

# Emergent Communication under Competition

Michael Noukhovitch  
Mila, Université de Montréal  
Montréal, Canada  
noukhovm@mila.quebec

Angeliki Lazaridou  
Deepmind  
London, UK  
angeliki@google.com

Travis LaCroix  
Mila, Université de Montréal  
Schwartz Reisman Institute, University of Toronto  
Montréal, Canada and Toronto, Canada  
lacroixt@mila.quebec

Aaron Courville  
Mila, Université de Montréal  
CIFAR Fellow  
Montréal, Canada  
aaron.courville@umontreal.ca

## ABSTRACT

The literature in modern machine learning has only negative results for learning to communicate between competitive agents using standard RL. We introduce a modified sender-receiver game to study the spectrum of partially-competitive scenarios and show communication can indeed emerge in a competitive setting. We empirically demonstrate three key takeaways for future research. First, we show that communication is proportional to cooperation, and it can occur for partially competitive scenarios using standard learning algorithms. Second, we highlight the difference between communication and manipulation and extend previous metrics of communication to the competitive case. Third, we investigate the negotiation game where previous work failed to learn communication between independent agents [6]. We show that, in this setting, both agents must benefit from communication for it to emerge; and, with a slight modification to the game, we demonstrate successful communication between competitive agents. We hope this work overturns misconceptions and inspires more research in competitive emergent communication. Code is available at <http://github.com/mnoukhov/emergent-compete>

## KEYWORDS

Emergent Communication; Multiagent Learning; Competitive Multiagent Learning

### ACM Reference Format:

Michael Noukhovitch, Travis LaCroix, Angeliki Lazaridou, and Aaron Courville. 2021. Emergent Communication under Competition. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

## 1 INTRODUCTION

Understanding the principles involved in the evolution of effective human communication is essential for the progression of artificial-intelligence research since following these same principles is likely our best path forward in developing effective communication strategies for use by interacting AI agents. Naturally, AI agents need a

common language to coordinate with one another and to communicate successfully with humans [49]. The emergence of communication protocols between learning agents has seen a surge of interest in recent years, but most work to date focuses on fully-cooperative agents that share rewards [16, 20, 33]. Work on partial or total conflict of interest between agents has been limited, with only negative results in modern ML literature [31]. Results suggest that when agents are not fully cooperative, they will fail to learn to use a communication channel effectively [6] without extra explicit incentives to do so [25]. This has contributed to a potential misconception that emergent communication (EC) is a purely cooperative pursuit [30]. Contrary to this view, theoretical results in classical game theory prove it is possible to use an existing cheap-talk communication channel to effectively coordinate actions even in situations where there is (partial) conflict of interest [15].

In this paper, we aim to reconcile these contradictory views and lay the groundwork for future research in competitive EC. We begin by reviewing the field of EC and previous work in competitive cases. Then, to study this problem in detail, we consider the simplest case of communication, where a sender must communicate a target to a receiver using a message. The receiver uses the message to guess the sender’s target [35]. The classic game is fully cooperative, with agents sharing reward proportional to the accuracy of the guess. To investigate whether meaningful communication can emerge between *competitive* agents, we introduce degrees of conflict between the sender and receiver using a parameter that smoothly controls the distance between agents’ targets. This allows us to study EC along a continuum between fully-competitive and fully-cooperative agents.

We present three key results for the field of EC and provide takeaways for each.

- (1) For the first time in modern EC, we find evidence that agents can learn to effectively communicate in competitive scenarios—given that their interests are at least partially aligned.
- (2) We investigate what qualifies as ‘communication’ and distinguish it from ‘manipulation’, leaning on previous literature in evolutionary game theory. In doing so, we extend the work of measuring cooperative communication [37] to the competitive case.
- (3) We go back to one of the first instances of competitive EC in Cao et al. [6] to reproduce and analyze their results. We bring

forward a strong argument to overturn their hypothesis that agents fail to learn to communicate. Instead, we propose that agents specifically learned to communicate ineffectively, and we highlight that both agents must be able to benefit from communication to ensure its feasibility.

Finally, we apply these three lessons to the negotiation game [6] where competitive agents previously failed to communicate. With a small modification to the game, we show that competitive agents learn to effectively communicate. We conclude with a view towards the future of the field and encourage future research to use the outlined takeaways to carefully build better models, environments, and metrics.

## 2 RELATED WORK

### 2.1 Emergent Communication

Emergent communication (EC) is the study of learning a communication protocol, typically in order to share information that is needed to solve some task [16]. Recent work has focused on deep learning [18] agents trained with reinforcement learning (RL) [47]. Investigations have remaining mostly in the realm of fully-cooperative games [10, 14, 32] with agents often communicating through discrete tokens from a vocabulary [20]. See Lazaridou and Baroni [31] for a full review of recent work.

Evaluating communicative success is often qualitative [20, 33] and varied [4]. Lowe et al. [37] investigate common issues in evaluation and find that improvement over a non-communication baseline is sufficient to demonstrate success in *non-situated* environments [49] where agents can only transfer information to each other through communication actions. They also propose a metric for situated games, Causal Influence of Communication (CIC), which measures the causal influence of one agent’s communication on another agent’s action. A similar metric is used as an auxiliary reward by Jaques et al. [25] and Eccles et al. [12] to improve performance in multiagent RL games. In both cases there is an underlying assumption that more influence means better communication; however, this assumption may not always hold in competitive environments.

### 2.2 Competitive Emergent Communication

We investigate the space of partially-competitive games, known as general-sum games, where there is some amount of cooperation and some amount of competition. Research on EC in these environments has been limited with many works using forms of competition but not actually having game-theoretic competitiveness. Singh et al. [44] claimed to learn communication using continuous signals (scalars) in mixed cooperative-competitive scenarios. However, their setup uses parameter sharing between opponents; their ‘mixed’ case is non-competitive (and implicitly cooperative); their competitive case is actually two stages—one fully cooperative and one fully competitive; and their result in the competitive scenario is essentially non-communication. Liang et al. [36] have two teams ‘competing’ to solve a task but in practice they are actually cooperative, helping each other bootstrap their communication, and not competing over the same reward.

There are two modern works that do look at EC in competitive games. Cao et al. [6] primarily focus on fully-cooperative agents that share a reward (‘prosocial’) but have a set of experiments with

independent agents that do not share a reward (‘selfish’). They find that independent agents playing a non-iterated, two-player negotiation game fail to learn to communicate effectively. They claim that agents that share a reward learn to communicate, whereas agents that do not may require an iterated game to learn effective communication. Jaques et al. [25] focused on many independent agents playing a multi-agent sequential social dilemmas (SSDs) [34] and propose an auxiliary influence reward to encourage communication and cooperation. They show that effective communication between competitive agents can be learned using their auxiliary reward, but results without it are weak.


We focus on the rigorous study of why and when communication is learned in competitive settings. In contrast to previous work, we seek not only to learn communication between competitive agents but also to explain sufficient conditions for it while accounting for many possible variables and confounding factors. Our work uses a two-player non-iterated game as in Cao et al. [6]; but, we delve further and evaluate the spectrum from fully cooperative to fully competitive. To do so, we use independent agents (‘selfish’ in Cao et al. [6] terminology) and a game environment where we can precisely control the level of competition. Compared to Cao et al. [6], we specify the level of competition, and probe agent behaviours as well as game dynamics to understand the underlying incentive structure. In contrast to Jaques et al. [25], we use a non-situated environment to guarantee no communication through the action space, we explicitly quantify our game’s level of competitiveness, and use a foolproof non-communication baseline. This work does not propose new architectures or learning rules but aims to critically examine existing beliefs, explain previous negative results, draw important distinctions, and recommend best practices for future work.

We follow previous work by comparing *learning algorithms* that are trained together as opposed to *learned agents* compared at test time. This is to avoid a significant issue of comparing EC agents—i.e., different protocols. Since an emergent protocol only has meaning between the agents that learned it together, comparing two learned agents at test time would require that they ad-hoc infer each others’ protocols. While ad-hoc play has made strides in fully cooperative environments [23], it has not been generally solved [1].

### 2.3 Sender-Receiver Games

Our experiments use the simple emergent-communication framework known as a sender-receiver game [35] (or ‘referential game’ [33]). In the classic game, the sender is given a target value to be communicated to the receiver via a message. The receiver receives the sender’s message and must decode it to predict the target value. Both players get the same reward: the negative of the receiver’s prediction error. In this fully-cooperative setting, players learn a protocol to transfer information as effectively as possible [45]. The messages between the sender and the receiver can be categorised as ‘cheap talk’: messages are costless, non-binding, non-verifiable and *may* affect the receiver’s beliefs [15].

Our work benefits greatly from Crawford and Sobel [9]—a seminal work in classical game theory. They study possible *fixed* communication equilibria under competition by giving the sender and receiver different targets and creating a conflict of interest. They



**Figure 1: Visual representation of the circular biased sender-receiver game. Agents are given targets  $T_R, T_S$  that are bias  $b$  apart. The receiver chooses an action  $a$  and agents receive the  $L_1$  losses  $L_1^r, L_1^s$**

perform a static analysis and prove the existence of Nash equilibria where the amount of information communicated is proportional to the alignment between the players’ interests; however, no informative equilibrium exists when interests diverge too significantly. However, static analysis only tells us which equilibria are possible, not necessarily which are feasible to achieve and maintain using learning algorithms (e.g. RL). LaCroix [29] provides a concrete case where static analysis fails to explain the properties of the game fully. Behaviour out of equilibrium can sometimes be surprising; for example, Wagner [48] proves that non-convergent dynamics can sustain partial information transfer even in a zero-sum signalling game.

As such, we do not characterise the static equilibria of our game; instead, we provide a *dynamic* analysis and show the *feasibility* of communication using standard learning rules in RL. We do not explicitly aim for equilibria [43] but look at the information transfer of communication protocols in flux (and therefore out of equilibrium). This is more in line with previous work in EC [6, 25] and allows us to give practical advice for training RL agents to communicate. Similarly, our work is related to the competitive sender-receiver games of Kamenica and Gentzkow [26] who make the assumption of rational, Bayesian agents whereas we work in field of modern EC which uses RL and deep learning [31].

### 3 THE CIRCULAR, BIASED SENDER-RECEIVER GAME

To investigate a range of competitive scenarios, we introduce a modified sender-receiver game with a continuous bias variable,  $b$ , that represents the agents’ conflict of interest, ranging from fully cooperative to fully competitive. The two players—the Sender ( $S$ ) and the Receiver ( $R$ )—have corresponding targets ( $T_s$  and  $T_r$ ), which are represented by angles on a circle that are  $b$  degrees apart:  $T_r = (T_s + b) \bmod 360^\circ$ .

The game starts with the sender’s target being sampled uniformly from the circle  $T_s \sim \text{Uniform}[0, 360]$ . The sender is given its target as input and outputs a categorical distribution  $S(T_s)$  over the vocabulary  $V$ . For its message, it stochastically samples a single token from the distribution  $m \sim S(T_s)$ . The receiver is given the message and deterministically outputs a scalar  $a = R(m)$  as its

action. The receiver does not see its own target so must rely on information from the sender’s message. The goal of each agent is to make the receiver’s action as close as possible to its *own* target value as both players get a loss proportional to the distance between the action and their respective targets. We choose to use an  $L_1$  loss which corresponds to the angle between the target and action so formally, the loss for each agent  $i$  is

$$L_1^i(a, T_i) = \min(|T_i - a|, 360^\circ - |T_i - a|)$$

Figure 1 gives an instance of this game; see Algorithm 1 in supplementary for the complete algorithm. It is evident that a game with  $b = 0^\circ$  is fully cooperative ( $L_1^r = L_1^s$ ) and a game with the maximum bias  $b = 180^\circ$  is fully competitive or constant-sum, a generalisation of zero-sum (see supplementary B.1 for the proof). All values in-between,  $b \in (0^\circ, 180^\circ)$ , represent the spectrum of partially cooperative/competitive *general-sum* games. This can be seen as an extension of the game from Crawford and Sobel [9], modified to cover the spectrum of cooperative/competitive scenarios.

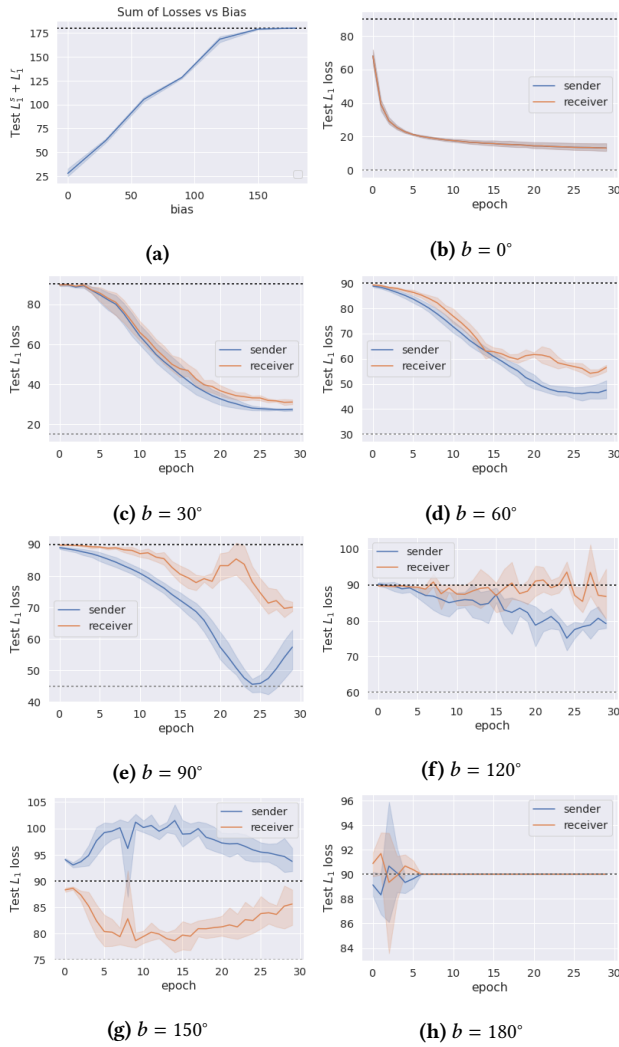
#### 3.1 Training Details

Both agents are implemented as MLPs with 2 hidden layers and ReLU [40] nonlinearities between layers. This was chosen as it is the smallest network that could reasonably solve the cooperative  $b = 0^\circ$  task (but we observed similar results with a 3 layer network as well). Optimization can be performed with either Gumbel-Softmax [24, 38]<sup>1</sup>, REINFORCE [50] (also known as the score function estimator [17]) or a combination of the two where the receiver is updated with backpropogation and the sender with REINFORCE [42] since the loss is differentiable with respect to the receiver. We choose the latter following previous work [7, 13] and Chaabouni et al. [8] who show it converges more robustly. We train both agents independently and for REINFORCE we use an added entropy regularisation term, as well as a baseline for variance reduction[50].

We train for 30 epochs of 250 batches, with batch size 64, and set the circumference of our circle to 36 (so that a loss of  $90^\circ$  corresponds to a value 9). Both agents are trained independently using Adam [27]. To evaluate, we use a fixed test set of 100 equidistant points  $\in [0^\circ, 180^\circ]$  and take the expectation over the messages using the sender’s distribution (as opposed to stochastically sampling a message during training). We do all hyperparameter searches with Orion [5], using random search with a fixed budget of 100 searches. We perform a hyperparameter search over both agents’ learning rates, hidden layer sizes, the vocabulary size, and entropy regularisation (when used). We always report results averaged over 5 random seeds, plotting the mean over seeds as a line and a shaded area indicating 95% confidence interval across seeds. Since competitive MARL may not converge we take the average over the last 10 epochs to capture agents’ final performance.<sup>2</sup> Hyperparameter search spaces and settings are described in supplementary A. Our model is implemented in PyTorch [41] and we use wandb [3] for logging.

<sup>1</sup>Since we train with independent agents that are competitive, we also thought Gumbel-Softmax would be inappropriate as it passes the receiver’s gradient to the sender. This would mean the sender is backpropogating through the receiver’s policy and is an unfair advantage.

<sup>2</sup>Our experiments generally never converge to one policy but have both agents dynamically maintain a final minimum loss nonetheless.



**Figure 2: (a) Lowest test loss ( $L_1^r + L_1^s$ ) for the circular biased game found with a hyperparameter search for each  $b \in [0, 30^\circ, 60^\circ, 90^\circ, 120^\circ, 150^\circ, 180^\circ]$ . (b-h) Individual agent losses for the best run for each bias  $b$ . We plot the baseline loss under non-communication ( $180^\circ$ ) as a dark dotted line. Loss lower than non-communication baseline indicates effective communication.**

## 4 COMMUNICATION IS PROPORTIONAL TO COOPERATION

### 4.1 Measuring Information Transfer

To evaluate the effectiveness of the communication emerged, we can simply look to the sum of agents'  $L_1$  losses. Under non-communication (or uninformative communication), we know that the receiver will just guess a point uniformly at random, so the expected loss for both agents is  $\mathbb{E}_{x \sim U(0,360)} [L_1^s(T_s, x)] = 90^\circ$ . Therefore, any error for either agent below  $90^\circ$  is evidence of information transfer [37].

Furthermore, since agents are not *situated* there is no other action space which agents can use to communicate, all information transfer must be through the communication channel [39]. Therefore, the lower  $L_1^r + L_1^s$ , the more information is passed through the learned protocol.

### 4.2 Emergent Communication for Different Levels of Competition

Both Cao et al. [6], Jaques et al. [25] found that independent agents could not effectively learn to communicate in a competitive scenario and argue this is supported by Crawford and Sobel [9]. Intuitively, though, the level of competition should affect whether communication is learned. To investigate, we use seven equidistant values of  $b \in [0^\circ, 30^\circ, 60^\circ, 90^\circ, 120^\circ, 150^\circ, 180^\circ]$  corresponding to different levels of competition. For each one, we do a hyperparameter search to find the lowest achievable sum of agents' errors, and empirically show whether communication can be learned under competition.

We report our results in Figure 2. The overall results in Figure 2a show that the lowest achieved sum of errors is proportional to the bias between agents. Given that our agents perform better than the non-communicative baseline, we can conclude that agents can learn to communicate effectively without any special learning rules or losses contrary to current literature in machine learning. This matches the theoretical results of Crawford and Sobel [9]; information transfer with communication is inversely proportional to the conflict of interest. We also confirm their theoretical result that effective communication is not possible past a certain bias – in our case  $b = 120^\circ$ , where information transfer becomes chaotic and negligible. We note that performance even in the fully cooperative  $b = 0^\circ$  case is not optimal because of the bottleneck of discrete communication and training instabilities of RL. We also show  $b = 180^\circ$ , the constant-sum game ( $L_1^s + L_1^r = 180^\circ$ ) in Figure 2h. Since the game is fully competitive, there is no incentive to communicate and all policies are uninformative after some initial instability. We show it here for completeness but, for brevity, we omit  $b = 180^\circ$  in future experiments.

Overall, this empirical result extends the theoretical work of Crawford and Sobel [9] to show that effective information transfer is not only theoretically possible, but also feasible to learn with the standard EC setup using RL. Furthermore, our results suggest that EC agents can effectively learn communication proportional to the level of competition.

## 5 MANIPULATION IS NOT COMMUNICATION

### 5.1 Information Transfer vs Communication

While we have found evidence of information transfer, does that necessarily mean our agents have *learned to communicate*? For example, if the sender were to be static and deterministic, the receiver could learn a perfect strategy to achieve 0 loss for itself. This would yield  $L_1^r + L_1^s = b$  which is optimal information transfer and therefore evidence of communication according to Lowe et al. [37]. But is this what we mean by *communication*?

To further explore this question, we would like to show how a different model can seem to have better communication under the current metric but be inconsistent with our goals. For our different model, we change to using continuous, scalar messages similar to

those used in a competitive setting in Singh et al. [44]. We make the sender’s message a real-valued scalar and appropriately change its the distribution to be a 1D Gaussian, for which we now learn the mean and variance (see Algorithm 2 in supplementary). As with the previous experiment, we choose the minimal architecture that does well on the  $b = 0^\circ$  case, which is again an MLP with two hidden layers where the sender has two separate heads on top, one for predicting each of the mean and variance of the Gaussian.

We re-run our experiments with continuous messages and report our results in Figure 3. The overall results in Figure 3a suggest that the protocols learned with continuous messages are all highly informative and better than those with discrete messages. However, looking at the best runs for highly competitive cases  $b = 120^\circ, 150^\circ$  in Figures 3f, 3g, we find that the receiver dominates the sender, performing near-optimally while the sender’s loss is usually close to the bias  $b$ . We can deduce that the receiver manages to guess its own target correctly. This is not advantageous to the sender, as it performs worse than if it were not communicating at all (dark dotted line). This suggests that the sender is not choosing to transfer information about the target but rather the receiver is reading it from the sender’s message despite the sender’s best efforts. Though it is information transfer, this does not resemble communication and requires us to consider literature in signalling [45] to delineate the differences.

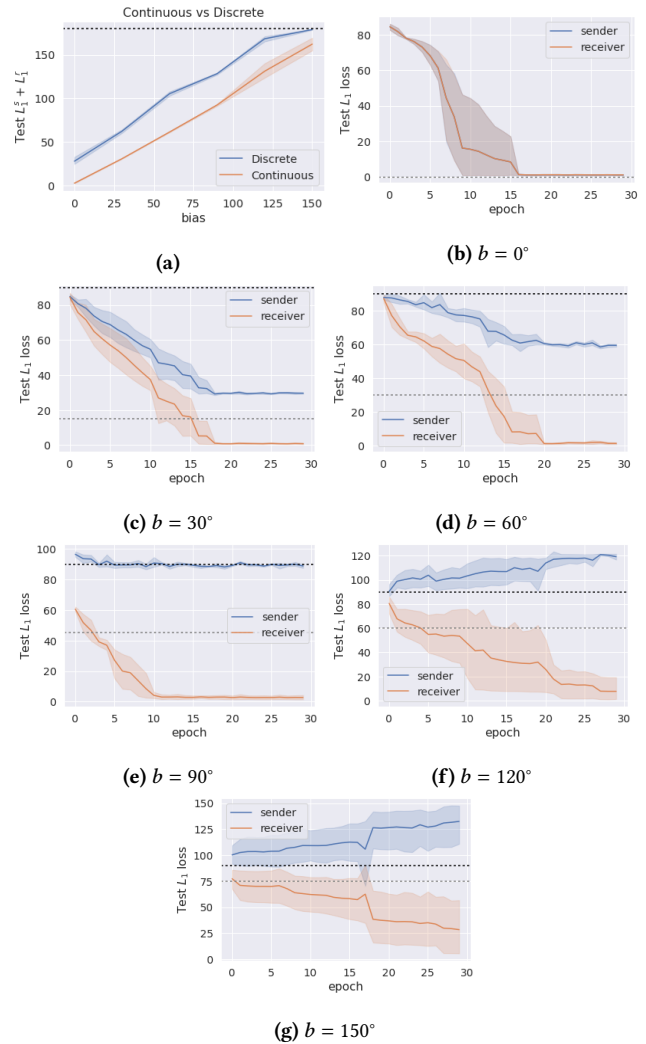
## 5.2 Communication vs Manipulation

The literature in signalling refers to information transfer as either *communication* [2] or *manipulation* of receivers by senders [11] or vice-versa [22]. The latter manifests as one agent dominating the other and is modelled as *cue-reading* or *sensory manipulation*. Communication, as opposed to manipulation, requires *both* agents to receive a net benefit [28], and this implies some degree of cooperation [35]. For fully-cooperative EC, previous metrics of joint reward [37], or causal influence [25, 37], are sufficient to find information transfer and can therefore be used as effective metrics. But for competitive scenarios, neither of these can distinguish between manipulation and cooperation [46] and using them as metrics (e.g. [25]) risks finding manipulation and confusing it for communication.

## 5.3 Measuring Competitive Communication

To accurately compare the inductive biases our models of discrete and continuous messages, we plot the results of all 100 hyperparameter searches instead of just choosing the best one. For each bias  $b$  we compare every hyperparameter configuration of our discrete and continuous message models using the  $L_1$  error for each agent. We report our results in Figure 4.

Since our focus is on the emergence of *communication*, we are looking for settings where both agents benefit: they perform better than either their fully-exploited losses ( $L_1^s < b$  and  $L_1^r < b$ ) or the loss under non-communication ( $L_1^s < 90^\circ$  and  $L_1^r < 90^\circ$ ). We see throughout that our discrete model is relatively fair. In contrast, the receiver with continuous messages dominates the sender. Points follow a general slope of  $L_1^s = L_1^r + b$  regardless of the bias  $b$ . In competitive scenarios  $b = 120^\circ, 150^\circ$  in Figure 4e,4f, the sender with continuous messages has very few runs that performs better than



**Figure 3: (a) Total test loss over biases  $b \in [30^\circ, 60^\circ, 90^\circ, 120^\circ]$  in the circular biased game for in a setup with continuous messages compared to discrete messages. (b-g) Individual agent reward for the best runs found for continuous messages for each bias. We plot the loss under non-communication ( $90^\circ$ , dark dotted line) and loss for an even split of the bias ( $b/2$ , light dotted line) as baselines. Loss higher than the non-communication baseline indicates manipulation.**

non-communication. All other runs clearly show manipulation by the receiver. This corroborates our intuition that this model with continuous messages is predisposed to finding manipulation, not communication.<sup>3</sup> On the other hand, our model with discrete messages has both agents perform similarly well and therefore biases towards communication. Though the receiver usually does better, we find that the difference between agents’ losses is not

<sup>3</sup>This model is only used to demonstrate how manipulation can be mistaken for communication. A different setup with continuous messages can learn strong non-manipulative communication (see supplementary C) and a discrete message system can learn manipulation (see certain runs in Figure 4e,4f)

overly large and indicative of widespread manipulation, like we see with continuous messages.

This demonstrates the importance of qualitatively evaluating our models’ biases but eventually, we need to choose the best hyperparameters and this requires a quantitative metric of competitive communication. Previous approaches[25] used the simple sum of agents’ losses but we propose the sum of squared losses  $L_2 = (L_1^s)^2 + (L_1^r)^2$  as a better heuristic metric. We can view our partially competitive scenario as having a *common-interest loss* ( $180^\circ - b$ ), in which both agents are fully cooperative, and a *conflict-of-interest loss* ( $b$ ), in which both agents are fully competitive. The sum of  $L_1$  losses optimises only for the common interest, whereas  $L_2$  prefers a more fair division of the conflict-of-interest loss in addition to optimising common interest (see supplementary B.2 for proof).

An important note is that this  $L_2$  metric need only be used for choosing the best hyperparameters. Within our game, we can still use the  $L_1$  (or another metric) as the loss function. As well, we stress that this metric is still just a heuristic and the best run found may still exhibit manipulation (as is the case in our setup continuous message), and it is important to verify that both agents cooperatively benefit from communication. Furthermore, given the extensive use of hyperparameter optimization in RL [21], we believe our qualitative hyperparameter plot is an effective way to find models with good biases towards communication. Finally, we note that we can extend other cooperative communication metrics [25, 37] to competitive scenarios by adding the condition that both agents benefit from communication.

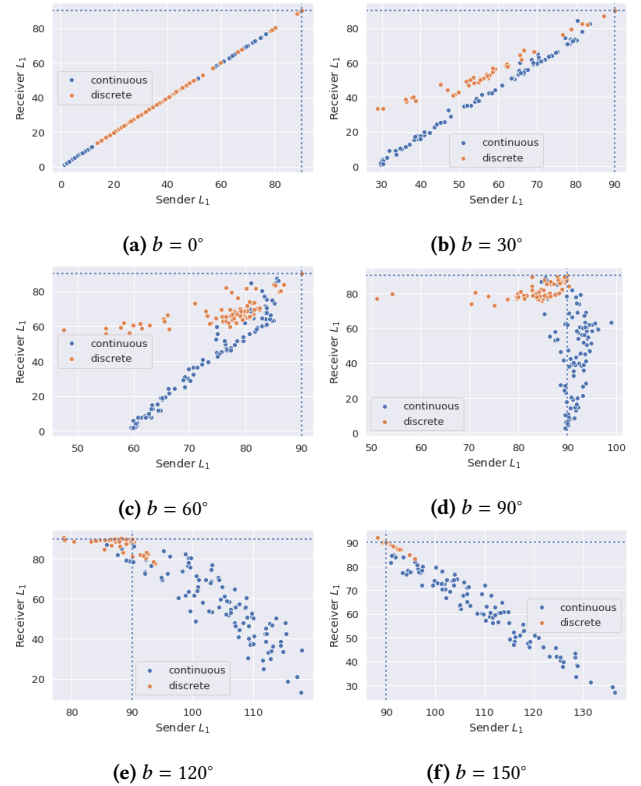
## 6 COMMUNICATION REQUIRES BOTH AGENTS TO POSSIBLY BENEFIT

### 6.1 The Negotiation Game

If both agents must benefit for communication to emerge, what happens if one agent is already at its optimal error? Cao et al. [6] proposed a negotiation game to study EC between agents  $A$  and  $B$  that are either fully-cooperative and collectively optimise total reward (‘pro-social’, reward-sharing) or competitive and independently optimise their own reward (‘selfish’). We study their ‘selfish’ setup<sup>4</sup> which they find fails to learn to communicate. They suggest the reason is either the game is too competitive [9] or not iterated. However, investigating more closely, the issue may be due to one agent being better off not communicating.

We describe the game briefly here; see the original paper for a full description. A pool of 3 types of items with quantity up to 5,  $i \in \{0, \dots, 5\}^3$  is initialised for a game between two agents. Each agent  $j$  is assigned a private utility  $u_j \in \{0, \dots, 10\}^3$  over each item and they negotiate with each other to split the items and try to gain a favourable split for themselves. At each turn,  $t$ , an agent creates a proposal for how to split up the objects  $p_j^t \in \{0, \dots, 5\}^3$ . If the other party accepts and the proposal is valid, agents are rewarded with  $R_j = p_j^t \cdot u_j$  and  $R_{1-j} = (i - p_j^t) \cdot u_j$ . If the proposal is invalid, both agents get 0 reward. If an agent refuses the proposal, it gives a counter-proposal and the game continues until either a proposal is accepted or the game goes past a limit  $N$  and both agents get

<sup>4</sup>‘Prosocial’ agents are fully-cooperative so they, by default, can learn to communicate optimally since communication is proportional to cooperation

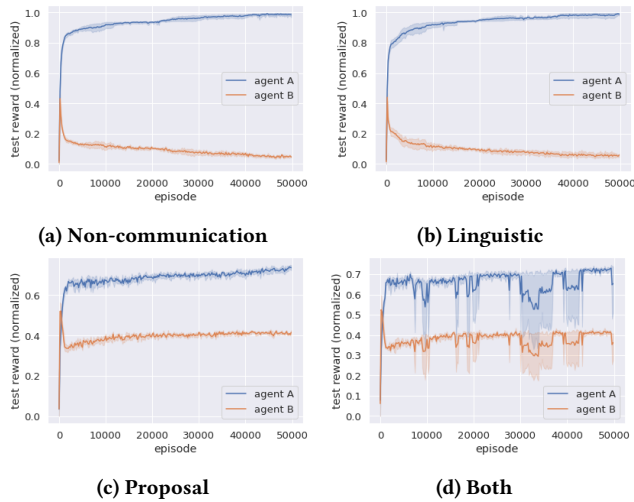


**Figure 4: Receiver vs Sender loss for all 100 hyperparameter runs for bias (a)  $0^\circ$  (b)  $30^\circ$  (c)  $60^\circ$  (d)  $90^\circ$  (e)  $120^\circ$  (f)  $150^\circ$  in the continuous and discrete message setups. We plot the baseline loss under non-communication ( $L_1 = 90^\circ$ ) as a dotted line. Loss higher than the non-communication baseline indicates manipulation. We report the loss averaged over 5 seeds.**


reward 0. Agent  $A$  always goes first. Lazaridou and Baroni [31] describe it as a modified, multi-turn version of the ultimatum game [19].

Cao et al. [6] experimented with four possible communication configurations. In the first case (‘non-communication’), agents do not communicate at all, so neither agent can see the other’s proposals. In the second case (‘proposal’), agents can see each other’s proposals. In the third case (‘linguistic’), agents can send each other a cheap talk message consisting of a sequence of tokens from a vocabulary. This corresponds to the EC case, and allows agents to learn to communicate either their proposal, utilities, or other information. The final configuration (‘both’) corresponds to agents having access to both the linguistic and the proposal channels of communication.

Cao et al. [6] find that for the ‘non-communication’ setting, agent  $B$  soundly dominates, taking nearly all the items. For the ‘linguistic’ setting, agents manage to split the item-pool relatively fairly but do not converge and exhibit chaotic dynamics. For the ‘proposal’ and ‘both’ settings, agents manage to converge to policies that split the item-pool perfectly evenly.



**Figure 5: Test reward per agent in our reproduction of the ‘selfish’ negotiation game [6] for a setup (a) without communication (b) with only the linguistic channel (c) with only the proposal channel (d) with both channels. We plot the mean and a 95% confidence interval over the seeds. Following the original paper, we report the reward divided by the total possible reward an agent could achieve.**



**Figure 6: The negotiation game using the masked-linguistic channel (a) reward per agent (b) percent of messages each agent unmasked**

## 6.2 Reproducing and Investigating

We reproduce the experiments and report our results in Figure 5. We largely manage to reproduce the results for the ‘proposal’ and ‘both’ settings, with the small difference that the reward is not perfectly split and agent *A* is now winning slightly. For the ‘non-communication’ setting our results are nearly identical except again, agent *A* is dominating instead of *B*. We believe this is reasonable as it matches results on the aforementioned ultimatum game [19] that gives the first-mover (*A*) an advantage. Stochastically varying the number of available turns  $N$ , as Cao et al. [6] propose, did not alleviate the first-mover advantage in our experiments.

However, we find noticeably different results for the ‘linguistic’ settings despite our best efforts.<sup>5</sup> Although Cao et al. [6] find that the game with a linguistic channel has both agents split the item

<sup>5</sup>Code for the original paper is not publicly accessible, so we contacted and worked with the original authors. We spent several weeks investigating after already having

pool equally, we find that the game is essentially unchanged from the ‘non-communication’ case. Our results in Figure 5b are nearly identical to those in Figure 5a.

## 6.3 Agents Learn Not To Communicate

To discover the reason for the discrepancy in results, we probe the behaviour of our agents. One possibility is that the game is too difficult to learn and there is an issue with optimization. To simplify, we use the linguistic policy to generating a binary mask  $m_j \in \{0, 1\}^3$  (instead of a sequence of tokens). We then set the agent’s linguistic message to the masked version of their proposal  $m_j \cdot u_j$ . Since proposals are grounded in the game (they determine the reward split), and learning a binary mask is much easier than a sequence with tokens from a vocabulary, we expect optimization issues to be greatly reduced. As well, this ‘masked-linguistic’ messaging system allows agents to explicitly hide or reveal their proposal to their opponent. In this way, we can probe whether the behaviour of an agent is to hide information or to share it.

We report our results in Figure 6 which shows that the reward (Figure 6a) with this channel is unchanged compared to the original ‘linguistic’ setup. This implies that agents are not having difficulty learning to communicate as it is incredibly easy to learn a mask  $[1, 1, 1]$  that would communicate their proposal precisely. Instead, looking at the percent of messages each agent unmasked (Figure 6b), it is clear agent *B* seems to be trying to communicate but agent *A* is specifically masking everything. Agent *A* explicitly chooses not to communicate their proposal.

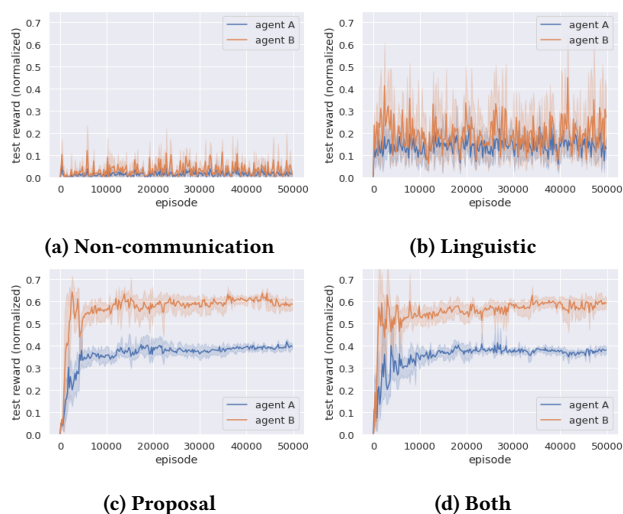
This is rational, because if agent *A* dominates in the case of non-communication, it is in their best interest to stay in that setup. By masking their proposal, they reduce this setting to the ‘non-communication’ setup as opposed to the ‘proposal’ setup where their reward is less. Cao et al. [6] suggest that agents *fail* to learn effective negotiation. Instead, we suggest that agent *A* specifically *chooses* to suppress any possible communication, preventing negotiation. This also suggests that our results in reproducing the negotiation game are more plausible as, in the original work, adding the cheap talk ‘linguistic’ channel causes the game to go from domination to a fair split.

From this we can infer that if an agent that has nothing to gain from communication, it should not be expected to learn to communicate. In particular, agents that dominate in a non-communicative setting (‘none’) but do worse when information is exchanged (‘proposal’) will explicitly learn not to communicate.

## 7 COMPETITIVE EMERGENT COMMUNICATION THROUGH NEGOTIATION

We have demonstrated three different results for competitive EC, but this inquiry started with the specific claim that independent agents could not learn to effectively communicate in the negotiation game [6]. We now examine the negotiation game from the perspective of our three takeaways. The game is not zero-sum as long as agents have different utilities for each item. Even though agents’ interests overlap, different utilities guarantee that the game

a working codebase but could not reproduce the results. All our code will be open-sourced and made available online.



**Figure 7: Reward in the negotiation game with termination through proposal agreement for setups (a) without communication (b) only linguistic channel (c) only proposal channel (d) both linguistic and proposal channels**

is general-sum so communication should be non-zero. An agent that manipulates instead of communicating would take all the items and this would be easy to spot. The only remaining issue is that agent A dominates in the non-communication case, making learning to communicate infeasible. By fixing the non-communication case, we should be able to learn to communicate in this game.

To alleviate the non-communication domination, we can change the conditions of termination to even the playing field. In the original work, agents explicitly output a termination signal that agrees to their opponent’s last proposal but agent B too easily learns to agree to whatever proposal it is given. We change the game to ignore the termination output and instead end the game when an agent outputs a proposal that is the same as their opponent’s last proposal,  $p_j^t == p_{1-j}^{t-1}$ . Thus, learning to terminate the game becomes more difficult and it is harder for agent A to take advantage of B. As well, it is more difficult to dominate in the non-communication case where you cannot communicate your proposal.

We report the results of using proposal termination in Figure 7. For ‘proposal’ and ‘both’, we see similar reward patterns to the original game except now agent B is slightly better than agent A. More importantly, though, is that the non-communication setting is *equally* badly for both agents and *both* are incentivised to communicate and improve their reward. With ‘linguistic’ channel agents successfully learn to communicate where they previous had not. Though the dynamics are slightly chaotic, the average reward for each agent is clearly above the non-communication baseline demonstrating effective information transfer. We also see both agents benefit from the communication so we can be sure this is not manipulation. The total reward utilization is low but this can be explained by the highly competitive nature of the negotiation game. Since agents always have overlapping interests, depending on the utilities each agent has for each item, the game may be much

closer to fully-competitive than fully-cooperative. Overall, with a small modification we have shown that even in the negotiation game, competitive agents can learn to communicate.

## 8 CONCLUSION

Importantly, we have demonstrated the first clear, positive result of EC under competition in modern machine learning using standard RL. We have expanded on conditions for competitively learning to communicate and presented three results. (1) We provide evidence that information transfer is proportional to cooperation, applying the theoretical results of Crawford and Sobel [9] to modern EC. (2) We highlight a subtle distinction between *information transfer*, *communication*, and *manipulation*, and demonstrated qualitative and quantitative approaches to differentiating the last two. We argue that you can extend metrics of cooperative EC [37] to competitive cases with the caveat that both agents must benefit from communication. This is especially useful for reflecting on the metric used to choose the best hyperparameters in general-sum games [25]. (3) We demonstrate that *both* agents must be able to benefit from communication for it to be feasible. To do so, we reproduced one of the first investigations into competitive EC, the negotiation game [6], and clarified their hypothesis for why communication did not emerge. Finally, we have combined our three takeaways to argue there is only a single issue preventing successful communication in the negotiation game. We have gone back to this first, modern negative result of competitive EC and with a small change, we have successfully learned to negotiate between competitive agents.

Modern EC can be seen through two overarching lines of research: understanding the properties of emergent languages and using EC to improve AI systems [31]. Research in both of these goals has been fruitful. However, it is usually limited to the case of fully-cooperative, reward-sharing agents which may be restricting to the fields and enforce a (false) perspective that communication must be fully cooperative. In contrast, other fields have already been using it as a tool for cooperation and coordination in partially-competitive environments [9, 15].

We believe that competitive EC can be useful to both avenues of research. In understanding emergent languages, competition can be another dimension to investigate which protocols emerge—e.g., how competition may regularise a community of speakers. Another line of research could investigate how more complex protocols change with the level of competition. For improving AI systems, competitive EC offer a more realistic perspective on human communication. For agents to understand that how human communication can lie or hide the truth, they must be at least partially competitive. And, for communication protocols to be resilient to lying, they may also need to be trained with competitive agents. We hope this paper invigorates research into competitive EC and highlights key takeaways for future work to progress this new and exciting field.

## ACKNOWLEDGMENTS

MN thanks Adrien Ali Taïga, David Krueger, Ryan Lowe, and reviewers for feedback, Hugh Perkins for initial code, and Compute-Canada for resources. MN thanks Microsoft Research, MN and AC thank CIFAR Canada AI Chair, and TL thanks Schwartz Reisman Institute at the University of Toronto for supporting this work.



## REFERENCES

- [1] Nolan Bard, Jakob N. Foerster, A. P. Sarath Chandar, Neil Burch, Marc Lanctot, H. Francis Song, Emilio Parisotto, Vincent Dumoulin, Subhodeep Moitra, Edward Hughes, Iain Dunning, Shibli Mourad, Hugo Larochelle, Marc G. Bellemare, and Michael H. Bowling. 2019. The Hanabi Challenge: A New Frontier for AI Research. *ArXiv abs/1902.00506* (2019).
- [2] Jeffrey A. Barrett and Brian Skyrms. 2017. Self-Assembling Games. *The British Journal for the Philosophy of Science* 68, 2 (2017), 329–353.
- [3] Lukas Biewald. 2020. Experiment Tracking with Weights and Biases. <https://www.wandb.com/> Software available from wandb.com.
- [4] Ben Bogin, Mor Geva, and Jonathan Berant. 2018. Emergence of communication in an interactive world with consistent speakers. *arXiv preprint arXiv:1809.00549* (2018).
- [5] Xavier Bouthillier, Christos Tsirigotis, François Corneau-Tremblay, Pierre Delaunay, Michael Noukhovitch, Reyhane Askari, Peter Henderson, Dendi Suhubdy, Frédéric Bastien, and Pascal Lamblin. 2019. Orion - Asynchronous Distributed Hyperparameter Optimization. <https://github.com/Epistimio/orion>.
- [6] Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z. Leibo, Karl Tuyls, and Stephen Clark. 2018. Emergent Communication through Negotiation. *CoRR abs/1804.03980* (2018). arXiv:1804.03980 <http://arxiv.org/abs/1804.03980>
- [7] Rahma Chaabouni, Eugene Kharitonov, Diane Bouchacourt, Emmanuel Dupoux, and Marco Baroni. 2020. Compositionality and Generalization In Emergent Languages. In *ACL 2020: 58th annual meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 4427–4442.
- [8] Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. 2019. Anti-efficient encoding in emergent communication. In *NeurIPS 2019 : Thirty-third Conference on Neural Information Processing Systems*. 6293–6303.
- [9] Vincent P. Crawford and Joel Sobel. 1982. Strategic Information Transmission. *Econometrica* 50, 6 (1982), 1431–1451. <https://doi.org/10.2307/1913390>
- [10] Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. 2019. TarMAC: Targeted Multi-Agent Communication. In *ICML*.
- [11] Richard Dawkins and John R. Krebs. 1978. Animal signals: Information or Manipulation? In *Behavioural Ecology*, J. R. Krebs and N. B. Davies (Eds.). Blackwell Scientific Publications, Oxford, 282–309.
- [12] Tom Eccles, Yoram Bachrach, Guy Lever, Angeliki Lazaridou, and Thore Graepel. 2019. Biases for Emergent Communication in Multi-agent Reinforcement Learning. In *NeurIPS 2019 : Thirty-third Conference on Neural Information Processing Systems*. 13111–13121.
- [13] Diane Bouchacourt Evgeny Kharitonov, Rahma Chaabouni and Marco Baroni. 2020. Entropy Minimization In Emergent Languages. In *ICML 2020: 37th International Conference on Machine Learning*.
- [14] Katrina Evtimova, Andrew Drodzov, Douwe Kiela, and Kyunghyun Cho. 2017. Emergent communication in a multi-modal, multi-step referential game. *arXiv preprint arXiv:1705.10369* (2017).
- [15] Joseph Farrell and Matthew Rabin. 1996. Cheap Talk. *The Journal of Economic Perspectives* 10, 3 (1996), 103–118.
- [16] Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*. 2137–2145.
- [17] Michael Fu. 2006. Chapter 19 Gradient Estimation. *Handbooks in Operations Research and Management Science* 13 (12 2006). [https://doi.org/10.1016/S0927-0507\(06\)13019-4](https://doi.org/10.1016/S0927-0507(06)13019-4)
- [18] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- [19] Werner Güth, Rolf Schmittberger, and Bernd Schwarze. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization* 3 (1982), 367–388. Issue 4.
- [20] Serhii Havrylov and Ivan Titov. 2017. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In *Advances in neural information processing systems*. 2149–2159.
- [21] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. 2018. Deep reinforcement learning that matters. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [22] Robert A. Hinde. 1981. Animal Signals: Ethological and Games-Theory Approaches are Not Incompatible. *Animal Behavior* 29 (1981), 535–542.
- [23] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. 2020. "Other-Play" for Zero-Shot Coordination. arXiv:2003.02979 [cs.AI]
- [24] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparameterization with Gumbel-Softmax. In *ICLR 2017 : International Conference on Learning Representations 2017*.
- [25] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Çağlar Gülçehre, Pedro A. Ortega, DJ Strouse, Joel Z. Leibo, and Nando de Freitas. 2018. Intrinsic Social Motivation via Causal Influence in Multi-Agent RL. *CoRR abs/1810.08647* (2018). arXiv:1810.08647 <http://arxiv.org/abs/1810.08647>
- [26] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.
- [27] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR abs/1412.6980* (2014).
- [28] John R. Krebs and Richard Dawkins. 1984. Animal Signals: Mind Reading and Manipulation. In *Behavioural Ecology*, J. R. Krebs and N. B. Davies (Eds.). Sinauer Associates, Sunderland, MA, 380–402.
- [29] Travis LaCroix. 2020. Communicative Bottlenecks Lead to Maximal Information Transfer. *Journal of Experimental and Theoretical Artificial Intelligence* 32, 6 (2020), 997–1014.
- [30] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Perolat, David Silver, and Thore Graepel. 2017. A Unified Game-Theoretic Approach to Multiagent Reinforcement Learning. In *NIPS*.
- [31] Angeliki Lazaridou and Marco Baroni. 2020. Emergent Multi-Agent Communication in the Deep Learning Era. (2020).
- [32] Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. 2018. Emergence of Linguistic Communication from Referential Games with Symbolic and Pixel Input. *CoRR abs/1804.03984* (2018). arXiv:1804.03984 <http://arxiv.org/abs/1804.03984>
- [33] Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2016. Multi-agent cooperation and the emergence of (natural) language. *arXiv preprint arXiv:1612.07182* (2016).
- [34] Joel Z. Leibo, Vinicius Flores Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-agent Reinforcement Learning in Sequential Social Dilemmas. *CoRR abs/1702.03037* (2017). arXiv:1702.03037 <http://arxiv.org/abs/1702.03037>
- [35] David Lewis. 1969. *Convention: A Philosophical Study*. Wiley-Blackwell, Oxford.
- [36] Paul Pu Liang, Jeffrey Chen, Ruslan Salakhutdinov, Louis-Philippe Morency, and Satwik Kottur. 2020. On Emergent Communication in Competitive Multi-Agent Teams. (2020), 735–743.
- [37] Ryan Lowe, Jakob N. Foerster, Y-Lan Boureau, Joelle Pineau, and Yann Dauphin. 2019. On the Pitfalls of Measuring Emergent Communication. In *AAMAS*.
- [38] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *ICLR 2017 : International Conference on Learning Representations 2017*.
- [39] Igor Mordatch and Pieter Abbeel. 2017. Emergence of Grounded Compositional Language in Multi-Agent Populations. *CoRR abs/1703.04908* (2017). arXiv:1703.04908 <http://arxiv.org/abs/1703.04908>
- [40] Vinod Nair and Geoffrey Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *ICML*, Vol. 27. 807–814.
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [42] John Schulman, Nicolas Manfred Otto Heess, Theophane Weber, and Pieter Abbeel. 2015. Gradient Estimation Using Stochastic Computation Graphs. In *NIPS*.
- [43] Yoav Shoham, Rob Powers, and Trond Grenager. 2003. Multi-Agent Reinforcement Learning: a critical survey. (2003).
- [44] Amanpreet Singh, Tushar Jain, and Sainbayar Sukhbaatar. 2018. Learning when to Communicate at Scale in Multiagent Cooperative and Competitive Tasks. In *ICLR*.
- [45] Brian Skyrms. 2010. *Signals: Evolution, Learning, & Information*. Oxford University Press, Oxford.
- [46] Brian Skyrms and Jeffrey A. Barrett. 2018. Propositional Content in Signals. <http://philsci-archive.pitt.edu/14774/>
- [47] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction* (second ed.). The MIT Press.
- [48] Elliott O. Wagner. 2012. Deterministic Chaos and the Evolution of Meaning. *British Journal for the Philosophy of Science* 63 (2012), 547–575.
- [49] Kyle Wagner, James A Reggia, Juan Uriagereka, and Gerald S Wilkinson. 2003. Progress in the simulation of emergent communication and language. *Adaptive Behavior* 11, 1 (2003), 37–69.
- [50] Ronald J. Williams. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine Learning* 8 (1992), 229–256.