

Evolving Compositionality in Evolutionary Language Games

José Fernando Fontanari and Leonid I. Perlovsky, *Senior Member, IEEE*

Abstract—Evolutionary language games have proved a useful tool to study the evolution of communication codes in communities of agents that interact among themselves by transmitting and interpreting a fixed repertoire of signals. Most studies have focused on the emergence of Saussurean codes (i.e., codes characterized by an arbitrary one-to-one correspondence between meanings and signals). In this contribution, we argue that the standard evolutionary language game framework cannot explain the emergence of compositional codes—communication codes that preserve neighborhood relationships by mapping similar signals into similar meanings—even though use of those codes would result in a much higher payoff in the case that signals are noisy. We introduce an alternative evolutionary setting in which the meanings are assimilated sequentially and show that the gradual building of the meaning-signal mapping leads to the emergence of mappings with the desired compositional property.

Index Terms—Complexity theory, game theory, genetic algorithms, simulation.

I. INTRODUCTION

THE CASE FOR the study of the evolution of communication within a multiagent framework was probably best made by Ferdinand de Saussure in his famous statement:

“language is not complete in any speaker; it exists only within a collectivity. . . only by virtue of a sort of contract signed by members of a community” [1].

Translated into the biological jargon, this assertion means that language is not the property of an individual, but the extended phenotype of a population [2]. More than one decade ago, seminal computer simulations were carried out to demonstrate that cultural [3] as well as genetic [4] evolution could lead to the emergence of ideal communication codes (i.e., arbitrary one-to-one correspondences between objects or meanings and signals), termed Saussurean codes, in a population of interacting agents. Typically, the behavior pattern of the agents was modeled by (probabilistic) finite-state machines. The work by

Hurford [3], in particular, set the basis of the Iterated Learning Model (ILM) for the cultural evolution of language, the typical realization of which consists of the interaction between two agents—a pupil that learns the language from a teacher [5]. In those studies, language is viewed as a mapping between meanings and signals. The communication codes that emerged from the agents’ interactions are, in general, noncompositional or holistic communication codes, in which a signal stands for the meaning as a whole. In contrast, a compositional language is a mapping that preserves neighborhood relationships—similar signals are mapped into similar meanings. If there is a nontrivial structure in both meaning and signal spaces then, in certain circumstances, making explicit use of those structures may greatly improve the communication accuracy of the agents. The emergence of compositional languages in the ILM framework beginning from holistic ones in the presence of bottlenecks on cultural transmission was considered a breakthrough in the computational language evolution field [5]–[7]. The aim of this contribution is to understand how compositional communication codes can emerge in an evolutionary language game framework [3], [4], [8], [9].

The way we introduce the structure of the signal space (i.e., the notion of similarity between signals) into the rules of the language game is through errors in perception: the signals are assumed to be corrupted by noise so that they can be mistaken for one of their neighbors in signal space [8]. Similarly, the structure of the meaning space enters the game by rewarding the agents that prompted by a signal, infer a meaning close to the meaning actually intended by the emitter. Of course, the reward for incorrect but close inferences must be smaller than that granted for the correct inference of the intended meaning (see [9] for a similar approach). Hence, the role played by noise in this context is similar to the role of the bottleneck transmissions in the ILM framework, since both make advantageous the exploration of the detailed structure of the meaning-signal mapping. In particular, we show that once a Saussurean communication code is established in the population, i.e., all agents use the same code, it is impossible for a mutant to invade, even if the mutant uses a better code, say, a compositional one. This is essentially the Allee effect [10], [11] of population dynamics that asserts that intraspecific cooperation might lead to inverse density dependence, resulting in the extinction of some (social) animal species when their population size becomes small. Of course, this effect is germane to the outcome of biological invasions involving such species. We note that most realizations of the ILM circumvent this difficulty by assuming that the population is composed of two agents only, the teacher and the pupil, and that the latter always replaces the former. However, according to de Saussure (see quotation above), this is not an acceptable framework for

Manuscript received December 15, 2005; revised May 11, 2006 and January 2, 2007. This work was supported in part by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under Grant FA9550-06-1-0202 and Grant FA8655-05-1-3031, and in part by CNPq and FAPESP under Project 04/06156-3. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon.

J. F. Fontanari is with the Instituto de Física de São Carlos, Universidade de São Paulo, São Carlos, São Paulo 13560-970, Brazil (e-mail: fontanari@ifsc.usp.br).

L. I. Perlovsky is with Harvard University, Cambridge, MA 02138 USA and also with the Air Force Research Laboratory, Hanscom AFB, MA 01731 USA (e-mail: Leonid.Perlovsky@hanscom.af.mil).

Digital Object Identifier 10.1109/TEVC.2007.892763

language. In addition, a bias toward compositionality is built in the inference procedure used by the pupil to fill in the gaps due to transmission bottlenecks, in which some of the meanings are not taught to the pupil. This bias towards generalization, together with cultural evolution, seems to be the key ingredients to evolve compositionality in the ILM framework.

Understanding as well as demonstrating how innovations that increase the expressive power of individuals can spread through a population is the essence of any evolutionary explanation to language evolution [9]. Accordingly, the solution we propose to the problem of evolving a compositional code in a population of agents that exchange signals with each other and receive rewards at every successful communication event is the incremental assimilation of meanings, i.e., the agents construct their communication codes gradually, by seeking a consensus signal for a single meaning at a given moment. Only after a consensus is reached, a novel meaning is permitted to enter the game. This sequential procedure, which dovetails with the classic Darwinian explanation to the evolution of strongly coordinated system, allows for the emergence of fully compositional codes, an outcome that we argue is very unlikely, if not impossible, in the traditional language game scenario in which the consensus signals are sought simultaneously for the entire repertoire of meanings.

II. MODEL

Here, we take the more conservative viewpoint that language evolved from animal communication as a means of exchanging relevant information between individuals rather than as a byproduct of animal cognition or representation systems (see, e.g., [12] and [13] for the opposite viewpoint). In particular, we consider a population composed of N agents who make use of a repertoire of m signals to exchange information about n objects. Actually, since the groundbreaking work of de Saussure [1], it is known that signals refer to real-world objects only indirectly as first the sense perceptions are mapped onto a conceptual representation—the meaning—and then this conceptual representation is mapped onto a linguistic representation—the signal. Here, we simply ignore the object-meaning mapping (see, however, [14] and [15]) and use the words object and meaning interchangeably. To model the interaction between the agents, we borrow the language game framework proposed by Hurford [3] (see also [8]) and assume that each agent is endowed with separate mechanisms for transmission (i.e., communication) and for reception (i.e., interpretation). More pointedly, for each agent we define a $n \times m$ transmission matrix P whose entries p_{ij} yield the probability that object i is associated with signal j , and a $m \times n$ reception matrix Q the entries of which, q_{ji} , denote the probability that signal j is interpreted as object i . Henceforth, we refer to P and Q as the language matrices. In general, the entries of these two matrices can take on any value in the range $[0,1]$ satisfying the constraints $\sum_{j=1}^m p_{ij} = 1$ and $\sum_{i=1}^n q_{ji} = 1$, in conformity with their probabilistic interpretation. In this contribution, however, we consider the case of binary matrices, in which the entries of Q and P can assume the values 0 and 1 only. There are two reasons for that. First, in the absence of errors in language learning, the evolutionary language game will eventually lead to binary transmission and reception matrices, regardless of

the values of m and n , and of the initial choice for the entries of those matrices [16]. Therefore, our restriction of the entry values to binary quantities has no effect on the equilibrium solutions of the evolutionary game. In addition, these deterministic encoders and decoders were shown to perform better than their stochastic variants [17]. Second, by assuming that the transmission and reception matrices are binary, we recover the synthetic ethology framework proposed by MacLennan [4], a seminal agent-based work on the evolution of communication in a population of finite state machines (see also [18]).

Although the reception matrix Q is, in principle, independent of the transmission matrix P , results of early computer simulations have shown that in a noiseless environment, the optimal communication strategy is the Saussurean two-way arbitrary relationship between an object and a signal, i.e., the matrices P and Q are linked such that if $p_{ij} = 1$ for some object-signal pair i, j , then $q_{ji} = 1$ [3]. These matrices are associated to the Saussurean communication codes introduced before, provided there are no correlations between the different rows of the matrix P , i.e., the assignment object-signal is arbitrary.

A. The Evolutionary Language Game

Given the transmission and reception matrices, the communicative accuracy or overall payoff for communication between two agents, say I and J , is defined as [3], [8], [19]

$$F(I, J) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^m \left(p_{ij}^{(I)} q_{ji}^{(J)} + p_{ij}^{(J)} q_{ji}^{(I)} \right) \quad (1)$$

from which we can observe the symmetry of the language game, i.e., both signaler and receiver are rewarded whenever a successful communication event takes place. By assuming such a symmetry, one ignores a serious hindrance to the evolution of language: passing useful information to another agent is an altruistic behavior [20], [21] that can be maintained in human societies thanks to the development of reciprocal altruism, in which unrelated individuals mutually benefit by exchanging the donor and the receiver roles multiple times [22]. However, the scarcity of empirical demonstrations of reciprocal altruism in nature, except for modern humans, motivated an alternative scenario for the evolution of language, namely, that human language evolved as a “mother tongue”—a communication system used among kin, especially between parents and their offspring [23].

In this contribution, we assume the validity of (1) and simply ignore the costs of honest signaling [20]. Hence, we take for granted the existence of special social conditions to foster reciprocal altruism among the agents or, alternatively, a mother tongue scenario in which the agents are related to each other. In this vein, it is interesting to note that although in the work by MacLennan [3] communication is defined following Burghardt [24] as “the phenomenon of one organism producing a signal that when responded to by another organism, confers some advantage to the signaler or his group” (see [25] for alternative definitions of communication), the actual implementation of the simulation rewards equally the two agents that take part in the successful communication event. In the case where only the receiver is rewarded, Saussurean communication fails to evolve [26].

Assuming, in addition, that each agent I interacts with every other agent $J = 1, \dots, N$ ($J \neq I$) in the population, we can immediately write down the total payoff received by I

$$F_I = \frac{1}{N-1} \sum_{J \neq I} F(I, J) \quad (2)$$

in which the sole purpose of the normalization factor is to eliminate the trivial dependence of the payoff measure on the population size N . Following the basic assumption of evolutionary game theory [27] this quantity is interpreted as the fitness of agent I . Explicitly, we assume that the probability that I contributes with an offspring to the next generation is given by the relative fitness

$$w_I = F_I / \sum_J F_J \quad (3)$$

which essentially implies that mastery of a public communication system adds to the reproductive potential of the agents [3].

There are several distinct ways to implement the language game. For instance, MacLennan [4] and Fontanari and Perlovsky [18] stick to the genetic algorithm approach (see, e.g., [28]) in which the offspring acquires both the transmission and reception matrices from its parent, assuming clonal or asexual reproduction. The offspring is identical to its parent except for the possibility of mutations that may alter a few rows of the language matrices. However, here we take a different viewpoint and reinterpret this genetic model within a learning context. We assume, in particular, that the offspring actually learns the language from its parent but that the learning is not perfect—there is a probability μ that the communication code it acquires is slightly different from its parent's. This very framework has been used to study the emergence of universal grammar and syntax in language [2], [29], [30].

An alternative learning scenario used by Nowak and Krakauer [8] assumes that the offspring adopt the language of its parent by sampling its response to every object k times. This approach makes sense only if the language matrices are not binary, though, as mentioned before, in the long run those matrices must become binary. For $k \rightarrow \infty$, the offspring is identical to its parent, which corresponds then to $\mu = 0$ in the previous learning scenario, whereas differences between parent and offspring arise in the case of finite $k > 1$. This sampling effect is qualitatively similar to the effect of learning errors in the scenario introduced before. For $k = 1$, already the first generation of offspring communicates through binary language matrices and so the sampling procedure is rendered ineffective. The reason is that a binary matrix P assigns each object to a unique signal (though this same signal can be used also for a distinct object), and so sampling the responses of the parent to the same object will always yield the same signal. As a result, the evolutionary process based on learning by sampling halts—the offspring become identical to their parents.

A similar but more culturally inclined approach is that followed by Hurford [3] and Nowak *et al.* [16]: instead of sampling the parent's responses, the offspring samples the responses of a certain number of agents in the population or even of the entire population. In this case, the hereditary component is lost since the offspring, in general, will not resemble its parent, and

so natural selection has no say in the outcome of the dynamics. In the case of Hurford [3], there is still a strong genetic component as the offspring inherits from its parent its strategy of inference. Similarly, the ILM for the cultural evolution of language (see [5] and [7] for reviews) in its more popular version consists of two agents only, the teacher and the pupil who learns from the teacher through a sampling process identical to that just described. The pupil then replaces the teacher and a new, tabula rasa pupil is introduced in the scenario. This procedure is iterated until convergence is achieved. In this case, the payoff (2) plays no role at all in the language evolutionary process and the stationary language matrices will depend strongly on the inference procedure used by the pupil to create a meaning/signal mapping from the teacher responses. Of particular interest for our purpose is the finding that compositional codes emerge in the case that the learning strategy adopted by the pupil supports generalization and that this ability is favored by the introduction of transmission bottlenecks in the communication between teacher and pupil. Such a bottleneck occurs when the learner does not observe the signal for some objects. This contrasts with the sampling effect mentioned before in which the learner observes the signals to every object. In this contribution, we study whether and in what conditions compositional codes emerge in an evolutionary language game.

B. The Meaning-Signal Mapping

As already pointed out, language is viewed as a mapping between objects (or meanings) and signals and compositionality is a property of this mapping: a compositional language is a mapping that preserves neighborhood relationships, i.e., nearby meanings in the meaning space are likely to be associated to nearby signals in signal space [5]. At first sight, this notion looks contradictory to the well-established fact that the relation between a word (signal) and its meaning is utterly arbitrary. For instance, as pointed out by Pinker [31],

“babies should not, and apparently do not, expect *cattle* to mean something similar to *battle*, or *singing* to be like *stinging*, or *coats* to resemble *goats*.”

In fact, Pettito demonstrated that the arbitrariness of the relation between a sign and its meaning is deeply entrenched in the child's mind [32]. On the other hand, sentences like *John walked* and *Mary walked* have parts of their semantic representation in common (someone performed the same act in the past) and so the meaning of these sentences must be close in the meaning space. Since both sentences contain the word *walked* they must necessarily be close in signal space as well. Following Pinker, we acknowledge a significant degree of arbitrariness at the level of word-object pairing. This might be a consequence of a much earlier (prehuman) origin of this mechanism, as compared with seemingly distinctly human mind mechanisms for sentence-situation pairing. From a mathematical modeling perspective, however, such a distinction is not essential for our purposes, since the signals (sentences or words) can always be represented by a single symbol—only the “distance” between them will reflect the complex inner structure of the signal space. For instance, suppose there are only two words that we represent, without lack of generality by 0 and 1. Hence, a binary sequence




Fig. 1. Example of a mapping meaning-signal for $n = m = 4$. The integers here may be viewed as labels for complex entities (e.g., sentences). The large circles indicate cyclic boundary conditions so that, e.g., signal 1 is 1 unit distant from signals 2 and 4. The code represented in the figure has compositionality $C = 1$.

or, equivalently, its decimal representation can represent any sentence in this language. Here, the relevant distance between two such sentences is the Hamming distance rather than the result of the subtraction between their labeling integers. This notion, of course, generalizes trivially to the case when the sentences are composed of more than two types of words.

For simplicity, in this paper, we consider the case where both signals and meanings are represented by integer numbers and the relevant distance in both signal and meaning space is the result of the usual subtraction between integers. Fig. 1 illustrates one of the $n \times m$ possible meaning-signal mappings. A quantitative measure of the compositionality of a communication code is given by the degree to which the distances between all the possible pairs of meanings correlates with the distance between their corresponding pairs of signals [7]. Explicitly, let Δm_{ij} be the distance between meanings i and j , and Δs_{ij} the distance between the signals associated to these two meanings. Introducing the averages $\Delta \bar{m} = \sum_{(ij)} \Delta m_{ij} / p$ and $\Delta \bar{s} = \sum_{(ij)} \Delta s_{ij} / p$, where the sum is over all distinct pairs $p = n(n-1)/2$ of meanings, the compositionality of a code is defined as the Pearson correlation coefficient [7]

$$C = \frac{\sum_{(ij)} (\Delta m_{ij} - \Delta \bar{m})(\Delta s_{ij} - \Delta \bar{s})}{\left[\sum_{(ij)} (\Delta m_{ij} - \Delta \bar{m})^2 \sum_{(ij)} (\Delta s_{ij} - \Delta \bar{s})^2 \right]^{1/2}} \quad (4)$$

so that $C \approx 1$ indicates a compositional code and $C \approx 0$ an unstructured or holistic code. This definition applies only to codes that implement a (not necessarily arbitrary) one-to-one correspondence between meaning and signal.

Strictly, here we do not address directly the emergence of compositionality, defined as the property that the meaning of a complex expression is determined by the meanings of its parts and the rules used to combine them. Rather, we focus on the emergence of structured communication codes, which preserve the topology of the meaning-signal mapping, in that similar meanings are associated with similar signals and *vice versa*. It seems that an important aspect of joint evolution of compositional cognition and compositional language is their evolution along with structural metric (or approximately metric) spaces

of cognition and meaning. In this contribution, we assume that a metric space exists, and explore the consequences for the emergence of compositionality. The connection between structured and compositional meaning-signal mappings can be made explicit if we consider an artificial scenario for which there is a prescription to derive the meaning of the whole given the meaning of the elementary parts. (Such prescription is clearly ruled out in real language since context and previous knowledge play a crucial role in our understanding of any situation.) In this case, the distance between any two composite meanings could be inferred by comparing their components and, consequently, by introducing a metric in the meaning space.

Our approach ties in with the view that properties of language such as compositionality are emergent characteristics of the explosion of semantic complexity occurred during hominid evolution [33]. Semantic complexity means not only a large number of cognitive categories (meanings) but also an increase in their perceived interrelationships, which are inherent properties of the topology of the meaning space. In fact, the number of objects for which a person has separate words is not too large: a recent estimate suggests a vocabulary of around 60,000 base words for well-educated adult native speakers of English [34]. This is not a very big number, and so it is reasonable to assume that object-word associations can be learned from examples, one by one. The number of situations that are combinations of objects, on the other hand, is larger than the number of all elementary particle events in the history of the Universe. This supports a need for the assumption of compositionality in language. As hinted in [33], a natural avenue to study the evolution of complex features of language (e.g., compositionality) is the increase of the complexity of the meaning space, which is exactly the approach we offer in this contribution.

C. Errors in Perception

So far as the communicative accuracy introduced in (1) is concerned, the structures of the meaning and signal spaces are irrelevant to the outcome of the evolutionary language game: the total population payoff is maximized when all agents adopt a code that implements a one-to-one correspondence between meanings and signals. Such a code is, of course, described by any one of the $n!$ permutation language matrices. The fact that ultimately all agents adopt the same communication code is a general result of population genetics related to the effect of genetic drift on a finite population [35]. To permit the structures of the meaning and signal spaces to play a role in the evolutionary game and so to break the symmetry among the permutation matrices so as to favor the compositional codes, we must introduce a new ingredient in the language game, namely, the possibility of errors in perception [8]. In fact, it is reasonable to assume that in the earlier stages of the evolution of communication the signals were likely to be noisy and so they could be easily mistaken for each other. The relevance of the structure of the signal space becomes apparent when we note that the closer two signals are, the higher the chances that they are mistaken for each other. This aspect of the model can be described by an agent-independent $m \times m$ confusion matrix E , the entries of which e_{ij} yield the probability of signal j being observed as signal i due to corruption by noise [8], [9].

To introduce the structure of the meaning space in the language game, we note first that (1) has a simple interpretation in the case of binary, but not necessarily permutation, language matrices: both signaler and receiver are rewarded with 1/2 unity of payoff whenever the receiver interprets correctly the meaning of the emitted signal. Otherwise, there is no reward to any of the two parts, no matter how close the inferred meaning is from the correct one. This gives us a clue as to how to modify the model in order to take into account the meaning structure—just ascribe some small reward value to both agents if the inferred meaning is close to the intended one. In fact, giving value to decisions which are not the best ones is a common assumption in decision and game theory [36] and seems to be consistent with what is actually observed in nature since, clearly, not every misinterpretation is equally harmful [9]. Consider for instance the Vervet monkey alarm calls [37]: misinterpreting a snake alarm for a leopard one, and hence running to a tree instead of standing up and looking in the grass, is clearly much better than misinterpreting it for an eagle call.

Following Nowak *et al.* [8] and Zuidema [9], we can formalize the notion of meaning similarity by introducing another agent-independent matrix, the $n \times n$ value matrix V , so that v_{ij} yields the payoff attributed to an agent which infers meaning i when the actual meaning the signaler intended to transmit was j . Hence, the overall payoff for communication between agents I and J becomes [9]

$$F(I, J) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n v_{ij} \left[P^{(I)} \times (E \times Q^{(J)}) + P^{(J)} \times (E \times Q^{(I)}) \right]_{ij} \quad (5)$$

where \times stands for the usual matrix multiplication. Note that (1) is recovered in the case that both value and confusion matrices are diagonal.

In particular, here we will consider the simple case in which there is a nonzero probability $\varepsilon \in [0, 1]$ that a signal, say signal j , be mistaken for one of its nearest neighbors only, $e_{j-1,j} = \varepsilon/2$ and $e_{j+1,j} = \varepsilon/2$. Of course, the probability that a signal is not corrupted by noise is $e_{j,j} = 1 - \varepsilon$. If signal j is in the boundary, $j = 1$ or $j = m$, then we use the cyclic structure of the signal space to set $e_{0,1} = e_{m,1} = \varepsilon/2$ and $e_{m+1,m} = e_{1,m} = \varepsilon/2$. So, in the example of Fig. 1, signal 4 can be mistaken only for signals 3 or 1 with probability ε . Similarly, agents are rewarded only if the inferred meaning is one of the nearest neighbors of the intended meaning. For example, if the intended meaning is j , then the only nonzero entries of the value matrix V are $v_{j,j} = 1$, $v_{j+1,j} = r$, and $v_{j-1,j} = r$. Meanings in the boundary, $j = 1$ and $j = n$, are treated using the cyclic boundary conditions as explained for the signal space. Here, $r \in [0, 1]$ is a parameter that measures the advantage, in terms of payoff, of using a compositional communication code rather than a Saussurean one.

Together with the presence of noise, this last ingredient—nonzero reward for inferring a meaning close to the correct one—should favor, in principle, the emergence of compositional communication codes in an evolutionary game guided by Darwinian rules. In what follows, we will show that

the problem of evolving efficient communication codes within an evolutionary framework, whether in the presence or not of noise, is more difficult than previously realized [4], [16], [18]. This problem differs from usual optimization problems tackled with evolutionary algorithms in that the maximization of the average population payoff requires a somewhat coordinated action of the agents. It is of no value for an agent to exhibit the correct “genome” (i.e., the transmission and reception matrices) if it cannot communicate efficiently with the other agents in the population because they use different language matrices.

The emergent view of compositionality adopted here differs from the approach followed by Nowak *et al.* [29] to study the evolution of syntactic (or combinatorial) communication. In that work, the conditions at which syntax is advantageous over non-syntactic or holistic languages were determined, namely, when the number of required signals to express the relevant meanings exceeds some threshold value. (It should be noted that combinatorial communication has its disadvantages too, since it boosts the potential for deception [38].) However, the finding that the adoption of a particular communication code is better for the population, in that it yields a higher overall payoff, is no guarantee that such code will actually spread in the population. On the contrary, in this contribution we show that the Allee effect will prevent its spreading. Additional assumptions, such as the semantic continuity of incremental learning proposed here, seem to be necessary to guarantee the emergence of compositional codes.

III. POPULATION DYNAMICS

We assume that the offspring learn their languages from their parents. Were it not for the effect of errors during learning, which results in small changes in the language matrices, the offspring would be identical to their parents. Like mutations in the genetic setup, these learning errors allow for the variability of the agents, and thus for the action of natural selection.

We start with N agents (typically $N = 100$) whose binary language matrices are set randomly. Explicitly, for each agent and for each meaning $i = 1, \dots, n$, we choose randomly an integer $j \in \{1, \dots, m\}$ and set $p_{ij} = 1$ and $p_{ik} = 0$ for $k \neq j$. Similarly, for each signal $j = 1, \dots, m$, we choose an integer $i \in \{1, \dots, n\}$ and set $q_{ji} = 1$ and $q_{jk} = 0$ for $k \neq i$. This procedure guarantees that initially P and Q are independent random probability matrices. Note that, in general, they are not permutation matrices at this stage. To calculate the total payoff of a given agent, say agent I , we let it interact with every other agent in the population. At each interaction, the emitted signal can be mistaken for one of the neighboring signals with probability ε . According to (5), at each communication event (an interaction) agent I receives the payoff value 1/2 if the receiver guesses the intended meaning of the signal that I has emitted, the payoff value $r/2$ if the receiver guessing is one of the nearest neighbors of the intended meaning, and payoff value 0, otherwise. Of course, the receiver obtains the same payoff accrued to agent I . Once the payoffs or fitness of all N agents are tabulated, the relative payoffs can be calculated according to (3), and then used to select the agent that will contribute with one offspring to the next generation.

To keep the population size constant, we must eliminate one agent from the population. To do that we will use two strategies: 1) to choose the agent to be eliminated at random, regardless of its fitness value and 2) to use an elitist strategy which eliminates the agent with the lowest fitness value. In both cases, the recently produced offspring is spared from demise. The first selection procedure is Moran's model of population genetics [35]. Both procedures differ from the standard genetic algorithm implementation [28] in that they allow for the overlapping of generations, a crucial prerequisite for cultural evolution which may be relevant when learning is allowed. In practice, however, Moran's model does not differ from the parallel implementation in which the entire generation of parents is replaced by that of the offspring in a single generation. We define the generation time t as the number of generations needed to produce N offspring with the consequent elimination of the same number of agents.

Finally, to allow for the appearance of novel codes (or language matrices) in the population, changes are performed independently on the transmission and reception matrices of the offspring with probability $u \in [0, 1]$. Explicitly, the transmission matrix P is modified by changing randomly the signal associated to an also randomly chosen meaning with probability u . A similar procedure updates the reception matrix Q . Hence, the probability that the same offspring has its transmission and reception matrices simultaneously altered by errors is u^2 and the probability that it will differ somehow from its parent is $\mu = 1 - (1 - u)^2$. Henceforth, we will refer to μ as the probability of error in language acquisition.

To facilitate comparison between different evolutionary algorithms, we define a properly normalized average payoff of the population

$$G = \frac{1}{nN} \sum_{I=1}^N F_I \quad (6)$$

so that $G \in [0, 1]$. The maximum value $G = 1$ is reached for Saussurean codes in the case of noiseless communication.

In Fig. 2, we present the effect of the inaccuracy in language acquisition on the average payoff of the population for the simplest situation, namely, $\varepsilon = 0$ (the receiver always gets the original signal) and $r = 0$ (only inference of the correct meaning is rewarded). The results show a stark difference between the elitist and the usual evolutionary strategy regarding the form they are affected by learning errors. Whereas the performance of Moran's model is degraded for high error rates [39], reaching the payoff of random binary matrices for $\mu = 1$, the elitist strategy actually benefits from those errors and gets to the maximum payoff for the highest possible error rate. In fact, for small but nonzero values of the error rate, the communication accuracy of the elitist strategy is practically constant and starts to increase only after μ crosses some threshold value $\mu \approx 0.02$. The performance of Moran's model, on the other hand, indicates the existence of an optimum value of the learning error for which the communication accuracy is maximum. Longer runs do not show any significant change of the pattern illustrated in Fig. 2. What enables the elitist strategy to take advantage of errors is the overlapping of generations together with the immediate removal




Fig. 2. Normalized average payoff G of the population as function of the probability of error in language acquisition μ in the case of $N = 100$ agents communicating about $n = 10$ meanings using $m = 10$ signals. The evolution was followed until $t = 2 \times 10^3$ for the elitist strategy (\circ) and until $t = 10^4$ for Moran's model (Δ). The symbols represent the average of over 50 independent runs. The error bars are smaller than the symbol sizes. For $\mu = 0$, we find $G = 0.255 \pm 0.005$ for both strategies, whereas for random language matrices we find $G = 0.1 \pm 0.0001$. The other parameters are $\varepsilon = r = 0$. The search space is the $m^n \times n^m$ space spanned by the two independent binary probability matrices P and Q .

of unfit agents from the population. This combination prevents the accumulation of inefficient agents in the population and the consequent degradation of the communication performance observed in Moran's model. Moreover, by eliminating the agent that performs worse in the language game, the elitist strategy adds an extra kick to the selective pressure towards better communication codes, in addition to the fitness regulation of offspring production described in (3).

The reason the elitist strategy can guide the population to a regime of practically perfect communication accuracy even in the presence of a constant flux of inefficient mutants ($\mu = 1$) is that a defective offspring, though spared from demise at birth, will almost certainly be purged from the population in the next step. We recall that a single generation comprises N such generation/elimination steps. In this scheme, the population can maintain at most a single defective agent, thus resulting in a reduction of the maximum normalized payoff by a factor $1/nN$. In view of the remarkable effectiveness of the elitist strategy to maximize the communication accuracy of the population, in what follows we will present the results for that strategy only.

Fig. 3 presents the average communication accuracy for 100 independent runs (populations) in a generic case in which the parameters ε and r , which couple the dynamics with the distances in the signal and meaning spaces are nonzero. Now, since the communication between any two agents is affected by noise, we must adopt a slightly different procedure to evaluate the payoff of the entire population. As before, we follow the evolutionary dynamics (i.e., the differential reproduction and learning-with-error procedures) until $t = 2 \times 10^3$, then we store the language matrices of all N agents. Keeping these matrices fixed, we evaluate the average population payoff in 100 contests. A contest is defined by the interaction between all pairs of agents in the population. Actually, according to (5),




Fig. 3. Normalized average payoff for the elitist (o) strategy at $t = 2 \times 10^3$ for 100 independent sample runs of the evolutionary dynamics. These results are compared with that of a fully compositional code (solid line) and of Saussurean codes (x). The parameters and search space are the same as in Fig. 2 with $\mu = 1$, except that now we have included a pressure for compositionality: the signals are corrupted with probability $\varepsilon = 0.2$ and the ratio between the payoffs for inferring a close and the correct meaning is $r = 0.25$. The optimal, compositional code yields $G \approx 0.85$ and the typical payoff of a Saussurean code is $G \approx 0.80$.

each interaction comprises two communication attempts, since any given agent first plays the role of the emitter and then of the receptor. Hence, a contest amounts to $N(N - 1)$ communication events. Of course, in the noiseless case ($\varepsilon = 0$), the payoff obtained would be the same in all contests. The procedural changes are needed to average out the effects of noise. For instance, in a single interaction two perfectly compositional codes could perform worse than two holistic codes if, by sheer chance, the signals happen to be corrupted only during the interaction of the compositional codes. To avoid such spurious effects the payoffs resulting from the interactions between any two agents are averaged out over 100 different interactions.

For the purpose of comparison, in Fig. 3 we also present the results for a population of agents carrying the same perfectly compositional code ($C = 1$), as well as for a similarly homogeneous population of agents carrying identical Saussurean codes. These are control populations that in contrast to the elitist populations, do not evolve. In the absence of noise, these control populations would reach the maximum allowed payoff, $G = 1$. We note that a perfectly compositional code is not a Saussurean code, in the sense that the one-to-one mapping between meaning and signals is not arbitrary. The elitist strategy seems to face great difficulties even to find a Saussurean code, as compared with the performance in the noiseless case (see Fig. 2) for instance, not to mention to find the optimum, perfect compositional code. Actually, in the presence of noise, the performance of the Saussurean code seems to pose an upper limit to the performance of the elitist strategy by acting as an attractor to the evolutionary dynamics.

It is instructive to calculate the average payoff G_c of a population composed of identical agents carrying a perfect compositional code. Consider the average payoff received by a given agent, say I , in a very large number of interactions with one of its siblings, say J . When I plays the signaler role its average




Fig. 4. Compositionality of the code carried by the agent with the highest payoff in the runs shown in Fig. 3. The compositionality of the perfect compositional code is, by definition, $C = 1$. There is a slight tendency to compositionality in the codes produced by the elitist (o) strategy as compared with those of the Saussurean codes (x).

payoff is $(1 - \varepsilon) \times 1/2 + \varepsilon \times r/2$, which, by symmetry, happens to be the same average payoff I receives when it plays the receiver role. Since all agents are identical, the expected payoff of any agent equals that of the population. Hence

$$G_c = 1 - \varepsilon(1 - r). \quad (7)$$

We can repeat this very same reasoning to derive the average payoff G_S of a homogenous population of Saussurean codes. In this case, by playing the signaler, I receives the average payoff $(1 - \varepsilon) \times 1/2 + \varepsilon \times 2/(n - 1) \times r/2$, where the factor $2/(n - 1)$ accounts for the fact that the reward $r/2$ is obtained only if the inferred meaning is one of the two neighbors of the correct meaning. This reasoning is valid for $n > 2$ only, since for $n = 2$ each meaning has a single neighbor, and so there is no difference between Saussurean and compositional codes. Taking into account the payoff received by I when playing the receiver yields

$$G_S = 1 - \varepsilon + \frac{2\varepsilon}{n - 1}r \quad (8)$$

for $n > 2$. Note that $G_c > G_S$ for $n > 3$. Similarly to the case $n = 2$, the Saussurean codes for $n = 3$ are compositional codes because of the cyclic boundary conditions in the meaning space. In Fig. 4, we show the compositionality of the code carried by the agent with the largest payoff value in each of the runs used to generate the data of Fig. 3. Although there is a slight tendency to compositionality in the codes produced by the elitist strategy, it is fair to say that the pressure to generate compositional code has not worked as expected, despite the clear advantage of such codes given the conditions of the experiment (see Fig. 3). As pointed out, the reason for that might be that the Saussurean codes act as barriers (local maxima) from which the evolutionary dynamics cannot escape, thus impeding it from reaching a perfect compositional code (global maximum).

The results depicted in Fig. 3 expose clearly the failure of the language evolutionary framework to produce efficient communication codes when the receiver must interpret noisy signals. To rule out the possibility that the cause of such failure




Fig. 5. Average payoff resulting from 100 independent runs of the noisy evolutionary language game with the search space restricted to permutation matrices (\circ) as a function of the pressure for compositionality. The error bars are smaller than the symbol sizes. The upper straight line is the function $G_c = (1 + r)/2$ that yields the average payoff of a perfect compositional code and the lower straight line is $G_S = 0.5 + 0.11r$ that yields the average payoff of a Saussurean code (see (7) and (8)). The parameters are $\varepsilon = 0.5$, $\mu = 0.9$, $N = 100$, and $n = m = 10$.

was the initial unlikely decoupling between production and interpretation, in the following, we will restrict the search space to that of Saussurean codes. Hence, for any agent, the transmission matrix P is a permutation matrix and the reception matrix Q has entries given by $q_{ji} = 1$ if $p_{ij} = 1$ and 0 otherwise (Q is also a permutation matrix). The initial population is composed of N agents adopting distinct Saussurean codes. To guarantee that all new codes generated by mutations stay within our search space, we modify the mutation procedure so that with probability μ the signal associated to a randomly chosen meaning, say i , is exchanged with the signal associated to another randomly chosen meaning, say k . This corresponds to the interchange of the rows i and k of the transmission matrix. The reception matrix is then updated accordingly. The sole genetic strategy we use in the forthcoming simulations is the elitist one, in which the worst performing agent is replaced by the offspring of the agent chosen by rolling the fitness wheel.

In Fig. 5, we show the results of the experiments with the evolutionary search restricted to the space of permutation matrices. The procedure we use here was the same as that employed to draw Figs. 3 and 4: after the evolutionary dynamics has settled to an equilibrium (i.e., all agents are using the same communication code, except for single temporary mutants), the resulting homogeneous population is then left to interact for 100 contests and the average payoff is recorded. However, instead of exhibiting the payoff obtained in the 100 independent runs as in Fig. 3, we exhibit in Fig. 5 only the average payoff calculated over those runs. Hence, to obtain each data point of this figure we need to generate a set of data similar to that used to draw Fig. 3. We choose as the independent variable the ratio between the payoffs for inferring a neighbor of the correct meaning and the correct meaning (r), which can be interpreted also as a selective pressure for evolving compositional codes. For the sake




Fig. 6. Average compositionality of the 100 evolved communication codes (\circ) whose payoffs are exhibited in Fig. 5, as well as of the same number of Saussurean codes (\times). The compositionality of a perfect compositional code is $C = 1$ by definition. The linear fitting of the average compositionality of the evolved codes yields a slope of ≈ 0.43 .

of comparison, Fig. 5 also shows the average payoffs of perfect compositional and random Saussurean codes.

The results in Fig. 5 indicate that for $r = 0$, the performance of the communication codes, regardless of whether random, compositional or evolved, are identical. Explicitly, in this case, we find $G = 1 - \varepsilon$ for any one-to-one mapping. Since the search space is now restricted to the space of permutation matrices, it is not a surprise that the payoffs of the Saussurean codes serve as lower bounds to those of the evolved codes. This trivial finding should not be confused with the unexpected result exhibited in Fig. 3, that the payoffs of the Saussurean codes serve as upper bounds to the payoffs of the evolved codes when the search space is enlarged to cover all binary language matrices. The results in Fig. 5 show clearly that, despite the fact that compositionality can greatly improve the communication payoff of the population (see upper straight line in that figure), the evolved codes fall short of taking full advantage of the structure of the meaning-signal space to cope with the noise in the communication. As a result, the evolved codes are far from the optimal, perfect compositional codes, although they fare better than the Saussurean codes. Fig. 6 explains the reason for that: the evolutionary dynamics actually succeeded to produce partially compositional codes, thus reducing the deleterious effects of noise.

It is interesting that the payoffs of the Saussurean codes increase when the pressure for compositionality increases [see Fig. 5 and (8)], although they remain largely noncompositional in average (see Fig. 6). The key to the explanation of this result is found in Fig. 4, where we can see that half of the samples of the random Saussurean codes exhibit a positive value of the compositionality, which is then associated to a payoff value greater than $1 - \varepsilon$ ($= 0.8$ in that case), while the representatives of the other half have a payoff of $1 - \varepsilon$ at worst. It is clear that the resulting average payoff must be an increasing function of r .

The reason that the evolutionary dynamics failed to produce perfect compositional codes, despite their obvious advantage to




Fig. 7. The evolution of the fraction f of agents carrying a perfect compositional code in an experiment in which they compete against agents carrying a Saussurean code. The parameters are $\varepsilon = 0.5$, $r = 0.25$, $N = 100$, and $n = m = 10$. The initial population is set so that (from top to bottom) $f = 0.8$, 0.5 , 0.42 , 0.419 , and 0.2 .

cope with noisy signals, is that once a nonoptimal communication code has become fixed (or even almost fixed) in the population, mutants carrying better codes cannot invade. In fact, those mutants will most certainly do badly when communicating with the resident agents and, as a result, will quickly be removed from the population. As pointed out, this is essentially the Allee effect of population dynamics.

The task faced by the evolutionary algorithm here is of an essentially different nature from that tackled in typical optimization problems in which the fitness of an agent is frequency independent. In such a case, a fitter mutant can always invade the resident population. To stress this phenomenon, Fig. 7 illustrates the competition between a fraction f of agents carrying (the same) perfect compositional code and a fraction $1 - f$ of agents carrying (the same) Saussurean code. This simulation is implemented using the elitist procedure described before, except that learning errors are not allowed, so that at any time an agent can carry only one of the two types of codes set initially. Alternatively, Fig. 7 can be interpreted as the competition between two different strategies: the perfect compositional and the holistic strategies. We can easily estimate the minimum fraction f_m of perfect compositional codes above which this strategy dominates the population. It is simply

$$\frac{f_m}{1 - f_m} = \frac{G_S}{G_c} \quad (9)$$

with G_c and G_S given by (7) and (8), respectively. For the parameters of Fig. 8, this estimate yields $f_m \approx 0.46$, which within statistical errors, is in very good agreement with the single run experiment described in the figure. Repetition of this experiment using Moran's model rather than the elitist strategy leads to the same result, except that the fixation of the winner strategy takes much longer—about 100 times longer than the fixation times exhibited in Fig. 7.

This simple analysis of the competition between suboptimal Saussurean codes and the optimal compositional codes lends support to our previous conclusion that compositional codes do

not evolve within the usual language evolutionary game framework because the evolutionary dynamics is very likely to get trapped in the local maxima—the Saussurean codes.

IV. INCREMENTAL MEANING ASSIMILATION

What we have been trying to do up to now is to evolve in a single shot a communication code that associates each of the n meanings (or objects) to one of the m signals available in the repertoire of the agents. As pointed out, in the case that the meaning-signal mapping has a nontrivial underlying structure, the optimal association is not completely arbitrary in the sense that in the presence of noise some codes (i.e., the perfect compositional codes) result in a much better communication accuracy than codes that implement an arbitrary one-to-one correspondence between meaning and signals (Saussurean codes). The results of the previous simulations lead us to conclude that it is very unlikely, if not impossible, that evolution through natural selection alone could take advantage of the structure of the meaning-signal space to produce the optimal, perfect compositional codes.

The outcome would be very different, however, if the task posed to the population were to reach a consensus on the signals to be assigned to the meanings in a sequential manner. In other words, let us consider the situation in which each agent has m signals available (here we set $m = 10$) and the population needs to communicate about a single meaning, say $i = 1$. The search space is reduced then to the space of the $1 \times m$ permutation matrices. (We restrict the search space to that of permutation matrices, for simplicity.) Once the consensus is reached (i.e., the signal assigned to meaning $i = 1$ is fixed in the population), a new meaning is presented and the population is then challenged to find a consensus signal for that meaning. The procedure is repeated until each of the $n = m$ meanings are associated to a unique signal.

In the case of structured meaning-signal mappings, the order of presentation of meanings to the population plays a crucial role on the outcome of this strategy, which we term sequential meaning assimilation. In particular, success is guaranteed only if the novel meaning is a neighbor of the previously presented meaning (e.g., $i = 2$ or $i = N$ in the case the previous assimilated meaning was $i = 1$). In this case, the question is whether the population will reach a consensus on a signal that is also a neighbor of the signal assigned to the previous meaning. Curve (a) of Fig. 8 shows that this scheme works neatly, and yields a fully compositional code provided that $\varepsilon \neq 0$ and $r \neq 0$. We note that when the number of assimilated meanings is less than the size of the repertoire of signals m , the payoff of the sequential assimilation scheme [curve (a)] falls below the average payoff a fully compositional code (dashed horizontal line), because until all meanings are presented, the codes produced by that scheme cannot take full advantage of the topology of the meaning and signal spaces. The following example explains the reason this is so. Consider the situation in which two meanings were assimilated, say $i = 1, 2$ and the signals assigned to them were $j = 6, 7$, respectively. The agents will receive no reward if the corrupted signals become 5 or 8 (we recall that $m = 10$ in this experiment), since at this point there are no meanings associated to these altered signals. In contrast, reward is always




Fig. 8. Average payoff of the population when the task is to produce consensus signals to n meanings presented sequentially at the time intervals $\Delta t = 100$. In curve (a), the new meaning is a neighbor of the previous one, whereas in curve (c), the order of presentation of the meanings is random. The result for the usual batch algorithm, in which all meanings are presented simultaneously, is shown in curve (b). The dashed horizontal line indicates the average performance of perfect compositional codes. The parameters are $\varepsilon = 0.5$, $r = 0.5$, $N = 100$, and $n = m = 10$.

guaranteed for the fully formed compositional code since, by definition, all meanings are assimilated at the very outset in this case. Of course, as seen in Fig. 8, this “surface” effect is attenuated as more meanings are assimilated. The fact that the final payoff of the single run displayed in curve (a) ends up being greater than the (theoretical) average payoff of the perfect compositional code is simply a statistical fluctuation. Curve (c) in Fig. 8 illustrates the failure of the sequential presentation scheme when the order of presentation of meanings is random. In fact, if the meanings are presented in an arbitrary order, say $i = 3$ after $i = 1$, then there is no selection pressure to prevent that the signal assigned to $i = 3$ be one of the neighbors of the signal associated to $i = 1$. Eventually, when the meaning $i = 2$ is presented this optimal signal will be unavailable to the agents, precluding thus the emergence of a compositional code. Finally, we note that the incremental learning scheme would work all the same if the repertoire of meanings were left fixed and the signals were presented one by one.

The proposed solution to the evolution of compositional codes in an evolutionary language game framework could be questioned, because it relies on the assumption that the new meanings entering the population repertoire must be closely related to the already assimilated meanings. However, this seems to be the manner in which the perceptual systems work during categorization: new meanings are usually hierarchically related to the assimilated ones and this could be, for instance, the reason for Zipf’s law of languages [40], [41]. In fact, as pointed out in [33], the hierarchical structure of language may be caused by our perception of reality, rather than the other way around. The case for a hierarchically organized world was made by Simon [42]:

“On theoretical grounds we could expect complex systems to be hierarchies in a world in which complexity had to evolve from simplicity.”

In addition, the evidence that nouns are easily changed into verbs (e.g., ship-shipped, bottle-bottled) [43] illustrates the same type of continuity in the signal space as well.

In any event, our solution is in line with the traditional Darwinian explanation to the evolution of the so-called irreducibly complex systems. Although the evolutionary game setting failed to evolve perfect compositional codes when the task was to produce a meaning-signal mapping by assimilating all meanings simultaneously, that setting proved successful when the meanings were created gradually.

V. CONCLUSION

Saussure’s notion of language as a contract signed by members of a community to set arbitrarily the correspondence between words and meanings leads to unexpected obstacles to the evolution of efficient communication codes in the evolutionary language game framework. In fact, the fixation of a communication code in a population is a once-for-all decision—it cannot be changed even if a small fraction of the population acquires a different, more efficient code (see Fig. 7). The situation here is similar to the evolutionary stable strategies of game theory [27], the escape from which is only possible if all players change their strategies simultaneously. Since such concerted, global changes are not part of the rules of the language game, there seems to be no way for the population to escape from nonoptimal communication codes.

In fact, languages evolve. A branch of linguistics named glottochronology (the chronology of languages) suggests the rule of thumb that languages replace about 20% of their basic vocabulary every 1000 years [44]. The abovementioned difficulty of changing the communication code is not in the replacement of old signals by new ones, but in the assignment of different meanings to old signals and *vice versa*. Of course, this would not be an issue if the evolutionary language game could lead the population to the optimal code (a perfectly compositional code, in our case); our simulations have shown that it always gets stuck in one of the local maxima that plague the search space. To point out this difficulty was, in fact, the main goal of the present contribution.

Our view of compositionality as the evolutionary stage following the settlement of simpler, unstructured communication codes, and the search for a continuous path connecting these two stages, led us to the same type of difficulties researchers working on a similarly elusive problem—the origin of life—have been struggling with for more than three decades [39]. For example, although the coordinated work of distinct genes is germane to the emergence of cells, it is still not clear how such an assemblage could be formed and maintained starting from selfish genes (see [45] for a review). In that sense, by exposing the obstacles to explain compositionality from an evolutionary perspective, our work follows the same research vein that led to the present understanding of prebiotic evolution.

The solution we put forward to this conundrum is a conservative one—we cannot explain the emergence of the entire meaning-signal mapping that displays the required compositional property via natural selection, but it is likely that the mapping was formed gradually with the addition of one meaning

at each time. This gradual procedure, that we term incremental meaning creation, leads indeed to fully compositional codes (see Fig. 8). It would be interesting to verify whether alternative, less conservative solutions such as the spatial localization of the agents, less than perfect metrics in meaning space, or the structuring of the population by age could lead to the dissolution of the language contract and so open an evolutionary pathway to more efficient communication codes.

REFERENCES

- [1] F. de Saussure, *Course in General Linguistics*. New York: McGraw-Hill, 1966, Translated by Wade Baskin.
- [2] M. A. Nowak, N. L. Komarova, and P. Niyogi, "Computational and evolutionary aspects of language," *Nature*, vol. 417, pp. 611–617, 2002.
- [3] J. R. Hurford, "Biological evolution of the Saussurean sign as a component of the language acquisition device," *Lingua*, vol. 77, pp. 187–222, 1989.
- [4] B. J. MacLennan, "Synthetic ethology: An approach to the study of communication," in *Artificial Life II*, C. G. Langton, C. Taylor, J. Doayne Farmer, and S. Rasmussen, Eds. Reading, MA: Addison-Wesley, 1991, vol. X, SFI Studies in the Sciences of Complexity, pp. 631–658.
- [5] K. Smith, S. Kirby, and H. Brighton, "Iterated learning: A framework for the emergence of language," *Artif. Life*, vol. 9, pp. 371–386, 2003.
- [6] H. Brighton, "Compositional syntax from cultural transmission," *Artif. Life*, vol. 8, pp. 25–54, 2002.
- [7] H. Brighton, K. Smith, and S. Kirby, "Language as an evolutionary system," *Phys. Life Rev.*, vol. 2, pp. 177–226, 2005.
- [8] M. A. Nowak and D. C. Krakauer, "The evolution of language," *Proc. Nat. Acad. Sci. USA*, vol. 96, pp. 8028–8033, 1999.
- [9] W. Zuidema, "Optimal communication in a noisy and heterogeneous environment," in *Proc. 7th Eur. Conf. Arti. Life (ECAL), Advances in Artificial Life*, W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, and J. Ziegler, Eds., 2003, vol. 2801, Lecture Notes in Artificial Intelligence, pp. 553–563.
- [10] W. C. Allee, *Animal Aggregations. A Study in General Sociology*. Chicago, IL: Univ. Chicago Press, 1931.
- [11] F. Courchamp, T. Clutton-Brock, and B. Grenfell, "Inverse density dependence and the Allee effect," *Trends Ecol. Evol.*, vol. 14, pp. 405–410, 1999.
- [12] D. Bickerton, *Language & Species*. Chicago, IL: Univ. Chicago Press, 1990.
- [13] I. Ulbaek, "The origin of language and cognition," in *Approaches to the Evolution of Language*, J. R. Hurford, M. Studdert-Kennedy, and C. Knight, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1998, pp. 30–43.
- [14] L. Steels, "Perceptually grounded meaning creation," in *Proc. 2nd Int. Conf. Multi-Agent Syst.*, M. Tokoro, Ed., 1996, pp. 338–344.
- [15] J. F. Fontanari and L. I. Perlovsky, "Meaning creation and modeling field theory," in *Proc. Int. Conf. Integr. Knowl. Intensive Multi-Agent Syst.*, C. Thompson and H. Hexmoor, Eds., 2005, pp. 405–410.
- [16] M. A. Nowak, J. B. Plotkin, and D. C. Krakauer, "The evolutionary language game," *J. Theor. Biol.*, vol. 200, pp. 147–162, 1999.
- [17] J. B. Plotkin and M. A. Nowak, "Language evolution and information theory," *J. Theor. Biol.*, vol. 205, pp. 147–159, 2000.
- [18] J. F. Fontanari and L. I. Perlovsky, "Evolution of communication in a community of simple-minded agents," in *Proc. Int. Conf. Integr. Knowl. Intensive Multi-Agent Syst.*, C. Thompson and H. Hexmoor, Eds., 2005, pp. 285–290.
- [19] D. K. Lewis, *Convention: A Philosophical Study*. Cambridge, MA: Harvard Univ. Press, 1969.
- [20] R. Dawkins and J. R. Krebs, "Animal signals: Information or manipulation?," in *Behavioural Ecology: An Evolutionary Approach*, J. R. Krebs and N. B. Davies, Eds. Oxford, U.K.: Blackwell, 1978, pp. 282–309.
- [21] J.-L. Dessalles, "Altruism, status and the origin of relevance," in *Approaches to the Evolution of Language*, J. R. Hurford, M. Studdert-Kennedy, and C. Knight, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1998, pp. 130–147.
- [22] R. L. Trivers, "The evolution of reciprocal altruism," *Quart. Rev. Biol.*, vol. 46, pp. 35–57, 1971.
- [23] W. T. Fitch, "Kin selection and mother tongues: A neglected component in language evolution," in *Evolution A of Communication Systems: A Comparative Approach*, K. Oller and U. Griebel, Eds. Cambridge, MA: MIT Press, 2004, pp. 275–296.
- [24] G. M. Burghardt, "Defining communication," in *Communication by Chemical Signals*, J. W. Johnston, Jr., J. D. G. Moulton, and A. Turk, Eds. New York: Appleton-Century-Crofts, 1970, pp. 5–18.
- [25] M. D. Hauser, *The Evolution of Communication*. Cambridge: MIT Press, 1996.
- [26] M. Oliphant, "The dilemma of Saussurean communication," *BioSystems*, vol. 37, pp. 31–38, 1996.
- [27] J. M. Smith, *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge Univ. Press, 1982.
- [28] M. Mitchell, *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press, 1996.
- [29] M. A. Nowak, J. B. Plotkin, and V. A. A. Jansen, "The evolution of syntactic communication," *Nature*, vol. 404, pp. 495–498, 2000.
- [30] M. A. Nowak, N. L. Komarova, and P. Niyogi, "The evolution of universal grammar," *Science*, vol. 291, pp. 114–118, 2001.
- [31] S. Pinker, *The Language Instinct*. Baltimore, MD: Penguin, 1994, p. 152.
- [32] L. A. Petitto, "Language in the prelinguistic child," in *Language Acquisition: Core Readings*, P. Bloom, Ed. Cambridge, MA: MIT/Bradford Press, 1994.
- [33] P. T. Schoenemann, "Syntax as an emergent characteristic of the evolution of semantic complexity," *Minds and Machines*, vol. 9, pp. 309–346, 1999.
- [34] J. Aitchinson, *Words in the Mind: An Introduction to the Mental Lexicon*. Oxford, U.K.: Blackwell, 1994.
- [35] W. J. Ewens, *Mathematical Population Genetics*, 2nd ed. New York: Springer, 2004.
- [36] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.
- [37] R. M. Seyfarth, D. L. Cheney, and P. Marler, "Monkey responses to three different alarm calls: Evidence of predator classification and semantic classification," *Science*, vol. 210, pp. 801–803, 1980.
- [38] M. Lachmann and C. T. Bergstrom, "The disadvantage of combinatorial communication," *Proc. R. Soc. Lond. B*, vol. 271, pp. 2337–2343, 2004.
- [39] M. Eigen, "Self-organization of matter and the evolution of biological macro-molecules," *Naturwissenschaften*, vol. 58, pp. 465–523, 1971.
- [40] H. A. Simon, "On a class of skew distribution functions," *Biometrika*, vol. 42, pp. 425–440, 1955.
- [41] L. B. Levitin, B. Schapiro, and L. I. Perlovsky, "Zipf's law revisited: Evolutionary model of emergent multiresolution classification," in *Proc. Conf. Intell. Syst. Semiotics*, Gaithersburg, MD, 1996, vol. 1, pp. 65–70.
- [42] H. A. Simon, *The Sciences of the Artificial*. Cambridge, MA: MIT Press, 1996.
- [43] P. J. Hopper and S. A. Thompson, "The discourse basis for lexical categories in universal grammar," *Language*, vol. 60, pp. 703–752, 1984.
- [44] C. Renfrew, *Archaeology and Language*. London, U.K.: Pimlico, 1998.
- [45] J. M. Smith and E. Szathmáry, *The Major Transitions in Evolution*. San Francisco, CA: Freeman, 1995.



José Fernando Fontanari received the Ph.D. degree in physics from the University of São Paulo (USP), São Paulo, Brazil, in 1988

He is a Professor of Theoretical Physics at USP. He worked as a Postdoctoral Researcher at the California Institute of Technology in J. Hopfield's group in 1989. The author of more than 80 papers in international journals with selective editorial policy has served as a Member of the Editorial Board of *Network: Computation in Neural Systems* from 1990–1992 and currently serves on the Editorial

Board of *Physics of Life Reviews* and *Theory in Biosciences*. His research focuses on the application of concepts and analytical tools from physics to problems in biology, in particular, population dynamics and evolutionary theory.

Dr. Fontanari was elected Fellow of the Institute of Physics (U.K.) in August 2004.



Leonid I. Perlovsky (M'85–SM'95) is a Visiting Scholar at Harvard University, Cambridge, MA, and Principal Research Physicist and Technical Advisor at the Air Force Research Laboratory, Hanscom AFB, MA. He is Program Manager for DOD Semantic Web program and leads several research projects. From 1985 to 1999, he served as Chief Scientist at Nichols Research, a \$0.5B high-tech DOD contractor leading the corporate research in intelligent systems, neural networks, sensor fusion, and target recognition. He served as

Professor at Novosibirsk University and New York University, and participated as a principal in commercial startups developing tools for text understanding, biotechnology, and financial predictions. His company predicted the market crash following 9/11 a week before the event, detecting activities of Al Qaeda traders, and later helped SEC looking for these guys. He delivered invited keynote plenary talks and tutorial lectures worldwide, published more than 250 papers, 7 book chapters, and authored a monograph *Neural Networks and Intellect* (Oxford University Press, 2001) (currently in the 3rd printing). *The Knowledge Instinct* is being published in 2007 by Basic Books.

Dr. Perlovsky received the IEEE Distinguished Member of Boston Section Award and the International Neural Network Society Gabor Award. He serves as Associate Editor for the IEEE TRANSACTIONS ON NEURAL NETWORKS, Editor-at-Large for *Natural Computations*, and Editor-in-Chief for *Physics of Life Reviews*. He organizes conferences on Computational Intelligence, and Chairs the IEEE Boston Computational Intelligence Chapter.