Emergent Communication under Varying Sizes and Connectivities

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recent advances in deep neural networks allowed artificial agents to derive their own emergent languages that promote interaction, coordination, and collaboration within a group. Just as we humans have succeeded in creating a shared language that allows us to interact within a large group, can the emergent communication within an artificial group converge to a shared, agreed language? This research provides an analytical study of the shared emergent language within the group communication settings of different sizes and connectivities. As the group size increases up to hundreds, agents start to speak dissimilar languages, but the rate at which they successfully communicate is maintained. Remarkably, we also show an evidence of the shared language becoming more generalizable at describing new referents with increased group size. We observe the emergence of different dialects when we restrict the group communication to have local connectivities only. Finally, we provide an optimization result of group communication graphs when the number of agents one can communicate with is restricted.

1 Introduction

Communication plays a vital role for us humans to interact with one another [Austin, 1975, Allwood, 1976, Linell, 2009]. Through communication, we acquire new information about the environment and other people around us, which further facilitates collaboration or competition and allows us to coordinate our future decisions and actions better.

Over the past decades, numerous research outcomes have succeeded in mimicking human behavior and making artificial agents equipped with interactive communication abilities [Batali, 1998, Cangelosi and Parisi, 2002, Christiansen and Kirby, 2003, Steels, 2003, Wagner et al., 2003]. The communication in the artificial system is emergent in a way that there are no pre-specified usage rules, semantics, or syntax, all of which are formed by the artificial agents upon their necessities [Lazaridou et al., 2017]. Not only does empowering agents with their own communicative abilities set an important milestone in interactive AI systems, it also brings us steps closer to the world where machines cooperate with humans [Mikolov et al., 2016, Crandall et al., 2018, Lazaridou and Baroni, 2020].

In human societies, hundreds of millions of people use the shared, agreed communication protocols to interact. As the group becomes more complex with increased size and connectivity, communication becomes more and more important for us to coordinate with one another. Fortunately, we humans excel at using languages for large-group coordination and problem-solving [Tomasello, 2010, Lupyan and Bergen, 2016]. However, in the AI system counterpart, the question of how the emergent communication will function under different group settings is understudied. As the group becomes more complex, will emergent communication become more effective for the system in promoting group coordination, as in human society, or will it lose its functionality? Will the emergent communication still converge to a shared, agreed language? In the worst case, efficient interaction among agents

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

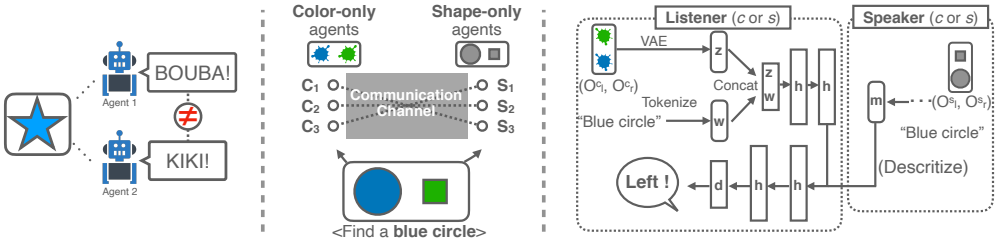

Figure 1: Left: A scenario in which two agents speak different languages to describe the same object. Middle: Bi-partite group communication using a variant of the referential game, involving color-only agents c and shape-only agents s. Right: Communication model architecture between speaker and listener agents. Both color-only agent c and shape-only agent s can be speakers and listeners.

through the emergent communication can be hampered because different agents will start to use different languages to describe the same observation (the left-most panel of Figure 1).

This research provides an analytical study on emergent communication in artificial groups of different sizes and connectivities. Without injecting architectural bias into the environment or agents that could favor or hamper communication, our goal is to investigate the effect of emergent communication and analyze how agents' general agreement on the communication protocol changes under varying group settings. To study this question, we construct a unique setting within which communication among agents is indispensable for achieving their shared goals. Communication is their only means to interact and compensate for their limited access to the environment. Furthermore, to make the communication among agents symmetric, we divide agents into two groups and allow an agent to communicate with another agent in a different group through a bi-partite communication graph. This symmetric setting allows all agents to function as both speakers and listeners and suppresses potential unfavored bias that might arise in the asymmetrical setting that potentially distracts our analyses on the convergence of a shared language [Bouchacourt and Baroni, 2019].

To quantitatively evaluate the functional aspect of communication, we look at (a) the language similarity or message agreement of agents within the same group when describing the same referents and (b) the rate at which an arbitrary pair of agents in the bi-partite communication group succeed in communication and achieve the shared goal.

To this end, our contributions and findings are as follows: (1) We propose a new variant of the referential game [David, 1969, Batali, 1998, Lazaridou et al., 2017, 2018] that involves two groups of agents with limited perceptual capabilities. Our construction serves the exact same purpose and motivation as the original referential game. With our variation, however, agents are now *symmetric* in terms of how much they perceive, speak, and listen. This functional symmetry prevents agents from overfitting a certain aspect of the communicative behavior (e.g., speaking) and producing ad-hoc communication techniques [Bouchacourt and Baroni, 2019, Kottur et al., 2017, Lowe et al., 2017]. Moreover, our construction is simple in a way that adding new degree of complexity and investigating in a group setting straightforward. (2) We find that as the size of the communication graph scales up to hundreds, the language used in the group becomes less similar amongst agents. However, the communication success rate is maintained. In addition, we analyze the extent to which the languages of emergent group communications generalize beyond the training experience. Remarkably, we find that when agents communicate to solve a relatively complex task, their communication generalizes *better* as the group size *increases*. (3) We observe that when agents only get to communicate with nearby agents, a particular form of communication peculiar to agents' locality emerges. At the same time, distant pairs of agents that have not interacted during training can still communicate. (4) We present and analyze the *optimized* communication graph in terms of overall communication success rate when the node outdegrees are fixed. Compared to the fully-connected all-to-all graph, the link density decreases by 81.25% but the communication success rate drops just by 5.89%. With random, unoptimized communication links, the success rate drops by over 40% with the same link density. Throughout the optimization process, we observe the emergence of *hub* nodes taking up a large share of the total communication channels and connecting the majority of other agent nodes, as we see in the real-world scale-free social networks [Barabási and Albert, 1999]. We release the source-code of the game setup and experiments for replicability and future analyses.

2 Related work

There have been successive attempts to let agents develop their own communication protocols through emergent communication [Batali, 1998, Cangelosi and Parisi, 2002, Christiansen and Kirby, 2003, Steels, 2003, Wagner et al., 2003], followed by more recent approaches that are based on deep neural networks [Foerster et al., 2016, Sukhbaatar et al., 2016, Hausknecht, 2016, Havrylov and Titov, 2017]. [1] The idea is applied and extended to a variety of collaborative [Lowe et al., 2020, Hu and Foerster, 2020, Jaques et al., 2019, Kajić et al., 2020, Mordatch and Abbeel, 2018] and competitive [Singh et al., 2019, Jaques et al., 2019] environments. Recently, advances have been made for the emergent communication to be operationalized at scale; to prevent agents from being "lazy" and let them generate meaningful signals [Sunehag et al., 2018]; to effectively schedule the communication in a limited bandwidth [Kim et al., 2019]; to foster conditions that encourage compositional and easier-to-teach languages [Li and Bowling, 2019]; to understand other agents' intention, belief, and point of view, i.e., theory of mind [Hu and Foerster, 2020]; to win a combat against another group [Singh et al., 2019]. Orthogonally, there has been attempts to make the interactive communication more dynamic and more human-like by increasing the complexity of the environment [Jaques et al., 2019, Das et al., 2019], making the communication span multiple turns [Evtimova et al., 2018, Jorge et al., 2016]. These lines of research focus on increasing the efficacy functionality of the emergent communication by suggesting more advanced perceptual, navigational, and generative AI capabilities.

On the other hand, our work falls into the category of an analytical study of understanding the emergent communication from various perspectives. Recent studies find properties of emergent language that sets itself aside from [Chaabouni et al., 2019, Kottur et al., 2017] or are shared with [Resnick et al., 2020, Baroni, 2020] natural human languages and seek for different possibilities to benefit existing NLP models [Lee et al., 2018, Lazaridou et al., 2020]. Bouchacourt and Baroni [2019] and Cao et al. [2018] studies various conditions under which the emerged language can be consonant and agreed as apposed to giving rise to idiolects or agent identity playing a crucial factor in interactive systems. Section 5.4 of our research shares some commonalities with the work of Graesser et al. [2019], where they report the emergence of linguistic continuums at the intersections of different linguistic communities. Our work is distinctive in a way that we do not assign agents into hard clusters of communities, and therefore, we observe finer granularities of linguistic continuums at agent-levels rather than at community-levels.

3 Game design

We design a variant of the referential game [David, 1969, Batali, 1998, Lazaridou et al., 2017, 2018] (the middle panel of Figure 1). The purpose of the game is identical to that of the original game; we test whether an artificial agent can *perceive* a given observation, *speak* through the emergent protocol about the observation according to the instruction, *listen* and make the right decision. Throughout the game, the message from the speaker agent should convey crucial information about the observation, namely, the position of the target object, for the game to succeed. The difference is in the *symmetry* of the game setting, with all agents having the equal degree of speaking and listening. As previously investigated, if an agent is confined in a certain communicative behavior only, there is a concern that it develops *ad-hoc* communication techniques that overfit the given environment without forming a general communication protocol [Kottur et al., 2017, Bouchacourt and Baroni, 2018, 2019, Lowe et al., 2017]. The simple yet symmetric setting allows us to separate the underlying traits of emergent group communication from the potential bias that arise due to agents overfitting the environment.

In each game trial, two agents are presented with multiple pixel-based objects: a target, distractors, and instruction. The goal is to choose the correct target object for the description in the instruction. Crucially, communication is indispensable for the agents to correctly find the target since each of the two agents has its own *limitation* in the perception: One agent can only recognize different shapes, and the other can only differentiate colors. Therefore, the information of the object retrieved by the agents from their observations is mutually exclusive, but at the same time is guaranteed to recover the full state information when aggregated. Under this setting, the agents communicate based on their limited observations to compensate for their deficiencies and correctly choose the target object described in the instruction.

[1] The work by Lazaridou and Baroni [2020] surveys the nascent literature.

Agent and action Each of the two non-embodied agents is chosen from two distinctive *groups*. We denote agents that can only perceive colors (i.e., the *color-only* agents) by c, and the corresponding group by $\mathcal{G}^c$, such that $c \in \mathcal{G}^c$. Similarly, the *shape-only* agents and the corresponding group are denoted by s and $\mathcal{G}^s$. The number of agents of both groups is set to be the same, i.e., $|\mathcal{G}^c| = |\mathcal{G}^s| = N$. We consider a situated environment, where the action space $\mathcal{A}$ is divided into the disjoint communication action space $\mathcal{A}^m$ and decision-making action space $\mathcal{A}^e$, such that $\mathcal{A}^m \cup \mathcal{A}^e = \mathcal{A}$ and $\mathcal{A}^m \cap \mathcal{A}^e = \emptyset$. Of color-only and shape-only agents, one is randomly selected and act as a *speaker*; it takes its observations, i.e., objects and instruction, as inputs and sends a *message* $m \in \mathcal{A}^m$ to the other agent as an output, which can be either discrete or continuous. Next, the other agent is randomly selected and act as a *listener*. This agent takes the message along with the observations and makes a *decision* $d \in \mathcal{A}^e$. The decision is about choosing the position of the correct target object. Thus, the size of the decision making action space $|\mathcal{A}^e|$ is equal to the number of objects in the observation. Finally, the communication between agents (c, s) is *successful* if the decision d finds the target object.

State and observation A state s is a 3-tuple $s = (o_t, O_d, w)$ with the target object o_t, the distractor objects O_d, and the instruction w. Each object is of different colors and shapes [2], and the instruction tells of choosing an object with certain color and shape. Therefore, for n and m different colors and shapes, there can be $n \times m$ varying instructions. We represent the instruction using a fixed-sized vector by concatenating n and m-dimensional one-hot vectors. Observations comprise objects and instructions.

Communication graph A bipartite *communication graph* $\mathcal{G}$ is formed between $\mathcal{G}^c$ and $\mathcal{G}^s$ as multiple pairs (c, s) communicate, i.e., form a set of directed *communication links* $(c, s), (s, c) \in \mathcal{E}$, and establish communication protocols.

Message agreement and communication success rate We introduce message agreement and communication success rates used to evaluate language convergence. First, let a and a' be two agents randomly sampled from the same color-only or shape-only group, i.e., $a, a' \in \mathcal{G}^c$ or $a, a' \in \mathcal{G}^s$, and let m^a and $m^{a'}$ be the messages generated by the agents from the same observation $o \in \mathcal{O}$. The message agreement of a certain group, which is either color-only or shape-only, in a communication graph is measured by the expected similarity between two messages m^a and $m^{a'}$: $\mathbb{E}_{a,a' \sim \mathcal{G}', o \sim \mathcal{O}}[f(m^a, m^{a'})|o]$, where $f(\cdot, \cdot)$ is a similarity measure and $\mathcal{G}' \in \{\mathcal{G}^c, \mathcal{G}^s\}$. Next, the communication success rate of a communication graph is measured as: $\mathbb{E}_{c \sim \mathcal{G}^c, s \sim \mathcal{G}^s, o \sim \mathcal{O}}[r(c, s, o)]$ with r being an identity function that returns 1 if agents successfully find the target object.

4 Group communication model

We describe how agents communicate with others within $\mathcal{G}$ their speaking and listening networks and explain the training procedure. In a communication graph, the selected color-only agent $c \in \mathcal{G}^c$ and shape-only agent $s \in \mathcal{G}^s$ with $(c, s) \in \mathcal{E}$ or $(s, c) \in \mathcal{E}$ communicate via fixed-sized messages and make collective decisions. The network architecture is illustrated in the right panel of Figure 1. The communication algorithm that summarizes the overall procedure can be referenced in the Appendix B.

Speaker In each trial of the game, one color-only or shape-only agent $c \in \mathcal{G}^c$ or $s \in \mathcal{G}^s$ is selected from its group and serves as a speaker, generating a message m. The speaker agent starts by processing a raw pixel-based observation input, consisting of multiple objects with different shapes or colors, into a fixed-sized encoding using variational autoencoder (VAE) [Kingma and Welling, 2014, Rezende et al., 2014]. [3] Each encoded vector is then concatenated with the instruction vector and is used as an input to the speaker network. The speaker network is a feedforward neural network and outputs a fixed-sized message m.

Listener The listener agent processes the observation in the same manner as the speaker network. However, rather than outputting the message, the listener network's intermediate layer is concatenated by the message generated from the speaker agent. The expanded intermediate layer is then used as an

[2] We also differentiate the size, rotation, and angle of objects and vary the color codes within the same color to diversify the state-space, as further detailed in the Appendix A.

[3] Each agent trains VAE separately so that no information is shared at this stage.

input for the upper layers of the network. It outputs the decision action d of size $|\mathcal{A}^e|$, i.e., predicted position of the target object.

Training During training, we iteratively sample communication links between two agents c and s, and apply backpropagation, with binary cross-entropy as the loss function over the training batch. We adopt centralized training and decentralized execution regime [Foerster et al., 2016]. Therefore, when two agents c and s communicate during training, the gradients flow through the communication channel from one agent's listening network to another speaking network. However, agents do not have access to the actual inner architecture, the parameter values of other agents, and no gradient flow during execution. Also, agents within $\mathcal{G}^c$ and $\mathcal{G}^s$ are completely decentralized. Message discretization is achieved by the Gumbel-softmax relaxation [Jang et al., 2017, Maddison et al., 2017]. Finally, we ensure that agents c and s communicate *enough* to learn the communication protocol and be fairly compared but not too much to prevent the generated messages from overfitting the training data by adopting the early stopping rule with the validation set. Refer to the Appendix C for the detailed architectural settings and training results.

5 Experiments

We investigate how different configurations of communication graph affect language convergence in group communication. Specifically, we answer the following questions: 1. **Group size:** In all-to-all communication, i.e., when all agents communicate with each other, how does the size of the communication group $N = |\mathcal{G}^c| = |\mathcal{G}^s|$ affect the message agreement, communication success rate, and generalizability? 2. **Group connectivity:** In neighbor-to-neighbor communication, i.e., when agents only communicate with nearby agents, how successful can a distant pair of agents – who has not exchanged messages during training – communicate? How similar do their languages become? 3. **Optimization:** Can we structuralize the optimal communication graph in terms of communication success rate when agents can only establish a limited number of communication links?

5.1 Experimental setup

Game configuration Agents observe 10 objects with 1 target object and 9 distractors. The size of the decision-action space $|\mathcal{A}^e|$ is set to 10. For both discrete and continuous messages, we set the dimensionality of $m \in \mathcal{A}^m$ to 10. Discrete messages are binary vectors using the Gumbel-softmax relaxation. Continuous messages are real-valued vectors. [4] Objects observations are RGB pixel-based images with sizes 32×320. Each image contains 10 objects with 6 different colors (red, green, blue, cyan, magenta, yellow) and 5 different shapes (ellipse, triangle, quadrilateral, pentagon, hexagon). Instruction is an 11-dimensional vector that concatenates 6 and 5-dimensional one-hot vectors for shapes and colors. Example object observations are in the Appendix A.

Data We generate 128,000 observations for a training set and 12,800 observations for test and validation sets. The early stopping rule is applied for the validation set, and the reported results are calculated using the test set. Finally, the model parameter settings are detailed in the Appendix C.

5.2 Preliminary results

The default success rate is 0.236 in the absence of communication, i.e., when a color-only or shape-only agent single-handedly tries to find the correct target object. The upper bound of the communication success rate is 0.972; there are cases of having one or more distractors with the same shape and color as the target object. In this case, even with perfect communication, the agents are forced to select one randomly. In addition to this, there are different *types* of messages that can be identified combinatorially with respect to the observation space; different agents generate similar messages if the objects in the images and the instructions are similar. Visualizations of different types of messages can be found in the Appendix D.

[4] We can straightforwardly restrict the range of the message values by e.g., applying the softmax or sigmoid function to the message. In our experiments, findings and statements made are consistent with and without such restrictions.

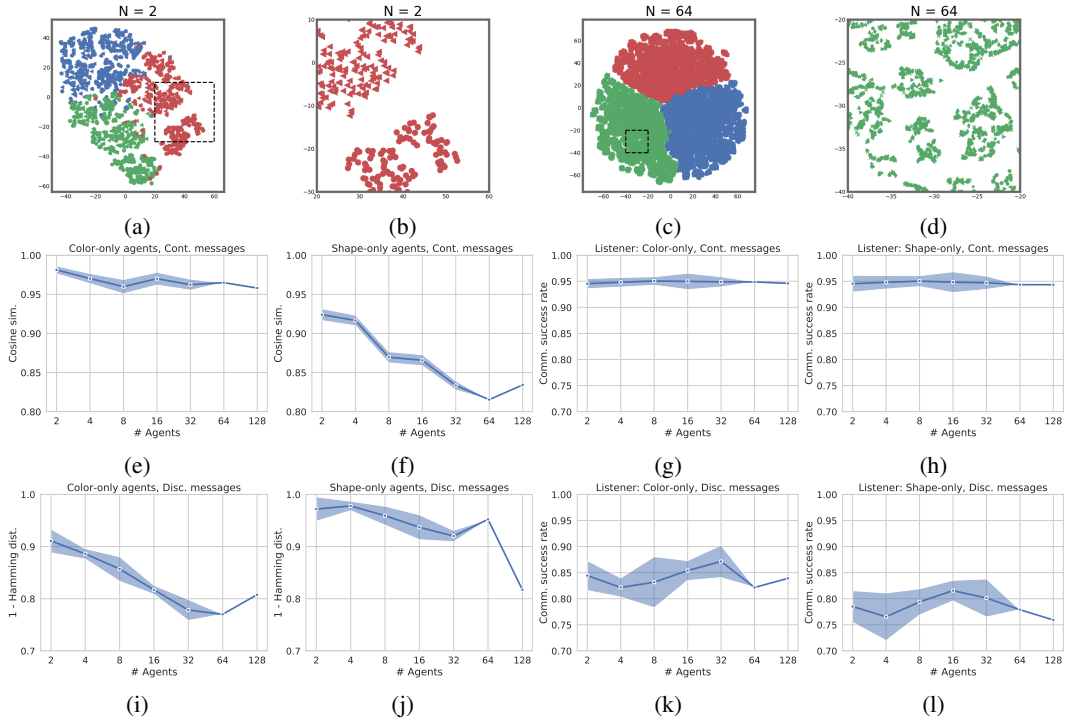

Figure 2: All-to-all communication. (a)-(d): t-SNE plots of messages. (b) and (d) are the magnified portions of (a) and (c) respectively. N is the number of agents and each point is a single communication message. Colors represent different types of messages identified by the authors, and markers represent different agents. As N grows, *what* is in the message becomes more relevant than *who* the speaker is for identifying the message. (e)-(l): Similarities of messages sent from different agents with same observation and communication success rates of the communication graph with respect to the number of agents N in the group. (e)-(h) and (i)-(l) are the results when agents use continuous and discrete messages, respectively. As N grows, the messages sent by different agents become less similar, but the communication success rate is maintained.

5.3 Group size

We answer the first question regarding the all-to-all communication, where all agents in $\mathcal{G}^c$ communicate with all agents in $\mathcal{G}^s$. The top panels of Figure 2 illustrate the t-SNE [Maaten and Hinton, 2008] plots of the messages generated, with N being the size of $\mathcal{G}^c$ and $\mathcal{G}^s$ ($|\mathcal{G}| = 2N$). We only show the color-only agents' results with a continuous message setting due to the space limit. [5] Each dot represents a single message, with different markers representing the messages generated by different agents. Different colors represent different *types* of messages identified by the authors, e.g., we group one type of messages by criteria: the target object is "blue" and it is at the "3rd" position in the observation. [6] As the group size increases, the messages tend to gather up with respect to their types (panels (a) and (c)), and within each cluster, the identities of the agents who transmit the messages become more difficult to be separated (panels (b) ad (d)). As N increases, *what* is entailed in the message becomes more relevant when identifying the message than *who* is sending the message.

Second, from the subfigures (e), (f), (i), (j) of figure 2, the message agreement calculated by the cosine similarity decreases as the group size N increases, for both color-only and shape-only groups and both continuous and discrete-message settings. On the other hand, in the subfigures (g), (h), (k), (l) of figure 2, the communication success rate does not increase nor decrease as N increases. Even though m^c and $m^{c'}$ become dissimilar as N increases, the rate of the color-only agents c and c' successfully communicating with the shape-only agents $\forall s \in \mathcal{G}^s$ does not decrease.

[5] Similar results can be obtained from the other configurations, e.g., shape-only agents with discrete messages.

[6] For the visualization purpose, we only display three types of messages, but there are more than three types of messages that we can identify.

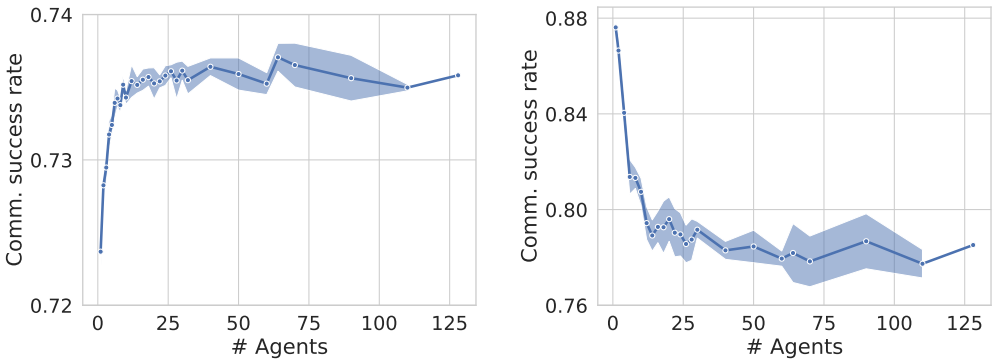

(a) Complex task

(b) Simple task

Figure 3: Communication success rate when agents communicate about observations not seen during training. For complex tasks the communication becomes more generalizable as the group size increases, while for the simple task it becomes less generalizable.

Overall, as N increases, agents start to use different languages to describe similar observations while maintaining the communication success rate. Furthermore, agents' languages becoming more similar does not directly translate to the increased communication success rate.

Generalizability under varying group sizes How well does the emergent protocol generalize to new observations that were not shown during the training? How does the size of the communication graph affect such generalization?

To assess generalization, we diversify the object shapes and colors in the test set. Thus, each color is sampled from the larger spectrum of color codes, and more degrees of freedom is added to the shapes while maintaining the label sets and the number of instructions. For example, an object with the color red can look less red in the extended test set. Refer to the Appendix A for the detailed settings.

Subfigure 3a shows the communication success rates of all-to-all communication graphs with varying sizes N for the extended test set. Since the observations are diversified and not seen during the training experience, the overall communication success rates decrease from 0.95 to less than 0.74. Remarkably, when the graph size increases, the communication success rate calculated using the extended test set also increases, suggesting that the learned communication protocol generalizes better to new observations when a larger number of agents are involved in the communication. Note that the increased generalizability is *not* due to agents' increased perceptual abilities. In fact, the number of observations provided to agents during training is maintained or even decreases depending on the early stopping criteria.

On the other hand, the subfigure 3b shows the generalization results when the task is much simpler with reduced action and observation spaces; the number of decision-action is reduced from 10 to 2; the number of colors and shapes for objects is reduced from 6 to 3 and 5 to 3 respectively. In this case, we observe that the emergent communication generalizes worse as the group size increases.

In summary, when agents are restricted with small observation and action spaces, their communication generalizes poorly as the size of the communication graph increases; in complex environments, the emergent protocol generalizes better as the group size increases.

5.4 Group connectivity

In neighbor-to-neighbor communication the color-only agents $\mathcal{G}^c = \{c_1, ..., c_N\}$ and the shape-only agents $\mathcal{G}^s = \{s_1, ..., s_N\}$ communicate with nearby-agents only, i.e., c_i only communicates with $\{s_{i-j}, ..., s_{i+j}\}$ and vice versa, with j being the maximum communication range. Thus, $j = 0$ signifies one-to-one communication.

Figure 4 illustrates the neighbor-to-neighbor communication results with N set to 64 and varying the maximum communication range j. The dots in the top panels represent the messages generated by the color-only agents after having similar observations, i.e., when the object explained in the instruction

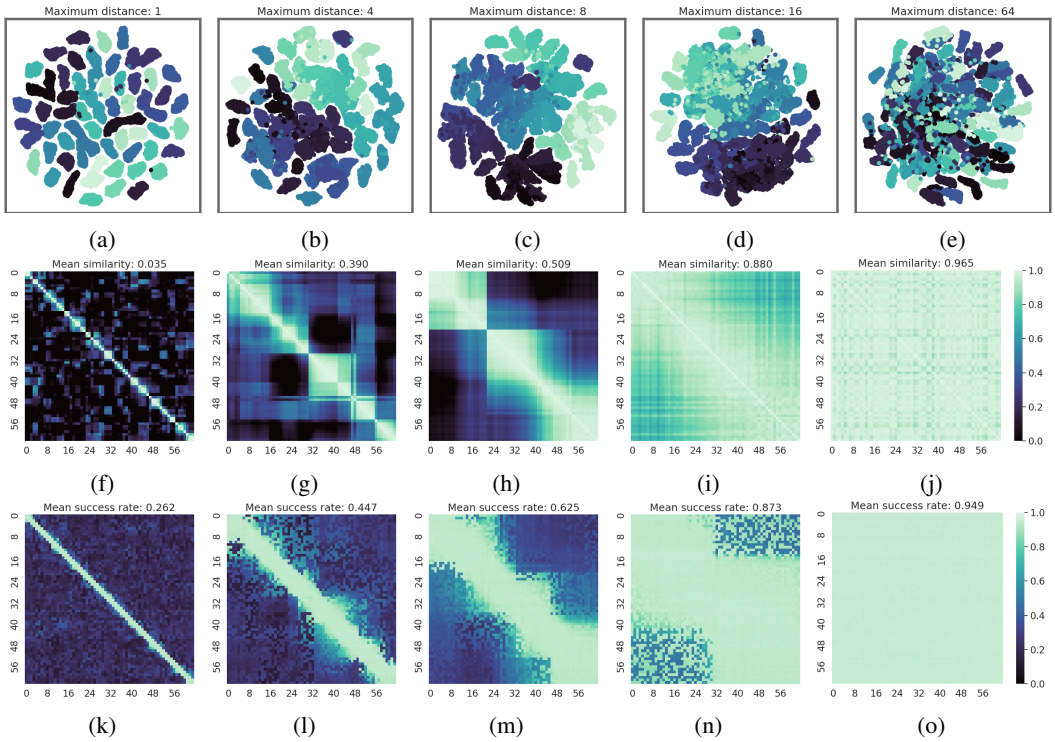

Figure 4: Neighbor-to-neighbor communication. (a)-(e): Color represents locality: lighter colors at the top of the bipartite graph and darker colors at the bottom. As agents communicate with larger number of nearby agents, groupings of communications that are peculiar to their locality, e.g., *dialects* are formed. At the same time, their messages become more similar ((f)-(j)) and the communication success rate increases ((k)-(o)), even between non-neighbors (off-diagonal parts).

Table 1: Communication graph optimization. We restrict the number of outgoing edges for each agents. The numbers are the communication success rates plus minus one standard error over five different trials.

	No Link	All-to-one	N-to-N	Rand. Links	Optimal
Success Rate	0.236	0.384 ± 0.031	0.454 ± 0.017	0.578 ± 0.154	$\mathbf{0.894 \pm 0.021}$

is on the left-most side of the image. [7] Different colors represent different localities of the agents; as i increases in $\mathcal{G}^c = \{c_1, ..., c_i, ..., c_N\}$, the color becomes darker, from light green to black.

When j is small ($j = 1$), multiple separate clusters form, and the dots of the same colors are scattered, suggesting that the agents' communication patterns are distinctive and do not depend too much on the locality. As j increases (j = 16), the groupings of messages emerge according to agents' locality; a particular form of communication is shared by agents at the top, middle, and bottom of the graph. Simultaneously, as j grows, the messages of the pairs of color-only agents that are far away from each other become more similar (the middle panels, off-diagonal parts; the row and column indices indicate the identities of color-only agents). Moreover, the communication success rate between the pair of agents that has never communicated during training increases (the bottom panels, off-diagonal parts; the row indices indicate the localities of the color-only agents and the columns are the localities of the shape-only agents).

As the maximum communication range increases, a particular form of communication peculiar to agents' locality, which is analogous to a *dialect* in human languages, emerges, and the agents' overall language gradually converges.

[7] Again, similar results can be retrieved for different configurations.

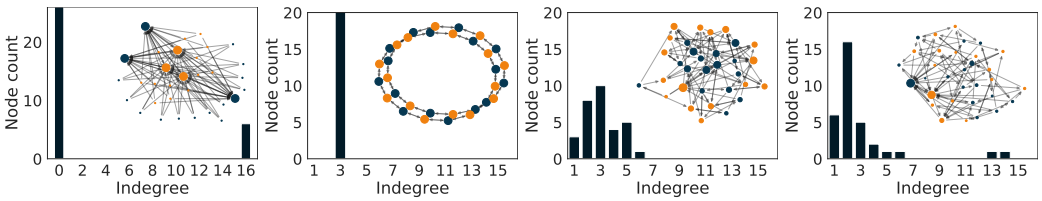

(a) All-to-one (b) N-to-N (c) Random (d) Optimized

Figure 5: Different communication graphs and the respective in-degree distributions.

5.5 Communication graph optimization

In the previous subsection, we observe that in the sparse setting a pair of agents located far away struggles to communicate successfully and become speaking dissimilar languages. In this regard, we aim to find the *optimal* communication graph that maximizes the communication success rate when all agents share the same limited number of outgoing communication links. With $N = 16$, we restrict the number of outgoing links to 3. A color (shape)-only agent can form communication links with any shape (color)-only agents but the total number of outgoing links from the agent always has to be 3. Therefore, the graph density remains constant at 0.1875 $(3/16)$ throughout optimization. We optimize the link structure using covariance matrix adaptation evolution strategy (CMA-ES) [Hansen and Ostermeier, 2001, Hansen, 2006, 2016] with population size 32. For each link structure setting, we use the averaged score from 4 different initializations to stabilize the optimization process. Refer to the Appendix E for additional details regarding the optimization process.

Table 1 shows the communication success rate of the optimized and comparison communication graphs. All except for no link have the same graph density of 0.1875. The visualizations of the graph structures along with their indegree histograms are shown in Figure 5. All-to-one and neighbor-to-neighbor (N-to-N) are two extreme cases. In All-to-one graph, only 3 randomly selected agents in $\mathcal{G}^c$ and $\mathcal{G}^s$ have the in-degrees of $N = 16$, while the rest of the agents have no incoming links. On the contrary, in neighbor-to-neighbor graph, all agents have the same in-degrees and out-degrees of 3. The communication success rate is less than 0.5 for both graphs. Also, when the links are formed randomly, the mean communication success rate is 0.578. Remarkably, we are able to achieve the communication success rate of 0.894 after the optimization. In all-to-all communication, the communication success rate is 0.950. The communication success rate decreases just by 5.89% while the communication density is dropped by 81.25%.

The panel (d) of Figure 5 show the visualization of the optimized communication graph; while most agents have in-degrees ranging from 1 to 4, one agent in each group $\mathcal{G}^c$ and $\mathcal{G}^s$ that has in-degree greater than 12. These agents act as the *hub* nodes, increasing the overall communication success rates. In random graphs (panel (c) of Figure 5), there is no such node with high centrality value.

6 Conclusion & discussion

In this work, we investigated whether the emergent language can converge in different configurations of the communication groups. We observed that when the number of group size increases, the language spoken by different agents becomes more discordant while the overall communication success rate is maintained. This is intriguing since the agents speaking similar languages does not necessarily translate to them communicating successfully. We found that in low-density communication graphs, distant pairs can still communicate with partially shared communication protocols, and further demonstrated the optimization of communication link structure. We find our result on linguistic generalization to be particularly enlightening since it directly evidences the benefit of emergent communication involving larger number of participants rather than concerns. One possible way to explain the better generlization under larger graphs is *over-parameterization* of communication networks, although we withhold from making a definitive statement as the theoretical reasoning on generalization of neural networks is yet to be conclusive [Neyshabur et al., 2015, Zhang et al., 2017, Nagarajan and Kolter, 2019].

The analyses on temporally extended and complex environment with embodied agents would be a desirable next step. Hill et al. [2020] focused on various environmental factors that lead to generalization in an embodied system. Overcoming the existing problem of multi-agent RL, such as the lazy-agent problem [Sunehag et al., 2018], the message scheduling issue [Kim et al., 2019], and joint exploration problem [Lowe et al., 2017], would also be an interesting future direction.

References

Jens Allwood. *Linguistic communication as action and cooperation*. University of Göteborg. Department of Linguistics, 1976.

John Langshaw Austin. *How to do things with words*, volume 88. Oxford university press, 1975.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286 (5439):509–512, 1999.

Marco Baroni. Linguistic generalization and compositionality in modern artificial neural networks. *Philosophical Transactions of the Royal Society B*, 375(1791):20190307, 2020.

John Batali. Computational simulations of the emergence of grammar. *Approach to the Evolution of Language*, pages 405–426, 1998.

Diane Bouchacourt and Marco Baroni. How agents see things: On visual representations in an emergent language game. In *EMNLP*, 2018.

Diane Bouchacourt and Marco Baroni. Miss tools and Mr fruit: Emergent communication in agents learning about object affordances. In *ACL*, 2019.

Angelo Cangelosi and Domenico Parisi. *Simulating the evolution of language*. Springer, 2002.

Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Leibo, Karl Tuyls, and Stephen Clark. Emergent communication through negotiation. In *ICLR*, 2018.

Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. Anti-efficient encoding in emergent communication. In *NeurIPS*, 2019.

Morten Christiansen and Simon Kirby. *Language evolution.* Oxford University Press, 2003.

Jacob Crandall, Mayada Oudah, Fatimah Ishowo-Oloko, Sherief Abdallah, Jean-François Bonnefon, Manuel Cebrian, Azim Shariff, Michael Goodrich, Iyad Rahwan, et al. Cooperating with machines. *Nature communications*, 9(1):1–12, 2018.

Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. Tarmac: Targeted multi-agent communication. In *ICML*, 2019.

Lewis David. Convention: a philosophical study, 1969.

Katrina Evtimova, Andrew Drozdov, Douwe Kiela, and Kyunghyun Cho. Emergent communication in a multi-modal, multi-step referential game. In *ICLR*, 2018.

Jakob Foerster, Yannis Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *NeurIPS*, 2016.

Laura Graesser, Kyunghyun Cho, and Douwe Kiela. Emergent linguistic phenomena in multi-agent communication games. In *EMNLP*, 2019.

Nikolaus Hansen. The CMA evolution strategy: a comparing review. In *Towards a new evolutionary computation*, pages 75–102. Springer, 2006.

Nikolaus Hansen. The CMA evolution strategy: a tutorial. *arXiv preprint arXiv:1604.00772*, 2016.

Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.

Matthew Hausknecht. *Cooperation and communication in multiagent deep reinforcement learning*. PhD thesis, University of Texas at Austin, 2016.

Serhii Havrylov and Ivan Titov. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In *NeurIPS*, 2017.

Felix Hill, Andrew Lampinen, Rosalia Schneider, Stephen Clark, Matthew Botvinick, James McClelland, and Adam Santoro. Environmental drivers of systematicity and generalization in a situated agent. In *ICLR*, 2020.

Hengyuan Hu and Jakob Foerster. Simplified action decoder for deep multi-agent reinforcement learning. In *ICLR*, 2020.

J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3): 90–95, 2007.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-softmax. In *ICLR*, 2017.

Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, Dj Strouse, Joel Leibo, and Nando De Freitas. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *ICML*, 2019.

Emilio Jorge, Mikael Kågebäck, Fredrik Johansson, and Emil Gustavsson. Learning to play guess who? and inventing a grounded language as a consequence. In *NeurIPS Workshop Track*, 2016.

Ivana Kajić, Eser Aygün, and Doina Precup. Learning to cooperate: Emergent communication in multi-agent navigation. *arXiv preprint arXiv:2004.01097*, 2020.

Daewoo Kim, Sangwoo Moon, David Hostallero, Wan Ju Kang, Taeyoung Lee, Kyunghwan Son, and Yung Yi. Learning to schedule communication in multi-agent reinforcement learning. In *ICLR*, 2019.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014.

Diederik Kingma and Max Welling. Auto-encoding variational Bayes. In *ICLR*, 2014.

Satwik Kottur, José Moura, Stefan Lee, and Dhruv Batra. Natural language does not emerge 'naturally'in multi-agent dialog. In *EMNLP*, 2017.

Angeliki Lazaridou and Marco Baroni. Emergent multi-agent communication in the deep learning era. *arXiv preprint arXiv:2006.02419*, 2020.

Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. Multi-agent cooperation and the emergence of (natural) language. In *ICLR*, 2017.

Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. Emergence of linguistic communication from referential games with symbolic and pixel input. In *ICLR*, 2018.

Angeliki Lazaridou, Anna Potapenko, and Olivier Tieleman. Multi-agent communication meets natural language: Synergies between functional and structural language learning. In *ACL*, 2020.

Jason Lee, Kyunghyun Cho, Jason Weston, and Douwe Kiela. Emergent translation in multi-agent communication. In *ICLR*, 2018.

Fushan Li and Michael Bowling. Ease-of-teaching and language structure from emergent communication. In *NeurIPS*, 2019.

Per Linell. *Rethinking language, mind, and world dialogically*. Information Age Publishers, 2009.

Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *NeurIPS*, 2017.

Ryan Lowe, Abhinav Gupta, Jakob Foerster, Douwe Kiela, and Joelle Pineau. On the interaction between supervision and self-play in emergent communication. In *ICLR*, 2020.

Gary Lupyan and Benjamin Bergen. How language programs the mind. *Topics in cognitive science*, 8(2):408–424, 2016.

Andrew Maas, Awni Hannun, and Andrew Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML*, 2013.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *JMLR*, 9(Nov): 2579–2605, 2008.

Chris Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*, 2017.

Tomas Mikolov, Armand Joulin, and Marco Baroni. A roadmap towards machine intelligence. In *CICLing*, 2016.

Igor Mordatch and Pieter Abbeel. Emergence of grounded compositional language in multi-agent populations. In *AAAI*, 2018.

Vaishnavh Nagarajan and Zico Kolter. Uniform convergence may be unable to explain generalization in deep learning. In *NeurIPS*, 2019.

Vinod Nair and Geoffrey Hinton. Rectified linear units improve restricted Boltzmann machines. In *ICML*, 2010.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *ICLR workshop track*, 2015.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035. Curran Associates, Inc., 2019. URL `http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf`.

Cinjon Resnick, Abhinav Gupta, Jakob Foerster, Andrew Dai, and Kyunghyun Cho. Capacity, bandwidth, and compositionality in emergent language learning. In *AAMAS*, 2020.

Danilo Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.

Amanpreet Singh, Tushar Jain, and Sainbayar Sukhbaatar. Learning when to communicate at scale in multiagent cooperative and competitive tasks. In *ICLR*, 2019.

Luc Steels. Evolving grounded communication for robots. *Trends in cognitive sciences*, 7(7):308–312, 2003.

Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Learning multiagent communication with backpropagation. In *NeurIPS*, 2016.

Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Leibo, Karl Tuyls, et al. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *AAMAS*, 2018.

Michael Tomasello. *Origins of human communication*. MIT press, 2010.

Kyle Wagner, James Reggia, Juan Uriagereka, and Gerald Wilkinson. Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69, 2003.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See section 6
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Appendix G
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We provide it in the supplementary material for the full replication of the results shown in the paper.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See subsection 5.1 and Appendix C.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We report stand errors. See section 5.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix F
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Supplementary material
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] The main data we are using are polygon images, not involving human data.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]