

An Infection-Based Mechanism in Large Convention Spaces

Norman Salazar, Juan A. Rodriguez-Aguilar, and Josep Ll. Arcos

IIIA, Artificial Intelligence Research Institute
CSIC, Spanish National Research Council
{norman, jar, arcos}@iiia.csic.es

Abstract. Regulating the behavior of autonomous agents is necessary to solve coordination problems and minimize conflicts in multi-agent systems (MAS). It is well known that in practice centralized approaches are not viable to accomplish this. Thus, distributed regulating mechanisms, such as mechanisms for the emergence of social conventions, are highly needed. Nevertheless, existing studies have not focused on determining how the size of the convention space may influence the emergence of conventions. To that end in this paper we apply a mechanism for the distributed, dynamic emergence of social conventions, to a problem with a large convention space: finding a common vocabulary (lexicon) for the agents of a MAS that allows them to perfectly communicate with neither ambiguity nor inconsistencies. Therefore, we empirically show that the mechanism can cope with large convention spaces.

1 Introduction

Regulating the behavior of autonomous agents in multi-agent systems (MAS) to improve its overall performance and effectiveness is a current subject of interest. In particular, to solve coordination problems and minimize conflicts. It is well known that centralized techniques that depend on global knowledge have become a less than viable approach to accomplish this. Therefore, distributed mechanisms have become highly desirable. In particular, those that coordinate the agents in a MAS through social conventions.

It has been argued that the space of alternative behaviors (or convention space) is an important factor in the outcome of convention problems [8]. In other words, the number of possible conventions the agents can establish may influence the emerging convention(s) (if any). Therefore, mechanisms must be able to cope with convention spaces of different sizes.

In [15, 14], Salazar et al. propose, with some success, an evolutionary computing mechanism based on the notion of social contagion to dynamically emerge social conventions, the so-called *infection-based mechanism* (IBM). Nevertheless, their experiments are limited to study cases where the space of conventions agents can reach is somewhat small. With this in mind, this paper focuses on validating the IBM by evaluating it in large convention spaces. To this aim,

we have selected language conventions as our case study domain because: (i) it provides large convention spaces; and (ii) it is a relevant problem for MAS.

In multi-agent systems (MAS), communication is a key factor for agents to successfully interact with each other. In particular, when agents rely on explicit communication, a *shared language or vocabulary* (i.e. communication system) is highly necessary. Nevertheless, in open, heterogeneous MAS, where no central authority exists, such language may not exist. Since no one enforces a common language, agents may have their own, limiting their successful interactions to agents with a similar or the same language (if any exists).

In such MAS, agents may use different terms to refer to the same concept, or may use the same term to refer to different concepts, creating ambiguities in their communications. Therefore, a mechanism that allows agents to distributedly reach language conventions (consensus) that improve their communications is highly desirable. Furthermore, in an open MAS, establishing conventions with an *offline* process may not be reliable. Because, the MAS conditions can change with time (e.g. the number of agents, their objectives, the environment). Hence, the need for a mechanism that allows agents to reach language conventions at the same time they normally operate to achieve their (individual) goals.

From the social sciences perspective, language establishment is a highly studied topic [2, 11]. It has been argued that languages are established as a form of a social convention, thus the relationship between a word and a concept is dependent on the interactions between individuals. Several studies have addressed the modeling of such interactions as language games between individuals [12] [4], with various levels of success. These games model language construction at a purely semiotic level, i.e. they neglect the semantic relationships between models and symbols. A common game to study conventionalization is the *naming game* [18]. This game focuses on the interactions of speakers and listeners that try to find names for objects to understand each other. Thus, the aim of the naming game is to study how a common lexicon (vocabulary) is established in a society. Thus, for our purposes we apply this notions from a MAS perspective.

With the purpose of studying different types of large convention spaces, we evaluate two different cases. Firstly, the one created by having the same number of names and concepts (objects). This scenario, is likely to exhibit ambiguity (also called specificity) because there is a high probability that agents assign the same name to different objects. Thus, it represents a behavior space with small number of desired conventions (with respect to the number of possible ones). Secondly, if the number of names is (relatively) much larger than the number of objects, ambiguity may be less likely to occur, but conventions may be harder to reach, since the space of possibilities is larger.

Several studies show that the social structure of a population affects how a language emerges [13, 6, 10]. This motivates that we further explore how different complex networks, such as small-world [19] and scale-free [1], as underlying topologies of our MAS may influence the adaptation mechanism.

In this paper, as stated above, we propose applying the infection-based mechanism as a distributed adaptation mechanism to engineer the emergence of lexi-

conventions in a MAS. Firstly, we model agents' communication interactions as a particular type of naming game. To that end, we base ourselves on the model proposed in [9]. Next, we use the IBM to provide lexicon adoption and adaptation. Lastly, we empirically show that thanks to such distributed adaptation mechanism, agents in an open MAS distributedly manage to reach a common lexicon. Furthermore, the emergent lexicon provides the MAS with a so-called, *perfect communication system* [9]. Consequently, this also shows the usefulness of the infection-based mechanism in large convention spaces.

Additionally, we analyze the robustness of the approach in dynamical settings, by allowing new agents to join a MAS at any time and by dynamically changing its interaction topology. Incoming agents are equipped with their own lexicon, which may be different to the one the agents in the MAS have already agreed upon (if any). Thus, we observe that despite changes in the population our adaptation mechanism leads the agents in an open MAS to adapt to the existing lexicon conventions or not.

To summarize, the contributions of this paper are: (1) we show how to successfully apply the IBM to emerge communication in an open MAS; (2) we propose a MAS communication model based on the naming-game; (3) we empirically show that the IBM can cope with large convention spaces; and (4) we empirically show that thanks to IBM a consensus on a single, globally-shared lexicon is feasible despite the interaction topology and dynamic changes.

The paper is organized as follows. Section 2 characterizes an open MAS and defines the communication model agents employ to interact. Section 3 briefly reviews the infection-based mechanism [14] and discusses its application to our problem. Section 4 presents some empirical evaluations. Finally, in section 5 we draw some conclusions and set paths to future research.

2 The Communication Model

In this section we define and characterize the communication model that the agents employ to interact between themselves. We based the model on a well studied one, the *naming game*. Therefore, the first part of the section focuses on describing such game, as well as its shortcomings (from our perspective) if used in MAS. Whereas, the second part proposes a communication model to overcome the identified shortcoming.

The naming game [18] is one of the most used models for studying language evolution. This game consists of two agents, a speaker who utters the name of an object, and a listener who must understand it. Thus, a game is successful if the listener can understand the speaker. Moreover, the game is considered as adaptive if speaker and hearer manage to increase the success of their communications through repeated interactions.

Even though this kind of game is mostly used in linguistics to understand the principles of language evolution [12], it can be highly useful in open MAS for agents to agree upon the lexicon to share. Therefore, in a broad sense it can be regarded as a model for ontology sharing [17]. Nevertheless, currently the most

common formulation of the naming game presents some impractical characteristics to make it useful for open MAS. First of all, it allows agents to create any word to refer to a particular object, which may be unrealistic in MAS since the number of concepts to name can most likely be bounded beforehand. Secondly, it allows the existence of multiple words to refer to the very same object (synonymy), which may cause ambiguities or inconsistencies in the communication between agents. Moreover, the predominant naming-game formulation makes no distinction between the communication model and the communication development (language acquisition) algorithm, (i.e. they are inter-wove).

To take into account the above-mentioned issues, we propose a communication model based on the one described by De Jong et al. in [7]. De Jong’s model borrows some of the notions of the naming game and defines them for a MAS. Moreover, it makes a distinction between the interaction model and the communication development algorithm. Nevertheless, it still considers word creation. Therefore, we propose to replace word creation with a word selection (from a finite set), similar to a not commonly used variation of the naming game presented in [16]. Henceforth, we shall refer to this communication model as the *name-matching game*, whose mechanics we describe in what follows.

We shall consider an open MAS composed of a set of autonomous agents, Ag . No central authority exists to rule the agents and agents only work with local knowledge. Each agent, $ag_i \in Ag$ knows a set of concepts (be them, for instance, object, topics, actions) O , which it employs to communicate with the other agents in the MAS. Some or all concepts can be shared between different agents. We also consider that all agents share a finite set of words, W , which they employ to refer to the concepts they use. Thus, agents interact with each other by exchanging messages composed of words from W . To facilitate communication among agents, each agent has a lexicon, $L_i : O \rightarrow W$, which assigns an external representation to the concepts it needs to employ. Thus, each agent uses its own lexicon L_i to find the appropriate word that represents the concept about which it wants to communicate. Moreover, we restrict the lexicon in such a manner that only one entry per concept is permitted. Hence, it is not possible to assign more than one word per concept (synonymy). Finally, the decoding function, $D_i : W \rightarrow 2^O$, is used to translate a given word to its related concept.

The convention space of this problem, C , stands for the set that contains all possible agent language conventions. For the sake of simplicity we shall measure its size as the number of possible lexicons to which all agents may agree to at one point, namely $|C| = |W|^{|O|}$. Notice that this is a lower bound since it only considers conventions in which all agents share the same lexicon.

Additionally, we consider that interactions among agents in a MAS are restricted by an interaction topology. We model an interaction topology as a graph (Ag, E) , where $E \subseteq Ag \times Ag$, whose edges correspond to relationships (neighborhoods) between agents. If $(ag_i, ag_j) \in E$, then ag_i and ag_j are neighbors, and thus they can interact with each other. Since the kind of MAS we consider is open (agents join or leave at will), interaction topologies may change with time. Specifically, interactions between agents are pair-wise. Each interaction is a com-

munication between an agent playing the role of *speaker*, $s \in Ag$, and another one playing the role of *hearer*, $h \in Ag$, relating to a certain concept, $o \in O_s$.

Given some interaction topology, each agent uses the words in its lexicon to build messages that exchanges with its neighboring agents. The recipient of an agent’s message may understand a message or not. This directly depends on the degree of agreement on the lexicons of sender and receiver. Overall, the higher the agreement on lexicons, the higher the number of successful interactions (and hence the lower the amount of misunderstandings).

In a MAS context, this interaction framework models various communication situations, in particular petitions. For example, an agent s (speaker) requiring an object available to an agent h (hearer), requesting a service or task, sending instructions, etc. Within this setting, communication is successful if agent s obtains the object or service it requires; or if it perceives that agent h soundly performs the requested task or some instruction. In other words, the game is successful if both agents can match the same word to the same concept.

To summarize, the mechanics of the game that we propose are as follows: (1) Agent s selects a concept, $o_s \in O_s$; (2) Agent s uses its lexicon, L_s , to find the word, w , that refers to o_s ; (3) Agent s communicates w to agent h ; (4) Agent h uses its decoding function, D_h , to interpret w into a concept $o_h \in O_h$; (5) Agent h responds according to its understanding of o_h ; and (6) The game is successful if s is satisfied by h ’s response (i.e. if $o_s = o_h$).

Our aim will be that agents achieve a so-called perfect communication system [7], where the lexicon mappings between words and concepts are one-to-one. Thus, as with synonymy, a desirable lexicon should not exhibit polysemy (i.e. same word for multiple concepts). The presence of polysemy increases the possibility of *ambiguosness* in the message interpretation of the hearer agent. If a particular agent’s lexicon, assigns the same word to two (or more) concepts, at the moment of decoding a message relating to this word, the hearer will have trouble deciding which concept the word refers to. In MAS communications, the specificity of a word quantifies the degree to which it identifies a single concept (the higher the specificity the less ambiguous the word). Thus, from here on we shall measure the specificity of a lexicon as the percentage of words in the lexicon with specificity equal to one. Therefore, a lexicon with 100% of specificity represents a lexicon with one-to-one mappings guaranteeing perfect communication.

Now observe that the ratio between the number of available words ($|W|$) and concepts ($|O|$) to the agents in the MAS depicts scenarios with different degrees of specificity. Hence, if $|W| < |O|$ we obtain games where full understanding (a perfect communication system) is impossible because ambiguity is unavoidable (a 100% lexicon specificity cannot occur). If $|W| = |O|$, ambiguity is likely to happen, but lexicons with a 100% specificity are feasible. Thus, for a large enough number of concepts, the resulting convention space C is large, namely $|C| = |W|^{|O|}$. Nevertheless, only a small number of desirable conventions exist. Finally, when $|W| > |O|$ we obtain games where the likelihood to present ambiguity is low, but where misunderstandings are possible because of different lexicons

naming the same object with different words. Moreover, if $|W|$ is considerably large then the resulting convention space is also considerable large.

Notice that the proposed communication model solves to some extent the above-mentioned impracticalities by: i) bounding the number of available words to a particular set; ii) preventing synonymy through lexicon restrictions; and iii) decoupling the communication development algorithm from the communication model. Therefore, the remaining issue to deal with is polysemy. We tackle this issue with the aid of the communication development algorithm, whose aim will be to emerge a common lexicon with high specificity. The next section presents the algorithm we employ to accomplish this aim.

3 IBM for Communication Development

As noticed above, our main goal is to engineer the emergence of desirable lexicon conventions in an open, dynamic MAS. To succeed in this endeavor, we must guarantee that the agents in the MAS converge to a common lexicon. Not only that, because we pursue that convergence occurs despite changes in the agent population caused by the openness of a MAS. Therefore, we aim at endowing a MAS (section 2) with a distributed, adaptive mechanism that ensures a continuous convergence to a common lexicon despite changes in the agent population, hence promoting the development of a common communication system.

In what follows we propose to apply the infection-based mechanism described by Salazar et al. in [15] as a mechanism that promotes communication development. Furthermore, we also discuss how to deploy the infection-based model in an open MAS since it is an issue which is not tackled in [15].

The IBM is a distributed evolutionary algorithm that allows agents in a MAS to self-regulate through the collective emergence of social conventions. It is based on the social phenomenon of social-contagion [5], which relates to the spreading of behaviors/knowledge between individuals akin to an infectious disease.

However, the IBM considers the notion of *positive infection*: agents with *good behaviors / knowledge* that help improve the social welfare become more infectious to spread their behaviors (knowledge). In the context of our problem, an agent whose lexicon is highly unambiguous (close to or no polysemy) and/or highly agrees with its peers' is considered to have a *good lexicon* because it leads to successful communications. Therefore, the agent is more likely to spread his lexicon (infect other agents). An agent whose lexicon shows either a low degree of agreement with its peers' or a considerable number of ambiguous words, is bound to lead the agent to unsuccessful communications. Hence, this agent's lexicon can be regarded as a *bad lexicon*, that should not spread and be positively infected by some agent with a good lexicon instead of being spread. The following subsection details how the infection-based mechanism was implemented.

3.1 Implementation

Here we propose to deploy the infection-based mechanism in an open MAS by embedding an infection-based (adaptation) module within each agent as de-




Fig. 1. Infection-based module embedded into agents.

depicted by figure 1 (squared boxes within agents labeled as IBM). These adaptation modules collaborate to run a distributed evolutionary algorithm (detailed below) to continuously improve their lexicons. Within this setting, the infection-based mechanism operates as follows. Each adaptation module (IBM) continuously monitors the success of the interactions of the agent where it is embedded, (figure 1). Recall that such interactions occur between agents linked by some interaction topology. Periodically, adaptation modules synchronize to fire a cooperative evolutionary process aimed at improving their agents' lexicons. More precisely, adaptation modules record the results of a fixed number of interactions, namely language games (such as described in section 2). We refer to such number as the *incubation time* ($t_{incubation}$). This parameter bounds the time period for adaptation modules to monitor the success that the lexicons in use bring to their agents.

Once the incubation time expires, each infection-based module starts its adaptation process. It employs two strategies to improve the lexicon of each agent: *spreading* it among the agent's neighbors; and *innovating* it to prevent stagnation. As outlined above, the mechanism gives agents a higher infectious power the higher their success in their interactions. In the context of our problem, agents with a good lexicons (low ambiguity and already adopted by other agents) should have a higher chance of infecting other agents (via their lexicons), thus following the *survival of the fittest* approach.

In order to assess an agent's lexicon success, each adaptation module requires an **evaluation function**. This function allows to value the results of the language games an agent has been engaged in during the incubation time. Moreover, each adaptation module also requires a secondary communication channel. Such channel needs to be specialized for the adaptation module, i.e. independent of the (primary) communication channel used by each agent (see figure 1). Through this secondary channel, adaptation modules realize the spreading (infection attempts) by exchanging information related to their lexicons and their usefulness in preventing misunderstandings.

Henceforth, although the operations described below are run inside agents' adaptation modules, for the sake of simplicity we refer to the agents hosting them. Then, since each agent receives infection attempts from multiple neighbors, a **selection process** is ran by the agent to determine the infecting peer, based on its fitness (namely, on the evaluation of the peer's lexicon). The selec-

tion operator is implemented as detailed in [15] by adapting the roulette selection in the classic GA literature [3] to make it decentralized.

Once an agent has chosen a peer, it runs an **infection process** (with probability $p_{infection}$) is ran, to have its lexicon injected with part of the infecting agent’s lexicon (both lexicons are represented as genes). The infection process is implemented as a classic crossover recombination. The classic crossover (*single-cut crossover*) randomly selects a cut point in the parents’ gene sequences to exchange their genes and produce two new individuals. Consider a contagious agent and an agent to infect as two parents. Instead of creating child individuals, an infection operator combines the genes of both parents. Furthermore, there is no restriction on the number of agents each agent can infect (per iteration), but no agent can be infected twice. Therefore, the fittest agents enjoy more opportunities to spread.

Finally, each agent also runs an **innovation process** that randomly changes (with probability $p_{mutation}$) the word assigned to a concept with one word out the set of words W . The infection-based mechanism is outlined in algorithm 1. Importantly, our algorithm runs distributedly: each agent decides whether to infect or mutate based on its local knowledge. The interactions to which algorithm 1 refers are language games (line 2). As mentioned above, agents in a MAS continuously engage in language games until the incubation time expires. Thereafter, all agents locally start their evolutionary processes. Once this process finishes, agents resume their interactions.

<pre> 1: repeat 2: Let the agent interact for time $t_{incubation}$; 3: $ag.evaluate()$; 4: $ag.sendInfectionAttempts()$; 5: $ag' \leftarrow ag.selection()$; 6: $ag.infection(ag', p_{infection})$; 7: $ag.innovation(p_{innovation})$; 8: until MAS stops </pre>
--

Algorithm 1: Infection-based Algorithm.

Notice that a very important feature of the infection-based mechanism is that it allows to interweave agents’ interactions to achieve their goals with their lexicon adaptation. In other words, adaptation occurs at the same time that agents operate.

4 Empirical Evaluation

We hypothesize that the adaptation mechanism detailed in section 3 can be applied to emerge conventions in MAS with large convention spaces. To that end, we use communication development, namely the agreement on lexicon conventions, as our experimental domain. This domain, has the potential of exhibiting an interesting large convention space (depending on the number of words and concepts). Therefore, we shall consider our hypothesis as solved if the IBM allows a MAS to emerge a global (near-) perfect communication system.

Given an open MAS whose agent communication interactions are modeled as a language game, we shall consider the mechanism as successful if it allows: (i) agents in the MAS to reach lexicon convention(s) with a high level of specificity, for different word/concept ratios and under the most common interaction topologies; and (ii) conventions can be maintained despite changes in the agent population and in the underlying interaction topology.

At the aim of validating these hypothesis, we designed various sets of experiments to empirically evaluate the mechanism for communication development (described in section 3) under different conditions of a particular MAS. Next, in subsection 4.1 we describe the interaction topologies that we employed. In subsection 4.2 we detail our empirical settings, and in the rest of subsections we present and discuss our empirical results.

4.1 Interaction Topologies

It has been argued that the social distribution of individuals is an important factor in the evolution of languages [13, 6, 10]. This distribution is modeled in our MAS by the underlying interaction topology. Thus, in order to empirically analyze the potential of the infection-based mechanism as a tool for lexicon evolution we chose the following interaction topologies:

Small-world These networks present the small-world phenomenon, in which nodes have small neighborhoods, and yet it is possible to reach any other node in a small number of hops. This type of networks are *highly-clustered* (i.e. have a high clustering coefficient). Formally, we note them as $W_V^{k,p}$, where V is the number of nodes, k the average connectivity, i.e., the average size of the node’s neighborhood, and p the re-wiring probability. We used the Watts & Strogatz model [19] to generate these networks.

Scale-free These networks are characterized by having a few nodes acting as highly-connected hubs, while the rest of them have a low connectivity degree. Scale-free networks are *low-clustered* networks. Formally we note them as $S_V^{k,-\gamma}$, where V is the number of nodes and its degree distribution is given by $P(k) \sim k^{-\gamma}$, i.e. the probability $P(k)$ that a node in the network connects with k other nodes is roughly proportional to $k^{-\gamma}$.

Notice that we discard to consider random networks because they rarely appear in actual-world networked systems.

4.2 Experimental Settings

Each *experiment* consists of 50 discrete event simulations, each one running up to 120000 time-steps (ticks). Each simulation runs with 1000 agents using one of the underlying topologies defined in section 4.1. At the beginning of each simulation, each agent uploads a random lexicon. During each simulation, at each time-step agents interact through communication, as defined in sections 2, with a randomly selected neighbor. The interactions occur by agents randomly choosing some concepts to send.




Fig. 2. Results of the name-matching game

The individual understanding, evaluation function, of each agent is measured as the number of times it has engaged in a successful communication as a speaker. This measure is reset after each incubation period in the infection-based algorithm (algorithm 1), namely once the interaction period is over.

We generate interaction topologies for the simulations as small-world and scale-free networks by setting the following parameters: $W_{1000}^{<10>,0.1}$ and $S_{1000}^{<10>,-3}$. The clustering coefficients of the topologies are 0.492 and 0.056 respectively. Notice that we generate a new interaction topology per simulation.


As to the parameters of the infection-based mechanism, we set them as follows: $p_{infection} = 0.65$, $p_{innovation} = 5 \times 10^{-5}$, and $t_{incubation} = 10$. This setting of the incubation time means that we require a low number of interactions before adapting lexicons. In other words, we can consider a continuous adaptation.

In order to observe the effect of the IBM over a MAS we probe simulations in two ways. On the one hand, to measure whether a lexicon convention is adopted, we observe the number of agents that share each lexicon per tick. We shall refer to the lexicon shared by the largest number of agents as the *dominant* lexicon. On the other hand, we also observe at every tick the *quality* of such lexicon. Given a lexicon its quality is determined by its specificity, namely the percentage of words that represent a single concept. For both dominant lexicons and specificity, we aggregate the measures obtained after 50 simulations using the inter-quartile mean.

We designed different sets of experiments to empirically validate our initial hypothesis regarding the validity of the infection based mechanism. The first set, analyzed in section 4.3, aims at showing that the IBM can guide a MAS to emerge a lexicon with high specificity in a scenario where the likelihood of ambiguity is high ($|W| = |O|$). Next, in section 4.4 we also test the case when the likelihood to present ambiguity is low (but misunderstandings are feasible), namely when $|W| > |O|$. Finally, in section 4.5 we study a dynamic setting where both the agent population and interaction topology change over time to test the robustness of the IBM.

4.3 Matching Game with same Words and Concepts

The aim of this section is to show that the infection-based mechanism (section 3) can guide a MAS, whose interactions are modeled as a name-matching game, to



(a) Small-world

Fig. 3. Results of a particular small-world simulation

emerge a lexicon with high specificity. Furthermore, we particularly focus on the case where the existing number of words is equal to the number of objects ($|W| = |O|$). Recall that this scenario is likely to promote the existence of lexicons with low-specificity because given a limited number of words an agent can probably assign the same word to more than one concept.

We set the number of concepts to 10, and the set of words, W , is also composed of 10 different words ($|C| = 10^{10}$). Thus, to prevent ambiguity a lexicon must manage to match each one of the ten concepts to a different word.

Figure 2 shows the evolution of the percentage of agents sharing the dominant lexicon along with the specificity of the lexicon. We observe that, on average, for a small-world topology the dominant lexicon convention exhibits a smooth and slow growth. Nevertheless, at least 80% of the agents reach a consensus regarding their lexicon. The figure also shows that the specificity of the dominant lexicon quickly reaches 90%, which means that one of the words matches two concepts. However, we observe that in the long run there is a trend towards total specificity.

As to scale-free networks, a dominant lexicon shared by a large number of agents emerges in a fast and sharp manner. This dominant lexicon almost immediately encompasses 90% of the agents with a 90% specificity. This state is maintained for some time, but at some point (around 60000 ticks) reaches beyond 95% of the population and the quality or the lexicon improves.

Moreover, observing a particular simulation provides some interesting insights. Figure 3 shows one of the 50 simulations performed for the small-world topology. In this plot, the transition towards a lexicon with 100% specificity is clearer. If we analyze what happens before the lexicon transition, we observe that the agent population sharing the dominant lexicon decreases (around 30000 ticks). This occurs because a lexicon with maximum specificity appears and starts pulling members out of the dominant lexicon at that point. By the 35000 time-step it is able to overtake the previously dominant lexicon and rapidly reaches a global consensus (almost all the agent population shares it).

Overall, the results of these experiments show that the IBM emerges a perfect communication system for more than 90% of the population. It is called perfect because: i) all the words in the dominant lexicon have a high specificity; ii) once established the lexicon becomes consistent overtime; and iii) the agents' communicative interactions are always successful.




Fig. 4. Results of the 2-names-matching game with $|W| = 100$ and $|O| = 10$.

4.4 Matching Game with Many Words

In the previous section we focused on experiments with high potential ambiguity, where there were as many concepts as words. Now we turn our attention to the other extreme of the spectrum, where the number of words is significantly greater than the number of concepts. In this case, the likelihood of ambiguity decreases when the number of words increases. This happens because assigning the same word to more than one concept is less probable. In other words, the number of lexicons with 100% specificity increases.

For this set of experiments we set $|W| \in \{20, 100\}$, whereas we only used 10 concepts ($|C| = 20^{10}$ and $|C| = 100^{10}$). As to the extreme case ($|W| = 100$) for a small-world topology, we observe in figure 4(a) that two (dominant) lexicon conventions with total specificity rapidly appear (< 100 ticks). One of them starts gaining members rapidly ($\sim 89\%$) while the other one appears to stabilize with a small percentage of the population ($\sim 11\%$). Since the second (small) agent group has a lexicon with 100% specificity and its communications are most likely occurring between its members, it can withstand infections coming from the dominant group. However, because the topology is a small-world scenario, members of this group may need to communicate with agents sharing the dominant lexicon. Hence, around 1000 ticks, time at which the dominant lexicon is shared by most of the population, the membership of the second group starts swaying towards the biggest group out of necessity. As times goes on (beyond what we show in the plot), the dominant lexicon settles in 98 % of the population. On scale-free topology see figure 4(b), we observe a more straightforward behavior. Almost immediately (~ 2000 ticks) a lexicon convention with total specificity is shared by the whole population, swaying over any secondary lexicon intending to rise.

As to the results with $|W| = 20$, on the small-word topology we observe the same behavior as the one presented in section 4.3 for $|W| = 10$. However, there are significant improvements regarding scale-free networks because there is a larger number of lexicons with high specificity.

4.5 Dynamic Population

In previous sections we showed, through multiple scenarios, the usefulness of the infection-based mechanism to engineer the emergence of lexicon conventions.




Fig. 5. Evolution of a dominant lexicon for a dynamic agent population over a scale free topology.

Nevertheless, all scenarios were somewhat static, in the sense that the MAS did not change with time. Hence, the purpose of our last experiments is to show the robustness of the IBM in a MAS that changes with time.

Thus, we model the open dynamics of a MAS by allowing the agent population and their neighborhoods to change over time. In practice, environment changes are achieved by dynamically changing the network topology. Hence, we proceeded as follows: 1) we create a scale-free network interaction topology up to certain number of agents; 2) we let agents interact over the interaction topology; and 3) after l simulation ticks, we introduce new agents along with their neighborhoods. Finally, we implemented the open dynamics of a MAS by interweaving the Barabasi-Albert (BA) scale-free network generation algorithm [1] with the MAS simulation. In other words, we ran the MAS and the BA algorithm at the same time.

We set the parameters for the experiments as follows. The starting MAS employed a scale-free underlying topology of 400 agents ($S_{400}^{<10>,-3}$). Every 400 iterations, 20 new agents with random lexicons joined the agent population ($l = 400$). The MAS continued growing until reaching the $S_{5746}^{<10>,-3}$ topology (5746 agents).

On average (inter-quartile mean), we can report that IBM helps agents rapidly converge to a lexicon, and incoming agents promptly join the dominant convention. Nevertheless, this initial dominant lexicon presented some ambiguity that diminished with time. Thus, at some point a lexicon with total specificity was found. We illustrate this claim with a particular example.

Figure 5 shows a simulation (a single run instead of the aggregation) where a dominant lexicon is quickly established. Notice that the dominant lexicon grows at almost the same rate as the population (left-hand y-axis). However, the dominant lexicon has an initial lexicon specificity percentage of $\sim 70\%$ (right-hand y-axis). Observe that as time goes on the lexicon specificity starts to improve. Furthermore, notice that improvements in the quality of the global convention lexicon (specificity) seems to be usually preceded by a decrease in the convention size (e.g. ~ 10000 , ~ 15000 and ~ 60000 ticks). These downward spikes mark the agent population's transition to better lexicons. Nevertheless, observe that some of the spikes do not improve the unambiguity. This also marks a change in the current dominant lexicon, but one that did not improve the specificity.

5 Conclusions and future work

In this paper we proposed the use of a recently studied distributed evolutionary algorithm as a viable tool to emerge conventions in multi-agent systems with large convention spaces.

We accomplish this by tackling communication development in MAS, which is a known problem. Through the infection-based mechanism it is possible to engineer the emergence of lexicon conventions in open heterogeneous MAS. The infection-based mechanism allows agents to distributedly converge to a lexicon with high specificity that guarantees the success of their communications, i.e. a perfect communication system. Even more, through this approach lexicon consensus are attainable despite the presence of dynamic changes in the MAS. In particular, changes in the agent population and its underlying topology (e.g. available communication links between the agents).

We apply the infection-based mechanism described in [15] as a communication development mechanism. Moreover, we make use of language games as a model for communication interactions. Hence, the agents' normal interactions are valued as successful or failed. Through these two elements our approach allows agents to find and improve lexicon conventions while they normally interact in the MAS. Nevertheless, notice that different communication models can be employed without changing the IBM.

We ran several experiments to provide empirical evidence of the capabilities of the proposed approach under different (reasonable) circumstances. Furthermore, because it is known that language emergence is influenced by the social structures, our experiments took this into account. Namely, through the use of the most common complex networks as underlying interaction topologies.

On the one hand, the results show that we can direct the emergence of a global lexicon consensus regardless of the topology. Nevertheless, the clustering of the topology affects the time it takes to reach such consensus. Therefore, in contrast with other approaches [6], the IBM reaches a single-vocabulary consensus, in the form of a globally shared lexicon. Moreover, this is accomplished in a largely populated multi-agent systems.

On the other hand, the results also helped to gauge the effect intrinsic to the relation in the number of words and concepts. When the number of words and concepts is equal, some difficulty exists in establishing an unambiguous lexicon, but conventions between agents may be reached faster. However, as the number of words increases, finding a lexicon with 100% specificity becomes somewhat easier, but reaching a common to most agents becomes harder. Nevertheless, despite this it is still possible to agree on a single lexicon global consensus.

Our final experiments showed that our adaptation mechanism can reach an unambiguous lexicon convention even if the agent population in the MAS and its interaction topology are constantly changing. Therefore, our approach shows to be valid even in highly dynamic scenarios.

To summarize, we empirically show that the IBM can reach desirable global conventions even when the convention space is large. Furthermore, this can also

be accomplished when the number of desired convention is considerably a small fraction of the convention space.

Finally, as to future work we plan to look to more dynamic settings. For instance, scenarios where agents can introduce new concepts at run-time. Moreover, we also plan to study the effect produced by each agent having a different set of available words.

References

1. R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47, 2002.
2. R. Axelrod. The dissemination of culture: A model with local convergence and global polarization. *The Journal of Conflict Resolution*, 41(2):203–226, 1997.
3. T. Bäck. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, 1996.
4. A. Baronchelli, M. Felici, V. Loreto, E. Caglioti, and L. Steels. Sharp transition towards shared vocabularies in multi-agent systems. *Journal of Statistical Mechanics: Theory and Experiment*, 06:P06014, 2006.
5. R. Burt. Social contagion and innovation: Cohesion versus structural equivalence. *American J. of Sociology*, 92:1287–1335, 1987.
6. L. Dall’Asta, A. Baronchelli, A. Barrat, and V. Loreto. Nonequilibrium dynamics of language games on complex networks. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 74(3):036105, 2006.
7. E. D. de Jong and L. Steels. A distributed learning algorithm for communication development. *Complex Systems*, 14:315–334, 2003.
8. B. DeVylder. *The Evolution of Conventions in Multi-Agent Systems*. PhD thesis, Artificial Intelligence Lab Vrije Universiteit Brussel, 2007.
9. E. D. D. Jong and L. Steels. A distributed learning algorithm for communication development. *Complex Systems*, 14:2003, 2003.
10. A. Kalampokis, K. Kosmidis, and P. Argyrakis. Evolution of vocabulary on scale-free and random networks. *Physica A: Statistical Mechanics and its Applications*, 379(2):665–671, Jun 2007.
11. D. Lewis. *Convention*. Cambridge: Harvard University Press., 1969.
12. V. Loreto and L. Steels. Social dynamics: Emergence of language. *Nature Physics*, 3:758–760, November 2007.
13. A. Puglisi, A. Baronchelli, and V. Loreto. Cultural route to the emergence of linguistic categories. *Proceedings of the National Academy of Sciences*, pages 0802485105+, June 2008.
14. N. Salazar, J. A. Rodriguez-Aguilar, and J. L. Arcos. An infection-based mechanism for self-adaptation in multi-agent complex networks. In S. Brueckner, P. Robertson, and U. Bellur, editors, *Proc. of the Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, pages 161–170, 2008.
15. N. Salazar, J. A. Rodriguez-Aguilar, and J. L. Arcos. Infection-based self-configuration in agent societies. In *Proc. of the 2008 GECCO conference companion on Genetic and evolutionary computation*, pages 1945–1952. ACM, 2008.
16. L. Steels. Self-organizing vocabularies. In C. G. Langton and K. Shimohara, editors, *Artificial Life V*, pages 179–184, Nara, Japan, 1996.
17. L. Steels. The origins of ontologies and communication conventions in multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 1:169–194, 1998.
18. L. Steels. Grounding symbols through evolutionary language games. In *Simulating the Evolution of Language*, chapter 10, pages 211–226. Springer Verlag, 2002.
19. D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, June 1998.