Structure and Variation of Signaling Conventions in Scale-free Networks

Roland Mühlenbernd

Eberhard Karls Universität Tübingen

Abstract. The *signaling game* is a game-theoretical model that can depict the dynamics enabling the emergence of semantic meaning that establishes as convention among members of a society. The signaling game alone describes a communicative situation, but by applying it as a repeated game and combining it with update dynamics, a path of the evolution of semantic meaning can be simulated. In this work I will i) combine the repeated signaling game with the update dynamics *innovative reinforcement learning*, and ii) conduct it on a society of agents arranged in *scale-free networks*. The simulation runs show not only that multiple regional conventions emerge and stabilize, but also reveal the way these regions are arranged and interact with each other.

Keywords: repeated signaling game, innovative reinforcement learning, scale-free networks, regional conventions, emergence of semantic meaning

1 Introduction

Semantic meaning of a language's expressions is in the least frequent cases a result of agreements among a speech community, but rather a product of regularities in communicative behavior. These regularities form semantic conventions, which – once emerged – are stable to a specific degree.[1] To formalize this process of the emergence of semantic meaning, Lewis invented a game-theoretical model, the *signaling game* [13]. But he took only a first step by working out the stability conditions for semantic meaning. Subsequent work analyzed the possible paths that lead to conventions of semantic meaning by considering a repeated version of the signaling game, combined with simple update mechanisms to adjust subsequent behavior of the game's participants [20].

As a matter of fact, most studies with signaling games consider only two-players setups, and therefore neglect the fact that an essential feature of a convention might be its emergence inside a population that contains more than two members. In the last decade, a number of studies addressed the question of how conventions of semantic meaning arise in realistic population structures; by applying repeated signaling games between connected agents placed in network

[1] Of course, forces of language change might shift the manifestation of semantic conventions, in form and in meaning as well, but generally they reveal at least a temporal stability.

structures: from lattice structures [27][14] to small-work networks [24][16]. All these studies revealed that for a multitude of circumstances a number of different regional conventions emerges. Nevertheless, more detailed analyses of how these conventions position themselves and interact with each other in a social network structure are still missing, and this study wants to make such a contribution.

In the last two decades a quite large number of studies emerged that incorporate the analysis of the behavior of agents that communicate via the *naming game* in diverse population structures to reach semantic conventions [21, 22, 4]. This game is in many aspects quite similar to a signaling game, since in both games there is a sender that encodes an information state with a message, and a receiver that decodes the message with an interpretation state.[2] One crucial difference between the signaling game and the standard naming game is the fact that the latter one has a built-in update rule: i.a. if communication is successful by sending a message m, the receiver keeps only the appropriate interpretation for m in his memory, and eliminates all possible alternative interpretations. In contrast, the signaling game does not have a built-in update rule, and the one applied in this study – reinforcement learning – does not eliminate alternatives, but only decreases their probability to be chosen in subsequent rounds.

Another line of research analyzed the emergence of conventions and norms by applying a repeated *coordination game* combined with reinforcement learning in social network structures [1, 19, 23]. These studies differ at least in one fundamental aspect: by applying a coordination game, they analyze conventions in terms of coordinated acting, not conventions in terms of communication.[3] Therefore, those studies' resulting phenomenon is a *behavioral convention*, whereas in this study a *semantic convention* is the phenomenon under examination.[4] Finally, different semantic conventions in terms of so-called *signaling systems* can be compared in degrees of similarity, a property that is also generally missing in behavioral conventions.

With this study I want to contribute to the line of research that deals with the question of how semantic conventions emerge and position themselves in a society. I use repeated signaling games in combination with an update mechanism that is called *reinforcement learning* [18]. Applying reinforcement learning on repeated signaling games is a popular combination in this field [20][5][6]. For my purposes I use a specific version of this learning model that is called *innovative reinforcement learning* [20][2]. As a previous study showed, this model supports the alternating dynamics of interacting regional conventions of a society [17]. In this article I will i) conduct simulation runs on so-called *scale-free networks*, and ii) analyze the way regional conventions are structured and vary, and especially

[2] In the literature of the naming game, information/interpretation states are often referred to as *concepts*, messages as *words*.

[3] Admittedly, a signaling game's basis is the utility table of a simple coordination game, and it is fair to say that both players' goal is to coordinate. But they do it in a *mediate* way by sending a message, whereas in a simple coordination game players try to coordinate in an *immediate* way.

[4] Another difference between coordination game and signaling game is, that the former is a symmetric, the latter an asymmetric game.

how neighboring regions interact with each other. The results will give new insights not only into the mechanisms of language evolution, but also e.g. into the mechanisms of language contact and its influence on semantic shift.

The article is divided in the following way: in Section 2, I will introduce some basic notions of repeated signaling games, innovative reinforcement learning and network theory. In Section 3, I will present the simulation experiments, their results and appropriate analyses of the data. In Section 4, I will give a final conclusion and discuss open questions.

2 Signaling Games, Learning and Networks

This section will give a coarse technical and theoretical background of the important concepts of this article: the signaling game, innovative reinforcement learning and network theory.

2.1 Signaling Games

A signaling game $SG = \langle \{S, R\}, T, M, A, Pr, U \rangle$ is a game played between a sender S and a receiver R. T is a set of information states, M is a set of messages and A is a set of interpretation states. $\Pr(t) \in \Delta(T)$ is a probability distribution over T and describes the prior probability that a state t is topic of communication.[5] $U : T \times A \to \mathbb{R}$ is a utility function that determines how well interpretation state and information state both correspond to each other.

A play of the game looks as follows: the so-called invisible player, call it destiny or nature, picks one of the states $t \in T$ and puts it into the sender's mind.[6] The sender wants to communicate this state by choosing a message $m \in M$ that she signals to the receiver. At this point the receiver knows the message that has been signaled, whereas the information state is concealed to him. In face of the given message m the receiver picks an interpretation state $a \in A$ with the purpose to match as good as possible with information state t; or in other words: to maximize the utility $U(t, a)$. Since utilities are identical to both participants, also the sender has an interest that the receiver decodes successfully. All in all, both interlocutors have an interest to perform an act of perfect information transmission, of perfect communication.

In this study I only consider signaling games that i) have a flat probability function, $\forall t \in T : Pr(t) = \frac{1}{|T|}$, and ii) have the same number of information states and interpretation states, thus $|T| = |A|$. Furthermore, each information state has a corresponding interpretation state, that is marked by the same index, c.f. t_i corresponds to a_i. Finally, this correspondence is valued by the utility function: if both states correspond to each other, the utility for both players is

[5] $\Delta(X) : X \to \mathbb{R}$ denotes a probability distribution over random variable X.

[6] Note that in this kind of game the function Pr determines how probable nature picks an information state; i.o.w. the probability of state t being in the sender's mind is always $Pr(t)$.

$$L_1: \begin{tikzcd} t_1 \arrow[r] & m_1 \arrow[r] & a_1 \\ t_2 \arrow[r] & m_2 \arrow[r] & a_2 \end{tikzcd} \qquad L_2: \begin{tikzcd} t_1 \arrow[dr] & m_1 \arrow[dr] & a_1 \\ t_2 \arrow[ur] & m_2 \arrow[ur] & a_2 \end{tikzcd}$$

Fig. 1: The two signaling systems of the 2×2-game.

a positive value $\alpha \in \mathbb{R}$; if not, it is zero or a negative value $\beta \in \mathbb{R}$. This kind of signaling game is called $n \times k$-game and defined as follows:

Definition 1 **($n \times k$-game).** *A $n \times k$-game is a signaling game SG with:* $|T| = |A| = n$, $|M| = k$, $\forall t \in T : Pr(t) = 1/|T|$ *and* $U(t_i, a_j) = \begin{cases} \alpha > 0 \text{ if } i = j \\ \beta \leq 0 \text{ else} \end{cases}$

Note that messages are initially meaningless in this game, since the utility function rewards a correspondence between information state and interpretation state, whereas the chosen message does not affect the reward. In other words, each message can be used to communicate a corresponding state pair $\langle t_i, a_i \rangle$. The meaningfulness of a message m is stated if it communicates a particular state; and this meaningfulness can arise from regularities in behavior. Behavior is defined in terms of strategies. A *behavioral sender strategy* is a function $\sigma : T \to \Delta(M)$, and a *behavioral receiver strategy* is a function $\rho : M \to \Delta(A)$. A behavioral strategy can be interpreted as a single agent's probabilistic choice.

If the game per se does not entail meaningfulness of messages, then what circumstances can tell us that a message is attributed with a meaning? The answer is: this can be indicated by the combination of sender and receiver strategies, called strategy profile: a message has a meaning between a sender and a receiver, if both participants use *pure strategies* that constitute a specific *isomorphic strategy profile.*

Note that for a 2×2-game, there are exactly two such strategy profiles, as depicted in Figure 1. Here in profile L_1 the message m_1 has the meaning of state pair $\langle t_1, a_1 \rangle$ and message m_2 has the meaning of state pair $\langle t_2, a_2 \rangle$. For profile L_2 it is exactly the other way around. Lewis called such strategy profiles *signaling systems* [13], which have interesting properties: they i) define a meaning of a message, ii) ensure perfect communication, iii) are *Nash equilibria* over expected utilities [8], and iv) are *evolutionary stable states* [25], and therefore attractors under many prominent evolutionary dynamics (e.g. replicator dynamics) [11].

A signaling system shows how semantic meaning can be expressed by participants' communicative behavior: a message has a meaning, if sender and receiver communicate according to a signaling system. But it does not explain, how participants come to a signaling system in the first place, by expecting that messages are initially meaningless. To understand the evolutionary paths that might lead from a meaningless to a meaningful message, it is necessary to explore the process that leads from participants' arbitrary communicative behavior to a behavior that constitutes a signaling system. Such a process can be simulated by repeated signaling games, where the participants' behavior is guided by *update dynamics.* One popular dynamics is called *reinforcement learning* [20][5][6].

2.2 Reinforcement Learning

Reinforcement learning can be delineated by an urn model, known as *Pólya urns* [18]. A single urn models a behavioral strategy, in the sense that the probability of making a particular decision is proportional to the number of balls in the urn that correspond to that choice. By changing the content of an urn after each access, an agent's behavior can be gradually adjusted. Such an urn model can be combined with a repeated signaling game in the following way: the sender of the game has an urn $\mho_t$ for each state $t \in T$, which contains balls for different messages $m \in M$. Similarly, the receiver of the game has an urn $\mho_m$ for each message $m \in M$, which contains balls for different states $a \in A$.

The behavioral strategy of sender and receiver as well can be described in terms of *response rules* that describe a participant's situation before making the next move. To be concrete, let's designate i) the number of balls of type m in sender urn $\mho_t$ with $m(\mho_t)$, and the overall number of balls in sender urn $\mho_t$ with $|\mho_t|$; and ii) the number of balls of type a in receiver urn $\mho_m$ with $a(\mho_m)$, and the overall number of balls in receiver urn $\mho_m$ with $|\mho_m|$. The resulting *sender response rule* σ and *receiver response rule* ρ is given in equations 1 and 2.

$$\sigma(m|t) = \frac{m(\mho_t)}{|\mho_t|} \qquad (1) \qquad\qquad \rho(a|m) = \frac{a(\mho_m)}{|\mho_m|} \qquad (2)$$

The actual learning dynamics is realized by changing the urn content dependent on the communicative success: if communication via t, m and a is successful, the number of balls of type m in urn $\mho_t$ and the number of balls of type a in urn $\mho_m$ each is increased by the positive utility value $\alpha \in \mathbb{R}$. In this way successful communicative behavior is more probable to reappear in subsequent rounds, successful behavior is reinforced.

This mechanism can be intensified by *lateral inhibition*: if communication via t, m and a is successful, not only will the number of ball type m in urn $\mho_t$ be increased, but also will the number of all other ball types $m' \in M \setminus \{m\}$ be decreased by $\gamma \in \mathbb{R}$. Similarly, for the receiver. Franke and Jäger [9] introduced the concept of *lateral inhibition* for reinforcement learning in signaling games and showed that it leads the system more speedily towards pure strategies.

Furthermore, *negative reinforcement* can be used to punish unsuccessful behavior. It changes urn contents in the following way: if communication via t, m and a is unsuccessful, the number of balls of type m in the sender urn $\mho_t$, and the number of balls of type a in the receiver urn $\mho_m$ each will be decreased by negative utility value $\beta \in \mathbb{R}$. In this way unsuccessful communicative behavior is less probable to reappear in subsequent rounds.

Note that reinforcement learning might have the property to slow down the learning effect: if the total number of balls in an urn increases over time, but the rewarding value α is a fixed value, then the learning effect mitigates. A way to prevent learning from slowing down is to keep the overall number of balls $|\mho|$ on a fixed value $\Omega \in \mathbb{R}$ by scaling the urn content appropriately after each round of play. Such a setup is a variant of so-called Bush-Mosteller reinforcement [7].

All in all, a reinforcement learning setup for a signaling game can be captured by $RL = \langle \alpha, \beta, \gamma, \Omega, \phi \rangle$, where α is the reward value, β the punishment value, γ the lateral inhibition value and Ω determines the constant urn size. Additionally, ϕ is a function that defines the initial urn settings. This standard account can be refined to an account that includes a mechanism for the invention of new messages, and is correspondingly called *innovative reinforcement learning.*

2.3 Innovative Reinforcement Learning

The basic idea of innovative reinforcement learning stems from Skyrms [20] and works as follows: each sender urn contains, next to the balls for each message, an additional ball type, which Skyrms calls *black ball.* Whenever the sender draws a black ball from one of her urns, she sends a new message that was never sent before. In other words, the sender invents a new message. Experiments with this setup were made for 2-players games [2] and multi-agent accounts[17].

The second study [17] used a setup with negative reinforcement and lateral inhibition. It showed that larger populations will probably never find an agreement (or at least need an unmanageable amount of runtime). The reason is given by the fact that the number of possible messages is verbatim unlimited and populations end up in a chaos of a never ending production of new messages. This phenomenon was shown even for a little community of 6 agents. But it could also be shown that by limiting the possible message set, this problematic nature of a never ending chaos can be avoided [17]. In such a game, when a sender draws a black ball, she doesn't send a completely new message, but sends a randomly chosen message from the given message set. By stating a message set that is substantially larger than the number of states, the game keeps its innovative nature, but avoids runtime problems.[7]

2.4 Network Structure and Network Properties

To ensure that a network structure resembles a realistic interaction structure of human populations, it should have *small-world* properties: i) a short *characteristic path length*, and ii) a high *clustering coefficient* [26].[8] Additionally, most often human networks display a third property, namely to be *scale-free*: the frequency of agents with ever larger numbers of connections roughly follows a power-law distribution. In this sense I consider a special kind of a scale-free network, which is both scale-free and has small-world properties [3]. This network type is constructed by a *preferential attachment* algorithm that takes two parameters m that controls the network density, and p that controls the clustering coefficient [10]. In my experiment I used a scale-free network with 500 nodes, $m = 2$ and $p = .8$, which ensures small-world properties.

[7] Note that to limit the message number supports the finding of a global convention. This is no surprise, since the less messages are given, the less signaling systems are possible: for an $n \times k$-game there are ${}^{k!}/_{(k-n)!}$ possible signaling systems.

[8] For the definition of these network properties I refer to Jackson's *Social and Economic Networks* [12], Chapter 2.

In the analysis of the simulation experiments I perused particular network properties, which I divide in *local properties* that describe characteristics of a node, and *global properties* that describe characteristics of a connected sub-network. There were two local network properties I was interested in: *degree centrality* that is determined by the number of a node's connections; and *individual clustering* that is determined by the number of connections among a node's neighborhood.[9] Furthermore, there are two global network properties: *transitivity* that can be interpreted as a measure of global clustering; and the *clustering coefficient* that is the average value of individual clustering among all nodes of a sub-network, and can be seen as a measure of local clustering.[8]

3 Simulation Experiment and Results

My experiment contained 10 simulation runs, each is arranged as follows: agents are placed in a social network structure; per simulation step they communicate by playing a signaling game with each of their connected neighbors and update their behavior by reinforcement learning. The concrete settings are as follows:

- *network structure*: a scale-free network with 500 agents (Holme-Kim algorithm [10] with $m = 2$ and $p = .8$)
- *signaling game*: a 3×9-game
- *reinforcement learning*: innovative reinforcement learning with $\alpha = 1$, $\beta = 1$, $\gamma = 1$, $\Omega = 20$, ϕ: all ball types are equiprobably distributed
- *stop condition*: reaching 100,000 simulation steps

The essential goal of my simulation runs is to analyze the structure and variation of what I will call a *convention region*. A convention region is a group of agents, which i) has only members that have learned the same convention, and ii) constitutes a connected sub-graph in the network.

3.1 The Structure of Regional Conventions

A basic result of all simulation runs was as follows: a number of different convention regions emerged. This result was already observed in previous studies with similar settings [27][14][24][16]. Figure 2 depicts the number of convention regions for the first 20,000 simulation steps, each line is the result for one of the ten simulation runs. As observable, the number of regions initially increases to a high value, and then decreases to a specific value n, and oscillates around this value. Depending on the particular simulation run, this value n is between 10 and 30. All simulation runs revealed that the number of regions continues oscillating around this value, up to the maximum number of 100,000 simulation steps. The fact that the number of convention regions has no long-term tendency of decreasing or increasing reveals a *long-term stability*, whereas the high amplitude of oscillation reveals a *short-term reactivity*.

[9] A node's neighborhood is the set of all nodes it is connected to.

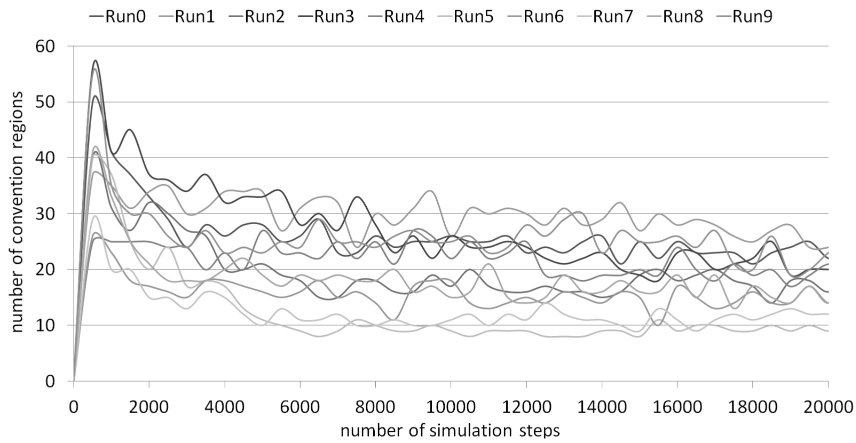

Fig. 2: The number of regional conventions over the first 20,000 simulation steps for all 10 simulation runs.

Another result of the simulation runs was the following: the global communicative success[10] also oscillates around a specific value between .75 and .95, depending on the simulation run. Apparently, despite the high number of multiple conventions, communication works quite well in this society. The reasons are i) that all members of their own convention region communicate perfectly successfully with each other, and ii) that neighboring regions often have similar conventions that guarantee at least partial success at the borders. To get an impression of such a diverse society, Figure 3a shows a resulting scale-free network with around 20 different convention regions, indicated by different colors.

It turned out that the conventions regions that evolved inside a network after 100,000 simulation steps differentiate highly in their size.[11] In each simulation run maximally one really large region emerged, and always a couple of medium size regions, and a high number of small regions. Figure 3b depicts the number of emerged convention regions for specific intervals of sizes, over all 10 simulation runs. As observable, the number of convention regions decreases with their size; c.f. while there evolved more that 50 convention regions with a size smaller than 50, there evolved only two convention regions with a size larger than 250.

A further, quite surprising result is the fact that there is a negative correlation between the size of a region and the degree centrality of its agents. This is clearly observable in Figure 4a, where each data point represents an agent's region size (x-axis) and her degree centrality (y-axis). This analysis reveals a negative Pearson correlation of -.405. Furthermore, by conducting the same analysis and considering only each region's agent with maximal degree centrality (as depicted in Figure 4b) the negative correlation is even stronger, namely -.634.

[10] The global communicative success is simply the utility value averaged over all conducted interactions at one simulation step,.

[11] A size of a region is determined by its number of members.

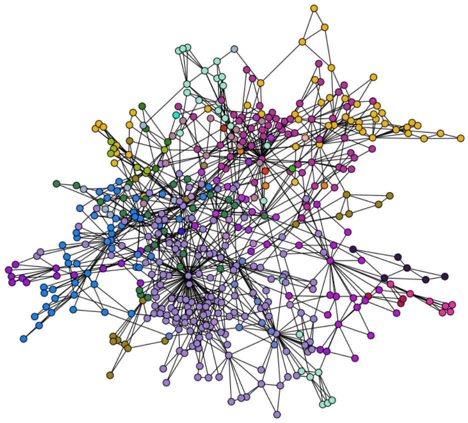

(a) exemplary scale-free network

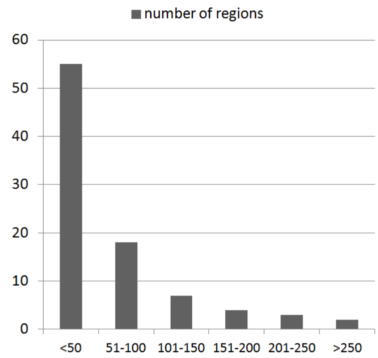

(b) histogram of region sizes

Fig. 3: Scale-free network with 500 nodes and around 20 convention regions (Figure 3a) and the number of convention regions for different intervals of size over all 10 simulation runs (Figure 3b).

This result says that agents with a high degree centrality are members of a small convention region, while large regions contain rather agents with a low degree centrality. This is a surprising result, since intuitively one would expect an outcome that is exactly the other way around; one would expect agents with a lot of neighbors of the same convention rather in a large convention region, and vice versa. One possible explanation is that agents with a high degree centrality tend to be more innovative and frequently create their own small society. A more detailed analysis that might explain this phenomenon is under study and will be part of a subsequent study.

3.2 The Local Interaction of Regional Conventions

Another issue of examination was the way convention regions change over time. As a basic result it turns out that most of the time regions have a rather fixed size that oscillates with a small deviation. But from time to time a process happens that I will call *spatial acquisition*: a convention region acquires a part of a neighboring region. Furthermore, all simulation runs reveal that spatial acquisition only happens between regions whose conventions have a high similarity and differ only in one used message. This result gives rise to the more general assumption that a particular degree of similarity between conventions is a necessary condition for realizing spatial acquisition at all. Further formal analyses are beyond the scope of this study and are left to future research.

In a more detailed analysis I wanted to figure out if particular global network properties of neighboring regions are supportive for the emergence of spatial acquisition among them. For that purpose I extracted all events that indicated a spatial acquisition between two convention regions and measured the transitivity

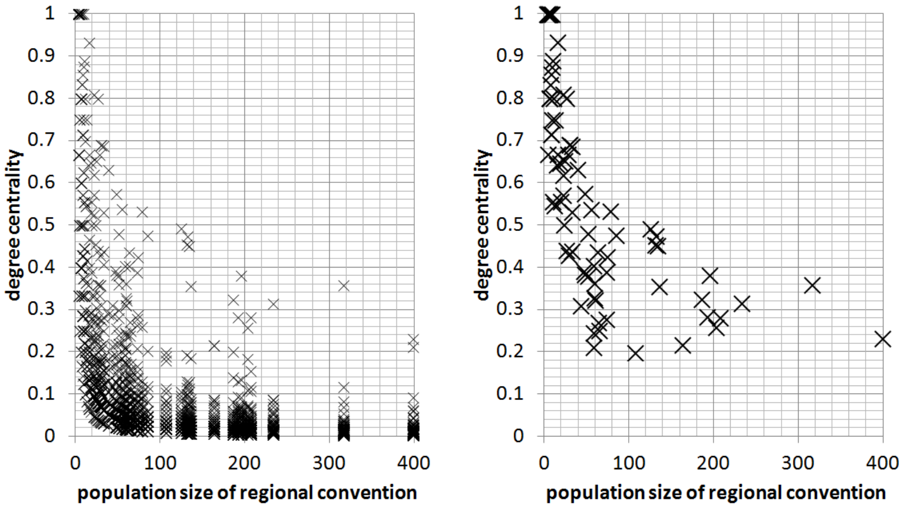

(a) all agents (b) agents with max. DC

Fig. 4: Data plot of agent's degree centrality (y-axis) in comparison to their region's size (x-axis). Figure 4a depicts all agent's data points, Figure 4b only data points of agents with maximal degree centrality in their region.

value and clustering coefficient[12]of the *occupying force* and the *retreating force*[13] just before the spatial acquisition event started. Furthermore, I also computed the transitivity value and clustering coefficient of the occupied territory right after the spatial acquisition event.

It turned out that there doesn't seem to be any systematic difference between the clustering coefficient of occupying and retreating force, whereas the transitivity values reveal a systematic distribution: the occupying force has always a lower value than the retreating force and the occupied territory (Figure 5a and 5b shoe the results of 5 exemplary spatial acquisition events). This is an indicator for the fact that areas of high transitivity are easier to occupy. In conclusion, while it is generally assumed that i) a high degree of local clustering - the clustering coefficient - is an indicator for stability and supports conventions to maintain, the result of my analysis shows that ii) a high degree of global clustering - transitivity - supports the possibility for conventions to spread and is therefore also susceptible to spatial acquisition by neighboring regions.

[12] The values are normalized to an expected average value for a given population size.

[13] For a spatial acquisition event with two participating convention regions, the occupying force is the one with increasing population size, whereas the retreating force is the one with decreasing population size.

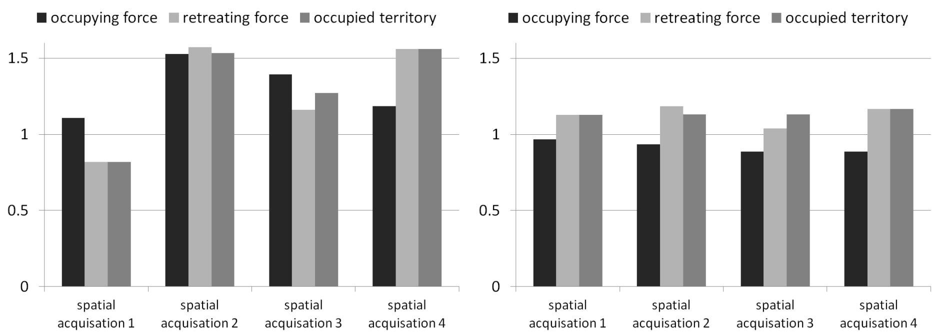

(a) clustering coefficient (b) transitivity

Fig. 5: The clustering coefficient (5a) and transitivity value (5b) for different language regions of four different spatial acquisition events. The black bars each represent the value of the occupying force, the light gray of the retreating force; both just before the acquisition has started. The dark gray bars each represent the value of the occupied territory after acquisition.

4 Conclusion

In this study I conducted simulation runs of 500 agents that i) are placed on a scale-free network, ii) play repeated signaling games with connected neighbors, and iii) update behavior according to innovative reinforcement learning. A general result in all runs was the emergence of convention regions. In further examination steps I analyzed the way regions i) are arranged, and ii) interact.

The first examination step revealed quite a counterintuitive result: agents with a high degree centrality are basically found in small convention regions, whereas large regions entail generally agents with a quite small degree centrality. As a first suggestion, this phenomenon seems to be an artifact of the specific learning dynamics and its possibility for innovation. The second examination step points to the fact that a territory with a high transitivity value is easy to occupy, while a low transitivity value supports preservation. Since this result follows from a small number of observed events, the assumption has to be tested for more robustness in subsequent studies with more data.

References

1. Airiau, S., Sen, S., Villatoro, D.: Emergence of Conventions through Social Learning – Heterogeneous Learners in Complex Networks. Autonomous Agents and Multi-Agent Systems 28(5), 779–804 (2014)
2. Alexander, J M., Skyrms, B., Zabell, S.L.: Inventing New Signals. Dynamic Games and Applications 2.1, 129–145 (2012)
3. Barabási, A.-L., Réka A.: Emergence of Scaling in Random Networks. Science 286, 509-512 (1999).

4. Baronchelli, A., Felici, M., Loreto, V., Caglioti, E., Steels, L.: Sharp Transition towards Shared Vocabularies in Multi-Agent Systems. Journal of Statistical Mechanics: Theory and Experiment 6, P06014 (2006)
5. Barret, J.A.: The Evolution of Coding in Signaling Games. Theory and Decision 67, 223–237 (2009)
6. Barret, J.A., Zollman, K.J.S.: The Role of Forgetting in the Evolution and Learning of Language. Journal of Experimental and Theoretical Artificial Intelligence 21.4, 293–309 (2009)
7. Bush, R., Mosteller, F.: Stochastic Models of Learning. New York: John Wiley & Sons (1955)
8. Crawford V.P., Sobel J.: Strategic Information Transmission. Econometrica 50, 1431–1451 (1982)
9. Franke, M., Jäger, G.: Bidirectional Optimization from Reasoning and Learning in Games. Journal of Logic, Language and Information 21.1, 117–139 (2012)
10. Holme, P., Kim, B.J.: Growing Scale-free Networks with Tunable Clustering. Physical Review E 65.2, 0261071-0261074 (2002).
11. Huttegger S.M.: Evolution and the Explanation of Meaning. Philosophy of Science 74, 1–27 (1998)
12. Jackson, M.O.: Social and Economic Networks. Princeton: Princeton University Press (2008)
13. Lewis, D.: Convention. A Philosophical Study. Harvard University Press, Cambridge (1969)
14. Mühlenbernd, R.: Learning with Neighbours: Emergence of Convention in a Society of Learning Agents. Synthese 183.S1, 87–109 (2011)
15. Mühlenbernd, R.: Convention and Innovation in Social Networks. Manuscript.
16. Mühlenbernd, R., Franke, M.: Signaling Conventions: Who Learns What Where and When in a Social Network? Proceedings of EvoLang IX, 242–249 (2012)
17. Mühlenbernd R., Nick J.D.: Language Change and the Force of Innovation. Pristine Perspectives on Logic, Language, and Computation - Lecture Notes in Computer Science 8607, 194–213 (2014)
18. Roth, A., Erev, I.: Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. Games and Economic Behaviour 8, 164–212 (1995)
19. Sen, S., Airiau, S.: Emergence of Norms through Social Learning. IJCAI 2007, 1507–1512 (2009)
20. Skyrms, B.: Signals: Evolution, Learning & Information. Oxford University Press, Oxford (2010)
21. Steels, L.: Self-Organizing Vocabularies. Artificial Life, 197–184 (1996)
22. Steels, L.: Grounding Symbols through Evolutionary Language Games. Simulating the Evolution of Language, 211–226 (2002)
23. Villatoro, D., Sabater-Mir, J., Sen, S.: Social Instruments for Robust Convention Emergence. IJCAI 2011, 420–425 (2011)
24. Wagner, E.: Communication and Structured Correlation. Erkenntnis 71.3, 377–393 (2009)
25. Warneryd K.: Cheap Talk, Coordination, and Evolutionary Stability. Games and Economic Behaviour 5, 532-546 (1993)
26. Watts D.J., Strogatz S.H.: Collective Dynamics of Small-World Networks. Nature 393, 440–442 (1998)
27. Zollman, K.J.S.: Talking to Neighbors: The Evolution of Regional Meaning. Philosophy of Science 72.1, 69–85 (2005)