

Miss Tools and Mr Fruit: Emergent communication in agents learning about object affordances

Diane Bouchacourt¹ and Marco Baroni^{1,2}

¹Facebook A.I. Research

²ICREA

{diane, mbaroni}@fb.com

Abstract

Recent research studies communication emergence in communities of deep network agents assigned a joint task, hoping to gain insights on human language evolution. We propose here a new task capturing crucial aspects of the human environment, such as natural object affordances, and of human conversation, such as full symmetry among the participants. By conducting a thorough pragmatic and semantic analysis of the emergent protocol, we show that the agents solve the shared task through genuine bilateral, referential communication. However, the agents develop multiple idiolects, which makes us conclude that full symmetry is not a sufficient condition for a common language to emerge.

1 Introduction

The advent of powerful deep learning architectures has revived research in simulations of language emergence among computational agents that must communicate to accomplish a task (e.g., Jorge et al., 2016; Havrylov and Titov, 2017; Kottur et al., 2017; Lazaridou et al., 2017; Lee et al., 2017; Choi et al., 2018; Evtimova et al., 2018; Lazaridou et al., 2018). The nature of the emergent communication code should provide insights on questions such as to what extent comparable functional pressures could have shaped human language, and whether deep learning models can develop human-like linguistic skills. For such inquiries to be meaningful, the designed setup should reflect as many aspects of human communication as possible. Moreover, appropriate tools should be applied to the analysis of emergent communication, since, as several recent studies have shown, agents might succeed at a task without truly relying on their communicative channel, or by means of *ad-hoc* communication techniques

overfitting their environment (Kottur et al., 2017; Bouchacourt and Baroni, 2018; Lowe et al., 2019).

We contribute on both fronts. We introduce a game meeting many desiderata for a natural communication environment. We further propose a two-pronged analysis of emerging communication, at the *pragmatic* and *semantic* levels. At the pragmatic level, we study communicative acts from a functional perspective, measuring whether the messages produced by an agent have an impact on the subsequent behaviour of the other. At the semantic level, we decode which aspects of the extra-linguistic context the agents refer to, and how such reference acts differ between agents. Some of our conclusions are positive. Not only do the agents solve the shared task, but genuine bilateral communication helps them to reach higher reward. Moreover, their referential acts are meaningful given the task, carrying the semantics of their input. However, we also find that even perfectly symmetric agents converge to distinct idiolects instead of developing a single, shared code.

2 The fruit and tools game

Our game, inspired by Tomasello’s (2014) conjecture that the unique cognitive abilities of humans arose from the requirements of cooperative interaction, is schematically illustrated in Fig. 1. In each episode, a randomly selected agent is presented with instances of two tools (*knife*, *fork*, *axe*...), the other with a fruit instance (*apple*, *pear*, *plum*...). Tools and fruits are represented by property vectors (e.g., *has a blade*, *is small*), with each instance characterized by values randomly varying around the category mean (e.g., an apple instance might be smaller than another). An agent is randomly selected to be the first to perform an action. The game then proceeds for an arbitrary number of *turns*. At each turn, one of




Figure 1: Our game. One agent receives a fruit, another two tools. Each agent sends a message in turn, until an agent ends the episode by choosing a tool. The agents are rewarded if the tool choice is optimal given the fruit.

the agents must decide whether to pick one of the two tools and stop, or to continue, in which case the message it utters is passed to the other agent, and the game proceeds. Currently, for ease of analysis, messages are single discrete symbols selected from a vocabulary of size 10, but extension to symbol sequences is trivial (although it would of course complicate the analysis). As soon as an agent picks a tool, the game ends. The agents receive a binary reward of 1 if they picked the better tool for the fruit at hand, 0 otherwise. The best choice is computed by a utility function that takes into account the interaction between tool and fruit instance properties (e.g., as in Fig. 1, a tool with a round edge might be particularly valuable if the fruit has a pit). Utility is relative: given a peach, the axe is worse than the spoon, but it would be the better tool when the alternative is a hammer.

Here are some desirable properties of our setup, as a simplified simulation of human interactions. The agents are fully symmetric and cannot specialize to a fixed role or turn-taking scheme. The number of turns is open and determined by the agents. In pure signaling/referential games (Lewis, 1969), the aim is successful communication itself. In our game, reward depends instead on tool and fruit affordances. Optimal performance can only be achieved by jointly reasoning about the properties of the tools and how they relate to the fruit. Humans are rewarded when they use language to solve problems of this sort, and not for successful acts of reference *per se*. Finally, as we use commonsense descriptions of everyday objects to build our dataset (see below), the distribution of their properties possesses the highly skewed characteristics encountered everywhere in the human environment (Li, 2002). For example, if the majority of fruits requires to be cut, a knife is intrinsically more useful than a spoon. Note that the agents do not have any a priori knowledge of the tools utility. Yet, baseline agents are able to dis-

cover context-independent tool affordances and already reach high performance. We believe that this scenario, in which communication-transmitted information complements knowledge that can be directly inferred by observing the world, is more interesting than typical games in which language is the only information carrier.

Game ingredients and utility We picked 16 tool and 31 fruit categories from McRae et al. (2005) and Silberer et al. (2013), who provide subject-elicited property-based commonsense descriptions of objects, with some extensions. We used 11 fruit and 15 tool features from these databases to represent the categories. We rescaled the elicitation-frequency-based property values provided in the norms to lie in the $[0, 1]$ range, and manually changed some counter-intuitive values. An object instance is a property vector sampled from the corresponding category as follows. For binary properties such as *has a pit*, we use Bernoulli sampling with p equaling the category value. For continuous properties such as *is small*, we sample uniformly from $[\mu - 0.1, \mu + 0.1]$, where μ is the category value. We then devised a function returning an utility score for any fruit-tool property vector pair. The function maps properties to a reduced space of abstract functional features (such as *break* for tools, and *hard* for fruits). Details are in Supplementary Section 1. For example, an apple with *is crunchy*=0.7 value gets a high *hard* functional feature score. A knife with *has a blade*=1 gets a high *cut* score, and therefore high utility for the *hard* apple. Some features, e.g., *has a handle* for tools, have no impact on utility. They only represent realistic aspects of objects and act as noise. Our dataset with full category property vectors will be publicly released along with code.

Datasets We separate the 31 fruit categories into three sets: in-domain (21 categories), validation and transfer (5 categories each). The in-domain set is further split into train and test parti-

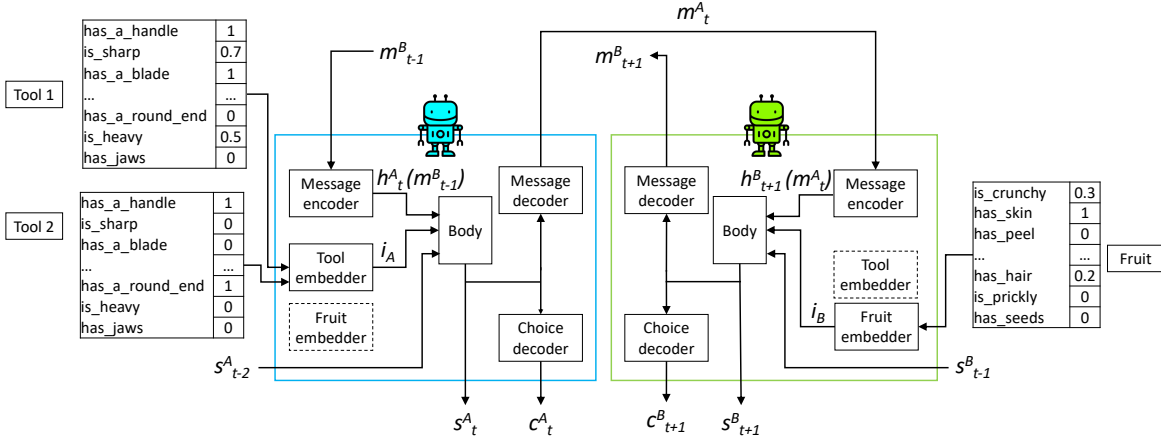


Figure 2: Two turns of dialogue. Dashed boxes are not used in this game episode due to the agent roles (the blue/left agent is Tool Player, the green/right one is Fruit player). The flow is explained in detail in the text.

tions. We train agents on the in-domain train partition and monitor convergence on the validation set. We report test performance on the in-domain test partition and on the transfer set. For example, the *peach* category is in-domain, meaning that distinct peach instances will be seen at training and in-domain testing time. The *nectarine* category is in the transfer set, so nectarine instances will only be seen at test time. This scheme tests the generalization abilities of the agents (that can generalize to new fruits since they are all described in the same feature space). We generate 210,000 in-domain training samples and 25,000 samples for the other sets, balanced across fruits and tools (that are common across the sets).

Game dynamics and agent architecture At the beginning of a game episode, two neural network agents A and B receive, randomly, either a pair of tools ($tool_1, tool_2$) (always sampled from different categories) or a *fruit*. The agent receiving the tools (respectively the fruit) will be Tool Player (respectively Fruit Player) for the episode.¹ The agents are also randomly assigned positions, and the one in position 1 starts the game. Figure 2 shows two turns of the game in which A (blue/left) is Tool Player and in position 1. The first turn is indexed $t = 0$, therefore A will act on even turns, B on odd turns. At game opening, each agent passes its input (tool pair or fruit) through a linear layer followed by a tanh nonlinearity, resulting in embedding i_A (resp. i_B). Then, at each turn t , an agent, for example A, receives the message m_{t-1}^B from agent B, and accesses its own previous internal state s_{t-2}^A (we refer to “memory” the addition

¹Agents must learn to recognize the assigned role.

of the agent’s previous state). The message m_{t-1}^B is processed by a RNN, and the resulting hidden state $h_t^A(m_{t-1}^B)$ is concatenated with the agent previous internal state s_{t-2}^A and the input embedding i_A . The concatenated vector is fed to the Body module, composed of a linear layer followed by tanh. The output of the Body module is the new A state, s_t^A , fed to Message and Choice decoders.

The Message decoder is an RNN with hidden state initialized as s_t^A , and outputting a probability distribution $p(m_t^A | s_t^A)$ over possible A messages. At training time, we sample a message m_t^A ; at test time we take the most probable one. The Choice decoder is a linear layer processing s_t^A and outputting a softmax-normalized vector of size 3. The latter represents the probabilities $p(c_t^A | s_t^A)$ over A’s possible choices: (i) $c_t^A = 0$ to continue playing, (ii) $c_t^A = 1$ to choose tool $tool_1$ and stop (iii) $c_t^A = 2$ to choose tool $tool_2$ and stop. Again, we sample at training and argmax at test time. If $c_t^A = 0$, the game continues. B receives message m_t^A , its previous state s_{t-1}^B and input embedding i_B , and it outputs the tuple $(m_{t+1}^B, c_{t+1}^B, s_{t+1}^B)$ etc., until an agent stops the game, or the maximum number of turns $T_{\max} = 20$ is reached.

When an agent stops by choosing a tool, for example $tool_1$, we compute the two utilities $U(tool_1, fruit)$ and $U(tool_2, fruit)$. If $U(tool_1, fruit) \geq U(tool_2, fruit)$, that is the best tool was chosen, shared reward is $R = 1$. If $U(tool_1, fruit) < U(tool_2, fruit)$ or if the agents reach T_{\max} turns without choosing, $R = 0$.² During learning, the reward is back-propagated

²We also tried directly using raw or normalized scalar utilities as rewards, with similar performances.

with Reinforce (Williams, 1992). When the game starts at $t = 0$, we feed the agent in position 1 a fixed dummy message m^0 , and the previous states of the agents s_{t-2}^A and s_{t-1}^B are initialized with fixed dummy s^0 . In the no-memory ablation, previous internal states are always replaced by s^0 . When we block communication, agent messages are replaced by m^0 . Supplementary Section 2 provides hyperparameter and training details.

3 Measuring communication impact

Message Effect is computed on single turns and uses causal theory (Pearl et al., 2016) to quantify how much what an agent utters impacts the other, compared to the counterfactual scenario in which the speaking agent said something else.

Consider message m_t^A uttered by A at turn t . If $c_t^A = 0$ (that is, A continues the game), m_t^A is processed by B, along with s_{t-1}^B and i^B . At the next turn, B outputs a choice c_{t+1}^B and a message m_{t+1}^B drawn from $p(c_{t+1}^B, m_{t+1}^B | s_{t+1}^B)$. B’s state s_{t+1}^B is deterministically determined by $m_t^A, c_t^A, s_{t-1}^B, i^B$, so we can equivalently write that c_{t+1}^B and m_{t+1}^B are drawn from $p(c_{t+1}^B, m_{t+1}^B | m_t^A, c_t^A, s_{t-1}^B, i^B)$. Conditioning on c_t^A, s_{t-1}^B, i^B ensures there are no confounders when we analyze the influence from m_t^A (Pearl et al., 2016). Supplementary Figure A1 shows the causal graph supporting our assumptions. We will not from here onwards write the conditioning on c_t^A, s_{t-1}^B, i^B explicitly.

We define $z_{t+1}^B = (c_{t+1}^B, m_{t+1}^B)$. We want to estimate how much the message from A, m_t^A , influences the next-turn behaviour (choice and message) of B, z_{t+1}^B . We thus measure the discrepancy between the conditional distribution $p(z_{t+1}^B | m_t^A)$ and the marginal distribution $p(z_{t+1}^B)$ not taking m_t^A into account. However, we want to assess agent B’s behaviour under *other possible received messages* m_t^A . To do so, when we compute the marginal of agent B’s $p(z_{t+1}^B)$, we intervene on m_t^A and draw the messages from the *intervention distribution*. We define $\tilde{p}(z_{t+1}^B)$, the marginal computed with counterfactual messages m_t^A , as:

$$\tilde{p}(z_{t+1}^B) = \sum_{m_t^A} p(z_{t+1}^B | m_t^A) \tilde{p}(m_t^A) \quad (1)$$

where $\tilde{p}(m_t^A)$ is the intervention distribution, different from $p(m_t^A | s_t^A)$. If at turn t , A continues the game, we define the *Message Effect* (ME) from agent A’s message m_t^A on agent B’s choice and

Given the message m_t^A from agent A.

- 1 Sample K pairs $z_{t+1,k}^B \sim p(z_{t+1,k}^B | m_t^A)$.
- 2 Sample J counterfactuals $m_{t,j}^A \sim \tilde{p}(m_t^A)$.
- 3 For each $z_{t+1,k}^B$, do

$$\tilde{p}(z_{t+1,k}^B) = \sum_{j=1}^J p(z_{t+1,k}^B | m_{t,j}^A) \tilde{p}(m_{t,j}^A).$$
- 4 Return

$$\text{ME}_t^{A \rightarrow B} = \frac{1}{K} \sum_{k=1}^K \log \frac{p(z_{t+1,k}^B | m_t^A)}{\tilde{p}(z_{t+1,k}^B)}.$$

Algorithm 1 Computation of $\text{ME}_t^{A \rightarrow B}$.

message pair, z_{t+1}^B as:

$$\text{ME}_t^{A \rightarrow B} = \text{KL} \left(p(z_{t+1}^B | m_t^A) \parallel \tilde{p}(z_{t+1}^B) \right) \quad (2)$$

where KL is the Kullback-Leibler divergence, and $\tilde{p}(z_{t+1}^B)$ is computed as in Eq. 1. This allows us to measure how much the conditional distribution differs from the marginal. Algorithm 1 shows how we estimate $\text{ME}_t^{A \rightarrow B}$. In our experiments, we draw $K = 10$ samples $z_{t+1,k}^B$, and use a uniform intervention distribution $\tilde{p}(m_t^A)$ with $J = 10$. This kind of counterfactual reasoning is explored in depth by Bottou et al. (2013). Jaques et al. (2018) and Lowe et al. (2019) present related measures of causal impact based on the Mutual Information (MI) between influencing and influenced agents. We discuss in Supplementary Section 3 possible issues with the MI-based approach.

Bilateral communication Intuitively, there has been a proper dialogue if, in the course of a conversation, each agent has said at least one thing that influenced the other. We operationalize this through our *bilateral communication* measure. This is a binary, per-game score, that is positive only if in the game there has been at least one turn with significant message effect in *both directions*, i.e., $\exists t, t' \text{ s.t. } \text{ME}_t^{A \rightarrow B} > \theta$ and $\text{ME}_{t'}^{B \rightarrow A} > \theta$. We set $\theta = 0.1$.³

³We considered setting θ to (i) the average ME returned by untrained agents, but this led to a threshold extremely close to 0, and (ii) the average of the agents’ ME values, but this counterintuitively penalized pairs of agents with high overall communication influence.

4 Results

We first confirm that the agents succeed at the task, and communication improves their performance. Second, we study their *pragmatics*, looking at how ablating communication and memory affect their interaction. Finally, we try to interpret the *semantics* of the agents’ messages.

4.1 Performance and pragmatics

We report mean and standard error of the mean (SEM) over successful training seeds.⁴ Each agent A or B can be either F (Fruit Player) or T (Tool Player) and in position 1 or 2, depending on the test game. We measure to what extent Tool Player influences Fruit Player ($ME^{T \rightarrow F}$) and *vice versa* ($ME^{F \rightarrow T}$). Similarly, we evaluate position impact by computing $ME^{1 \rightarrow 2}$ and $ME^{2 \rightarrow 1}$. We average ME values over messages sent during each test game, and report averages over test games. Note that we also intervene on the dummy initialization message used at $t = 0$, which is received by the agent in position 1. This impacts the value of $ME^{2 \rightarrow 1}$. If the agent in position 1 has learned to rely on the initialization message to understand that the game is beginning, an intervention on this message will have an influence we want to take into account.⁵ Similarly, in the no-communication ablation, when computing ME values, we replace the dummy fixed message the agents receive with a counterfactual. Finally, we emphasize that the computation of ME values does not interfere with game dynamics and does not affect performance.

Both communication and memory help Table 1 shows that enabling the agents to communicate greatly increases performance compared to the no-communication ablation, both with and without memory, despite the high baseline set by agents that learn about tool usefulness without communicating (see discussion below). Agents equipped with memory perform better than their no-memory counterparts, but the gain in performance is smaller compared to the gain attained from communication. The overall best performance is achieved with communication *and* memory. We also see that the agents generalize well, with transfer-fruit performance almost matching that on in-domain fruits. Next, we analyze in detail

⁴That is, training seeds leading to final validation performance above 85%.

⁵Conversely, we ignore the messages agents send when stopping the game, as they are never heard.

the impact of each factor (communication, memory) on agent performance and strategies.

No-communication, no-memory We start by looking at how the game unfolds when communication and agent memory are ablated (top left quadrant of Table 1). Performance is largely above chance ($\approx 50\%$), because, as discussed in Section 2, some tools are intrinsically better on average across fruits than others. Without communication, the agents exploit this bias and learn a strategy where (i) Fruit Player never picks the tool but always continues the game and (ii) Tool Player picks the tool according to average tool usefulness. Indeed, Tool Player makes the choice in more than 99% of the games. Conversation length is 0 if Tool Player starts and 1 if it is the second agent, requiring the starting Fruit Player to pass its turn. Reassuringly, ME values are low, confirming the reliability of this communication score, and indicating that communication-deprived agents did not learn to rely on the fixed dummy message (e.g., by using it as a constant bias). Still, we observe that, across the consistently low values, Fruit Player appears to affect Tool Player significantly more than the reverse ($ME^{F \rightarrow T} > ME^{T \rightarrow F}$). This is generally observed in all configurations, and we believe it due to the fact that Tool Player takes charge of most of the reasoning in the game. We come back to this later in our analysis. We also observe that the second player impacts the first more than the reverse ($ME^{2 \rightarrow 1} > ME^{1 \rightarrow 2}$). We found this to be an artifact of the strategy adopted by the agents. In the games in which Tool Player starts and immediately stops the game, we can only compute ME for the Tool/position-1 agent, by intervening on the initialization. The resulting value, while tiny, is unlikely to be exactly 0. In the games where Fruit Player starts and Tool Player stops at the second turn, we compute instead two tiny MEs, one per agent. Hence, the observed asymmetry. We verified this hypothesis by removing single-turn games: the influence of the second player on the first indeed disappears.

Impact of communication The top quadrants of Table 1 show that communication helps performance, despite the high baseline set by the “average tool usefulness” strategy. Importantly, when communication is added, we see a dramatic increase in the proportion of games with bilateral communication, confirming that improved performance is not due to an accidental effect of adding

	Metric	No communication		With communication	
		In	Transfer	In	Transfer
No memory	Av. perf. (%)	84.83 ± 0.09	84.0 ± 0.11	96.9 ± 0.32	94.5 ± 0.37
	$ME^{F \rightarrow T}$	0.133* ± 0.01	0.14* ± 0.01	5.0* ± 0.39	5.0* ± 0.36
	$ME^{T \rightarrow F}$	0.05 ± 0.02	0.030 ± 0.01	3.9 ± 0.38	3.3 ± 0.30
	$ME^{1 \rightarrow 2}$	0.066 ± 0.00	0.067 ± 0.01	3.9 ± 0.29	3.7 ± 0.26
	$ME^{2 \rightarrow 1}$	0.12* ± 0.02	0.10* ± 0.01	5.0* ± 0.38	4.7* ± 0.33
	Bi. comm. (%)	1.4 ± 0.31	1.3 ± 0.40	88 ± 2.49	89 ± 2.24
	Av. conv. length	0.508 ± 0.01	0.52 ± 0.01	2.16 ± 0.08	2.21 ± 0.10
	T chooses (%)	99.4 ± 0.63	99.6 ± 0.56	85 ± 2.09	83 ± 2.56
With memory	Av. perf. (%)	88.5 ± 0.11	87.7 ± 0.16	97.4 ± 0.12	95.3 ± 0.16
	$ME^{F \rightarrow T}$	0.11* ± 0.01	0.13* ± 0.01	3.0* ± 0.29	2.8* ± 0.24
	$ME^{T \rightarrow F}$	0.064 ± 0.01	0.071 ± 0.01	1.8 ± 0.22	1.8 ± 0.21
	$ME^{1 \rightarrow 2}$	0.085 ± 0.01	0.10 ± 0.01	2.4 ± 0.29	2.3 ± 0.22
	$ME^{2 \rightarrow 1}$	0.093 ± 0.01	0.103 ± 0.01	2.4 ± 0.22	2.4 ± 0.21
	Bi. comm. (%)	3.8 ± 0.61	4.6 ± 0.68	78 ± 2.55	78 ± 2.65
	Av. conv. length	1.50 ± 0.06	1.46 ± 0.06	2.7 ± 0.11	2.7 ± 0.11
	T chooses (%)	87.3 ± 1.34	85.8 ± 1.48	81 ± 2.94	81 ± 3.00

Table 1: Test performance and pragmatic measures mean and SEM in different settings. “Av. perf.” (average performance) denotes % of samples where best tool was chosen, “Bi. comm.” denotes % of games with bilateral communication taking place. “Av. conv. length” is average conversation length in turns. “T chooses” denotes % of games ended by Tool Player. Values of ME with an asterisk * are statistically significantly higher than their reverse (e.g. $ME^{F \rightarrow T} > ME^{T \rightarrow F}$). Best “Av. perf.” and “Bi. comm.” in bold.

a new channel (Lowe et al., 2019). ME and average number of turns also increase. Fruit Player is the more influential agent. This effect is not due to the artifact we found in the no-communication ablation, because almost all conversations, including those started by Tool Player, are longer than one turn, so we can compute both $ME^{F \rightarrow T}$ and $ME^{T \rightarrow F}$. We believe the asymmetry to be due to the fact that Tool Player is the agent that demands more information from the other, as it is the one that sees the tools, and that in the large majority of cases makes the final choice. Supplementary Table A4 shows that the gap between the influence of the Fruit Player on the Tool player and its reverse is greater when the Fruit Player is in position 2. This, then, explains $ME^{2 \rightarrow 1} > ME^{1 \rightarrow 2}$ as an epiphenomenon of Fruit Player being more influential.

Is memory ablation necessary for communication to matter? An important observation from previous research is that depriving at least one agent of memory might be necessary to develop successful multi-turn communication (Kottur et al., 2017; Cao et al., 2018; Evtimova et al., 2018). This is undesirable, as obviously language should not emerge simply as a surrogate for mem-

ories of amnesiac agents. The performance and communicative behaviours results in the bottom right quadrant of Table 1 show that, in our game, genuine linguistic interaction (as cued by ME and bilateral communication scores) is present even when both agents are equipped with memory. It is interesting however to study how adding memory affects the game dynamics independently of communication. In the bottom left quadrant, we see that memory leads to some task performance improvement for communication-less agents. Manual inspection of example games reveals that such agents are developing turn-based strategies. For example, Tool Player learns to continue the game at turn t if $tool_1$ has a round end. At $t + 1$, Fruit Player can use the fact that Tool Player continues at t as information about relative tool roundness, and either pick the appropriate one based on the fruit or continue to gather more information. In a sense, agents learn to use the possibility to stop or continue at each turn as a rudimentary communication channel. Indeed, exchanges are on average longer when memory is involved, and turn-based strategies appear even with communication. In the latter case, agents rely on communication but also on turn-based schemes, resulting in lower ME val-

Messages	Fruit	Tool 1	Tool 2
Both	37 ± 1.70	31 ± 1.21	24 ± 1.07
F	37 ± 1.75	23.3 ± 0.66	16.7 ± 0.51
T	14.1 ± 0.79	32 ± 1.17	25 ± 1.04
Stats. (%)	5.786 ± 0.00	8.76 ± 0.01	7.682 ± 0.01

Table 2: Semantic classifier % accuracy mean and SEM over successful training seeds.

ues and bilateral communication compared to the no-memory ablation. Finally, the respective positions of the agents in the conversation no longer impact ME ($ME^{1 \rightarrow 2} \approx ME^{2 \rightarrow 1}$). This might be because, with memory, the starting agent can identify whether it is at turn $t = 0$, where it almost always chooses to continue the game to send and receive more information via communication. Intervening on the dummy initialization message has a lower influence, resulting in lower $ME^{2 \rightarrow 1}$.

4.2 Conversation semantics

Having ascertained that our agents are conducting bidirectional conversations, we next try to decode what are the *contents* of such conversations. To do this, we train separate classifiers to predict, from the message exchanges in successful in-domain test game, what are Fruit, Tool 1, Tool 2 in the game.⁶ Consider for example a game in which fruit is *apple* and tools 1 and 2 *knife* and *spoon*, respectively. If the message-based classifiers are, say, able to successfully decode *apple* but not *knifelspoon*, this suggests that the messages are about the fruit but not the tools. For each prediction task, we train classifiers (i) on the whole conversation, i.e., both agents’ utterances (*Both*), and (ii) on either Player’s utterances: Fruit (*F*) or Tool only (*T*). For comparison, we also report accuracy of a baseline that makes guesses based on the train category distribution (*Stats*), which is stronger than chance. We report mean accuracy and SEM across successful training seeds. Supplementary Section 5 provides further details on classifier implementation and training.

The first row of Table 2 shows that the conversation as a whole carries information about any object. The second and third show that the agents are mostly conveying information about their respective objects (which is very reasonable), but also, to a lesser extent, but still well above baseline-level,

⁶We focus on the in-domain set as there are just 5 transfer fruit categories. We also tried predicting triples at once with a single classifier, that consistently reached above-baseline but very low accuracies.

about the other agent’s input. This latter observation is intriguing. Further work should ascertain if it is an artifact of fruit-tool correlations, or pointing in the direction of more interesting linguistic phenomena (e.g., asking “questions”). The asymmetry between Tool 1 and 2 would also deserve further study, but importantly the agents are clearly referring to both tools, showing they are not adopting entirely degenerate strategies.⁷

We tentatively conclude that the agents did develop the expected semantics, both being able to refer to all objects in the games. Did they however developed *shared* conventions to refer to them, as in human language? This would not be an unreasonable expectation, since the agents are symmetric and learn to play both roles and in both positions. Following up on the idea of “self-play” of Graesser et al. (2019), after a pair of agents A and B are trained, we replace at test time agent B’s embedders and modules with those in A, that is, we let one agent play with a copy of itself. If A and B are speaking the same language, this should not affect test performance. Instead, we find that with self-play average game performance drops down to 67% and 65% in in-domain and transfer test sets, respectively. This suggests that the agents developed their own idiolects. The fact that performance is still above chance could be due to the fact that the latter are at least partially exchangeable, or simply to the fact that agents can still do reasonably well by relying on knowledge of average tool usefulness (self-play performance is below that of the communication-less agents in Table 1). To decide between these interpretations, we trained the semantic classifier on conversations where A is the Fruit Player and B the Tool Player, testing on conversations about the same inputs, but where the roles are inverted. The performance drops down to the levels of the Stats baseline (Supplementary Table A5), supporting the conclusion that non-random performance is due to knowledge acquired by the agents independently of communication, and not partial similarity among their codes.

5 Related work

Games Among the long history of early works that model language evolution between agents (e.g.

⁷We experiment with single symbol messages (and multi-turn conversation) but using longer messages we could potentially witness interesting phenomena such as the emergence of compositionality. We leave this exploration for future work.

Steels, 2003; Brighton et al., 2003), Reitter and Lebiere (2011) simulate human language evolution with a Pictionary type task. Most recently, with the advent of neural network architectures, literature focuses on simple referential games with a sender sending a single message to a receiver, and reward depending directly on communication success (e.g., Lazaridou et al., 2017; Havrylov and Titov, 2017; Lazaridou et al., 2018). Evtimova et al. (2018) extend the referential game presenting the sender and receiver with referent views in different modalities, and allowing multiple message rounds. Still, reward is given directly for referential success, and the roles and turns of the agents are fixed. Das et al. (2017) generalize Lewis’ signaling game (Lewis, 1969) and propose a cooperative image guessing game between two agents, a question bot and an answer bot. They find that grounded language emerges without supervision. Cao et al. (2018) (expanding on Lewis et al., 2017) propose a setup where two agents see the same set of items, and each is provided with arbitrary, episode-specific utility functions for the object. The agents must converge in multi-turn conversation to a decision about how to split the items. The fundamental novelty of our game with respect to theirs is that our rewards depend on consistent, realistic commonsense knowledge that is stable across episodes (hammers are good to break hard-shell fruits, etc.). Mordatch and Abbeel (2018) (see also Lowe et al., 2017) study emergent communication among multiple (> 2) agents pursuing their respective goals in a maze. In their setup, fully symmetric agents are encouraged to use flexible, multi-turn communication as a problem-solving tool. However, the independent complexities of navigation make the environment somewhat cumbersome if the aim is to study emergent communication.

Communication analysis Relatively few papers have focused specifically on the analysis of the emergent communication protocol. Among the ones more closely related to our line of inquiry, Kottur et al. (2017) analyze a multi-turn signaling game. One important result is that, in their game, the agents only develop a sensible code if the sender is deprived of memory across turns. Evtimova et al. (2018) study the dynamics of agent confidence and informativeness as a conversation progresses. Cao et al. (2018) train probe classifiers to predict, from the messages, each agent util-

ity function and the decided split of items. Most directly related to our pragmatic analysis, Lowe et al. (2019), who focus on simple matrix communication games, introduce the notions of positive signaling (an agent sends messages that are related to its state) and positive listening (an agent’s behaviour is influenced by the message it receives). They show that positive signaling does not entail positive listening, and commonly used metrics might not necessarily detect the presence of one or the other. We build on their work, by focusing on the importance of mutual positive listening in communication (our “bilateral communication” measure). We further refine the causal approach to measuring influence they introduce. Jaques et al. (2018) also use the notion of causal influence, both directly as a term in the agent cost function, and to analyze their behaviour.

6 Discussion

We introduced a more challenging and arguably natural game to study emergent communication in deep network agents. Our experiments show that these agents do develop genuine communication even when (i) successful communication *per se* is not directly rewarded; (ii) the observable environment already contains stable, reliable information helping to solve the task (object affordances); and (iii) the agents are not artificially forced to rely on communication by erasing their memory. The linguistic exchanges of the agents are not only leading to significantly better task performance, but can be properly pragmatically characterized as dialogues, in the sense that the behaviour of each agent is affected by what the other agent says. Moreover, they use language, at least in part, to denote the objects in their environment, showing primitive hallmarks of a referential semantics.

We also find, however, that agent pairs trained together in fully symmetrical conditions develop their own idiolects, such that an agent won’t (fully) understand itself in self play. As convergence to a shared code is another basic property of human language, in future research we will explore ways to make it emerge. First, we note that Graesser et al. (2019), who study a simple signaling game, similarly conclude that training single pairs of agents does not lead to the emergence of a common language, which requires diffusion in larger communities. We intend to verify if a similar trend emerges if we extend our game to larger agent

groups. Conversely, equipping the agents with a feedback loop in which they also receive their own messages as input might encourage shared codes across speaker and listener roles.

In the current paper, we limited ourselves to one-symbol messages, facilitating analysis but greatly reducing the spectrum of potentially emergent linguistic phenomena to study. Another important direction for future work is thus to endow agents with the possibility of producing, at each turn, a sequence of symbols, and analyze how this affects conversation dynamics and the communication protocol. Finally, having shown that agents succeed in our setup, we intend to test them with larger, more challenging datasets, possibly involving more realistic perceptual input.

Acknowledgments

We thank Rahma Chaabouni, Evgeny Kharitonov, Emmanuel Dupoux, Maxime Oquab and Jean-Rémi King for their useful discussions and insights. We thank David Lopez-Paz and Christina Heinze-Deml for their feedback on the causal influence of communication. We also thank Francisco Massa for his help on setting up the experiments.

References

- Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis X. Charles, D. Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14:3207–3260.
- Diane Bouchacourt and Marco Baroni. 2018. How agents see things: On visual representations in an emergent language game. In *Proceedings of EMNLP*, pages 981–985, Brussels, Belgium.
- Henry Brighton, Simon Kirby, and Kenneth Smith. 2003. *Situated cognition and the role of multi-agent models in explaining language structure.*, volume 2636, pages 88–109. Springer-Verlag GmbH.
- Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z. Leibo, Karl Tuyls, and Stephen Clark. 2018. Emergent communication through negotiation. In *Proceedings of ICLR*.
- Edward Choi, Angeliki Lazaridou, and Nando de Freitas. 2018. Compositional obverter communication learning from raw visual input. In *Proceedings of ICLR Conference Track*, Vancouver, Canada. Published online: <https://openreview.net/group?id=ICLR.cc/2018/Conference>.
- Abhishek Das, Satwik Kottur, José M. F. Moura, Stefan Lee, and Dhruv Batra. 2017. Learning cooperative visual dialog agents with deep reinforcement learning. In *2017 IEEE International Conference on Computer Vision (ICCV)*.
- Katrina Evtimova, Andrew Drozdov, Douwe Kiela, and Kyunghyun Cho. 2018. Emergent communication in a multi-modal, multi-step referential game. In *Proceedings of ICLR Conference Track*, Vancouver, Canada. Published online: <https://openreview.net/group?id=ICLR.cc/2018/Conference>.
- Laura Graesser, Kyunghyun Cho, and Douwe Kiela. 2019. Emergent linguistic phenomena in multi-agent communication games. *CoRR*.
- Serhii Havrylov and Ivan Titov. 2017. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In *Proceedings of NIPS*, pages 2149–2159, Long Beach, CA, USA.
- 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.
- Emilio Jorge, Mikael Kågebäck, and Emil Gustavsson. 2016. Learning to play Guess Who? and inventing a grounded language as a consequence. In *Proceedings of the NIPS Deep Reinforcement Learning Workshop*, Barcelona, Spain. Published online: <https://sites.google.com/site/deeprlnips2016/>.
- Satwik Kottur, José Moura, Stefan Lee, and Dhruv Batra. 2017. Natural language does not emerge ‘naturally’ in multi-agent dialog. In *Proceedings of EMNLP*, pages 2962–2967, Copenhagen, Denmark.
- Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. 2018. Emergence of linguistic communication from referential games with symbolic and pixel input. In *Proceedings of ICLR Conference Track*, Vancouver, Canada. Published online: <https://openreview.net/group?id=ICLR.cc/2018/Conference>.
- Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017. Multi-agent cooperation and the emergence of (natural) language. In *Proceedings of ICLR Conference Track*, Toulon, France. Published online: <https://openreview.net/group?id=ICLR.cc/2017/conference>.
- Sang-Woo Lee, Yu-Jung Heo, and Byoung-Tak Zhang. 2017. Answerer in questioner’s mind for goal-oriented visual dialogue. <https://arxiv.org/abs/1802.03881>.
- David Lewis. 1969. *Convention*. Harvard University Press, Cambridge, MA.

- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? End-to-end learning of negotiation dialogues. In *Proceedings of EMNLP*, pages 2443–2453, Copenhagen, Denmark.
- Wentian Li. 2002. Zipf’s law everywhere. *Glottometrics*, 5:14–21.
- Ryan Lowe, Jakob Foerster, Y-Lan Boureau, Joelle Pineau, and Yann Dauphin. 2019. Measuring emergent communication is tricky. In *Proceedings of AAMAS*, Montreal, Canada. In press.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *NIPS*, pages 6382–6393.
- Ken McRae, George Cree, Mark Seidenberg, and Chris McNorgan. 2005. Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37(4):547–559.
- Igor Mordatch and Pieter Abbeel. 2018. Emergence of grounded compositional language in multi-agent populations. In *AAAI*, pages 1495–1502. AAAI Press.
- Judea Pearl, Madelyn Glymour, and Nicholas Jewell. 2016. *Causal Inference in Statistics: A Primer*. John Wiley & Sons.
- David Reitter and Christian Lebiere. 2011. How groups develop a specialized domain vocabulary: A cognitive multi-agent model. *Cognitive Systems Research*, 12:175–185.
- Carina Silberer, Vittorio Ferrari, and Mirella Lapata. 2013. Models of semantic representation with visual attributes. In *Proceedings of ACL*, pages 572–582, Sofia, Bulgaria.
- Luc Steels. 2003. Evolving grounded communication for robots. *Trends in cognitive sciences*, 7:308–312.
- Michael Tomasello. 2014. *A Natural History of Human Thinking*. Harvard University Press, Cambridge, MA.
- Ronald Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.