

# EMERGENT COMMUNICATION IN HUMAN-MACHINE GAMES

**Nicolo’ Brandizzi, Luca Iocchi**

Department of Computer, Automation and Management Engineering, Sapienza University of Rome,  
 via Ariosto 25 Roma 00185, Italy, IT  
 {brandizzi,iocchi}@diag.uniroma1.it

## ABSTRACT

In this paper, we present how Emergent Communication (EmeCom) literature can be used to characterize a subset of human communication ability using recurrent features. EmeCom does not directly require human involvement; however, we advocate that the entire field spawned from a desire to achieve human-machine interaction and that explicitly considering human-AI interaction is an important research direction. Our goal is to bring emergent communication researchers together by demonstrating that different approaches can be regarded as subsets of the same problem. Also, we recognize the importance of incorporating other aspects of human interaction into this field.

## 1 INTRODUCTION

Communication between humans and machines is regarded as a prerequisite to achieving artificial general intelligence (AGI); thus, researchers have been interested in developing AIs that can communicate with humans using natural language.

Recent technological advances and attention-based architectures Vaswani et al. (2017) made transformers the preferred solution for human-machine conversation, GPT-3 being the most prominent example Brown et al. (2020). Although transformers can achieve high performances in human communication systems, they learn through statistical inferences on pre-existing languages, which does not align with natural language emergence in humans. This training pipeline fails to capture the essence of communication which, in humans, emerged as a means to achieve a common goal rather than being the objective itself.

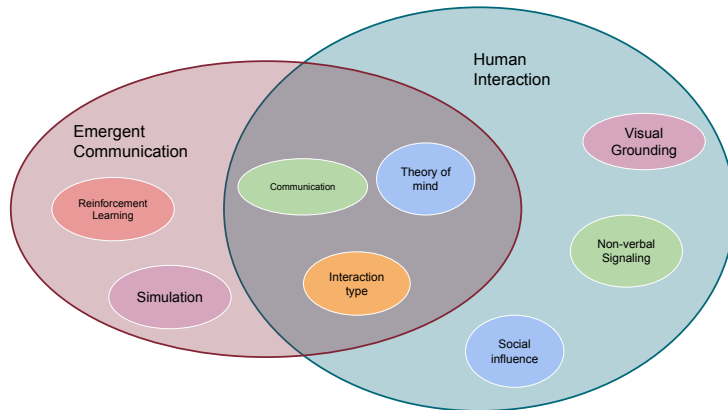


Figure 1: Framework division showing how different topics in EmeCom draw inspiration from a more general human interaction set.

This learning misalignment gave rise to the field of *Emergent Communication* (EmeCom) Wagner et al. (2003) to narrow the gap between machine and human language emergence. EmeCom aims to develop language through interaction between artificial agents, which can perceive the environment through observation and change its state by performing specific actions. Two fields heavily inspire this pipeline in artificial intelligence, Multi-Agent System (MAS) and Reinforcement Learning (RL) which we consider two fundamental pillars of human society and mind, respectively.

Games hold a meaningful impact on the learning process. Indeed, it has been shown how playing is a vital part of learning Spinka et al. (2001) that allows animals to develop the physical and psychological skills to handle unexpected events in which they experience a loss of control. Works such as Allee et al. (1949) emphasize how cheetahs cubs learn cooperation and social skills during playing, which will later be employed for hunts. Moreover, Tahmores (2011) identifies playing as fundamental for the proper development of social skills in children, and Kirriemuir & McFarlane (2004) investigates the literature on educational video games. Furthermore, RL finds its main application in games; thus, EmeCom is heavily based on game frameworks.

This paper aims to define common characteristics of EmeCom and highlight how each of them models a different aspect of human-human interactions. By specifying common characteristics, we can differentiate between several architectures and connect the literature based on our assumption.

**Paper Structure** We focus on two main categories that we believe hold fundamental similarities with human interactions: *theory of mind*, Sec. 3, where agents are aware of other intelligent entities in the environment and actively try to model their behavior, and *interaction types*, Sec. 2, defining the possible configuration between agents in a shared environment. Due to space limitation, we omitted another fundamental aspect of EmeCom: the *game environment* itself, which can be divided in *communication-goal*<sup>1</sup> and *communication-aid*<sup>2</sup>. Finally, we provide a complete table of referenced papers, Table 1 with a correlated plot of papers published by year Fig. 2.

## 2 INTERACTION TYPE

In this section, we focus on the types of interaction among players. Although real-world interactions are complex, our goal is to split them into fixed categories and study the contribution of each one.

We must first introduce the difference between internal and external interaction types. Given a set  $A_N$  of  $N$  agents, we define a team  $T_m \subseteq A_N$  as a subset of agents and we denote with  $\{T_1, \dots, T_k\}$  a partition of all the agents in  $k$  disjoint teams with  $\sum_{i=1, \dots, k} |T_i| = N$ .

Given a partition of agents (i.e., a set of teams), we define a game to be cooperative when any negotiation between teams lead to a sub-optimal outcome  $S_{so}$  for all the party involved. While each team actively tries to reach an optimal state  $S_o$ , trough negotiation it can avoid the worst outcome  $S_p$  where  $S_o < S_{so} < S_p$ . On the other hand, a zero-sum game competitive setting always implies a worst outcome for all teams but one. In the latter, negotiation becomes ineffective, and each team must prevail over the others.

Now we can introduce the concept of *internal cooperation*, where all the agents of one team are cooperative with one another. This is a necessary condition for the emergence of communication to arise. Indeed, to the best of our knowledge, there are no works in EmeCom tackling internal competitiveness. On the other hand, *external cooperation* is defined as two or more disjoint teams of agents sharing a common goal, thus being incentivized to cooperate. Differently, *external competitiveness* arises in zero-sum games, where one team’s victory implies a pessimal outcome for all other teams. In these cases, negotiation is not useful.

**Internal Cooperation** As discussed, cooperativeness is a fundamental aspect for communication to emerge Smith (2010); Nowak & Krakauer (1999). As shown in Table 1, most of the treated works deal with internal cooperation settings. Most notably, Cao et al. (2018) studies how pro-social agents favor cheap talks and achieve better results than selfish ones. On the same line, Graesser

<sup>1</sup>Comprehending those games where communication is the objective itself, such as Referential games (Ref) and Task and Talk (TnT).

<sup>2</sup>Where communication assists the agents in achieving a goal

et al. (2019) suggests how intricate language evolution can emerge from simple social interactions between agents.

**External Cooperation** External cooperation occurs when collaboration is extended between teams of agents and can be split into two types. The first one is the standard instance for external cooperation. Given a set of disjoint teams, the procedure is to initialize all the agents at the same time and let them interact with the environment, and each other Evtimova et al. (2018); Lowe et al. (2017). The peculiarity of this method is how every agent is initialized simultaneously and has no prior experience with the environment. Differently, when teams are pre-trained with one another before being introduced to other teams, this is referred to as population learning. This setting is heavily influenced by the study of human nature, specifically by language development between different cultures Briscoe (2002). In fact, the number of agents can influence the performance directly. It has been shown that a population of agents generalizes better than a pair. Moreover, the interaction between populations of agents can emerge a language that is easier to teach and understand Li & Bowling (2019); Lowe et al. (2019); Fitzgerald (2019).

**External Competitiveness** On the other hand, Liang et al. (2020) show how competitiveness can develop better communication protocols and improve general performance. In particular, they speculate how competition among a set of agents leads to the development of a communication protocol that prioritizes compositionality, performance, and convergence. It is worth noting that Liang et al. work with cooperative groups of competitive agents, thus making this setting more akin to human experience. The first three categories deal with the spatial component of human connection. The number of agents/teams in the environment can be utilized to characterize the differences between them. On the other hand, humans live in a temporal dimension in which information and experiences are passed down from one generation to the next. This temporal aspect drives an extra interaction type known as iterative learning, which is opposed to previously addressed ones.

**Iterative Learning** The concept of iterative learning is orthogonal to the population one. While this category does not strictly refer to an interaction type, we believe it captures a fundamental aspect of human interactions; thus, we decided to include it in this section. Indeed iterative learning is a cultural transmission where a population of agents passes their knowledge to a new one, and the loop repeats indefinitely. In the field of linguistics, iterative learning has been directly compared to a bottleneck in the learning system, which allows for generalization Nowak & Krakauer (1999). In Kirby et al. (2014) the authors summarize the approaches taken to study the emergence of natural language with an iterative point of view. They report various laboratory experiments in which iterative learning caused the shift towards languages that are consistent with prior biases. The same line of research was adopted by Scott-Phillips & Kirby (2010), where the author proposes a homonymy<sup>3</sup> filter to increase the structure and compositionality of the language through various generations.

On the same line, Ren et al. (2020) theorizes a correlation between compositionality, generalization, and topological similarity. For this reason, the authors build a training pipeline where artificial agents are trained on data generated by past versions of themselves and show how this advantage the emergence of language with high topological similarity.

These, and more works Dagan et al. (2021); Cogswell et al. (2019); Lu et al. (2020), show how language constructs passed down from generation to generation of learning agents lead to easier to learn languages. On the other hand, Guo et al. (2019) reports how iterative learning leads to the emergence of a compositional language; however, the results are highly dependent on input representation.

### 3 THEORY OF MIND

A fundamental propriety arising in the EmeCom field is directly derived from human behavior and can be seen as an attempt to model other agents' beliefs. Indeed, as humans, we tend to build a belief model of how other humans may react to certain stimuli and update it with each new observation Gopnik & Wellman (1992); Premack & Woodruff (1978). This behavior was first theorized in

<sup>3</sup>Relation between words that have the same form but different meanings (e.g., a pen: a small enclosure for animals; a female swan); common in natural languages.

Premack & Woodruff (1978) and named Theory of Mind (ToM), after the humans’ ability to represent the mental states of others. More recently, Rabinowitz et al. (2018) applied ToM to let artificial agents build a model of other agents’ observation and behavior alone.

The following paragraphs identify two main approaches to tackle the modeling problem: (i) *Agent’s modeling*, where the artificial agents are aware of others in the game and actively model their behavior to some extent; (ii) *Influencing others*, a direct extension of agent’s modeling, where agents also manipulate other agents’ behavior based on their objective.

**Agents’ Modeling.** Modeling other agents is a well-established line of work in the MARL community. Although modeling other humans as separated entities is natural for us, machines need specific formulations to approximate the same behavior.

When considering the original ToM, these formulations can leverage the similarities between agents’ belief systems. In this case, a communication message can be generated by maximizing an agent’s own understanding. This procedure takes the name of *obverter technique* and has shown effective for the emergence of compositional languages Choi et al. (2018); Bogin et al. (2018).

On the other hand, mental models can also be based on other agents’ actions and perceptions without assuming similar belief systems. For example, Raileanu et al. (2018) augment the agents’ policy with a prediction of other agents’ behavior and show that the agents can learn better policies using their estimate of the other players’ goals in both cooperative and competitive settings. However, this work considers environments where communication is not present.

Lowe et al. (2017) describe agents adapting to each other when trained in conjunction; this finding led to study agents who can reason about other agents and adjust the communication protocol accordingly Andreas & Klein (2016). On the same line, Rodriguez et al. (2019) describe agents modeling the conceptual understanding of other agents by switching partners with different proprieties; for example, a color-blinded agent will not benefit from a color-oriented description of an image. Moreover, Grover et al. (2018) show how to split the representation learning into two parts: a *generative* embedding simulates an agent’s policy, while a *discriminative* one distinguishes one agent from the other. The combination of these two representations yields the best results in two physics simulation environments, especially when the learning takes place in an offline manner<sup>4</sup>.

The above mentioned works show how agent modeling allows communication to quickly adapt and specialize to the task at hand. However, such specialization leads to languages that are difficult to interpret by humans.

In a multi-agent environment, whom to communicate with can be as important as the content of the message itself. For instance, Das et al. (2019) consider a MAS setting and build an architecture, TarMac<sup>5</sup>, based on soft-attention to let the agents choose their target. More precisely, each message is a continuous value split into two components: a value  $v$  which encodes the information that an agent wants to send, and a signature  $k$  multiplied with a query vector  $q$ . The dot product is normalized and used as a similarity measure between the sender signature and the receiver query. Intuitively, attention weights are high when both sender and receiver predict similar signature and query vectors. Following Lowe et al. (2017), the authors apply centralized critic training by temporal difference and decentralized execution with RMSProp. They report improved performances for their method, specifically how multiple rounds of communication lead to significantly higher performance than simply increasing message size. This finding opens a new horizon on the possibility to increase performances in a MAS since it does not require signal sharing for and to each agent in the system. Future research could also investigate the behavior of TarMac when agent communication is associated with a cost.

**Influencing Others.** The next natural step in this direction is given in Foerster et al. (2018) where the authors propose the Learning with Opponent Learning Awareness (LOLA) framework, which attempts to model the opponent’s policy and actively tries to influence it using policy gradient optimization. This then evolves into higher-order LOLA, where agents are aware of opponents trying to

<sup>4</sup>They still report an increase winning chance with online representation learning compared to no representation at all.

<sup>5</sup>Targeted multi-agent communication

influence them, leading to third-order derivatives, which are computationally expensive. This same approach can be applied in mixed human-robot settings, such as in Xie et al. (2020), where the author build an RL environment where the agent employs an encoder-decoder architecture to model the action of a human being. This model is then used to approximate the human policy to maximize the total discounted reward. While most of the treated literature concerns MAS with artificial agent only, just a handful of papers consider mixed human-robot teams Jaques et al. (2019); Finn et al. (2017); Hawkins et al. (2020); Xie et al. (2020). As a matter of fact, we believe that this line of work should be investigated more deeply in the field and all the above mentioned paper could be extended in the same way.

## 4 CONCLUSION

In this paper, we carried out an analysis of the current emergent communication literature. We defined two main proprieties, namely **theory of mind**, Sec 3, and **interaction type**, Sec 2, and summarized the literature tracing parallelism with human real-world interactions. *Theory of mind* takes direct inspiration from the human ability to model other intelligent entities as individuals separated from the environment. A natural extension is *influencing other* agents based on the previous modeling.

On the other hand, we introduce the concept of *internal and external interactions*, which happen between agents belonging to the same and different teams, respectively. This allows us to define a new kind of *external competitiveness* setting, seen as an environment where only one team can win while the others experience the worst outcome. Finally, we offer a complete list of the treated works, Table 1, where we correlate papers with the relative category they fall in, and a plot showing the publication year in Fig 2.

Our work can assist interested researchers in approaching this new and exciting field and connect experts in the area under a common objective. In fact, we promote the idea that EmeCom has spawned by the necessity for human-robot collaboration; thus, the literature can be seen as tackling diverse aspects of human-to-human interaction. Under this assumption, we aim to connect apparently diverse works under a unified line of research.

**Future Work** Most of the analyzed literature refers to multi-agent systems where the entities are artificial only; thus, the characteristics of human communication are approximated in a close setting. However, having a human-in-the-loop approach and, more explicitly, considering mixed human-robot teams is a promising research line. Finally, we would like to bring to attention other aspects of human interaction not typically discussed in EmeCom literature, such as visual grounding Antol et al. (2015) and nonverbal communication Breazeal et al. (2005).

## REFERENCES

- Warder Clyde Allee, Orlando Park, Alfred Edwards Emerson, Thomas Park, Karl Patterson Schmidt, et al. Principles of animal ecology. 1949.
- Jacob Andreas and Dan Klein. Reasoning about pragmatics with neural listeners and speakers. In Jian Su, Xavier Carreras, and Kevin Duh (eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pp. 1173–1182. The Association for Computational Linguistics, 2016. doi: 10.18653/v1/d16-1125. URL <https://doi.org/10.18653/v1/d16-1125>.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- Yoram Bachrach, Richard Everett, Edward Hughes, Angeliki Lazaridou, Joel Z Leibo, Marc Lanctot, Michael Johanson, Wojciech M Czarnecki, and Thore Graepel. Negotiating team formation using deep reinforcement learning. *Artificial Intelligence*, 288:103356, 2020.
- Ben Bogin, Mor Geva, and Jonathan Berant. Emergence of communication in an interactive world with consistent speakers. *CoRR*, abs/1809.00549, 2018. URL <http://arxiv.org/abs/1809.00549>.
- Nicolo’ Brandizzi, Davide Grossi, and Luca Iocchi. Rlupus: Cooperation through emergent communication in the werewolf social deduction game. *CoRR*, abs/2106.05018, 2021. URL <https://arxiv.org/abs/2106.05018>.
- Cynthia Breazeal, Cory D Kidd, Andrea Lockerd Thomaz, Guy Hoffman, and Matt Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *2005 IEEE/RSJ international conference on intelligent robots and systems*, pp. 708–713. IEEE, 2005.
- Ted Briscoe. *Linguistic evolution through language acquisition*. Cambridge University Press, 2002.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z. Leibo, Karl Tuyls, and Stephen Clark. Emergent communication through negotiation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=Hk6WhagRW>.
- Edward Choi, Angeliki Lazaridou, and Nando de Freitas. Compositional obverter communication learning from raw visual input. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=rknt2Be0->.
- Michael Cogswell, Jiasen Lu, Stefan Lee, Devi Parikh, and Dhruv Batra. Emergence of compositional language with deep generational transmission. *CoRR*, abs/1904.09067, 2019. URL <http://arxiv.org/abs/1904.09067>.
- Gautier Dagan, Dieuwke Hupkes, and Elia Bruni. Co-evolution of language and agents in referential games. In Paola Merlo, Jörg Tiedemann, and Reut Tsarfay (eds.), *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pp. 2993–3004. Association

- for Computational Linguistics, 2021. doi: 10.18653/v1/2021.eacl-main.260. URL <https://doi.org/10.18653/v1/2021.eacl-main.260>.
- Abhishek Das, Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE international conference on computer vision*, pp. 2951–2960, 2017.
- Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. Tarmac: Targeted multi-agent communication. In *International Conference on Machine Learning*, pp. 1538–1546. PMLR, 2019.
- Katrina Evtimova, Andrew Drozdov, Douwe Kiela, and Kyunghyun Cho. Emergent communication in a multi-modal, multi-step referential game. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=rJGZq6g0->.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *CoRR*, abs/1703.03400, 2017. URL <http://arxiv.org/abs/1703.03400>.
- Nicole Fitzgerald. To populate is to regulate. *CoRR*, abs/1911.04362, 2019. URL <http://arxiv.org/abs/1911.04362>.
- Jakob N. Foerster, Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with opponent-learning awareness. In Elisabeth André, Sven Koenig, Mehdi Dastani, and Gita Sukthankar (eds.), *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, pp. 122–130. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2018. URL <http://dl.acm.org/citation.cfm?id=3237408>.
- Alison Gopnik and Henry M Wellman. Why the child’s theory of mind really is a theory. 1992.
- Laura Graesser, Kyunghyun Cho, and Douwe Kiela. Emergent linguistic phenomena in multi-agent communication games. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pp. 3698–3708. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1384. URL <https://doi.org/10.18653/v1/D19-1384>.
- Aditya Grover, Maruan Al-Shedivat, Jayesh K. Gupta, Yuri Burda, and Harrison Edwards. Learning policy representations in multiagent systems. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1797–1806. PMLR, 2018. URL <http://proceedings.mlr.press/v80/grover18a.html>.
- Shangmin Guo, Yi Ren, Serhii Havrylov, Stella Frank, Ivan Titov, and Kenny Smith. The emergence of compositional languages for numeric concepts through iterated learning in neural agents. *CoRR*, abs/1910.05291, 2019. URL <http://arxiv.org/abs/1910.05291>.
- Robert X. D. Hawkins, Minae Kwon, Dorsa Sadigh, and Noah D. Goodman. Continual adaptation for efficient machine communication. In Raquel Fernández and Tal Linzen (eds.), *Proceedings of the 24th Conference on Computational Natural Language Learning, CoNLL 2020, Online, November 19-20, 2020*, pp. 408–419. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.conll-1.33. URL <https://doi.org/10.18653/v1/2020.conll-1.33>.
- Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Çağlar Gülçehre, Pedro A. Ortega, DJ Strouse, Joel Z. Leibo, and Nando de Freitas. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning*

- Research*, pp. 3040–3049. PMLR, 2019. URL <http://proceedings.mlr.press/v97/jaques19a.html>.
- Simon Kirby, Tom Griffiths, and Kenny Smith. Iterated learning and the evolution of language. *Current opinion in neurobiology*, 28:108–114, 2014.
- John Kirriemuir and Angela McFarlane. Literature review in games and learning. 2004.
- Satwik Kottur, José M. F. Moura, Stefan Lee, and Dhruv Batra. Natural language does not emerge ‘naturally’ in multi-agent dialog. In Martha Palmer, Rebecca Hwa, and Sebastian Riedel (eds.), *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pp. 2962–2967. Association for Computational Linguistics, 2017. doi: 10.18653/v1/d17-1321. URL <https://doi.org/10.18653/v1/d17-1321>.
- Fushan Li and Michael Bowling. Ease-of-teaching and language structure from emergent communication. In *Advances in Neural Information Processing Systems*, pp. 15851–15861, 2019.
- Paul Pu Liang, Jeffrey Chen, Ruslan Salakhutdinov, Louis-Philippe Morency, and Satwik Kottur. On emergent communication in competitive multi-agent teams. In Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith (eds.), *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS ’20, Auckland, New Zealand, May 9-13, 2020*, pp. 735–743. International Foundation for Autonomous Agents and Multiagent Systems, 2020. URL <https://dl.acm.org/doi/abs/10.5555/3398761.3398849>.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6379–6390, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/68a9750337a418a86fe06c1991a1d64c-Abstract.html>.
- Ryan Lowe, Abhinav Gupta, Jakob Foerster, Douwe Kiela, and Joelle Pineau. Learning to learn to communicate, 2019.
- Yuchen Lu, Soumye Singhal, Florian Strub, Olivier Pietquin, and Aaron C. Courville. Countering language drift with seeded iterated learning. *CoRR*, abs/2003.12694, 2020. URL <https://arxiv.org/abs/2003.12694>.
- Igor Mordatch and Pieter Abbeel. Emergence of grounded compositional language in multi-agent populations. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Martin A Nowak and David C Krakauer. The evolution of language. *Proceedings of the National Academy of Sciences*, 96(14):8028–8033, 1999.
- David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4):515–526, 1978.
- Neil C. Rabinowitz, Frank Perbet, H. Francis Song, Chiyuan Zhang, S. M. Ali Eslami, and Matthew Botvinick. Machine theory of mind. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4215–4224. PMLR, 2018. URL <http://proceedings.mlr.press/v80/rabinowitz18a.html>.
- Roberta Raileanu, Emily Denton, Arthur Szlam, and Rob Fergus. Modeling others using oneself in multi-agent reinforcement learning. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4254–4263. PMLR, 2018. URL <http://proceedings.mlr.press/v80/raileanu18a.html>.



- Yi Ren, Shangmin Guo, Matthieu Labeau, Shay B. Cohen, and Simon Kirby. Compositional languages emerge in a neural iterated learning model. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=HkePNpVKPB>.
- Rodolfo Corona Rodriguez, Stephan Alaniz, and Zeynep Akata. Modeling conceptual understanding in image reference games. In *Advances in Neural Information Processing Systems*, pp. 13155–13165, 2019.
- Thomas C Scott-Phillips and Simon Kirby. Language evolution in the laboratory. *Trends in cognitive sciences*, 14(9):411–417, 2010.
- Eric Alden Smith. Communication and collective action: language and the evolution of human cooperation. *Evolution and human behavior*, 31(4):231–245, 2010.
- Marek Spinka, Ruth C Newberry, and Marc Bekoff. Mammalian play: training for the unexpected. *The Quarterly review of biology*, 76(2):141–168, 2001.
- Aghajani Hashtchin Tahmores. Role of play in social skills and intelligence of children. *Procedia-Social and Behavioral Sciences*, 30:2272–2279, 2011.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Kyle Wagner, James A. Reggia, Juan Uriagereka, and Gerald Wilkinson. Progress in the simulation of emergent communication and language. *Adapt. Behav.*, 11(1):37–69, 2003. doi: 10.1177/10597123030111003. URL <https://doi.org/10.1177/10597123030111003>.
- Annie Xie, Dylan P. Losey, Ryan Tolsma, Chelsea Finn, and Dorsa Sadigh. Learning latent representations to influence multi-agent interaction. In Jens Kober, Fabio Ramos, and Claire J. Tomlin (eds.), *4th Conference on Robot Learning, CoRL 2020, 16-18 November 2020, Virtual Event / Cambridge, MA, USA*, volume 155 of *Proceedings of Machine Learning Research*, pp. 575–588. PMLR, 2020. URL <https://proceedings.mlr.press/v155/xie21a.html>.

Work	Theory of mind		Interaction Type				Game Environment		
	Modeling	Influencing	Int. Coop	Ext. Coop	Ext. Comp	Iter.	Com-goal	Com-aid	No-com
Jaques et al. (2019)	x	x	x	x	x			x	
Andreas & Klein (2016)	x	x	x				Ref		
Evtimova et al. (2018)	x		x				visual-Ref		x
Foerster et al. (2018)	x	x	x				Ref		
Hawkins et al. (2020)	x	x	x				Ref		
Rodriguez et al. (2019)	x		x						x
Xie et al. (2020)	x	x							
Rabinowitz et al. (2018)	x								
Brandizzi et al. (2021)			x		x			ww negotiation	
Cao et al. (2018)			x						
Raileanu et al. (2018)			x						x
Das et al. (2019)		x	x		x			navigation	
Choi et al. (2018)			x				Ref		
Bogin et al. (2018)			x				x	x	
Grover et al. (2018)			x					navigation	
Bachrach et al. (2020)			x						x
Das et al. (2017)			x				visual-TnT		
Fitzgerald (2019)			x				Ref		
Graesser et al. (2019)			x			x	Multi-Ref		
Kirby et al. (2014)			x			x	x		
Kottur et al. (2017)			x				TnT		
Li & Bowling (2019)			x			x	Ref		
Liang et al. (2020)			x		x		TnT	navigation	
Mordatch & Abbeel (2018)			x						
Guo et al. (2019)			x			x		x	
Ren et al. (2020)			x			x	Ref		
Cogswell et al. (2019)			x			x	TnT		
Dagan et al. (2021)			x			x	Ref		

Table 1: Complete list of treated works and related categories.

**Theory of Mind** refers to Sec. 3, where instances are either classified as *modeling* (other agents) and *influencing*. Some rows also distinct between the kind of modeling embedded into agents: trustworthiness is investigated in competitive settings where the nature of a received message is not clear a priori; speaker-selection looks into the efficiency of communication; imitation learning is concerned with belief modeling based on non communicative actions; obverter refers to the obverter technique.

**Interaction type** includes *internal/external cooperation external competitiveness*, and *iterative learning*, as defined in Sec. 2. External includes multiple teams of interacting agents, while internal refers to the dynamic of a single team.

Lastly, **game environment**, categorizes the setting based on the role of communication which can be the goal of the game, *com-goal*, of assistance, *com-aid*, or not present, *no-com*. *Com-goal* can be either a referential game, Ref, a Task and Talk, TnT, or a visual by-product of the two. *Com-aid* mainly differentiate to negotiation or navigation tasks.




Figure 2: Number of papers per year of publication