

The evolutionary language game: An orthogonal approach

Tom Lenaerts^{a,*}, Bart Jansen^b, Karl Tuyls^c, Bart De Vylder^b

^a*Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle, Université Libre de Bruxelles,
Avenue Franklin Roosevelt 50, 1050 Brussels, Belgium*

^b*Artificial Intelligence Lab, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium*

^c*Theoretical Computer Science Group, Limburgs Universitair Centrum, Universitaire Campus, Building D, 3590 Diepenbeek, Belgium*

Received 24 April 2004; received in revised form 9 February 2005; accepted 9 February 2005

Available online 23 March 2005

Communicated by Martin Nowak

Abstract

Evolutionary game dynamics have been proposed as a mathematical framework for the cultural evolution of language and more specifically the evolution of vocabulary. This article discusses a model that is mutually exclusive in its underlying principals with some previously suggested models. The model describes how individuals in a population culturally acquire a vocabulary by actively participating in the acquisition process instead of passively observing and communicate through peer-to-peer interactions instead of vertical parent–offspring relations. Concretely, a notion of social/cultural learning called the naming game is first abstracted using learning theory. This abstraction defines the required cultural transmission mechanism for an evolutionary process. Second, the derived transmission system is expressed in terms of the well-known selection–mutation model defined in the context of evolutionary dynamics. In this way, the analogy between social learning and evolution at the level of meaning–word associations is made explicit. Although only horizontal and oblique transmission structures will be considered, extensions to vertical structures over different genetic generations can easily be incorporated. We provide a number of simplified experiments to clarify our reasoning.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Language acquisition; Cultural evolution; Horizontal transmission; Evolutionary dynamics

1. Introduction

Words form the basic unit of a language. Humans use these words to identify things or actions in their environment. Using a word implies that the user also associates a particular meaning with that word. Wrong associations between words and meanings leads to misinterpretations which require corrections. These associations do not stand on their own. Vocabulary and more generally language are population-level phenomena which spread through cultural transmission systems. Questions concerning the minimal requirements of these transmission schemes in order to acquire

a language are pertinent to the better understanding of the actual dynamics. In order to answer these, and other, questions, we need well-founded models which examine the cultural evolution of language.

It is established that social learning plays a crucial role in the acquisition of language. Although, there is not a real consensus on the actual social learning mechanisms. [Boyd and Richerson \(1985\)](#) generally define social learning as the transmission of stable dispositions by teaching or imitation. Yet, the actual mechanics of the latter processes can differ. Two well-known models are the observational learning model and the operant conditioning model ([Rosenthal and Zimmerman, 1978](#)). The first model assumes that learning occurs by pure observations and that properties are acquired through statistical sampling of these observations. The second stresses the importance of a stimulus and the response when acting upon this stimulus. In the

*Corresponding author. Tel.: +32 2 650 27 12; fax: +32 2 650 27 15.

E-mail addresses: tom.lenaerts@ulb.ac.be (T. Lenaerts),
bartj@arti.vub.ac.be (B. Jansen), karl.tuyls@luc.ac.be
(K. Tuyls), bartdv@arti.vub.ac.be (B. De Vylder).

context of language acquisition observational learning has been examined in numerous situations. Our view on language acquisition belongs to the second type of models. The primary motivation for choosing this perspective is that we consider language learning to be functional, i.e. directed toward the communication of meaning.

The social learning scheme defines the general layout of the cultural evolutionary system, yet different transmission structures exist. Three alternative forms are often cited: vertical, oblique (role-model) and horizontal transmission (Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 1985). The work here focuses on the latter and some small comments will be made about oblique transmission. Horizontal transmission refers to the transmission between peers instead of parents and children (as in vertical transmission). Hence, there is no primary role, individuals can be either antagonist or protagonist in the social interaction. In oblique transmission, the roles are explicitly defined. An individual is either teacher or student. Moreover, oblique transmission implicitly assumes that among the teachers there is some degree of coherence in the language. The major difference with vertical transmission is that in both horizontal and oblique structures, the transmission can occur within one ‘genetic’ generation. Hence, there is a difference between ‘genetic’ and ‘cultural’ time. Boyd and Richerson (1985) refer to this situation as an asymmetric inheritance system. This difference may lead to conflicts since cultural transmission may favor other trait variants than genetic transmission. We will not explicitly discuss such conflicts here.

The previous discussion provides a combination of two underlying principals which are orthogonal with those used for previously designed models (Hurford, 1989; Oliphant and Batali, 1997; Nowak et al., 1999; Nowak and Komarova, 2001; Kirby and Hurford, 2002). Yet, as far as we are aware, none of them actually considers the suggested perspective. Nevertheless, large amount of evidence exists that both cornerstones have played a crucial rule in the origin and evolution of language. To clarify the major differences, take for instance the mathematical framework discussed by Nowak et al. (1999). First, as indicated earlier, the vocabulary in their model is acquired through observational learning, i.e. learning how to associate a word and a meaning without experiencing it oneself (Rosenthal and Zimmerman, 1978; Boyd and Richerson, 1985). In this approach, imitation of observed behaviors between communicating population members forms the primary mechanism for the student/child to acquire the language. Hence, the role of cognitive processes is restricted to an almost literal imitation of the shared lexicon. Second, the authors apply a cultural evolutionary model with vertical transmission between ‘genetic’ generations. The dynamics describe how through a process of

blending inheritance the lexicon of the different population members converges toward a shared one. In their context, the communication between the (cultural or genetical) parents of the same genetic generation was examined. Our aim is to provide an alternative mathematical framework that incorporates those features of cultural evolutionary system which are orthogonal to their model.

In Section 2 we will outline the basic model. Afterward, in Section 3 the cultural transmission scheme which describes how the associations between words and meanings are transmitted between peers is discussed. The dynamics of the transmission scheme will be analysed in Section 4. In Section 5, the dynamics of the evolutionary language game will be outlined and discussed. In this section, it will be shown that the cultural transmission scheme is equivalent to selection–mutation models discussed in the context of evolutionary game dynamics. Finally, the paper will be concluded in Section 6.

2. The model

2.1. The complete picture

The complete model consists of a population of individuals which possess a number of capabilities to acquire and communicate meaning. Here, it is assumed that each individual can perform a number of tasks: direct the others attention toward some objects in an environment, perceive these objects and assign meaning to them (Steels, 1995, 1999; Tomasello, 2003). The most primitive way to perform this first task is by pointing or some other gesture to manipulate the attention of the individual with whom one wants to communicate. The second and third task require that an individual maintains a set of meanings which are associated with different objects in the environment and a lexicon which collects the associations between the different words and meanings. Hence, the functional process consists of a combination of individual learning to discriminate objects in the environment and cultural transmission to communicate this meaning towards others. The current discussion we will only consider the latter since it captures the elements of the evolutionary process which we want to investigate.

Two final assumptions are made. First, we will, at each step of the discussion, always consider the communicative effects between two individuals first. Second, all individuals have homogeneous language skills. The initial simplification is introduced to capture the actual social learning process before making any generalisations towards populations. Yet, the extension towards populations is crucial due to the population-level consequences of cultural transmission. The latter

assumption is introduced as a simplification. All individuals use the same mechanisms to acquire the lexical information. Plasticity of the phenotype in terms of the social learning scheme and its actual parameters will be introduced, yet not within the same population. Such an extension towards heterogeneous populations will be made in future work where the combination of cultural and genetic transmission in one model is discussed.

2.2. The abstraction

Given these assumptions, we can define the population members. The state of the i th individual ϕ^i consists, on one hand, of a set of associations \mathbf{D}^i between objects from the environment (o_j) and features that discriminate an object from the other objects in the environment d_j^i

$$\mathbf{D}^i = \{(o_1, d_1^i), (o_1, d_2^i), \dots, (o_2, d_j^i), \dots, (o_j, d_l^i), \dots\}.$$

The set \mathbf{D}^i allows ambiguous pairing of discriminating features and objects. A possible example of these features could be the differences in color and shape of a set of cars. On the other hand, the state contains a lexicon \mathbf{L}^i which is again a set of associations between particular meanings (d_j^i) and words (w_l^i)

$$\mathbf{L}^i = \{(d_1^i, w_1^i), (d_1^i, w_2^i), \dots, (d_j^i, w_2^i), \dots, (d_k^i, w_l^i), \dots\}.$$

Again the set \mathbf{L}^i allows ambiguous pairing between meanings and words. Since the origin of the relation between objects and meanings is not investigated here, it will be assumed that each individual uses the same associations between objects and meanings. Hence, $\mathbf{D}^i = \mathbf{D}$ for all individuals i and moreover, the association set \mathbf{D} is finite and contains d elements.

Next, each individual has an initial finite collection of words W of size w which will be shared between all individuals. In other words, all words are given and none will be created. Although in practice unrealistic, this is introduced here to keep things simple. The number of word–meaning associations will be $n = w \times d$. Convention specifies that these associations between words and meanings are represented by dynamic lexicons. These lexicons are dynamic since their contents, i.e. associations between words and meanings, changes over time according to the use of the different words for different meanings during communication. Hence, words which are used infrequently may disappear over time.

To simulate this adaptive behavior, each entry (d_k^i, w_l^i) in lexicon \mathbf{L}^i is associated with a value v_{kl}^i . This value specifies a strength of the association between a particular meaning and word. It is assumed that each value v_{kl}^i is in the range $[0, 1]$. Hence, values v_{kl}^i close to 1 specify a strong relation and values v_{kl}^i close to 0 specify a weak relation. Hence, the lexicon \mathbf{L}^i can be defined as

a matrix with rows and columns specifying the strength of the associations between meanings and words:

$$\mathbf{L}^i = \begin{pmatrix} v_{11}^i & v_{12}^i & \dots & v_{1w}^i \\ v_{21}^i & v_{22}^i & \dots & v_{2w}^i \\ \dots & \dots & \dots & \dots \\ v_{d1}^i & v_{d2}^i & \dots & v_{dw}^i \end{pmatrix}.$$

Note that it is not assumed that \mathbf{L}^i is a probability matrix. In summary, the dynamically sufficient state of a language individual i is defined as

$$\phi^i = (\mathbf{D}, \mathbf{L}^i).$$

Given this definition, a framework for language evolution of an individual should express how the state changes using some *transformation laws* T , producing a trajectory in the state–space defined by the lexical matrices.

2.3. Transformation of the lexical matrix

As argued in the introduction, the transformation of the lexical matrix is determined by the social learning process and the cultural transmission structure. The social learning mechanism will determine which values will change and with what amount and the cultural transmission structure will determine whose values will change, i.e. both speaker and hearer or some other configuration.

The social learning process will be the (pointing-and-) *naming game* (Steels, 1996, 1999; Tomasello, 2003). The game occurs between two individuals in a shared environment. The goal is to learn a shared lexicon. In the game, each individual can play either of two roles: speaker or hearer. The actual role will influence the learning progress of each individual. Although the naming game itself is an extreme simplification of the actual process of language acquisition, it highlights the mechanics of the cultural transmission process which we try to understand. The cultural transmission structure is peer-to-peer interaction or horizontal transmission. Each individual participates as an equal partner in the interaction and they will learn from each other. To understand the dynamics of this cultural evolution system we need a well-founded mathematical framework. The following sections will describe how we derive the mathematical framework and relate it to standard evolutionary models.

3. The cultural transmission scheme

In this section we will first describe the game itself. We will always refer to the speaker with the subscript i and the hearer with subscript j . The naming game will either end in success or failure and based on this outcome, the

lexicon is updated. In Section 3.2 we will describe the updating scheme. A few simple examples will be given in Section 3.3.

3.1. The naming game

Assume, first, two individuals. The game occurs as follows: we know from Section 2.1 that both individuals are capable of determining a discriminative meaning (d_k^i and d_l^j , respectively) for the specific object in an environment. Yet, it is assumed that in the game only the speaker will perform this step (Steels, 1999). Assume further that each individual, speaker and hearer, respectively, has a lexical matrix \mathbf{L}^i and \mathbf{L}^j . The speaker utters a word w_a^i that is associated with the meaning d_k^i with the highest value v_{ka}^i .

This word is received by the hearer who in turn collects the different meanings associated with the word w_a^i . Given this collection of meanings, the hearer determines the meaning d_l^j which has the highest strength in relation to the received word, i.e. v_{lb}^j . If both meanings are the same, the game is a success. If not, the game fails. To inform the speaker about the success or failure of the game, the hearer utters the word (w_b^j) it associates with the specific object. As a result both individuals know whether or not the game succeeded and can act accordingly. This action is an update scheme that alters the different values in the lexical matrices of both speaker and hearer. The next section, describes the rules.

3.2. The updating scheme

Upon success or failure, both individuals update the values in their lexical matrices. Remember that w is the number of words in the set of words of each individual and d is the number of meanings in the set of meanings of each individual. Moreover, δ is a value in the interval $[0, 1]$ and specifies the amount of change used for updating the lexical matrices. Typically the value of $\delta = 0.1$. Given these parameters, the rules are:

- When the game is a success, the values v_{ka}^i and v_{lb}^j associated with the specific meaning–word combination in both individuals are reinforced with δ . Simultaneously, the speaker decreases the values corresponding to the associations between different words and the same meaning with a penalty value $\delta/(w-1)$. While the hearer decreases the values corresponding to the associations between different meanings and the same word with a penalty value $\delta/(d-1)$.
- When the game fails, the opposite changes are performed, i.e. the values v_{ka}^i and v_{lb}^j are decreased with δ , the speaker increases the values corresponding to associations between different words with the same

meaning using the value $\delta/(w-1)$ and the hearer increases the values corresponding to the associations between different meanings with the same word using the value $\delta/(d-1)$.

In both cases, the updating scheme will ensure that the values v_{kl}^i remain in the range $[0, 1]$ (this is ensured by reducing the value v_{kl}^i back to 1 when it exceeds 1 or increasing the value v_{kl}^i back to 0 when it is less than 0 after the update).

3.3. Simple examples

In order to assess the success of communication between two individuals, some measure is required. In the naming game, success occurs when the individuals can understand each other. Hence, one can count the number of successes and failures during a certain period of time over all the individuals. We define this success rate as the communicative *coherence* between individuals or within a population.

In the following examples, listed in Fig. 1, this success is used to assess the performance of two individuals. Depending on the settings of the parameters different things can be observed. Here, we observe the rate of convergence and whether the absorbing states of the game correspond to a state of optimal communicative coherence. Moreover, we examine whether these absorbing states are attained when w is not equal to d , i.e. homonymy and synonymy scenarios. The role of speaker or hearer is assigned randomly. As one can see when comparing the top-left and top-right plots in Fig. 1, the reinforcement value δ will determine the rate of convergence toward a shared lexicon.

In order for the naming game to be close to reality, the outcome of the game will need to show that homonymy is common and synonymy is rare. Homonyms are words which have different meanings yet the same form. The collection of homonyms consists of two major classes; coincidental and polysemic homonyms. The first kind are the result of chance and there is usually no relation between the different meanings. An example is ‘bark’ which can either refer to the skin of a tree or the sound of a dog. The second kind results from the historical development of different meanings for the same word. There usually is some semantic link between the two meanings. An example is the leg of either a table or a person.

In the context of the reduced naming game discussed here it is difficult to examine these issues since the number of words and meanings are fixed. Yet, one can observe whether the learning dynamics would allow for homonyms to maintain themselves. As one can see in the lower-left plot in Fig. 1, when $w < d$ the game will not converge to a maximal communicative coherence and will remain oscillating somewhere below this maximum. When one inspects the lexical matrices of the two

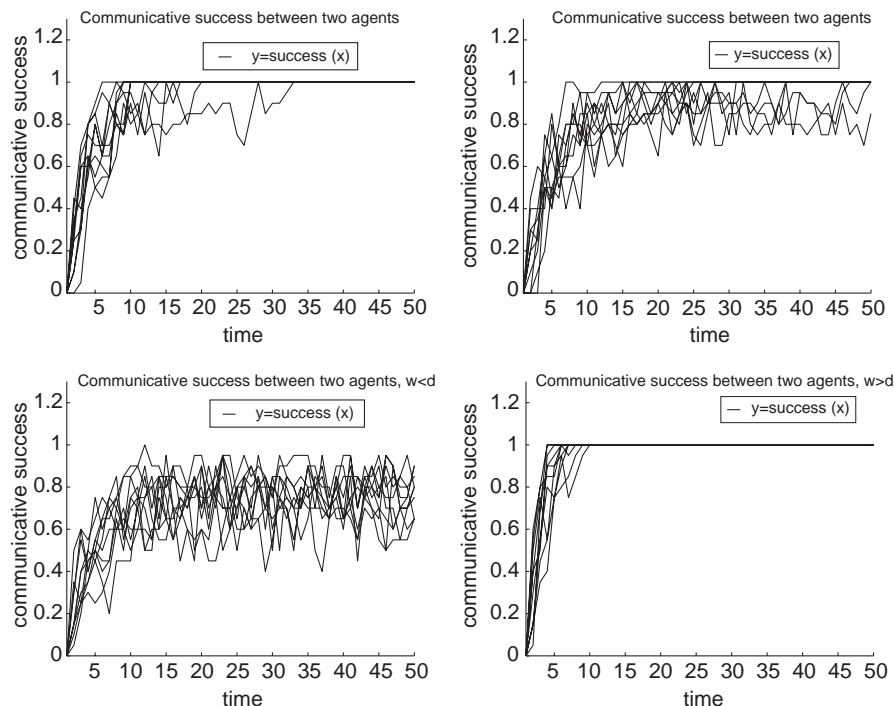


Fig. 1. *Communicative coherence between two individuals*: Plotting the communicative coherence of a two-individual naming game for the first 1000 iterations. The coherence is recorded for an interval of 20 iterations. The setup of the game consists of a set W of seven words and a set D of seven meanings. The learning occurs between two individuals whose role is determined randomly. Each plot shows the communicative coherence which measures how well two individuals understand each other. On each plot, 10 runs are shown. Other parameter settings are: $\delta = 0.1$ for all the figures except for the top-right one where $\delta = 0.05$. The two figures in the lower-left and lower-right show the performance of the naming game when the number of meanings differs from the number of words. In the figure in the lower-left, there are more meanings than words. This setup allows for one word to refer to different meanings (coincidental and polysemic homonyms). In the figure in the lower-right, there are more words than meanings. In this setup we allow different words with similar meaning (synonyms).

individuals there will always be multiple interpretations for some words. Moreover, these interpretations will differ between individuals resulting in lower communicative coherence.

Synonyms are words with similar meanings. As for the homonyms, one can verify whether under the learning dynamics of the naming game, synonyms survive. The general expectation is that they disappear. To validate this expectation a situation is constructed where synonyms can be created. This is the case when there are more words than meanings. In this experimental setup, the naming game will always converge to binary matrices where some words will not be used, i.e. some columns in the lexical matrices will contain all zeros. These results can be observed in the lower-right plot in Fig. 1. Hence, the experiments seem to show that synonyms disappear, yet polysemic words remain forever as long as no new words are introduced.

In Fig. 2, the same experiments are shown for a population of 10 individuals. At each iteration two individuals are selected randomly. It can be observed that similar observations apply. Homonyms remain in the population and synonyms disappear. Moreover, an initial observation can be made on the rate of convergence in relation to the number of individuals in

the population: the more individuals in the system, the more time it will take. This difference in convergence rate can be observed when comparing the two top-panels in Fig. 2. A more extensive analysis is required to provide a concrete proof on the relation between the parameters and the convergence behavior. Some previous work conducted by Kaplan (2000) can illuminate already some issues.

In general, the naming game belongs to the class of operant conditioning models of social learning. The stimulus that is provided here is determined by, on one hand, the manipulation of the individuals toward a joint attentional space and, on the other hand, the utterance of a word which the speaker associates with the word. The response is triggered by the hearer when it does not understand the speaker and can be either success or failure. From this response, both speaker and hearer derive the necessary information to alter their internal lexical beliefs. Rosenthal and Zimmerman (1978) discuss at length different situations where language is acquired through this approach. In that context the naming game provides a computational model that shows that language can indeed be learned in this way. Now in order to understand the dynamics, a mathematical framework is derived from the naming game.

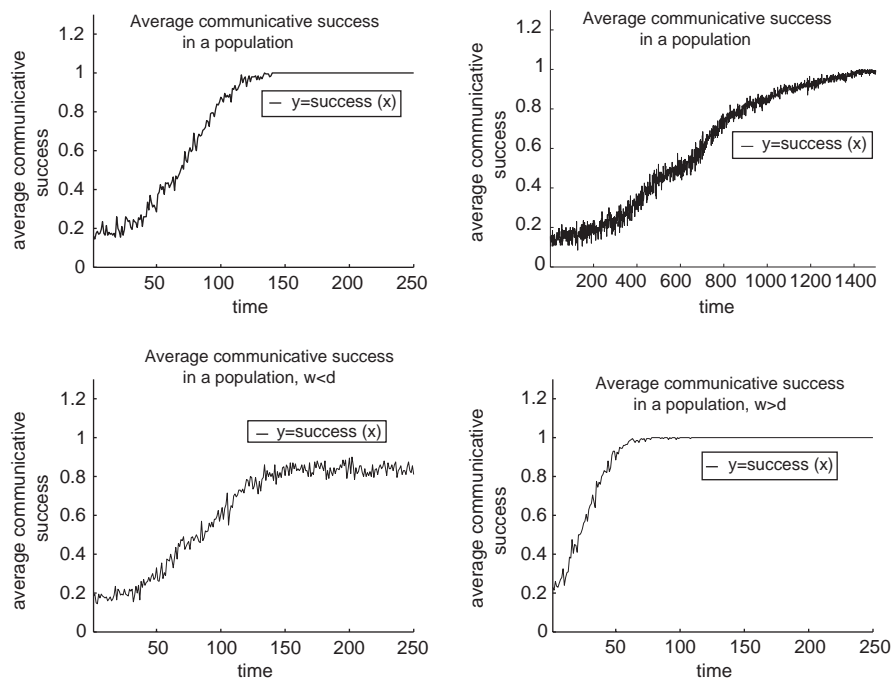


Fig. 2. *Communicative coherence between 10 individuals*: Plotting the average communicative coherence of a naming game with 10 individuals averaged over 10 runs. The coherence measure was recorded every 20 iterations. We performed 5000 iterations in total for each experiment except the one in the top-right panel where 30 000 were performed with 30 individuals. The setup of the game consists of a set W of seven words and a set D of seven meanings. Each plot shows the communicative coherence which measures how well two individuals understand each other. On each plot, the average of 10 runs is shown. For all plots, we assumed $\delta = 0.1$. The two figures in the lower-left and lower-right show the performance of the naming game when the number of meanings differs from the number of words. In the figure in the lower-left, there are more meanings than words. This setup allows for one word to refer to different meanings (coincidental and polysemic homonyms). In the figure in the lower-right, there are more words than meanings. In this setup we allow different words with similar meaning (synonyms).

4. Learning dynamics of the naming game

Given the naming game algorithm, one would like to express it in terms of some well-founded mathematical framework so it could give rise to theories on the dynamic behavior of such a system. The current discussion is limited to the explanation of the model. Theories on the kinds of absorbing states, whether they can be reached and whether they are stable are left for future work. In this section, we will use some concepts and models from learning automata theory (Narendra and Thathachar, 1989) to set up a mathematical model of the naming game.

To reach the final model some steps need to be taken: in Section 4.1, two normalized matrices will be introduced to facilitate the analysis. The normalization is introduced for reasons of mathematical simplicity, understandability and to make the relation to learning automata theory explicit. Moreover, normalization makes it possible to compare different lexicons when there is no upper limit on the strength value, i.e. when $v_{kl}^i \in [0, \infty]$. How the matrix is normalized, depends on the role of the individual in the game. Furthermore, to evaluate the success of each normalized matrix, a measurement is required. In Section 4.2 this issue will be discussed briefly. Afterward, the updating rules from

Section 3.2 will be expressed in terms of linear reinforcement schemes (Narendra and Thathachar, 1989). This immediately introduces a generalization of the naming game since different rules might be used. Using this formalization of the naming game, a number of simple experiments with different learning schemes are performed in Section 4.4. Finally, in Section 4.5 an example is provided on the use of a different transmission structure.

4.1. Normalization of the lexical matrix

In order to analyse the learning dynamics described in the previous discussion, the lexical matrix of an individual is normalised. At this point, the importance of the role in a cultural transmission process is made explicit. Depending on the role different normalizations can be performed. Remember that an individual can be either speaker or hearer in the naming game. When the individual is a speaker, L^i will be normalized according to the rows. This normalization transforms each row into a normalized vector where the sum of the value in a particular row equals one.

The values in each row might be interpreted as probabilities, yet this confuses things since, in the naming game, the speaker does not select some word

according to a probability distribution. One might change the naming game and select a word in a row according to the distribution defined by the normalized values. Yet, in initial experiments, which will not be discussed here, we have observed that this would result in bad behavior. Thus, the only meaning assigned to each value in a normalized row (or column) is that of strength. Moreover, the semantics of this strength is determined by its relation to the other strength values in the particular row or column, i.e. its order. Further experiments will be conducted to examine the difference between these interpretations in a more general framework.

When the individual plays the role of hearer, \mathbf{L}^i will be normalized according to the columns. This normalization transforms each column into a normalized vector associating one word with different meanings. The same reasoning as with the columns can be applied here. In summary, since there are two roles, two normalization matrices can be defined: \mathbf{P}^i and \mathbf{Q}^i . These matrices are derived from \mathbf{L}^i by normalizing, respectively, the columns or rows (assume that $k \in \{1, \dots, d\}$ and $l \in \{1, \dots, w\}$)

$$p_{kl}^i = \frac{v_{kl}^i}{\sum_{l=1}^w v_{kl}^i} \quad \text{or} \quad q_{kl}^i = \frac{v_{kl}^i}{\sum_{k=1}^d v_{kl}^i}. \quad (1)$$

Producing the normalized matrices:

$$\mathbf{P}^i = \begin{pmatrix} p_{11}^i & p_{12}^i & \dots & p_{1w}^i \\ p_{21}^i & p_{22}^i & \dots & p_{2w}^i \\ \dots & \dots & \dots & \dots \\ p_{d1}^i & p_{d2}^i & \dots & p_{dw}^i \end{pmatrix} \quad \text{or} \quad \mathbf{Q}^i = \begin{pmatrix} q_{11}^i & q_{12}^i & \dots & q_{1w}^i \\ q_{21}^i & q_{22}^i & \dots & q_{2w}^i \\ \dots & \dots & \dots & \dots \\ q_{d1}^i & q_{d2}^i & \dots & q_{dw}^i \end{pmatrix},$$

where

$$0 < p_{kl}^i < 1 \quad \text{and} \quad \sum_{l=1}^w p_{kl}^i = 1$$

for normalized matrix \mathbf{P}^i or where

$$0 < q_{kl}^i < 1 \quad \text{and} \quad \sum_{k=1}^d q_{kl}^i = 1$$

for normalized matrix \mathbf{Q}^i .

The normalized matrices can be restored to their original format by performing the inverse operation of Eq. (1), i.e. $((1)^{-1})$. Note that the previous sum of the row (or column) or some value has to be remembered to perform that operation otherwise different inverse calculations are possible. Given the transformation rule specified by Eq. (1) and its inverse, a general scheme for the stochastic language learning process can be formulated as can be seen in Fig. 3.

In Fig. 3 a lexical matrix \mathbf{L}^i is assumed at time step $(t-1)$. This matrix is transformed into a normalized matrix according to the role of the individual. This role is randomly assigned. At time step $(t-1)$, the role of the individual is hearer and hence the matrix \mathbf{Q}^i is derived using Eq. (1). Next, $\mathbf{Q}^i(t+1)$ is determined through one (or more) iteration(s) of the naming game which is defined by the updating scheme \mathcal{U} . This matrix is in turn transformed back into a new lexical matrix at the next time step. The transformation laws T , briefly mentioned at the end of Section 2, hence consist of three steps. As in Nowak et al. (1999), the laws of transformation T , assuming no noise, will lead toward absorbing states. This can be observed in the examples discussed later. Before discussing the central transformation rule, i.e. the updating scheme \mathcal{U} , we first provide a short motivation for our coherence measure that reflects the progress of the stochastic learning process.

4.2. Measuring the language success

In order to assess the success of communication between two individuals, a measure is required. In general, two approaches can be taken. First, one could use the actual communicative success between individuals as it occurs within a certain time interval. The success is determined by the procedure that selects two individuals to communicate and about what they will communicate. Such a measure, called language




Fig. 3. Stochastic learning process defined for naming game: \mathbf{L}^i refers to the lexical matrix and \mathbf{Q}^i (\mathbf{P}^i) refers to the normalized hearer (speaker) matrix. Eq. (1) and its inverse are used to determine the normalized matrices and \mathcal{U} is the updating scheme that changes the normalized values in a particular row or column.

coherence, was used in the experiments shown in Figs. 1 and 2. Second, one could introduce a measure which determines the expected communicative success for the next game. This value would be completely determined by the current state of the individuals.¹ In the current work, language coherence will be used to observe the dynamics of the language game.

4.3. Acquiring the vocabulary

The naming game describes an update scheme that stipulates how the lexical matrix is updated when the game succeeds or fails. Given the stochastic process defined in Fig. 3, this updating scheme \mathcal{U} (defined by the naming game) can also be formalized. The rules which are constructed here differ in correlation to the role of the individual and whether or not the game was a success.

Assume two individuals with state $\phi^i = (\mathbf{D}, \mathbf{L}^i)$ and $\phi^j = (\mathbf{D}, \mathbf{L}^j)$. The number of meanings and words are fixed, finite and given, i.e. the set of words W contains w elements and the set of meanings D contains d elements. Moreover, each individual has an update (or reinforcement) mechanism \mathcal{U} that updates the probabilities of either the speaker matrix \mathbf{P}^i (or \mathbf{P}^j) or the hearer matrix \mathbf{Q}^i (or \mathbf{Q}^j) depending on the roles of both individuals.

As specified in the naming game in Section 3, if individual ϕ^i is the speaker it selects randomly a meaning and selects the associated word relative to the assigned strengths. This word is transmitted to the hearer ϕ^j who determines which meanings are associated with the received word. From this collection of meanings the one with the highest strength is chosen. If both meanings from the individuals are the same the game succeeds, otherwise it fails. Success or failure at iteration t is represented by the binary variable $\beta(t)$. Thus $\beta(t)$ is considered to be the response of the environment when a particular word is uttered. In principle, this response does not have to be binary, yet the current work will be restricted to binary values for $\beta(t)$: success after t iterations is $\beta(t) = 0$ and failure is $\beta(t) = 1$.

Upon success, the association between that particular meaning and the word is reinforced and the other associations are decreased. When the game fails the opposite occurs. Hence a reward is assigned for success and a penalty for failure. Here, we assume that upon success the reinforcement value δ_s is used. Likewise, when the game fails, the value δ_f is used.

Given this information, a general reinforcement scheme can be defined for the speaker ϕ^i . When the role is speaker, a row in \mathbf{P}^i is updated. This means that the relation between a particular meaning and a set of

words is changed. The following equations provide the general algorithm which consists of two update rules for the speaker; one to update the probabilities of the selected word l (Eq. (2)) and one to update all the other words $h \neq l$ (Eq. (3)):

$$p_{kl}^i(t+1) = p_{kl}^i(t) + \delta_s(1 - \beta(t))(1 - p_{kl}^i(t)) - \delta_f\beta(t)p_{kl}^i(t) \quad (2)$$

when the word used to communicate was w_l^i for meaning d_k^i . Eq. (2) specifies that the normalized value p of the selected word–meaning association k, l is increased with $\delta_s(1 - p_{kl}^i(t))$ when the association was successful and is decreased with $\delta_f p_{kl}^i(t)$ when the association fails.

$$p_{kh}^i(t+1) = p_{kh}^i(t) - \delta_s(1 - \beta(t))p_{kh}^i(t) + \delta_f\beta(t)((w-1)^{-1} - p_{kh}^i(t)) \quad (3)$$

for all the other words w_h^i , with $l \neq h \in W$, who might also be associated with the meaning d_k^i , yet who are not used. Eq. (3) specifies that the normalized value p of all the other word–meaning associations k, h are decreased with $\delta_s p_{kh}^i(t)$ when the association k, l was successful and are increased with $\delta_f((w-1)^{-1} - p_{kh}^i(t))$ when the association k, l fails. A careful reader can observe that the value of $\delta_f((w-1)^{-1} - p_{kh}^i(t))$ can become negative. This is indeed the case when $(w-1)^{-1} < p_{kh}^i(t)$. The reasoning behind this effect is the following. The configuration where all p -values in a particular row have the same value is interpreted as the situation where, given a certain meaning, one does not know which word should be selected. Hence every word has equal chance. Only when the environment provides useful positive feedbacks can certain associations be preferred over others. Hence upon negative feedback, the state of the row is moved closer to a state where one does not know which word to select, i.e. all normalized strengths have the same value $(w-1)^{-1}$. This move toward this state of total uninformedness allows the system to explore alternative associations.

The same definitions can be specified for the hearer ϕ^j . In this case a column in the normalized matrix \mathbf{Q}^j is updated. This means that an association between a word and a collection of meanings is altered: the selected meaning is reinforced while the other meanings are degraded:

$$q_{kl}^j(t+1) = q_{kl}^j(t) + \delta_s(1 - \beta(t))(1 - q_{kl}^j(t)) - \delta_f\beta(t)q_{kl}^j(t) \quad (4)$$

when the word received during communication was w_l^j for meaning d_k^j .

$$q_{rl}^j(t+1) = q_{rl}^j(t) - \delta_s(1 - \beta(t))q_{rl}^j(t) + \delta_f\beta(t)((d-1)^{-1} - q_{rl}^j(t)) \quad (5)$$

for all the other meanings d_r^j , with $k \neq r \in D$, who might also be associated with the word w_l^j , yet who are not used. The interpretation of Eq. (4) and (5) is equivalent to that of the speakers update scheme.

¹An example of such a measure was introduced in Nowak et al. (1999): $F(\phi^i, \phi^j) = \frac{1}{2} \sum_{k \in D} \sum_{l \in W} p_{kl}^i q_{kl}^j + p_{kl}^j q_{kl}^i$. This measure is meaningless here since it assumes a different interpretation of the normalized matrices \mathbf{P} and \mathbf{Q} , i.e. probability distributions.

Eqs. (2)–(5) describe how the vocabulary is learned by each individual from both perspectives of hearer and speaker. The reinforcement scheme designed here for the naming game has been discussed and analysed before in the context of learning automata and more specifically variable structure stochastic automata (Bush and Mosteller, 1951, 1955; Narendra and Thathachar, 1989). This type of learning automata form a class of more general stochastic systems wherein state transition and action probabilities are updated by a reinforcement scheme at every time step. The specific scheme captured by the equations here is called a *linear reinforcement scheme*. Variations of this scheme are constructed by constraining the reinforcement values δ_s and δ_f in particular ways, e.g. when $\delta_s = \delta_f = \delta$ and $0 < \delta < 1$, the equations become a *linear reward–penalty scheme*. Another variation assumes for instance that $\delta_s > \delta_f > 0$. This update scheme is called a *linear reward- ϵ -penalty scheme* and this will be used in the following sections.

4.4. Simple examples and different learning schemes

Given the normalized matrices, the linear schemes and the success function discussed in Section 4.2, a collection of experiments can be discussed. The mathematical framework should produce the same effects as those observed in Fig. 1. Again, we assume two individuals

with similar specification as in the naming game examples. In Fig. 4 the experiments were conducted using the linear reward–penalty scheme. In Fig. 5 the same experiments were repeated using the linear reward- ϵ -penalty scheme. The exact settings of the reward δ_s and penalty δ_f can be found below the figures. When comparing the results from each scheme to the results obtained in Fig. 1, one can observe that both the reward–penalty and the reward- ϵ -penalty scheme have a similar qualitative behavior as the original naming-game.

In all experiments, both δ_s and δ_f determine the rate at which a lexicon is learned. This can easily be observed in the different figures when comparing the plots in the top row. For a high δ_s value, the shared lexicon is achieved faster than for low values. Moreover, in the homonymy scenarios in the different figures (lower-left plot) it can be observed that the learning scheme maintains homonyms in the vocabulary of the individual. Yet, in the context where synonyms are possible (lower-right), the learning dynamics will converge toward a binary matrix where some words will not be used. Hence, one can see that linear schemes from learning automata theory have a qualitative behavior which is equivalent to the standard naming game.

In Fig. 6, the same population experiment as in Fig. 2 was repeated using a linear reward- ϵ -penalty update scheme. As the number of individuals increases, so will

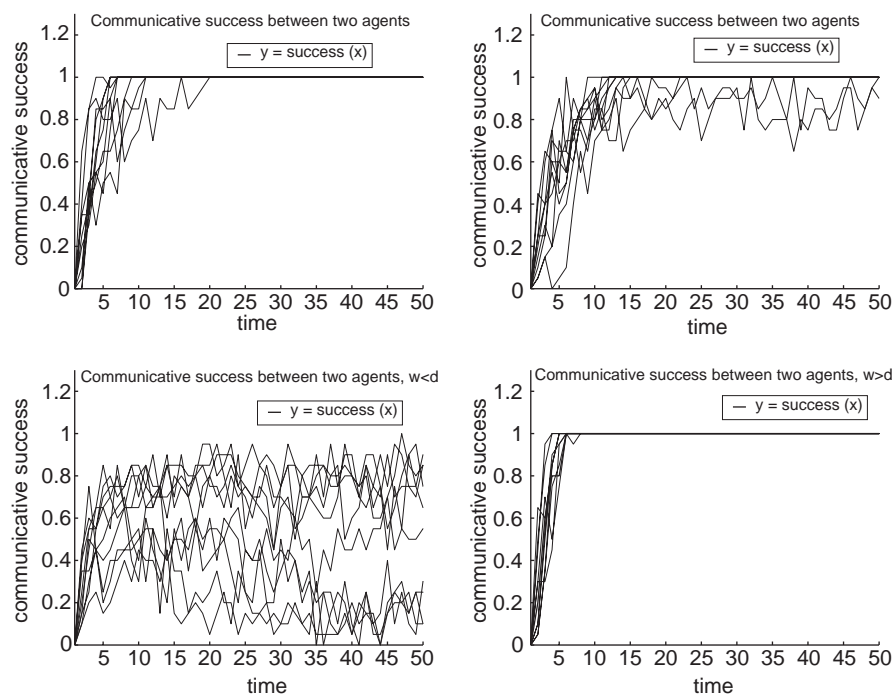


Fig. 4. *Communicative coherence using linear reward–penalty reinforcement scheme (two individuals)*: Plotting the communicative coherence of the formalized naming game with the linear reward–penalty reinforcement scheme. Again, the game consists of a set W of seven words and a set D of seven meanings ($w = d$). The learning occurs between two individuals whose role was determined randomly. The figure in the top-left corner shows 10 runs of the formal framework with $\delta_s = \delta_f = 0.1$. The same experiment with $\delta_s = \delta_f = 0.05$ is shown in the top-right figure. In both cases the game results in a binary matrix associating a particular meaning with a particular word. The only difference is the number of iterations required to reach the solution. The two figures in the lower-left and lower-right corner are homonymy ($w < d$) and synonymy ($w > d$) scenarios.

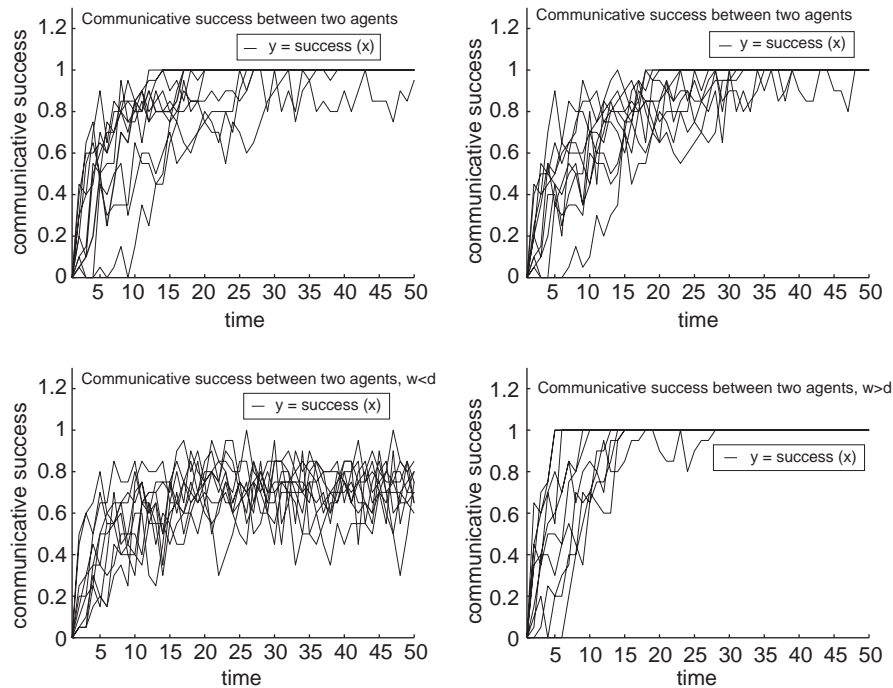


Fig. 5. *Communicative coherence using linear reward- ϵ -penalty reinforcement scheme (2 individuals)*: Plotting the communicative coherence of the formalized naming game with the linear reward- ϵ -penalty reinforcement scheme. Again, we use a set W of seven words and a set D of seven meanings ($w = d$). The figure in the top-left corner shows 10 runs of the formal framework with $\delta_s = 0.1$ and $\delta_f = 0.01$. The same experiment with $\delta_s = 0.05$ is shown in the top-right figure. In both cases the game results in a binary matrix associating a particular meaning with a particular word. The only difference is the number of iterations required to reach the solution. The two figures in the lower-left and lower-right corner are homonymy ($w < d$) and synonymy ($w > d$) scenarios.

the time to converge to a shared language. Moreover, the learning rate will play a crucial role on how easily this shared lexicon is attained (Narendra and Thathachar, 1989). Hence, a good balance between reward and penalty is necessary.

Remember from Section 4.1 that the actual learning dynamics of the naming game corresponds to a stochastic learning process whose general scheme is illustrated in Fig. 3. This general scheme consists of three steps. Firstly, the lexical matrix is transformed into a normalized matrix whose normalization depends on role of the individual, i.e. either speaker or hearer. Secondly, the update scheme \mathcal{U} changes the normalized matrix according to the outcome of the naming game. Finally, the normalized matrix is transformed back into the lexical matrix. In the previous sections, the updating scheme \mathcal{U} was formalized and some examples were provided to visualize the relation with the naming game experiments. The next logical step is to specify how all things fit together in the cultural transmission scheme T that connects two lexical matrices (see Fig. 3).

4.5. Transmission structure

In Section 2.3, we briefly highlighted that the goal of a mathematical framework is to understand the cultural transmission system T which determines how the strengths

in the individual's lexicon change over time. From Section 4.3, we have learned that the effect of social learning on the lexicon is determined by the role of the individual in the interaction and whether the interaction was successful or not. So far, it was assumed the individuals interacted through peer-to-peer interaction. This form of interaction assumes a horizontal transmission structure in the cultural inheritance system. As a consequence of this transmission structure, both speaker and hearer update their lexicons.

The current model does not have to be limited to horizontal inheritance structures. Other cultural inheritance system like oblique and vertical transmission structures can easily be incorporated. For instance, assume that a certain amount of individuals of the population will be assigned the role of teacher and the remaining set will be the students. It will be assumed that the teachers have acquired some coherence among themselves in their language. Since we assume that the teachers have coherence in their communication, they will not update their lexical matrices. Yet, alternative modes could add the fact that even the teachers still need to improve their communication skills among each other. The students will acquire the language from the teachers using the social learning scheme defined by the naming game. In other words, the students will always be hearer and use the updating rules associated the normalized matrix Q^j .

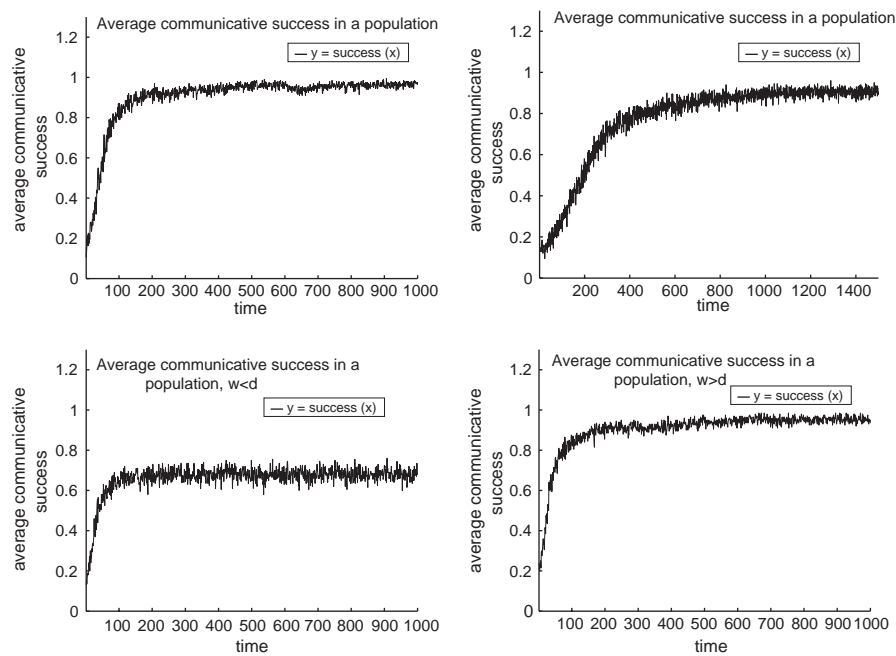


Fig. 6. *Communicative coherence using linear reward- ϵ -penalty reinforcement scheme (10 individuals)*: Plotting the average communicative coherence of a naming game with 10 individuals averaged over 10 runs. In this example, the reward- ϵ -penalty scheme was used. The success measure was recorded every 20 iterations, hence we performed 20 000 iterations in total except for the plot in the upper-right corner where the experiment was performed using 30 individuals and 30 000 iterations. The setup of the game consists of a set W of seven words and a set D of seven meanings. Each plot shows the communicative coherence which measures how well two individuals understand each other. On each plot, 10 runs are shown. For all plots, we assumed $\delta_s = 0.2$ and $\delta_f = 0.1$. The two figures in the lower-left and lower-right are again homonymy ($w < d$) and synonymy ($w > d$) scenarios.

In Fig. 7 an experiment using this simplified oblique transmission scheme is shown. Each plot in the figure shows the communicative coherence between the teacher and the students where the teachers are speakers and the students are hearers. Moreover, the communicative coherence among the students themselves is also measured. The general idea is that it is not enough to understand the teachers, each student should also understand his or her peers. In the top row of Fig. 7, one can see that communicative coherence is reached in all circumstances. Yet when the teachers have a vocabulary which contains either homonyms or synonyms the dynamics change. In the bottom-left, the teachers possess two words which can be used for different meanings. It can be observed that although students are able to understand their teacher, they are not capable of understanding each other. Moreover, they initially could until iteration 1200 where the language coherence starts to decrease. When examining the lexical matrix, this reason is that from there on, each student starts to use one of the homonyms exclusively for one of their meanings. Yet all possible configurations remain present in the entire population. Hence, differences in individual frequency can be observed where individuals of type A use one word for a particular meaning and individuals of type B use another word for that meaning. As a result, these two groups do not understand each other anymore.

A synonym scenario is visualized in the bottom-right plot of Fig. 7. Again the teachers are configured in such a way that they each possess the same two synonyms. As one can observe, the social learning process converges to a situation where the students understand their teachers completely. Yet, the communication among themselves is not optimal. This results from the fact that certain meanings cannot be learned. These results show some remarkable issues which require further investigations on themselves.

The experiments from Fig. 7 are just one way of examining oblique transmission. Again, extensions to this basic framework can be introduced by relaxing previous assumptions. Moreover, simulations can be performed using combinations of both horizontal and oblique transmission. Such a framework combining would allow teachers (or students) to communicate among themselves. In summary, each individual can have different roles depending on the interaction taking place. One can be parent or child, teacher or student, and peers all at the same time. All these roles will influence the actual vocabulary produced by the interactions.

Vertical transmission between genetic generations can also be introduced. In this context, a genotype could be introduced that specifies the genetically heritable properties of the individuals' language capabilities. The offspring can then inherit these capabilities in their original or mutated form. This would produce a model

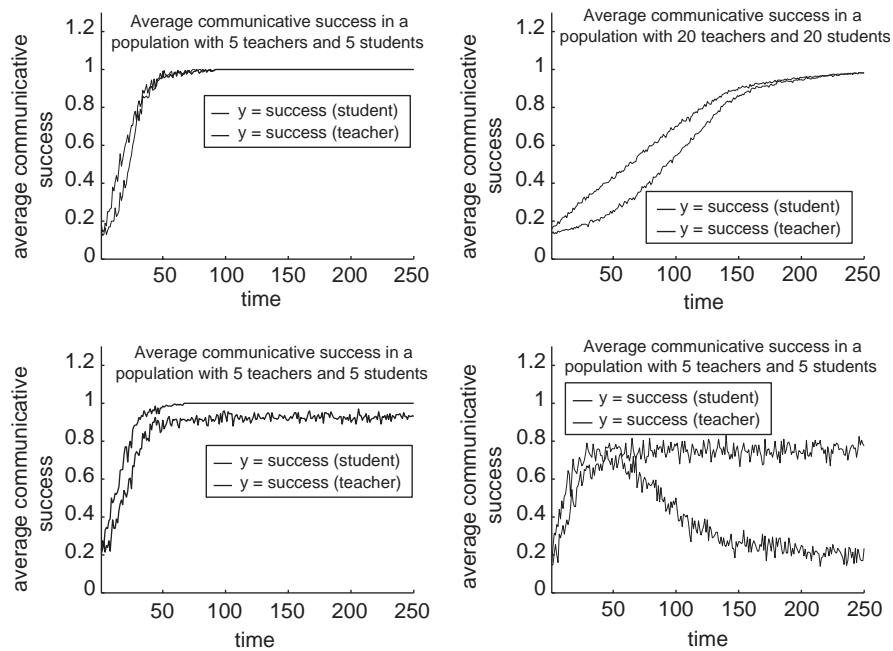


Fig. 7. *Communicative coherence using oblique transmission*: Plotting the average communicative coherence of a student population playing the naming game with a population of teachers averaged over 10 runs. In this example, the reward- ϵ -penalty scheme was used to learn the language using $\delta_s = 0.2$ and $\delta_f = 0.1$. Both the teacher and student populations consists of five individuals each. Yet, in the top-right figure, 20 students were used instead of five. Moreover, the game was played over 30 000 generations. Initially, all the parents have the same vocabulary, i.e. a randomly generated binary matrix. For the experiments in the top row, the matrices consist of seven words and seven meanings. The two figures in the lower-left and lower-right are again homonymy ($w < d$) and synonymy ($w > d$) scenarios. In all plots, the top line shows the language coherence between teachers and students and the lower line visualizes the language coherence among the students.

where both genetic and cultural transmission mechanisms can be studied as in for instance (Smith, 2004). For now, we limit ourselves to horizontal and oblique transmission schemes.

5. The evolutionary dynamics of language

In the previous section, an interpretation of the naming game in terms of learning automata theory was provided. Here, we discuss the relation between this learning theory and evolutionary dynamics. By drawing this relation, the evolutionary interpretation of cultural evolution of language is made explicit.

There has always existed an intuitive relation between learning and biological evolution. The origin of this relation stems from the fact that in both cases evolution is generally described as the gradual movement from mediocre or even bad solutions to better ones.

In the naming game, at the individual level, each individual constructs associations between meanings and words and vice versa. Which associations are considered highly relevant by the individual depend on the interaction experience of the individual with the other individuals it encounters. The learning process slowly changes the internal probabilities that associate words and meanings until some shared solution is found. From a biological perspective, the

lexicon can be considered to be a collection of ‘populations of word–meaning associations’ in the individual’s mind. These populations are formed each time step according to the role the individual is playing in the interaction. Hence a constant migration into subpopulations and regrouping in the complete population occurs. Learning corresponds to small incremental changes inside these subpopulations. It is this idea of small adaptive changes that corresponds to people’s perception of simple biological evolutionary processes. We referred to the processes at this level as an instance of social or cultural learning (Boyd and Richerson, 1985; Tomasello et al., 1993).

The analogy between learning and biological evolution can be made explicit through models of evolutionary game theory. The idea of defining this relation is not new. Børgers and Sarin (1997) specify a relation between the Cross learning model and replicator dynamics. They show that in the continuous time limit the learning model converges to the asymmetric version of the replicator dynamics from evolutionary game theory. By specifying this relation, the authors provide a non-biological interpretation of evolutionary game theory. It is this interpretation which explicitly links learning and evolutionary dynamics.

This result is not limited to the Cross learning model. Due to the relation of this model with the general theory of Learning automata (Tuyls et al., 2002) and with the Q-

learning model from Reinforcement Learning (Tuyls et al., 2003), these results are generally applicable. In the context of the models discussed here, the Cross learning model is a model similar to a learning process where no penalties are given when wrong associations are made.

Given these previous models of the relation between learning and evolution, the same thing can be done for the social learning scheme defined for the naming game. We will show that social learning scheme corresponds to a linear selection–mutation model as defined in the context of evolutionary game theory and evolutionary dynamics (Hofbauer and Sigmund, 1988).

5.1. Specification of the evolutionary model

Eqs. (2)–(5) can be used to derive the replicator dynamics for the speaker. A similar set of equations is required for the hearer and can in turn be used to derive the replicator dynamics of the hearer. To perform the task for the speaker, the expected change ($E[\Delta p_{kl}^i | p_k^i(t), q_l^j(t)]$) of a particular meaning–word association needs to be calculated. This expected change specifies how the strength of a word–meaning association changes from one generation to the next given a particular state description for both the speaker strengths and hearer strengths. Using the formula of the conditional expected value, the expected change in strength for the normalized value p of player i when choosing the association between meaning k and word l is determined by a combination of Eqs. (2), (3) and the conditional probability function that specifies that at time t the speaker i selects meaning k and the hearer j selects word l :

$$E[\Delta p_{kl}^i | p_k^i(t), q_l^j(t)] = p_{kl}^i(t+1) \times \sum_{h=1}^w \sum_{r=1}^d p_{kh}^i(t) q_{rl}^j(t).$$

After substitution of certain parts with Eqs. (2) and (3) and a number of derivation steps the following equation is obtained:

$$\begin{aligned} E[\Delta p_{kl}^i | p_k^i(t), q_l^j(t)] &= p_{kl}^i(t) \left(\sum_{r=1}^d \delta_s(1 - \beta(t)) q_{rl}^j(t) \right. \\ &\quad \left. - \sum_{h=1}^w p_{kh}^i(t) \sum_{r=1}^d \delta_s(1 - \beta(t)) q_{rl}^j(t) \right) \\ &\quad - \delta_f \beta(t) p_{kl}^i(t) \sum_{h=1}^w p_{kh}^i(t) \sum_{r=1}^d q_{rl}^j(t) \\ &\quad + \frac{\delta_f \beta(t)}{w-1} \sum_{h' \neq l}^w p_{kh'}^i(t) \sum_{r=1}^d q_{rl}^j(t), \end{aligned}$$

where k refers to the rows and l to the columns and w is the amount of words and d the amount of meanings. In order to provide a mapping of this equation to the replicator equation a number of substitutions need to be

made. These substitutions are merely syntactic sugar to clarify the relation with the replicator equation.

- $\delta_s(1 - \beta(t)) = \mathbf{A}_{ij}$: The update amount in case of success is considered to be the payoff that is received by the words and meanings upon successful interaction. The collection of all the \mathbf{A}_{ij} values defines the payoff matrix \mathbf{A} :

$$\mathbf{A} = \begin{pmatrix} 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ \delta_s(1 - \beta(t)) & \delta_s(1 - \beta(t)) & \dots & \delta_s(1 - \beta(t)) \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 0 \end{pmatrix}.$$

The resulting matrix \mathbf{A} is $w \times d$ -dimensional and we will refer to this matrix as the payoff matrix. Each entry in this matrix represents the payoff that certain word–meaning combination gets. Note that only the row is used which corresponds to the meaning selected by the speaker. All other entries are in this case irrelevant since they refer to other meanings which were not selected by the speaker. Furthermore, each value $\delta_s(1 - \beta(t))$ is only assigned when the communication is a success. This makes the payoff matrix an unusual one since all entries will become zero when the game fails. Yet it is required due to the distinction in how the update is performed between successful and unsuccessful interactions in the naming game.

- $\delta_f \beta(t) = \mu_{ij}$: The update amount in case of failure is considered to be the mutation rate from one type to another. The collection of all these rates corresponds to the mutation matrix \mathbf{U} :

$$\mathbf{U} = \begin{pmatrix} \delta_f \beta(t) & \delta_f \beta(t) & \dots & \delta_f \beta(t) \\ \delta_f \beta(t) & \delta_f \beta(t) & \dots & \delta_f \beta(t) \\ \dots & \dots & \dots & \dots \\ \delta_f \beta(t) & \delta_f \beta(t) & \dots & \delta_f \beta(t) \end{pmatrix}.$$

As with the payoff matrix, the matrix \mathbf{U} becomes zero when the communication was successful.

Given these syntactic substitutions, the following equation is now obtained:

$$\begin{aligned} E[\Delta p_{kl}^i | p_k^i(t), q_l^j(t)] &= p_{kl}^i(t) \left(\sum_{r=1}^d \mathbf{A}_{kr} q_{rl}^j(t) - \sum_{h=1}^w p_{kh}^i(t) \sum_{r=1}^d \mathbf{A}_{hr} q_{rl}^j(t) \right) \\ &\quad - \mu_{kl} p_{kl}^i(t) \sum_{h=1}^w p_{kh}^i(t) \sum_{r=1}^d q_{rl}^j(t) \\ &\quad + \frac{\mu_{kl}}{w-1} \sum_{h' \neq l}^w p_{kh'}^i(t) \sum_{r=1}^d q_{rl}^j(t). \end{aligned}$$

The different sums in this equation can be further simplified by writing them as matrix multiplication.

$$E[\Delta p_{kl}^i | p_k^i(t), q_l^j(t)] = p_{kl}^i(t)(e_k \mathbf{A} q_l^j(t) - p_k^i(t) \mathbf{A} q_l^j(t)) \quad (6)$$

$$- \mu_{kl} p_{kl}^i(t) \quad (7)$$

$$+ \frac{\mu_{kl}}{w-1} (1 - p_{kl}^i(t)). \quad (8)$$

This rewriting produces a replicator equation that describes the dynamics of meaning–word associations within the speakers mind. The symbol e_k refers to a row vector that contains all zeros except for a one at the k th position. The equation consists clearly of three parts:

Part (6) corresponds to a standard replicator equation as discussed in the context of evolutionary game theory. It describes the effects of selection on the frequency of a particular meaning–word combination relative to the average performance in that population. Note that the population here corresponds to a particular row from the lexicon when the individual's role is speaker.

Parts (7) and (8) correspond to the effect that mutation has on the frequency of a particular meaning–word association. Concretely Part (7) of the equation expresses the fact that meaning–word combinations can decrease in number and hence disappear. Part (8) of the equation expresses an increase in the frequency due to mutations of other meaning–word associations to this particular, one i.e. the association between meaning k and word l .

It can be observed further that the mutation part will only be active when the interaction, or better the communication, failed. Hence selection is active when the interaction succeeded. This separation is very logical since good associations should be exploited and bad associations require further exploration.

A similar equation can be derived for the hearer. Instead of using a payoff matrix \mathbf{A} , we use the payoff matrix \mathbf{B} .

$$\mathbf{B} = \begin{pmatrix} 0 & 0 & \dots & \delta_s(1 - \beta(t)) & \dots & 0 \\ 0 & 0 & \dots & \delta_s(1 - \beta(t)) & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \delta_s(1 - \beta(t)) & \dots & 0 \end{pmatrix}$$

in this matrix a certain column is filled with the value $\delta_s(1 - \beta(t))$. A similar derivation then produces:

$$E[\Delta q_{kl}^j | p_k^i(t), q_l^j(t)] = q_{kl}^j(t)(p_k^i(t) \mathbf{B} e_l - p_k^i(t) \mathbf{B} q_l^j(t)) \quad (9)$$

$$- \mu_{kl} q_{kl}^j(t) \quad (10)$$

$$+ \frac{\mu_{kl}}{d-1} (1 - q_{kl}^j(t)). \quad (11)$$

Here, the symbol e_l refers to a column vector that contains all zeroes except for a one at the l th position. This derivation describes for each probability q_{kl}^j of a certain column in the hearer matrix \mathbf{Q}^j how these values change over time. Again, one can observe a clear

separation into selection to exploit the use of good associations and mutation to explore for better associations between meanings and words.

Thus in general, both discrete replicator equations show how the frequency of each meaning–word association in the mind of each individual gradually changes over time. As can be observed, if the success of a specific association between a word and a meaning is higher than the average success of all associations between a particular meaning and the corresponding words of the speaker (or between a particular word and all applicable meanings in the hearer), then the frequency of this association will increase. Otherwise it will decrease. Hence, acquiring the vocabulary through a usage-based interaction as described by the naming game is equivalent to a multipopulation game theoretical model (Weibull, 1996) where, depending on the role, each individual's internal state corresponds to a set of populations of meaning–word associations. At each interaction, one of these subpopulations for each individual is used and will be updated according to the dynamics defined by both replicator equations.

Now what happens in the naming game? We know that the feedback from the environment $\beta(t)$ can only have the values zero in case of success and one in case of failure. These feedback values determine when selection or mutation are applied. On one hand, if the communication was successful ($\beta(t) = 0$), the data currently present in the lexicon is exploited through the selection equation (defined by Eqs. (6) and (9)). On the other hand, when the communication was unsuccessful ($\beta(t) = 1$), alternative meaning–word associations are introduced through some mechanism analogous to mutation (defined by the other equations). Furthermore, the specific values of the reward δ_s and the penalty δ_f determine the kind of linear reinforcement scheme. These values will have an effect on the qualitative behavior of the replicator equation. Note that it is not necessary for both individuals to have the same values for δ_s and δ_f .

When $\delta_f = 0$, we see that both selection–mutation models are reduced to a selection model. Since no variation is introduced when the interaction fails, the selection dynamics only act on the given data and can converge to local optima. Hence, whether or not it will find the lexicon which is shared with other individuals depends on the initial state of the lexicon. In other words, there is only exploitation of the known lexicon data when the interaction was successful and no exploration towards lexicons which are shared between all the individuals.

When $\delta_f > 0$ the mutation parts of the model become active but not at the same time as selection. As a consequence a better exploration of the space of lexicons becomes possible and it will be easier to find a shared lexicon. In Fig. 8, one can observe the evolutionary

dynamics of a linear reward- ε -penalty scheme. Other variations were tried but not shown here.

In summary, Both selection–mutation equations express the biological interpretation of the learning dynamics in the naming game. The evolutionary dynamics defined by the selection–mutation process, describe explicitly that the cultural evolution of language is a consequence of selecting those meaning–word combinations which are used most often and that over time, words can change their meaning.

In respect to the brief discussion on a vertical transmission structure in Section 4.5, one could assume that \mathbf{A} and \mathbf{U} are genetically determined for each individual and that offspring inherit this kind of information. Moreover, this would also mean that the values used in this matrices would not have to be the same. Then the genetic evolutionary system could produce both matrices which perform best given in the context of a particular social learning system.

5.2. Population-level consequences

As argued extensively by Boyd and Richerson (1985), since social learning causes phenotypic traits to be exchanged between individuals, social learning has population-level consequences. Thus, since language is

acquired, by individuals through actual communication, language is a population-level phenomenon.

Here, the population consist of individuals which are completely defined by their lexical matrix \mathbf{L}^i . Given the laws of transformation T , these matrices change and these changes define a trajectory in the state space defined by the matrices. Hence a population of individuals is described as a distribution over the possible lexical matrices. Moreover, the dynamics visualized in Fig. 3, describe how this distribution changes over time. In other words, how the average lexical configuration and the corresponding variance change over time.

As argued by Boyd and Richerson (1985) in their model, learning has two opposing effects on the variance of distribution of cultural variants in the population: first, it causes the individuals to change their cultural variants toward a common goal. Consequently, the variance in the population decreases. Second, errors made during learning increase the variance. Since we do not discuss the extensions here where the exchange of lexical information is noisy, this fact cannot be observed here. We leave this for later exploration. Nevertheless, the first effect is observable in the experiments. The cultural variants correspond to the lexical matrices of the different population members. The naming game causes the different individuals to modify their lexical

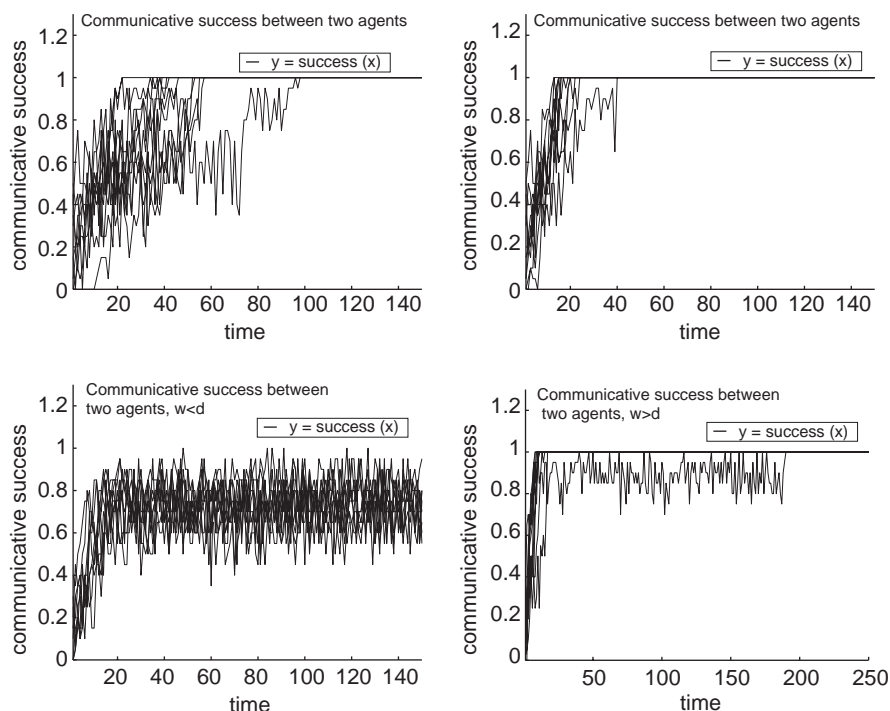


Fig. 8. *Communicative coherence of the replicator-dynamics (two individuals)*: Again, we use a set W of seven words and a set D of seven meanings ($w = d$). The figure in the top-left corner shows 10 runs of the formal framework with $\delta_s = 0.1$ and $\delta_f = 0.01$. The same experiment with $\delta_s = 0.2$ and $\delta_f = 0.05$ is shown in all other plots. In both figures at the top, the game converges to a solution with binary matrices. There is only a difference in the convergence rate. The two figures in the lower-left and lower-right corner are homonymy ($w < d$) and synonymy ($w > d$) scenarios. Again, homonymy remains while synonymy disappears.

variants toward a shared lexicon. Hence the variation in lexical information decreases. Moreover, it will disappear since the cultural inheritance process here is avoid of noise.

6. Conclusion

In this work, a mathematical framework for the cultural evolution of language was designed and examined. Starting from the cultural inheritance mechanism, i.e. the naming game, a formalization of this process was defined. We demonstrated that variations of this formalization could easily be made. These variations can be expressed in terms of parameters and the kind of inheritance mechanism used. For the different steps in the discussion we provided simple experiments which showed the actual effects between individuals and in populations of individuals. Moreover, the evolutionary dynamics of the cultural inheritance system was made explicit using a linear selection–mutation model. This model showed that change in the cultural variants in a population is due to a process of selection in favor of the most frequently used word–meaning associations and that new associations emerge through mutations between combinations.

Thus we can summarize that the contribution of this work is three-fold. First it provided a model of evolutionary language evolution which is orthogonal to previously defined ones. A combination of these models with this one would allow for a more meaningful exploration of mathematical models and the role of each social learning process. Second, this model highlights the crucial nature of the role, active or passive and speaker or hearer, of the individual in an interaction. Third, we explicitly relate existing evolutionary dynamical models with the research in the origin and evolution of language. In this way, we provide the necessary mechanisms to language researchers to construct meaningful models to reach their personal research goals.

The presented model provides an alternative perspective on the dynamics of language evolution without going into detail on the genetic inheritance mechanism. Although, a simplification in different aspects we think that this base model provides a first step toward a more complete picture of the entire process. For instance, other social learning mechanisms can be introduced or further examinations can be made of the incorporation of vertical transmission. Moreover, one can consider the phenotypic plasticity of each individual and how the actual parameters of the language acquisition device are genetically inherited in this mechanism. Apart from these obvious extensions, we feel that the next crucial step is to extend the current evolutionary model toward

a system where both meaning and language is culturally acquired.

Acknowledgements

When this work was performed, Tom Lenaerts was funded by FWO Flanders on an ESF-EUROCORES project titled “The Origin of Man, Language and Languages: the cultural self-organization of the cognitive; robotic and computational models”. Both Karl Tuyls and Bart Jansen are funded by a doctoral grant of the Institute for the Advancement of Scientific Technological Research Flanders (IWT). We thank the reviewer for providing extensive and useful commentary that allowed us to improve the quality of the article.

References

- Børgers, T., Sarin, R., 1997. Learning through reinforcement and replicator dynamics. *J. Econom. Theory* 77 (1), 1–14.
- Boyd, R., Richerson, P., 1985. *Culture and the Evolutionary Process*. The University of Chicago Press, Chicago, London.
- Bush, R., Mosteller, F., 1951. A mathematical model for simple learning. *Psychol. Rev.* 58 (5), 313–323.
- Bush, R., Mosteller, F., 1955. *Stochastic Models for Learning*. Wiley, New York.
- Cavalli-Sforza, L., Feldman, M., 1981. Cultural transmission and evolution. *Monographs in Population Biology*, vol. 160. Princeton University Press, Princeton, NJ.
- Hofbauer, J., Sigmund, K., 1988. *The Theory of Evolution and Dynamical Systems: Mathematical Aspects of Selection*. Cambridge University Press, Cambridge.
- Hurford, J., 1989. Biological evolution of the saussurean sign as a component of the language acquisition device. *Lingua* 77 (2), 187–222.
- Kaplan, F., 2000. L'émergence d'un lexique dans une population d'agents autonomes. These de doctorat, LIP6, Université Paris VI.
- Kirby, S., Hurford, J., 2002. The Emergence of Linguistic Structure: An overview of the Iterated Learning Model. Springer, London, pp. 121–148 (Chapter 6).
- Narendra, K., Thathachar, M., 1989. *Learning Automata: An Introduction*. Prentice-Hall International, Englewood Cliffs, NJ.
- Nowak, M.A., Komarova, N.L., 2001. Towards an evolutionary theory of language. *Trends Cognitive Sci.* 5 (7), 288–295.
- Nowak, M.A., Plotkin, J.B., Krakauer, D., 1999. The evolutionary language game. *J. Theor. Biol.* 200 (2), 147–162.
- Oliphant, M., Batali, J., 1997. Learning and the emergence of coordinated communication. *Newslett. Center Res. Lang.* 11 (1).
- Rosenthal, T., Zimmerman, B., 1978. *Social Learning and Cognition*. Academic Press, New York.
- Smith, K., 2004. The evolution of vocabulary. *J. Theor. Biol.* 228 (1), 127–142.
- Steels, L., 1995. Perceptually grounded meaning creation. In: Lesser, V. (Ed.), *Proceedings of the First International Conference on Multi-Agent Systems*. MIT Press, Cambridge MA.
- Steels, L., 1996. Emergent adaptive lexicons. *Proceedings of the Simulation of Adaptive Behavior Conference*.
- Steels, L., 1999. *The Talking Heads Experiment*, vol. 1: Words and Meanings. Special pre-edition for LABORATORIUM, Antwerpen.

- Tomasello, M., 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Harvard University Press, Cambridge, MA, London, England.
- Tomasello, M., Kruger, A., Ratner, H., 1993. Cultural learning. *Behav. Brain Sci.* 16 (3), 495–552.
- Tuyls, K., Lenaerts, T., Verbeeck, K., Maes, S., Manderick, B., 2002. Towards a relation between learning agents and evolutionary dynamics. In: *Proceedings of the Fourteenth Belgium–Netherlands Conference on Artificial Intelligence*. K.U. Leuven, Belgium, pp. 315–322.
- Tuyls, K., Verbeeck, K., Lenaerts, T., 2003. A selection–mutation model for q-learning in multi-agent systems. In: *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 693–700.
- Weibull, J., 1996. *Evolutionary Game Theory*. MIT Press, Cambridge, MA.