

# Mastering emergent language: learning to guide in simulated navigation

Anonymous EMNLP submission

## Abstract

To cooperate with humans effectively, virtual agents need to be able to understand and execute language instructions. A typical setup to achieve this is with a scripted teacher which guides a virtual agent using language instructions. However, such setup has clear limitations in scalability and, more importantly, it is not interactive. Here, we introduce an autonomous agent that uses discrete communication to interactively guide other agents to navigate and act on a simulated environment. The developed communication protocol is trainable, emergent and requires no additional supervision. The emergent language speeds up learning of new agents, it generalizes across incrementally more difficult tasks and, contrary to most other emergent languages, it is highly interpretable. We demonstrate how the emitted messages correlate with particular actions and observations, and how new agents become less dependent on this guidance as training progresses. By exploiting the correlations identified in our analysis, we manage to successfully address the agents in their own language.

## 1 Introduction

Developing intelligent agents that can communicate with humans in a cooperative way is a long-standing goal of AI (Wooldridge, 2009; Mikolov et al., 2016). Because supervised approaches are too static, not allowing agents to learn the interactive aspects of communication, current research on agent communication tends to focus on processes where language can naturally emerge (Lazaridou et al., 2017; Mordatch and Abbeel, 2017; Havrylov and Titov, 2017).

In this paper, we introduce an autonomous agent (the Guide) that uses discrete communication to interactively assist another agent (the Learner) in navigational and operational tasks. In contrast

with previous work where a Guide’s language is hardcoded (Co-Reyes et al., 2018), the communication protocol used by our Guide is fully emergent and does not require additional supervision to be trained. Nevertheless, the Guide speaks a language that preserves the desired properties of a manually programmed module: it is a language that generalizes compositionally across incrementally more difficult problems, and that is highly interpretable.

By analyzing the emergent communication protocols, we demonstrate strong correlations between messages and actions, and between messages and salient properties of the environment. Quantifying the causal influence of communication shows how Learners that are assisted by a Guide gradually become less dependent on the message channel. Moreover, we show how to interpret the ungrounded agent language, and even learn to ‘speak’ it by letting an agent follow pre-set trajectories described by discrete messages.

The paper makes three main contributions. First, we introduce an autonomous agent that can interactively guide another agent via emergent, discrete communication, speeding up the learning in a collection of sequential decision making problems. Second, we demonstrate that the developed language is general enough to be reused in levels different from the level it was developed on. Third, we perform extensive analysis on the observed communication, which enables us to interpret the agents’ language, and even to address them directly.

## 2 Related Work

In this section we refer to important related work on the main themes of the current research: following language instructions, imitation learning and emergent communication.

**Following language instructions** Much research has been dedicated to linguistic instruction following. Winograd (1972) tried to develop a program to interpret language directly, by hard-coding a great number of linguistic and physical regularities. More recently, Cangelosi et al. (2006) showed that robots equipped with neural networks are capable of learning action concepts from language instructions. Similar results were obtained by Tellex et al. (2011), Chen and Mooney (2011) and Artzi and Zettlemoyer (2013).

In recent years, several artificial environments have been proposed to develop and assess agents' instruction-following capacities (Hermann et al., 2017; Brodeur et al., 2017; Chaplot et al., 2018; Wu et al., 2018). Relevant research includes Mei et al. (2016), Yu et al. (2018), Bahdanau et al. (2018), Wang et al. (2016) and Williams et al. (2018).

In this work we use a synthetic gridworld with natural language instructions called 'BabyAI'. Chevalier-Boisvert et al. (2018) introduced the set-up and presented a baseline model. Co-Reyes et al. (2018) extended this model with a correction module, which provides additional linguistic feedback to an agent. This module, however, is fully hard-coded. We introduce a comparable mechanism that is entirely trained.

**Imitation learning** Most of the models in our research are trained using imitation learning, where the objective of an agent is to simulate the behavior of another agent with more experience or knowledge of a particular task. Forms of imitation learning have been successfully applied in the context of following language instructions by e.g. Hemachandra et al. (2015), Bahdanau et al. (2018) and Chevalier-Boisvert et al. (2018).

A range of imitation learning algorithms has been proposed in the literature, including Searn (Daumé et al., 2009), SMILe (Ross and Bag-nell, 2010) and DAgger (Ross et al., 2011). Inverse reinforcement learning can also be applied, as suggested by Russell (1998), Ng and Russell (2000) and Ziebart et al. (2008). We use behavioral cloning (Pomerleau, 1991), which treats expert trajectories as series of states labelled with actions that a new agent can learn to predict in a supervised setting.

**Emergent communication** In the current research, we study the emergence of communication

through discrete language in an artificial, multi-agent set-up. The emergence and evolution of language has received much attention in linguistics, game theory and cognitive science (Skyrms, 2010; Wagner et al., 2003; Crawford, 1998).

Experiments related to ours were conducted by Jorge et al. (2016) and Sukhbaatar et al. (2016), who showed that agents can autonomously evolve a communication protocol that helps them to play a game. Lazaridou et al. (2017) obtained similar results, but also investigated how to change the environment in order to interpret the emergent language more easily. Inspired by this research, Havrylov and Titov (2017) developed a referential game where agents learn to assist each other by sending variable-length sequences of discrete symbols. Mordatch and Abbeel (2017) showed that multi-agent communities can give rise to a grounded compositional language that helps speakers achieve their goals.

### 3 Approach

In this section we describe the BabyAI game, as well as the models used in the experiments and the training regime.

#### 3.1 BabyAI game

All our experiments are conducted using the BabyAI framework, introduced by Chevalier-Boisvert et al. (2018).<sup>1</sup> The platform lets virtual agents play games at a number of different levels. Each of these levels are combinations of a partially observable two-dimensional gridworld, and a mission presented as a language instruction. The challenges include various operational and navigational tasks. The environment contains several kinds of objects (e.g. balls, keys, boxes and doors) that come in different colors. See Figure 1 for a sample frame.

At each time step  $t$ , an agent receives as input the instruction string  $i_t$  (constant during a game), and the  $7 \times 7 \times 3$ -tensor  $o_t$  encoding the currently observable part of the surroundings. Her subsequent action  $a_t$  alters the state of the environment and thereby determines the next observation  $o_{t+1}$ .

In total, BabyAI has 19 levels with increasing degrees of complexity. In this paper, we consider six of these levels, which we consider to be a representative fragment of the game as a whole. See Table A.1 in the supplementary materials for

<sup>1</sup>Code available at <https://github.com/mila-udem/babyai>.




Figure 1: A screenshot of BabyAI level PutNextLocal with mission ‘put the green box next to the green ball’. The red arrow represents the agent and the shaded area the currently observable part of the environment.

an overview of the specific core skills required to solve each of them. The BabyAI levels are designed in such a way that competencies required to solve some of them are also needed in others. E.g., to complete GoToLocal an agent has to know how to navigate a single room with distractors, and this skill is also necessary to address PickupLoc. Thanks to this hierarchy, BabyAI is a suitable framework for studying curriculum learning.

### 3.2 Model

We consider two kinds of agents: *Learners* and *Guides*. The fully continuous Learners can approach a task alone, or be assisted by a Guide. A Guide is an agent that not only learns to solve a task, but also to emit discrete messages. These messages can be used to help a new Learner master a task more efficiently.<sup>2</sup>

**The Learner** The Learner architecture is as described by [Chevalier-Boisvert et al. \(2018\)](#). Instructions are presented to a GRU ([Cho et al., 2014](#)) and observations to a convolutional network. The results of these operations are processed together by a FiLM module ([Perez et al., 2018](#)) to obtain a joint representation of the command and the environment. A memory LSTM cell ([Hochreiter and Schmidhuber, 1997](#)) uses this representation, together with past FiLM outputs, to compute the tensor that is passed to a two-layered policy module. The action receiving the highest probability after applying softmax to the policy output is executed.

**The Guide** The Guide’s architecture is similar to the Learner’s, but differs in one crucial aspect: between the memory LSTM and the policy mod-

ule we introduce a bottleneck. Its function is to create the communication channel, where the information flowing from the input-processing layers has to be compressed in the form of a sequence of discrete tokens. These messages can later be sent to a Learner.

Generation of the messages takes place in a GRU decoder that receives the Guide’s memory LSTM output as input. To do so, continuous information has to be discretized. In order to maintain end-to-end differentiability, we use a straight-through Gumbel softmax estimator for this purpose ([Bengio et al., 2013](#); [Jang et al., 2017](#)). The resulting string of symbols can be passed to the policy module to predict an action after it has been encoded again by a GRU encoder.

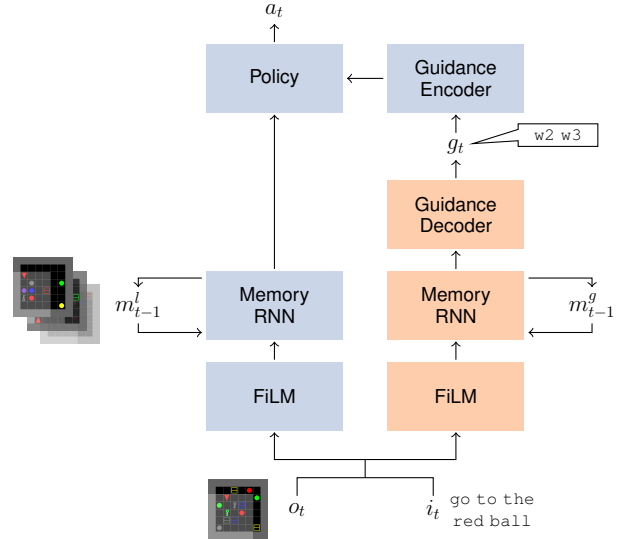


Figure 2: Schematic visualization of the model architecture. Included variables are  $o_t$ : state of the (observable part of the) environment,  $i_t$ : linguistic instruction,  $g_t$ : linguistic guidance,  $a_t$ : action,  $m_{t-1}^l$ : Learner’s memory state,  $m_{t-1}^g$ : Guide’s memory state. Subscript represents point in time. Blue modules are considered part of the Learner; red ones part of the Guide.

By first training a Guide and then pairing the experienced Guide with a new Learner, we investigate if and how agents can learn to make use of the available communication channel. The set-up used to test this is visualized in Figure 2. A new Learner is coupled with a pretrained Guide, and receives her discrete guidance at each time step. The Learner is provided with a new GRU encoder to interpret the guidance, and uses this information as additional input to her policy module.

<sup>2</sup>Code will be released upon acceptance.

### 3.3 Training

Because we use behavioral cloning as imitation learning algorithm, we effectively reduce training to a supervised setting. The Learner, the Guide and the combined Learner-Guide set-ups are all end-to-end architectures that receive observations and instructions as input, and whose parameters are updated by backpropagating a negative log likelihood loss function.

The agents are supposed to internalize the policy demonstrated by an expert agent, which requires that we first have such an expert at our disposal. [Chevalier-Boisvert et al. \(2018\)](#) implemented a heuristic bot that is able to solve all BabyAI levels. We do not use this hard-coded expert, but instead apply proximal policy optimization ([Schulman et al., 2017](#)) to train RL (Reinforcement Learning) experts. We do this because it makes our approach more generalizable, and because [Chevalier-Boisvert et al.](#) note that ‘demonstrations produced by [an RL] agent are easier for the learner to imitate’ than those produced by a symbolic bot. Therefore, our experiments follow three consecutive training stages: (1) Train RL expert (with Learner set-up) to generate demonstrations. (2) Train new Learner (as baseline) and new Guide separately on the RL expert data, by imitation learning. (3) Train new Learner + pretrained Guide (from stage 2) on the RL expert data, by imitation learning. In stage 3, the pretrained Guide’s weights are not frozen so they can be finetuned.

Stage 1 only has to be performed once to generate the expert data that we use in imitation learning. In all our experiments, the Guide emits messages of a fixed length of 2. There are three distinct tokens ( $w_0$ ,  $w_1$  and  $w_2$ ), so that there are 9 possible messages in total. For an overview of the hyperparameter settings we refer to Table A.2 in the supplementary materials.

## 4 Experiments and results

### 4.1 Intra-level guidance

First, we study the effect of the communication channel within single levels. We train the baseline model (the default Learner), and compare its performance over time with a Learner that receives messages from the best-performing pretrained Guide, which is itself being finetuned to coadapt with the Learner. We compare the models in terms of validation success rate, which is the percentage of validation episodes (out of 500) that

a model manages to complete successfully. Note that this metric differs from the accuracy, which is the percentage of individual actions that are predicted correctly.

We do not expect the guided Learner to achieve a higher success rate than the baseline Learner, because we already know that the Learner is in principle able to solve all levels when given enough training time. What we do hope to see is that a new Learner masters a task more quickly when aided by the messages of a pretrained Guide.

Results of three runs with three random seeds on all levels are shown in Figure 3. These plots clearly show that on each of the considered levels, the messages emitted by the pretrained Guide cause a great increase in the learning speed of a new agent.

### 4.2 Curriculum guidance

We assess if the Guide is also useful when we apply curriculum learning: are the messages emitted by a Guide pretrained at one level helpful for a Learner facing a more complex level? We consider transfer from GoToLocal to PickupLoc and PutNextLocal, and between PickupLoc and PutNextLocal (both directions).

Guides are trained at the mentioned base levels, and coupled with a new Learner at the related target levels. In all of these cases, we hypothesize that the messages sent by the Guide should be useful also at the target levels, even though the Guide never actually encountered them. This is because all of the levels concerned require an agent to be capable of navigating a single room with distractors (competencies ROOM, DISTR-BOX and DISTR of Table A.1 in the supplementary material). Hence, if the guidance messages are sufficiently general, they should be useful across base and target levels.

We compare performance of the models with different base and target levels with those for which base and target level coincide (discussed in the previous section). Additionally, we assess the performance at the target levels of a single Learner pretrained at the base levels. We expect a pretrained Learner to perform better than a new Learner assisted by a pretrained Guide, because in the former all parameters have been optimized for a related level, whereas in the latter the Learner can only profit from the pretraining by means of the discrete guidance messages.



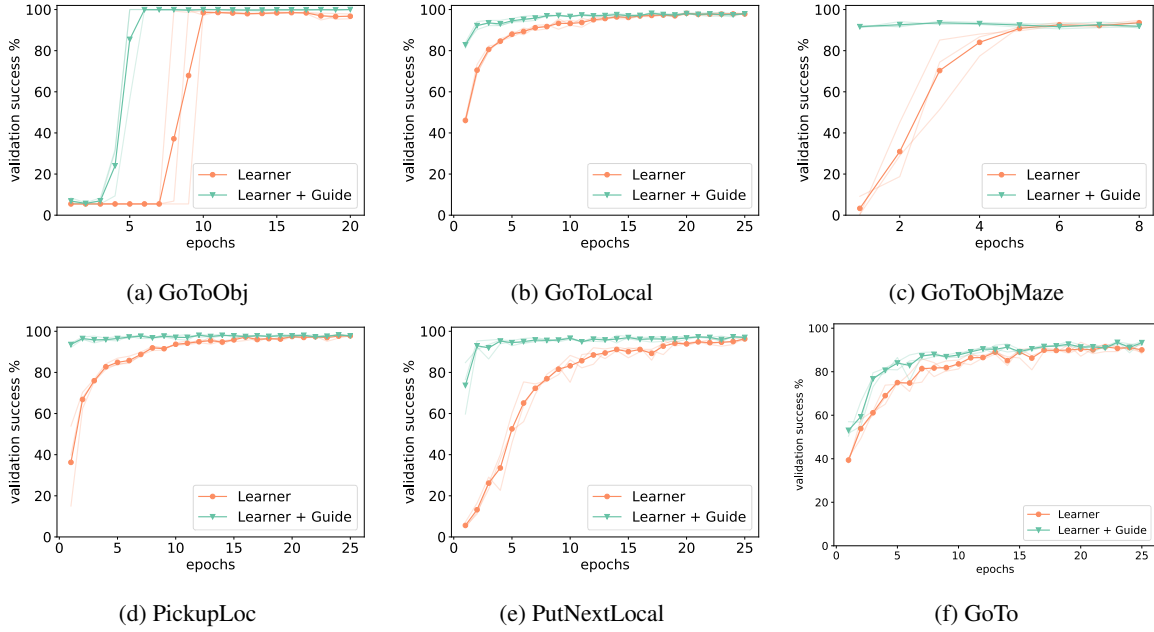


Figure 3: Development of validation success rate when training a single Learner and a Learner assisted by a pretrained Guide on the indicated BabyAI levels. Average over three runs; results per run are shown by shaded lines.

Figure 4 shows the results. In level PickupLoc, we see that pretraining the Guide at levels GoToLocal and PutNextLocal brings the learning curve close to the one of a learner with a Guide pretrained at PickupLoc itself. Surprisingly, the difference between the new Learners with a pretrained Guide and the pretrained Learners is not significant. In PutNextLocal, the results are more divergent, and the Learner with a Guide pretrained at the same level outperforms all other set-ups. Especially pretraining the Guide at PickupLoc accelerates learning, but not as much as pretraining the Learner at PickupLoc.

The observation that Guides can even be helpful to new Learners at levels for which they were not optimized suggests that their messages convey information that is relevant to the internalization of skills that generalize across levels.

### 4.3 Multi-level guidance

The previous experiment focused on the transfer of guidance from one level to another one. This transfer can be considered inductive because it relies on generalization to unseen levels. In the next experiment we also check for a deductive transfer: if a Guide is trained on a combination of levels, how useful is she when assisting a Learner trained from scratch on only one of these levels?

To investigate this, we train a single Guide on

GoToLocal, PickupLoc and PutNextLocal at the same time, by aggregating the training sets of these three levels. None of the model dimensions are changed, implying that this multi-level Guide has to address three different levels with the same number of parameters as a regular Guide. Because of this handicap, we hypothesize that the multi-level Guide will not accelerate the development of a new Learner as rapidly as a Guide specialized at a single level.

The results of this experiment are visualized in Figure 5. At all three levels the multi-level Guide turns out to be just as helpful to a new Learner as a Guide that was specifically optimized for a single level. This again suggests that the Guide does not produce messages aimed at solving specific levels, but rather at acquiring skills that are relevant in each of them.

## 5 Analysis

In this section we investigate the communication evolved by the agents in more detail.

### 5.1 Messages/actions correlations

We analyse the conditional probability distribution of actions given the messages produced by the pretrained or finetuned Guide. This reveals a strong correlation between messages and actions, as illustrated in Figure 6, visualizing the condi-

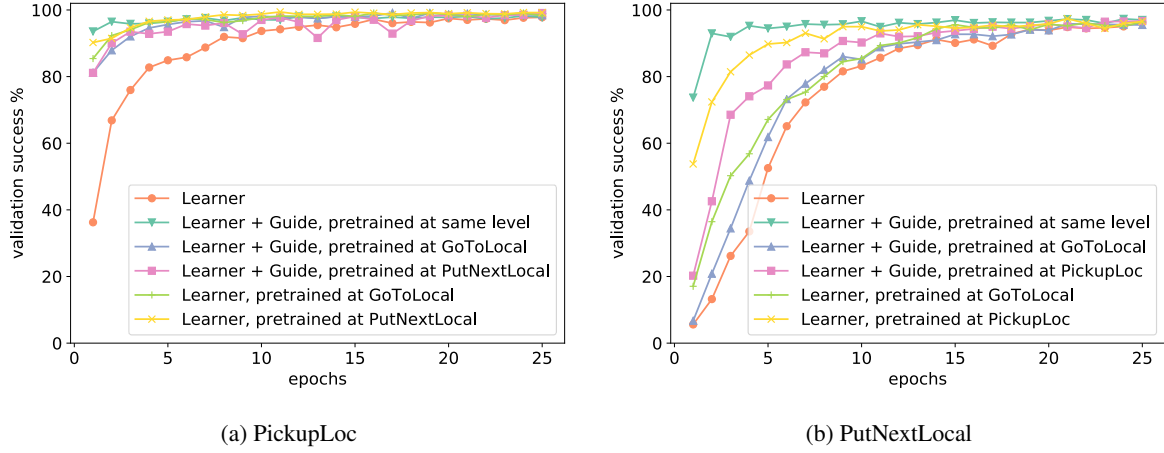


Figure 4: Development of validation success rate when training a single Learner from scratch, a Learner assisted by a pre-trained Guide, or a pre-trained Learner on several BabyAI levels.

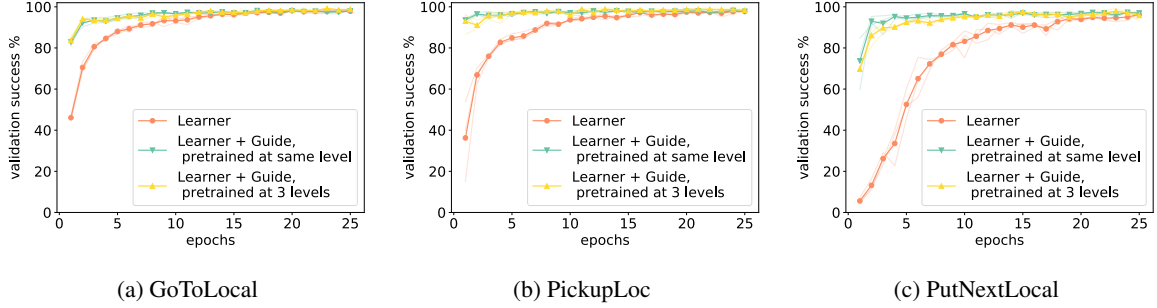


Figure 5: Development of validation success rate when training a single Learner, a Learner assisted by a Guide pretrained on a single level and a Learner pretrained on GoToLocal, PickupLoc and PutNextLocal simultaneously. Average over three runs; results per run are shown by shaded lines.

tional distribution of actions taken by the best-performing Guide, conditioned on her own messages, after training stage 2.

We show the results for level PutNextLocal, because here the agent has to perform most actions. Similar plots can be obtained at other levels. The distribution changes after pairing the Guide with a new Learner, and gradually becomes less spiky as learning progresses. We refer to the supplementary materials for Figure A.1, an illustration of this distribution in level PutNextLocal, and Figure A.2, which shows both distributions in level GoToObj.

## 5.2 Input/messages correlations

We now look at the conditional distribution of the Guide’s messages, given the observational input. To do so, we categorize the agent’s field of perception according to four salient observable facts: whether the target object lies directly ahead of the agent, on her right, on her left, or whether the target object is not currently visible.

In Figure 7, the results of the pretrained Guide

are shown for level GoToObj. They show that the Guide has a clear preference for a particular message when the object that has to be navigated towards lies ahead, on the left or on the right. Apparently, the guidance message distribution is conditioned on the observation of such situational circumstances. When the target object is invisible, the distribution is more spread out, indicating that the agent lets her policy depend on other factors in such cases.

Similar plots can be obtained at other levels, as shown for GoToLocal in Figure A.3 in the supplementary materials. Distributions at higher levels tend to be somewhat more diffuse, indicating a higher uncertainty about the messages to emit. This is probably due to the presence of distractors. Moreover, there must be many other observable facts in the environment that impact the distributions, and that are not included here.




Figure 6: Barplots visualizing the conditional distribution of actions taken in PutNextLocal by the best Guide after stage 2, given the generated messages, based on 500 sample episodes unseen during training. The bins in the individual barplots correspond to actions and the bar heights to probabilities. E.g. the top left barplot shows that the Guide performs action 3 with probability 1.0 after emitting message  $w_0 \ w_0$ .

### 5.3 Causal influence

In the analysis so far, we looked at statistical correlations. This does not tell us anything yet about causality: do the Guide’s messages actually *cause* behavioral patterns in the Learner? To quantify the impact that the guidance messages have on the actions taken by a new Learner, we implement and compute a metric introduced for this purpose by Lowe et al. (2019): Causal Influence of Communication (CIC).

CIC is intended to measure ‘positive listening’: the extent to which emergent communication influences an agent’s behavior. We consider one-step causal influence, c.q. the effect of the Guide’s message on the next action of the Learner. It is computed as:

$$\text{CIC} = \frac{1}{|T|} \sum_{t \in T} \sum_{m \in M} \sum_{a \in A} p_t(a, m) \log \frac{p_t(a, m)}{p_t(a)p_t(m)}, \quad (1)$$

where  $T$  is the set of ‘test games’ (c.q. the frames of the validation episodes),  $M$  the set of messages the Guide can send and  $A$  the set of actions the Learner can take. Each  $t \in T$  determines the probability distributions  $p_t(a)$ ,  $p_t(m)$  of Learner’s actions and Guide’s messages, as well as their joint distribution  $p_t(a, m)$ .

Figure 8 visualizes the development of the CIC metric in the set-up combining a new Learner with a pretrained Guide at all considered levels. The plots clearly show that CIC tends to decrease over

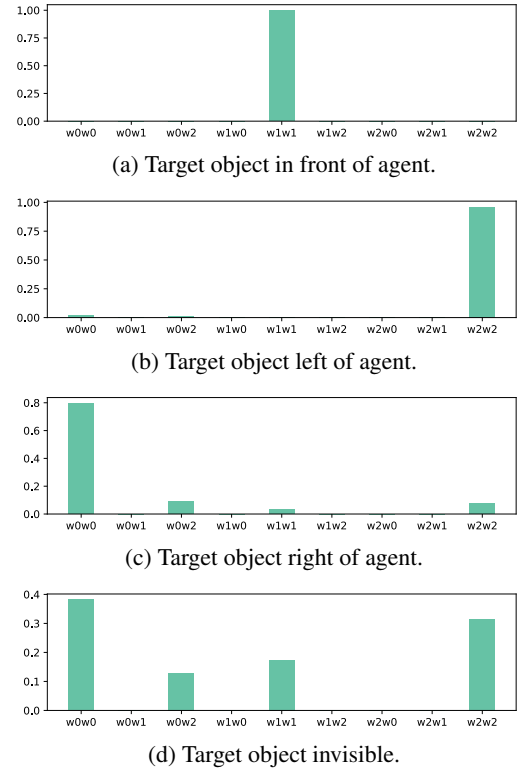


Figure 7: Barplots visualizing the conditional distribution of messages produced by the best Guide in stage 2 at level GoToObj, given the indicated observable events, based on 500 sample episodes unseen during training.

time, as illustrated by the linear trendlines. This indicates that the guidance messages have a high causal effect on the Learner’s actions in the early stages, and gradually become less important as training progresses and the Learner grows more confident. The only exception in Figure 8 is level GoToObj, where we see a positive correlation between training time and CIC. We note, however, that also at this level the trend is negative if more epochs after model convergence are considered.

### 5.4 Learning the agents’ language

The correlations between actions and messages taught us how to interpret the Guide’s language. This means that we do not only know how to make sense of witnessed utterances, but that we can also try to speak the language ourselves, encouraging the Learner to perform the actions that have proven to correlate with particular messages. To test this, we train a Learner with a pretrained Guide at levels GoToObj and GoToLocal, and identify the epoch where the CIC is highest. This should be the point in time where the Learner is most susceptible to

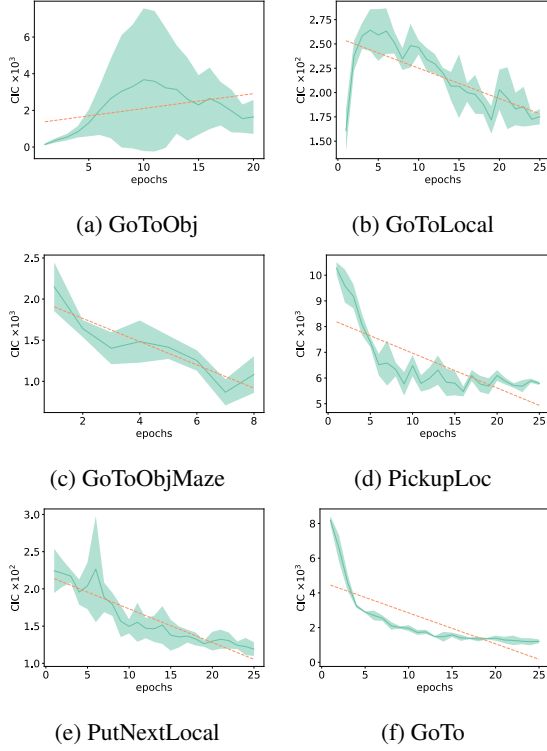
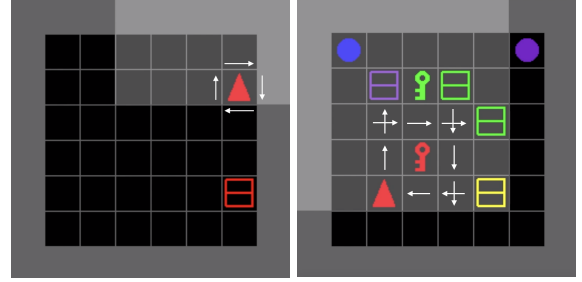


Figure 8: Development of CIC on 500 validation samples when training a Learner assisted by a pretrained Guide on the indicated BabyAI levels. Average over three runs; standard deviations are shown by shaded regions. Dashed orange lines show linear trends.

the message channel. We compute the correlations between guidance messages and performed actions at this stage, as in Section 5.1, to establish the expected correspondence between messages and actions. Next, we ‘hijack’ the Guide, and start sending our own messages to the Learner.

Communication success is assessed in two ways. First, we provide the Learner with a set of messages that should describe a pre-set trajectory (a ‘choreography’) in the grid, as per the computed message-action correlations. In particular, we try to tell the agent to turn around its own axis (the ‘pirouette’), and to describe a larger circular movement (the ‘waltz’). We check if the performed action sequences adhere to these choreographies. Second, we report quantitative results by randomly sending the Learner 500 messages that strongly correlate with specific actions. We then calculate the ‘obedience’: the percentage of messages that are followed by the expected action.

In GoToObj, the Learner successfully performs a pirouette and a waltz upon receiving the guidance messages that are supposed to describe these movements. See Figure 9a for the pirouette. This



(a) Pirouette in GoToObj (b) Waltz in PutNextLocal

Figure 9: Trajectories of Learners in indicated levels when told by means of messages to perform (a) a pirouette (c.q.  $4 \times w0 \ w0$ ) and (b) a waltz (c.q.  $4 \times w1 \ w2, w1 \ w2, w1 \ w1$ )

is also the case in GoToLocal. See Figure 9b for the waltz. In GoToObj we reach an obedience of 76%, and in GoToLocal one of 72%. These percentages show that the Learner does not blindly follow the guidance messages, even when the CIC is highest. The observational input always remains important, especially when random messages are sent so that the correlation between input and guidance is broken. However, in most cases the agent does what we tell her to. Hence, to some extent we have mastered the agent language.

## 6 Conclusion

We showed how neural agents can autonomously develop a discrete communication channel to help each other solve the BabyAI game. By applying curriculum learning, we demonstrated that the guidance messages encode information that is general enough to be useful in unseen levels. In the analysis, we saw that messages correlate with actions as well as with observations. Application of the CIC metric revealed that new agents gradually become more independent of the guidance messages they receive. Finally, we used the interpretability of the emerged language to successfully communicate with agents directly.

In future research we would like to extend the developed method to frameworks other than BabyAI. We also wish to investigate methods to break the strong correlation between messages and actions, and incentivize guidance of a more general nature. Moreover, we are interested in experimenting with the number of tokens, the message length and the learning method, to see if similar results can be obtained if we use reinforcement instead of imitation learning.



## References

- Yoav Artzi and Luke Zettlemoyer. 2013. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association for Computational Linguistics*, 1:49–62.
- Dzmitry Bahdanau, Felix Hill, Jan Leike, Edward Hughes, Pushmeet Kohli, and Edward Grefenstette. 2018. Learning to follow language instructions with adversarial reward induction. *arXiv preprint arXiv:1806.01946*.
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Simon Brodeur, Ethan Perez, Ankesh Anand, Florian Golemo, Luca Celotti, Florian Strub, Jean Rouat, Hugo Larochelle, and Aaron Courville. 2017. Home: A household multimodal environment. *arXiv preprint arXiv:1711.11017*.
- Angelo Cangelosi, Emmanouil Hourdakis, and Vadim Tikhonoff. 2006. Language acquisition and symbol grounding transfer with neural networks and cognitive robots. In *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, pages 1576–1582. IEEE.
- Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, and Ruslan Salakhutdinov. 2018. Gated-attention architectures for task-oriented language grounding. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- David L Chen and Raymond J Mooney. 2011. Learning to interpret natural language navigation instructions from observations. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. 2018. Babyai: First steps towards grounded language learning with a human in the loop. *arXiv preprint arXiv:1810.08272*.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- John D Co-Reyes, Abhishek Gupta, Suvansh Sanjeev, Nick Altieri, John DeNero, Pieter Abbeel, and Sergey Levine. 2018. Guiding policies with language via meta-learning. *arXiv preprint arXiv:1811.07882*.
- Vincent Crawford. 1998. A survey of experiments on communication via cheap talk. *Journal of Economic Theory*, 78(2):286–298.
- Hal Daumé, John Langford, and Daniel Marcu. 2009. Search-based structured prediction. *Machine learning*, 75(3):297–325.
- Serhii Havrylov and Ivan Titov. 2017. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2149–2159. Curran Associates, Inc.
- Sachithra Hemachandra, Felix Duvallet, Thomas M Howard, Nicholas Roy, Anthony Stentz, and Matthew R Walter. 2015. Learning models for following natural language directions in unknown environments. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5608–5615. IEEE.
- Karl Moritz Hermann, Felix Hill, Simon Green, Fumin Wang, Ryan Faulkner, Hubert Soyer, David Szepesvari, Wojciech Marian Czarnecki, Max Jaderberg, Denis Teplyashin, et al. 2017. Grounded language learning in a simulated 3d world. *arXiv preprint arXiv:1706.06551*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *Proceedings of the International Conference on Learning Representations*.
- Emilio Jorge, Mikael Kågebäck, and Emil Gustavsson. 2016. Learning to play guess who? and inventing a grounded language as a consequence. *CoRR*, abs/1611.03218.
- Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017. Multi-agent cooperation and the emergence of (natural) language. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Ryan Lowe, Jakob Foerster, Y-Lan Boureau, Joelle Pineau, and Yann Dauphin. 2019. On the pitfalls of measuring emergent communication.
- Hongyuan Mei, Mohit Bansal, and Matthew R Walter. 2016. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Tomas Mikolov, Armand Joulin, and Marco Baroni. 2016. A roadmap towards machine intelligence. In *Computational Linguistics and Intelligent Text Processing - 17th International Conference, CICLing 2016, Konya, Turkey, April 3-9, 2016, Revised Selected Papers, Part I*, pages 29–61.
- Igor Mordatch and Pieter Abbeel. 2017. Emergence of grounded compositional language in multi-agent populations.

- Andrew Y. Ng and Stuart Russell. 2000. Algorithms for inverse reinforcement learning. In *Proc. 17th International Conf. on Machine Learning*, pages 663–670. Morgan Kaufmann.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- D. A. Pomerleau. 1991. [Efficient training of artificial neural networks for autonomous navigation](#). *Neural Computation*, 3(1):88–97.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635.
- Stéphane Ross and Drew Bagnell. 2010. [Efficient reductions for imitation learning](#). In *AISTATS*, volume 9 of *JMLR Proceedings*, pages 661–668. JMLR.org.
- Stuart Russell. 1998. Learning agents for uncertain environments (extended abstract). In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, pages 101–103. ACM Press.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Brian Skyrms. 2010. *Signals: Evolution, Learning, and Information*.
- Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2016. [Learning multiagent communication with backpropagation](#). In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2244–2252. Curran Associates, Inc.
- Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R Walter, Ashis Gopal Banerjee, Seth Teller, and Nicholas Roy. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- Kyle Wagner, James A. Reggia, Juan Uriagereka, and Gerald S. Wilkinson. 2003. [Progress in the simulation of emergent communication and language](#). *Adaptive Behavior*, 11(1):37–69.
- Sida I Wang, Percy Liang, and Christopher D Manning. 2016. Learning language games through interaction. *arXiv preprint arXiv:1606.02447*.
- Edward C Williams, Nakul Gopalan, Mine Rhee, and Stefanie Tellex. 2018. Learning to parse natural language to grounded reward functions with weak supervision. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–7. IEEE.
- Terry Winograd. 1972. Understanding natural language. *Cognitive psychology*, 3(1):1–191.
- Michael J. Wooldridge. 2009. *An Introduction to MultiAgent Systems (2. ed.)*. Wiley.
- Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. 2018. Building generalizable agents with a realistic and rich 3d environment. *arXiv preprint arXiv:1801.02209*.
- Haonan Yu, Haichao Zhang, and Wei Xu. 2018. Interactive grounded language acquisition and generalization in a 2d world. *arXiv preprint arXiv:1802.01433*.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA.

## A Supplemental Materials

	required skills							
	ROOM	DISTR-BOX	DISTR	GOTO	PUT	PICKUP	LOC	MAZE
GoToObj	x							
GoToLocal	x	x	x	x				
GoToObjMaze	x							x
PickupLoc	x	x	x			x	x	
PutNextLocal	x	x	x		x			
GoTo	x	x	x	x				x

Table A.1: Considered BabyAI levels and skills required to solve them, from [Chevalier-Boisvert et al. \(2018\)](#). ROOM: navigating a  $6 \times 6$  room; DISTR-BOX: ignoring distracting grey boxes; DISTR: ignoring distracting objects of any kind; GOTO: understanding ‘go to’ instructions; PUT: understanding ‘put’ instructions; PICKUP: understanding ‘pick up’ instructions; LOC: understanding location expressed relative to initial agent position; MAZE: navigating a  $3 \times 3$  maze of randomly connected  $6 \times 6$  rooms.

# observation CNN output units	256
# memory LSTM units	2048
# instruction GRU units	256
word embedding size	256
guidance message length	2
guidance vocabulary size	3
batch size	128 (GoToObj, GoTo), 256 (GoToObjMaze), 512 (otherwise)
learning rate	$1 \cdot 10^{-4}$
optimizer	Adam ( $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 1 \cdot 10^{-5}$ )
# training instances	1K (GoToObj), 25K (GoToObjMaze), 50K (GoToLocal), 100K (otherwise)

Table A.2: Hyperparameter settings used in experiments.

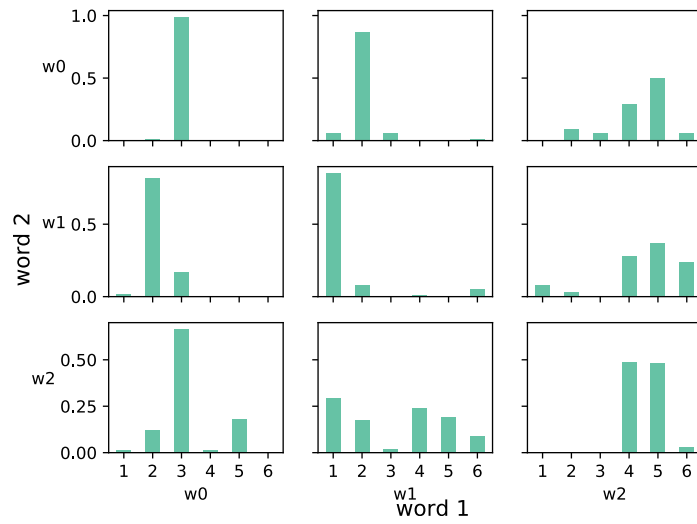




Figure A.1: Barplots visualizing the conditional distribution of actions taken in PutNextLocal by the best Learner after stage 3, given the Guide’s messages, based on 500 sample episodes unseen during training. The bins in the individual barplots correspond to actions and the bar heights to probabilities. E.g. the top left barplot shows that the Guide performs action 3 with probability close to 1.0 after emitting message  $w_0$ .






(a)




(b)


Figure A.2: Barplots visualizing the conditional distribution of actions taken in GoToObj by (a) the best Guide after stage 2 and (b) the best Learner after stage 3, given the Guide’s messages, based on 500 sample episodes unseen during training. The bins in the individual barplots correspond to actions and the bar heights to probabilities. E.g. the top left barplot in Figure (a) shows that the Guide performs action 2 with probability 1.0 after emitting message  $w_0$ .




(a) Target object in front of agent.



(b) Target object left of agent.



(c) Target object right of agent.



(d) Target object invisible.

Figure A.3: Barplots visualizing the conditional distribution of messages produced by the best Guide in stage 2 at level GoToLocal, given the indicated observable events, based on 500 sample episodes unseen during training.