

---

# A perspective on multi-agent communication for information fusion

---

**Homagni Saha**

Department of Mechanical Engineering  
Iowa State University  
Ames, IA 50011  
hsaha@iastate.edu

**Vijay Venkataraman**

Honeywell Aerospace  
Plymouth, MN 55441  
Vijay.Venkataraman@honeywell.com

**Alberto Speranzon**

Honeywell Aerospace  
Plymouth, MN 55441  
Alberto.Speranzon@honeywell.com

**Soumik Sarkar**

Department of Mechanical Engineering  
Iowa State University  
Ames, IA 50011  
soumiks@iastate.edu

## Abstract

Collaborative decision making in multi-agent systems typically requires a predefined communication protocol among agents. Usually, agent-level observations are locally processed and information is exchanged using the predefined protocol, enabling the team to perform more efficiently than each agent operating in isolation. In this work, we consider the situation where agents, with complementary sensing modalities must co-operate to achieve a common goal/task by learning an efficient communication protocol. We frame the problem within an actor-critic scheme, where the agents learn optimal policies in a centralized fashion, while taking action in a distributed manner. We provide an interpretation of the emergent communication between the agents. We observe that the information exchanged is not just an encoding of the raw sensor data but is, rather, a specific set of directive actions that depend on the overall task. Simulation results demonstrate the interpretability of the learnt communication in a variety of tasks.

## 1 Introduction

In this paper, we analyze communication protocols learnt by a team of agents equipped with complementary sensor modalities and tasked with a common goal. We call this “task based multi-modal decision making”, wherein agents learn to map their sensor measurements and the information communicated by other agents, directly into actions based on the common goal. In this setting, each agent has access only to its own sensor data but needs to rely on the communication with other agents to obtain task relevant information from that agent’s sensor modality. We present a way to interpret the emergent communication by visualizing this mapping into the agent’s action space. We find the communication that is learnt, within a reinforcement learning paradigm, is not only emergent, but is task dependent and adaptive to the size of the communication channel.

In relation to existing literature on learning for multi-agent systems our work borrows from the general framework of “Markov-games”, proposed in Lowe et al. [2017], Mordatch and Abbeel [2018], Foerster et al. [2016] and references therein. In particular, we consider (Lowe et al. [2017]) for our learning problem. Related to the emergent communication aspect, central to this work, we consider ideas from the literature on “multi-agent referential games” (Golland et al. [2010], Andreas and Klein [2016], Evtimova et al. [2017], Lazaridou et al. [2018]), where a (sender) agent communicates highly structured information (images and text) to a (receiver) agent which has to interpret what the other agent saw. However, here we are interested in the evolution of communication for

unstructured data under joint interactions using an actor-critic algorithm. Emergent communication was also studied in (Kottur et al. [2017], Cao et al. [2018]), however, we focus our analysis of emergent communication on the action space. Specifically, we project the learnt communication on the action space and visually analyze the results. This enables us to more clearly interpret the learnt communication. It is shown that powerful joint representations of the world can be encoded through task dependent communication which is easy to interpret under complementary sensing modality constraints.

## 2 Environment and tasks

For our experimental study, we consider a two-dimensional world with two agents and  $L$  landmarks. Each landmark has a color  $\{red, green, blue\}$  and shape  $\{triangle, circle, square\}$  property. Our agents have complementary sensing modalities: one of the agents, denoted as color-agent, can only observe the color of the landmarks and the other, denoted as shape-agent, can only observe the shape of the landmarks. We assume that both agents can measure their (relative) distance from all landmarks but cannot measure their distance from each other. At every discrete time step the agents take both physical movement actions (a unit movement in one of the four directions or stand still) and communicative actions, namely, broadcast a  $k$ -bit message. The communication message sent by one agent is received by the other in the next time step.

Each agent’s observation is a vector  $\mathbf{o}_i = [x_{i1}, y_{i1}, \mathbf{m}_{i1}, \dots, x_{iL}, y_{iL}, \mathbf{m}_{iL} | c_{i1}, \dots, c_{ik} | g_{i1}, \dots, g_{iL}]$ , where  $x_{ij}$  and  $y_{ij}$  denote the horizontal and vertical distances, respectively, between the  $i^{th}$  agent and the  $j^{th}$  landmark;  $\mathbf{m}_{ij}$  denotes a one hot encoding of the  $j^{th}$  landmark’s property as sensed by the  $i^{th}$  agent. For example,  $\mathbf{m}_{ij} \in \{[1, 0, 0], [0, 1, 0], [0, 0, 1]\}$  denotes the encoding of the  $j^{th}$  landmark’s color for the color agent or shape for the shape agent. The vector  $[c_{i1}, \dots, c_{ik}]$ , denotes the  $k$  bit word received by the  $i^{th}$  agent. Finally,  $[g_{i1}, \dots, g_{iL}]$  denotes one hot encoding of the target landmark properties provided to agent  $i$ . We base our study on the three collaborative tasks, described below.

**Task 1: Cross modal information exchange:** In this task, the map contains three landmarks. No two landmarks have the same shape or color. During each episode, agents and landmarks are placed randomly in the map. One of the landmarks is designated as “target” and the goal, for both agents, is to reach the designated target landmark. This target landmark’s property is indicated to the agents using  $[g_{i1}, \dots, g_{iL}]$  as described previously. Consider as example where the target landmark is a blue circle, if we were to provide to the color agent the color properties, it can trivially navigate to the target given that it has full knowledge of where the different colored landmarks are with respect to itself. To avoid this and to encourage communication, we pass the encoding corresponding to the shape of the target landmark to the color-agent (circle in this example) and vice versa for the shape agent. This creates a situation where the agents need to exchange information in order to successfully navigate to the right landmark.

**Task 2: Multi target consensus:** For this task, the map contains six landmarks each with shape and color properties and no two landmarks have the same set of properties. However, there can be two landmarks with the same color or shape and the target landmark is unique when both properties are considered. The goal is for both agents to move to the target landmark. The agent observation and action spaces are similar to the previous scenario, but here the property of the target landmark is specified in the agent’s own modality. Specifying the encoding for circle, as the target, to the shape agent does not trivially solve the problem as there can be two circles and co-ordination with the color agent is necessary to figure out which circle is blue and then move towards it. In this task, the agents need to reach a consensus on which is the target landmark by learning to reason over their observation spaces.

**Task 3: Collaborative localization:** Here the setup is similar to the information exchange task. However, no target landmark is specified as the goal is for the agents to meet with each other in the shortest possible time. Here a constant negative reward  $R_t$  supplied to the agents at each time-step to encourage meeting up fast. Here the agents must learn to estimate their relative position with respect to each other and then take actions to move closer.

Summarizing, all the above tasks share the following key challenging characteristics. Agents have (i) different sensing modalities; (ii) No knowledge about other agent’s sensor or position; (iii) No common world coordinate frame in their state space; (iv) a finite communication bandwidth.

**Reward structure and learning framework:** We primarily used three types of rewards:  $R_d = \sum_{i=1}^{i=n} \sqrt{x_{iT}^2 + y_{iT}^2}$ , where  $i$  is the agent number,  $x_{iT}$  and  $y_{iT}$  are the horizontal and vertical distances of  $i^{th}$  agent from the target landmark, and  $n$  is the number of agents. We define an instantana-

Task	k	Reward	Metric (%)	
			M <sub>1</sub>	M <sub>2</sub>
Information exchange	2	$R_d$	99	80
	3	$R_d$	100	89.5
	4	$R_d$	100	98.9
Multi target consensus	4	$R_d$	15.1	0.5
	4	$R_d + R_i$	91.8	61.6
Collaborative localization	3	$R_d + R_t$	80.5	100

Table 1: Performance metrics for the three test tasks with variations in reward structure and communication channels.

neous reward  $R_i = H$ , where  $H$  is a large number if at least one agent is touching the target at the current time step and 0 otherwise. For collaborative localization task,  $R_t$  is a constant penalty per time step and  $R_d$  is the inter agent distance. For all tasks, our reinforcement learning (RL) framework is based on the MADDPG algorithm (Lowe et al. [2017]), which relies on centralized training and decentralized execution, making it suitable for multi-agent problems. The core of MADDPG is an actor-critic scheme (Grondman et al. [2012]) that maintains a critic for each agent and the critics have access to actions (movement and communications) and rewards of all the agents. This helps with the problem of non-stationarity in multi-agent environments. In all experiments, we parameterize the output of both actors and the critics with a three layered fully connected network with ReLU activations. It must be noted that, although the agent state space allows for real numbers in the communication stream, the use of Gumbel-Softmax estimator (Jang et al. [2016]) transforms these into discrete valued messages. While  $2^k$  word variations are possible, we observe that the agents limit their vocabulary use to  $k + 1$  words. For 3 channels of communication, the word vocabulary was limited to  $w_0 = [0, 0, 0]$ ,  $w_1 = [1, 0, 0]$ ,  $w_2 = [0, 1, 0]$ ,  $w_3 = [0, 0, 1]$ . Details of the learning framework, training hyperparameters and reward curves are provided in Appendix A.

### 3 Results

In the following we evaluate agent performance using simple metrics and then provide interpretations of the emerged communication between the agents.

**Performance metrics:** We use two metrics to evaluate the performance of the agents in achieving the common goal. Let  $m_1$  denote the number of episodes in which at least one of the agent reaches the target landmark,  $m_2$  denote the number of episodes where both agents reach the target landmark, and  $N$  denote the total number of test episodes. We use  $N = 1000$ , let  $M_1 = m_1/N$  and  $M_2 = m_2/N$ . For the collaborative localization task  $m_2$  denotes the number of times the agents meet with each other, and  $m_1$  is the starting distance between agents divided by total distance travelled by both agents in the episode. As we mentioned,  $k$  is the number of communication bits available to the agent. Table 3 shows these metrics for the three different tasks.

**Emerg ed communication:** Table 3 shows that the agents are able to successfully complete the information exchange and collaborative localization tasks almost every time. In the more complex multi target consensus task the agents achieve a 61.6% success rate which is better than random chance of 25%. In order to better understand how the communication aid the agents, we devise a way to visualize this relation as follows. At every time step the  $i^{th}$  agent decides its actions based on its observations  $\mathbf{o}_i$  comprising of relative position to the landmarks, the word received from the other agent and target landmark (if applicable). For a given test case the target landmark is fixed. Then for every possible word that can be received, we can place the agent in a fixed position of the environment and query the learnt policy to find in which direction the agent would move. We can then repeat this for all possible agent positions and color code its preferential direction of motion at each location, obtaining a picture as shown in figure 1. As expected, we observe random motion in the beginning of the training. Over time the color agent learns to move to the blue triangle if the word uttered by the shape agent is  $[1, 0, 0]$ . Similarly the shape agent learns to move towards the blue circle for the same utterance by the color agent. Visualization for other word utterances are given in the Appendix, see figure 6. Note that both agents are either focusing on a blue or circular object for all word utterances as the given target in this example is a blue circle. It is remarkable that the agents are able to solve a complex map alignment and reasoning problem by directing each others actions through communication. Examples of the final learnt policy for the *information exchange* and *collaborative localization* tasks are shown in figure 2 left and right respectively. In the information exchange task, we observe that each unique word uttered by the color agent causes the shape-agent to move close to a specific shape of landmark irrespective of its current position in the map (e.g  $[1, 0, 0, 0]$  causes movement towards the circle). In the *collaborative localization* task each unique word uttered by the color agent causes the shape agent to move towards a focus point in the

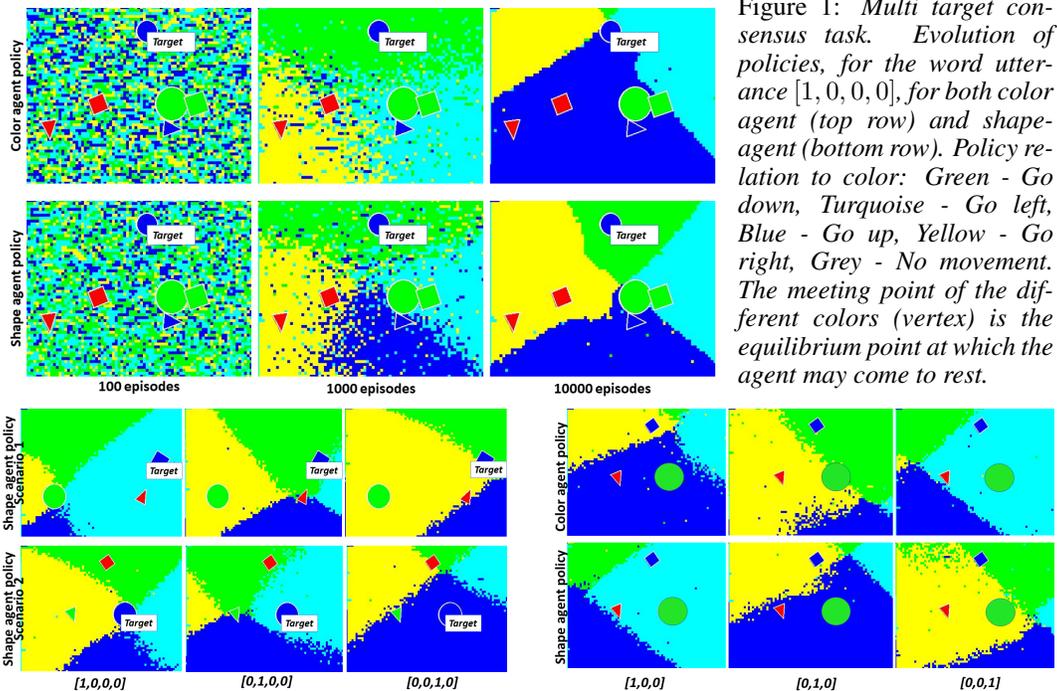


Figure 1: *Multi target consensus task. Evolution of policies, for the word utterance  $[1, 0, 0, 0]$ , for both color agent (top row) and shape agent (bottom row). Policy relation to color: Green - Go down, Turquoise - Go left, Blue - Go up, Yellow - Go right, Grey - No movement. The meeting point of the different colors (vertex) is the equilibrium point at which the agent may come to rest.*

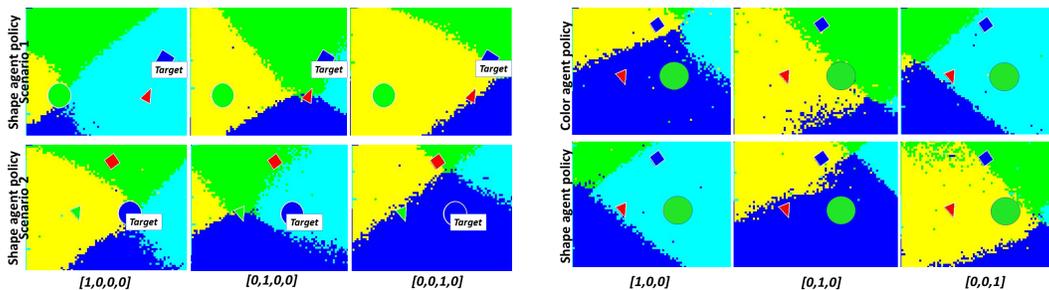


Figure 2: **Left**-Information exchange task. From left to right, the color-agent utters the words  $[1, 0, 0, 0]$ ,  $[0, 1, 0, 0]$ ,  $[0, 0, 1, 0]$  to the shape agent. Top and bottom rows represent different landmark configurations. **Right**-Collaborative localization task. From left to right, policy for the words  $[1, 0, 0, 0]$ ,  $[0, 1, 0, 0]$ ,  $[0, 0, 1, 0]$  are visualized for both color (top row) and shape agents (bottom row).

map and vice-versa. During the episode, both agents continuously change their utterance in order to force their partner to travel towards each other and meet in the shortest time possible.

**Effect of communication channels:** Changing the number of channels affects the learning capacity. In the information exchange task when 3 channels are provided agents associate their target landmarks as “topmost”, “leftmost”, or “bottom most” landmark in the map, depending on received communication. While this is a clever reference, as the “topmost” etc. landmark can easily be the same for the two agents as their reference axes are only translated from each other. However, this leads to a failure when the target landmark is located in the middle or is the “rightmost”. When using 4 channels agents can directly associate their targets to the property communicated to them by the other agent. In figure 2 (left), shape agent interprets  $[1, 0, 0, 0]$  from the color agent as a signal to go to a circle,  $[0, 1, 0, 0]$  to a triangle and  $[0, 0, 1, 0]$  to a square. This improves performance greatly.

## 4 Conclusion

We studied the application of multi agent reinforcement learning for task driven multimodal decision making. We analyzed the emergence of interpretable communication between agents and found that adaptive and non trivial communication protocols can be learned based on number of available communication channels and imposed reward structures. The size of the communication channel can be crucial in deciding the amount of information that is required to reconcile various modalities with each other and reward structures affect the nature of the learned communication. Visualizing emergent policies in the agent’s action spaces confirms that powerful joint representation of the world can be encoded through communication.

## 5 Acknowledgement

We are thankful to Shashank Shivkumar at Honeywell Aerospace for discussions related to this paper ranging from initial idea for the research through algorithm implementation and testing.

## References

- Jacob Andreas and Dan Klein. Reasoning about pragmatics with neural listeners and speakers. *arXiv preprint arXiv:1604.00562*, 2016.
- Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z Leibo, Karl Tuyls, and Stephen Clark. Emergent communication through negotiation. *arXiv preprint arXiv:1804.03980*, 2018.
- Katrina Evtimova, Andrew Drozdov, Douwe Kiela, and Kyunghyun Cho. Emergent communication in a multi-modal, multi-step referential game. *arXiv preprint arXiv:1705.10369*, 2017.
- Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2137–2145, 2016.
- Dave Golland, Percy Liang, and Dan Klein. A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 conference on empirical methods in natural language processing*, pages 410–419. Association for Computational Linguistics, 2010.
- Ivo Grondman, Lucian Busoniu, Gabriel AD Lopes, and Robert Babuska. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):1291–1307, 2012.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. Natural language does not emerge naturally in multi-agent dialog. *arXiv preprint arXiv:1706.08502*, 2017.
- Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. Emergence of linguistic communication from referential games with symbolic and pixel input. *arXiv preprint arXiv:1804.03984*, 2018.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pages 6379–6390, 2017.
- Igor Mordatch and Pieter Abbeel. Emergence of grounded compositional language in multi-agent populations. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.