothing

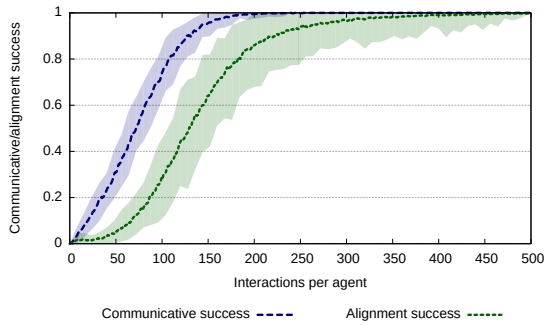

(a) Frequency Strategy

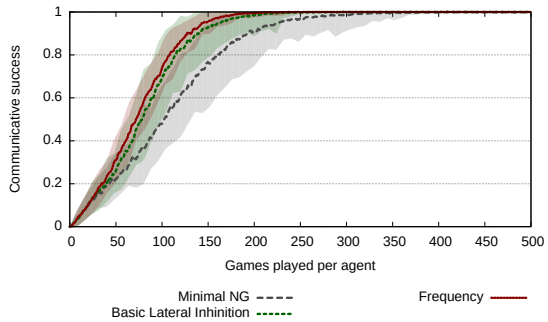

(b) Communicative Success

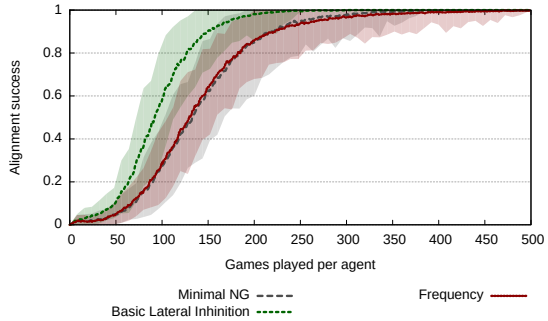

(c) Alignment Success

Figure 15: Experimental results for the Frequency Strategy. (a) Communicative and alignment success for the Frequency Strategy. (b) Comparison of Minimal NG, Basic LI and Frequency Strategy in terms of communicative success. (c) Comparison of Minimal NG, Basic LI and Frequency Strategy in terms of alignment success. (a), (b) and (c) population size: 50, object size: 5, total games: 10000, number of simulations: 100, error: 5 to 95 percentile. The Basic LI Strategy in (c) used parameters $\delta_{inc} = \delta_{dec} = 0.1$, $\delta_{inh} = 0.5$ and $s_{initial} = 0.5$.

more than update the frequency of the spoken word, but only for the listener. The reason only the listener updates is that the frequency should capture what is the preference of the other agents and his own should just be that of the others. If the agent would also increase the frequency as speaker he would, half of the time, amplify his own preference rendering him less adaptive. De Vylder (2007)(p. 139) delves deeper into this issue and compares the impact of hearer-only versus hearer-and-speaker alignment in Naming Games. These updating mechanics are also illustrated in Figure 14.

In the Frequency Strategy, just like in the Baseline NG Strategy, agents store and keep all names they ever encountered, which means they will also reach full communicative success as shown in Figure 15a. The population however reaches communicative success considerably faster because only the most frequent names will be used in production and so many invented names won't spread through the entire population. And indeed Figure 15b shows that communicative success with the Frequency Strategy much improves upon the Minimal NG Strategy, which itself improved on the Baseline NG Strategy. The reason for this lies in the fact that the Minimal NG Strategy deletes competing names in successful games, which might lead to a failure if one of the competing names is encountered later on. For alignment success we do not see the same gain as with communicative success. Alignment success remains below that of the Basic LI Strategy and almost coincides with that of the Minimal NG Strategy. Why alignment success between the Minimal NG Strategy and the Frequency Strategy is near identical remains an open question to me.

6 Conclusion

In this chapter I focused on the problem of form competition and the Naming Game that introduces it. The Naming Game is designed in such a way that it evokes only the conventionalization problem of name or word form competition. By equating meanings with holistic objects and perfect feedback the problem of meaning-related uncertainty is avoided. The only problem the agents need to solve is that of bootstrapping and maintaining a minimal set of names for a set of objects.

By presenting and comparing different strategies I have shown that multiple solutions exist to the problem of name or word form competition. Alignment, either by removing or inhibiting competitors or by counting frequencies, is crucial for a population to successfully play the Naming Game. All these successful strategies, except the Imitation Strategy, showed a fast transition from low levels of alignment to high levels of alignment. This sudden rise in the emergence of the communication system is also observed in the formation of basic human communication systems (Galantucci, 2005).

From the Imitation Strategy we learned that sometimes, less is more. By reducing the representational capabilities, the agents were capable of reaching full alignment because it gave rise to a winner takes all dynamics. Imitation agents can be seen as maximally adaptive, in that they immediately change

Strategy	Baseline	Imitation	Minimal NG	LI	Freq
(1) full CS	yes	yes	yes	yes	yes
(2) WFC solved	no	yes	yes	yes	yes
(3) scoring	no	no	no	yes	yes
(4) fast transition	N/A	no	yes	yes	yes
(5) competition damping	no	no	yes	yes	no

Table 1: Overview of Naming Game strategies. Depicts whether (1) strategies could reach full communicative success (CS), (2) could solve word form competition (WFC), (3) implemented a scoring mechanism, (4) showed a fast transition in alignment, and whether (5) they implemented an explicit damping of competitors.

their preference upon an exposure. Adaptivity, in the sense that agents need to change their name inventory based on their communicative encounters, is indeed a crucial ingredient for the emergence of a minimal naming system.

The limited representational capabilities of the Imitation agents came at the cost of slow alignment. The Minimal NG Strategy extends the abilities of agents in terms of both representation and processing. On the representational level it allows agents to maintain lists of competing names, which in turn allows the agents to be slightly less "adaptive" when confronted with a novel name. Agents can store the name as a competitor of the other names he has heard for that object. Only when a name is either heard os used successfully again will he remove its competitors. This change lead to considerably faster convergence both in terms of communicative as alignment success.

The Minimal NG Strategy is a special case of a more general type of strategies that rely on lateral inhibition. Successful words are awarded and competing words are damped. Such a scheme was implemented in a gradual way by adding a score to each name object pairing. With this added representational capability agents did not need to remove competing names so hastily as in the Minimal NG Strategy. Instead, based mainly on the inhibition parameter δ_{inh}, agents have a more gradual memory. It turned out that lower values for inhibition did not improve alignment success.

From these results it seems that adding a scoring mechanism yields only marginal improvements in terms of communicative and alignment success. Although this holds true for the Naming Game we will see in the following chapter that when other problems are introduced next to the problem of word form competition, slower update mechanics become increasingly beneficial and even crucial for alignment.

The last discussed strategy never removes or inhibits names and thus differs considerably from all previous extensions to the Baseline Strategy. It extends the baseline by adding the representational capability of keeping a frequency for each name object pair. This frequency only needs to be updated as a listener and can be done regardless of communicative success. In production a speaker following the Frequency Strategy chooses the most frequent corresponding name.

This strategy reaches full communicative and alignment success. While it does not improve alignment success compared to the Minimal NG Strategy it does improve communicative success. Because it does not implement a direct competition inhibition like the lateral inhibition strategies, alignment moves slightly slower. An overview of all strategies is given in Table 1.

References

Shun-ichi Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27:77–87, 1977. ISSN 0340-1200. URL `http://dx.doi.org/10.1007/BF00337259`. 10.1007/BF00337259.

A Baronchelli, L DallAsta, A Barrat, and V Loreto. The role of topology on the dynamics of the naming game. *The European Physical Journal Special Topics*, 143(1):233–235, 2007. URL `http://www.springerlink.com/index/10.1140/epjst/e2007-00092-0`.

Andrea Baronchelli. The minimal naming game: a complex systems approach. In Luc Steels, editor, *Experiments in Cultural Language Evolution*. John Benjamins, Amsterdam, 2012.

Andrea Baronchelli, Emanuele Caglioti, Vittorio Loreto, and Luc Steels. Complex systems approach to language games. In *Proceedings of the Second European Conference on Complex Systems ECCS'06*, Oxford (UK), 2006a.

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

Andrea Baronchelli, Claudio Castellano, and Romualdo Pastor-Satorras. Voter models on weighted networks. *Physical Review E.*, 83, 066117, 2011.

Anrea Baronchelli, Vittorio Loreto, and Luc Steels. In-depth analysis of the naming game dynamics: The homogeneous mixing case. *International Journal of Modern Physics C*, 19:785–812, 2008.

Richard A. Blythe. Generic modes of consensus formation in stochastic language dynamics. *J. Stat. Mech*, page P02059, 2009.

R. Bod, J. Hay, and S. Jannedy. *Probabilistic linguistics*. MIT Press, 2000.

Joan Bybee. The emergent lexicon. *CLS 34: The panels*, pages 421–435, 1998.

Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81:591–646, 2009.

Joachim De Beule and Bart De Vylder. Does language shape the way we conceptualize the world? In B. G. Bara, L. Barsalou, and M. Bucciarelli, editors, *the XVII Annual Conference of the Cognitive Science Society*, pages 559–564, 2005. ISBN ISBN 0-9768318-1-3.

Joachim De Beule, Bart De Vylder, and Tony Belpaeme. A cross-situational learning algorithm for damping homonymy in the guessing game. In Luis Mateus Rocha, Larry S. Yaeger, Mark A. Bedau, Dario Floreano, Robert L. Goldstone, and Allessandro Vespignani, editors, *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 466–472, Cambridge, MA, 2006. MIT Press.

Bart de Boer. Self organization in vowel systems. *Journal of Phonology*, 28: 441–465, 2000.

Bart de Boer and W. Tecumseh Fitch. Computer models of vocal tract evolution: an overview and critique. *Adaptive Behaviour*, 18:36–47, 2010.

B. De Vylder and K. Tuyls. How to reach linguistic consensus: A proof of convergence for the naming game. *Journal of Theoretical Biology*, 242(4): 818–831, October 2006.

Bart De Vylder. *The Evolution of Conventions in Multi-Agent Systems*. PhD thesis, Vrije Universiteit Brussel, Brussels, 2007.

Feng Fu and Long Wang. Coevolutionary dynamics of opinions and networks: From diversity to uniformity. *Physical Review E*, 78, 2008.

Bruno Galantucci. An experimental study of the emergence of human communication systems. *Cognitive Science*, 29(5):737–767, 2005.

H.K. Hartline, H.G. Wagner, and F. Ratliff. Inhibition in the eye of limulus. *J. Gen. Physiology*, 39:651, 1956.

Tom Lenaerts, Bart Jansen, Karl Tuyls, and Bart De Vylder. The evolutionary language game: An orthogonal approach. *Journal of Theoretical Biology*, 235 (4):566 – 582, 2005.

A. Lipowski and D. Lipowska. Bio-linguistic transition and baldwin effect in an evolutionary naming-game model. *Int. J. of Mod. Phys. C*, 19(3):399 – 408, 2008.

R.R. Liu, C.X. Jia, H.X. Yang, and B.H. Wang. Naming game on smallworld networks with geographical effects. *Physica A*, 388:3615–3620, 2009.

Q. Lu, G. Korniss, and B.K. Szymanski. Naming games in twodimensional and small-world-connected random geometric networks. *Physical Review E*, 77 (1:16111), 2008.

Pierre-Yves Oudeyer. The self-organization of speech sounds. *Journal of Theoretical Biology*, 233(3):435 – 449, 2005.

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

Luc Steels. Self-organizing vocabularies. In Christopher G. Langton and Katsunori Shimohara, editors, *Proceeding of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge, MA, 1996. MIT Press.

Luc Steels. The origins of ontologies and communication conventions in multi-agent systems. *Journal of Agents and Multi-Agent Systems*, 1(1):169–194, 1998.

Luc Steels. *The Talking Heads Experiment. Volume 1. Words and Meanings.* 1999. Special pre-edition for LABORATORIUM, Antwerpen.

Luc Steels. Language as a complex adaptive system. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J.J. Merelo, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature*, LNCS 1917, pages 17–28, Berlin, 2000. Springer-Verlag.

Luc Steels and Tony Belpaeme. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4):469–489, 2005.

Luc Steels and Frédéric Kaplan. Spontaneous lexicon change. In *Proceedings COLING-ACL 1998*, pages 1243–1250, San Francisco, CA, August 1998a. Morgan Kaufmann.

Luc Steels and Frédéric Kaplan. Stochasticity as a source of innovation in language games. In Christoph Adami, Richard K. Belew, Hiroaki Kitano, and Charles E. Taylor, editors, *Proceedings of the Sixth International Conference on Artificial Life*. MIT Press, 1998b.

Luc Steels and Frédéric Kaplan. Collective learning and semiotic dynamics. In Dario Floreano, Jean-Daniel Nicoud, and Francesco Mondada, editors, *Advances in Artificial Life. Proceedings of the Fifth European Conference, ECAL'99*, volume 1674 of *Lecture Notes in Computer Science*, pages 679–688, Lausanne, Switzerland, September 1999. Springer.

Luc Steels and Frédéric Kaplan. Bootstrapping grounded word semantics. In T. Briscoe, editor, *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, pages 53–73. Cambridge University Press, Cambridge, 2002.

Luc Steels and Martin Loetzsch. The grounded naming game. In Luc Steels, editor, *Experiments in Cultural Language Evolution*. John Benjamins, Amsterdam, 2012.

Luc Steels and Angus McIntyre. Spatially distributed naming games. *Advances in Complex Systems*, 1(4):301–323, January 1999. doi: 10.1142/S021952599800020X.

R. Sutton and A. Barto. Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88:135 – 170, 1981.

Sven Van Segbroeck, Steven de Jong, Ann Now, Francisco C Santos, and Tom Lenaerts. Learning to coordinate in complex networks. *Adaptive Behavior*, 18:416 –427, 2010.

Remi van Trijp. The emergence of semantic roles in Fluid Construction Grammar. In Andrew D.M. Smith, Kenny Smith, and Ramon Ferrer i Cancho, editors, *The Evolution of Language. Proceedings of the 7th International Conference (EVOLANG 7)*, pages 346–353, Singapore, 2008. World Scientific Press.

Paul Vogt. *Lexicon Grounding on Mobile Robots*. PhD thesis, Vrije Universiteit Brussel, 2000.

Paul Vogt and Hans Coumans. Investigating social interaction strategies for bootstrapping lexicon development. *Journal of Artificial Societies and Social Simulation*, 6(1), 1 2003.

Ludwig Wittgenstein. *Philophische Untersuchungen*. Suhrkamp, Frankfurt, 1967.