

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Characterizing tradeoffs between teaching via language and demonstrations in multi-agent systems

#### **Permalink**

<https://escholarship.org/uc/item/97j4r1nn>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

#### **Authors**

Yu, Dhara  
Goodman, Noah  
Mu, Jesse

#### **Publication Date**

2023

Peer reviewed

# Characterizing tradeoffs between teaching via language and demonstrations in multi-agent systems

Dhara Yu, Noah D. Goodman, Jesse Mu

Stanford University

{dharakyu, ngoodman, muj}@stanford.edu

## Abstract

Humans teach others about the world through *language* and *demonstration*. When might one of these modalities be more effective than the other? In this work, we study the factors that modulate the effectiveness of language vs. demonstration using multi-agent systems to model human communication. Specifically, we train neural network agents to teach via language or demonstration in a grounded communication task, manipulating 1) the inherent difficulty of the task and 2) the competence of the teacher. We find that teaching by demonstration is more effective in the simplest settings, but language is more effective as task difficulty increases, due to its ability to generalize more effectively to unseen scenarios. Overall, these results provide converging evidence for a tradeoff between language and demonstration as teaching modalities in humans, and make the novel predictions that demonstration may be optimal for easy tasks, while language enables generalization in more challenging settings.

**Keywords:** teaching; communication; reference games; computational modeling; neural networks

## Introduction

One of the defining characteristics of human collective intelligence is the ability for individuals to transmit information about the world to others through the act of teaching (Tomasello, 1999). While teaching takes many forms, two of the most important modalities are teaching by *language*, i.e. describing an underlying rule or desired behavior, or *demonstration*, i.e. directly exhibiting the desired behavior. For instance, an experienced chef might instruct a novice by providing a language explanation of a recipe, or demonstrate the process through step-by-step preparation.

When is it optimal to use language vs. demonstrations to teach others? Studies of child development suggest that teaching by demonstration may be a cognitively simpler modality: young children tend to teach with demonstration whereas older children and adults teach with language (Ellis & Rogoff, 1982; Strauss et al., 2002). On the other hand, language may be better at conveying richer generalizations, especially in tasks requiring more abstract reasoning (Chopra et al., 2019; Sumers et al., 2023; Tessler et al., 2020).

While one might conclude from this literature that language is a more effective medium than demonstrations for complex tasks, one confounding factor is human language depends on rich shared conventions and abstractions (Clark, 1996) that make it especially effective for teaching. Thus, the extent to which these tradeoffs exist for interlocuters learning conventions from scratch is unknown. Understanding these tradeoffs

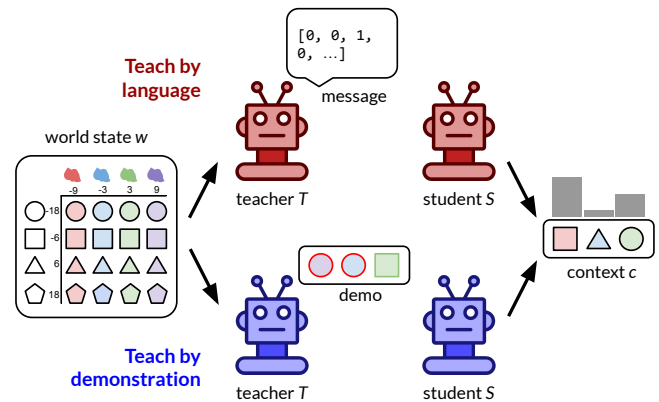


Figure 1: Overview of signaling bandits task in the language (red) and demonstration (blue) conditions. World state  $w$  defines the reward associated with each (shape, color) pair. In the language setting, the teacher produces a message represented as one-hot encodings of vocabulary tokens, while in the demonstration setting, the teacher samples an example context and indicates its selection for the context. The student in both conditions selects an item in the context.

would not only provide additional evidence for the *a priori* relative effectiveness of teaching modalities, but also inform theories of language evolution grounded in teaching (Laland, 2017): did language become a prominent form of teaching because it was already equipped with useful abstractions? Or could the need to teach others more complex tasks alone motivate language development?

It is difficult to study human teaching in the complete absence of shared conventions: even humans playing abstract and synthetic tasks in the literature (Hawkins et al., 2020) quickly learn linguistic conventions that map onto familiar shared concepts. Instead, we examine the tradeoffs between language and demonstration using multi-agent communication (Lazaridou & Baroni, 2020) as a model of human communication. Specifically, to study teaching in the absence of conventions, we train neural network agents that learn open-ended communication strategies from scratch to play a game known as the signaling bandits task (Sumers et al., 2021). Within these simulations, we study how the relative effectiveness of learning language- vs. demonstration-based teaching protocols varies based on 1)

difficulty of the task and 2) competence of the teacher.

Our experiments show that while agents can devise both language- and demonstration-based teaching protocols, there are clear tradeoffs in the effectiveness of the two mediums. When the task is relatively easy, demonstration agents outperform language agents, but as task difficulty increases, language agents improve in performance, while the ability of demonstration agents drops, corroborating the experimental literature that language is a more useful medium in more challenging settings. Finally, we vary the amount of training data shown to the teacher as a way to tweak teacher competence. The results show that language as a communication medium is more robust to teachers with low competence, with demonstration agents failing to learn in limited-data settings.

Together, our results provide converging evidence consistent with behavioral experiments (Chopra et al., 2019; Sumers et al., 2023) for the increased effectiveness of human language as a teaching modality for more difficult tasks, *even in the absence of conventions*. Our models further predict that i) demonstrations may be optimal for teaching simpler tasks, and ii) language enables more effective transfer of abstractions as teacher knowledge declines. These simulations support theories that human language may have emerged partially because it is a more effective teaching medium as agents acquire and teach more complex knowledge.

## Task

As a testbed to study the emergence of language- and demonstration-based teaching strategies, we use the signaling bandits task (Sumers et al., 2021) (Figure 1). This task generalizes the classic dyadic reference game (Lewis, 1969) to the multi-armed bandit setting (Sutton & Barto, 2018). We pose this as a game to agents in our world, by imagining a *teacher* who knows how to perform the task, and who then tries to teach a *student* how to perform it. Crucially, to succeed at this game the teacher must prepare the student for a variety of situations, unlike in reference games where the teacher only needs to communicate about one specific referring context (but see Mu and Goodman, 2021 for an example of a more generalized reference game). Formally, each game  $(c, w)$  is defined by a *context*  $c$  consisting of a set of objects  $c = (c_1, \dots, c_n)$ , with each object having a set of boolean features  $\phi(c_i) \in \{0, 1\}^K$ , and a *world state*  $w$ , which describes the reward associated with each feature. For the selected object  $c_a$  and world state  $w$ , the reward function  $R$  is linear over the object features:

$$R(c, w, a) = w^\top \phi(c_a). \quad (1)$$

Given the true world state  $w$  and a set of possible utterances  $\mathcal{U}$ , the teacher  $T$ 's job is to send an utterance  $u \in \mathcal{U}$  to the student  $S$  such that the student selects the object with the highest utility for some context  $c$ , revealed only to the student. Specifically, given  $\hat{u} \sim T(u | w)$  and  $\hat{a} \sim S(a | \hat{u}, c)$ , the reward for both agents is  $R(c, w, \hat{a})$ . In this work, we examine how different classes of utterances  $\mathcal{U}$ , namely language and demonstrations, facilitate success on this task.

## Models

To study the effects of different teaching modalities, we train neural teacher and student agents to play the signaling bandits game. Specifically, the teacher will be trained to both (1) play the game given the ground-truth world state, and (2) send an utterance to the student. In turn, the student will be trained to select the highest utility object given the teacher's message and some game context.

### Teacher

We can represent the teacher via the conditional distributions:

$$T_{\text{game}}(a | c, w) \propto \exp(f^T(w)^\top g^T(c_a)), \quad (2)$$

$$T_{\text{utt}}(u | w) = \prod_{i=1}^{|u|} p_{\text{RNN}}(u_i | u_{<i}; f(w)). \quad (3)$$

Here,  $T_{\text{game}}$  is the game-playing portion of the teacher, whose job is to select objects given a context and the known ground-truth world state. Specifically, the probability of a teacher selecting object  $c_a$  given a context  $c$  and world state  $w$  is equal to the dot product between learned representations of the world state  $f^T(w)$  and the corresponding object  $g^T(c_a)$ , normalized by a softmax decision rule (Sutton & Barto, 2018), where  $f$  and  $g$  are multi-layer perceptrons (MLPs). We train the teacher to play the signaling game itself to learn representations of world states  $f^T(w)$  and contexts  $g^T(c_a)$  useful for message generation, as well as provide a source of demonstration data for the demonstration agents (introduced later).

Second,  $T_{\text{utt}}$  describes the utterance-generating portion of the teacher, where  $p_{\text{RNN}}$  is the probability of utterance token  $u_i$  according to a Gated Recurrent Unit (GRU; Cho et al., 2014) recurrent neural network (RNN) decoder that conditions on the past messages up to time  $i$ ,  $u_{<i}$ , and the same representation of the world state  $f(w)$  used in  $T_{\text{game}}$ .

### Student

The student is represented similarly to the game-playing portion of the teacher agent, except instead of using the world state  $w$ , it conditions on the teacher's utterance  $u$ :

$$S(a | u, c) \propto \exp(f^S(u)^\top g^S(c_a)). \quad (4)$$

Again, the probability of object  $c_a$  is proportional to the dot product between representations of the utterance  $f^S(u)$  and object  $g^S(c_a)$ , similar to  $f^T$  and  $g^T$  in Equation 2.

### Objectives

Now, given a dataset of games  $D = \{(c_i, w_i)\}_{i=1}^M$ , we train our agents jointly via gradient descent to maximize the expected reward  $J(S, T)$  of both the teacher and student:

$$J_T(c, w) = \mathbb{E}_{\hat{a}^T \sim T_{\text{game}}(a|c,w)} [R(c, w, \hat{a}^T)], \quad (5)$$

$$J_{S,T}(c, w) = \mathbb{E}_{\hat{u} \sim T_{\text{utt}}(u|w), \hat{a}^S \sim S(a|\hat{u},c)} [R(c, w, \hat{a}^S)], \quad (6)$$

$$J(S, T) = \mathbb{E}_{c,w \sim D} [J_T(c, w) + J_{S,T}(c, w)]. \quad (7)$$

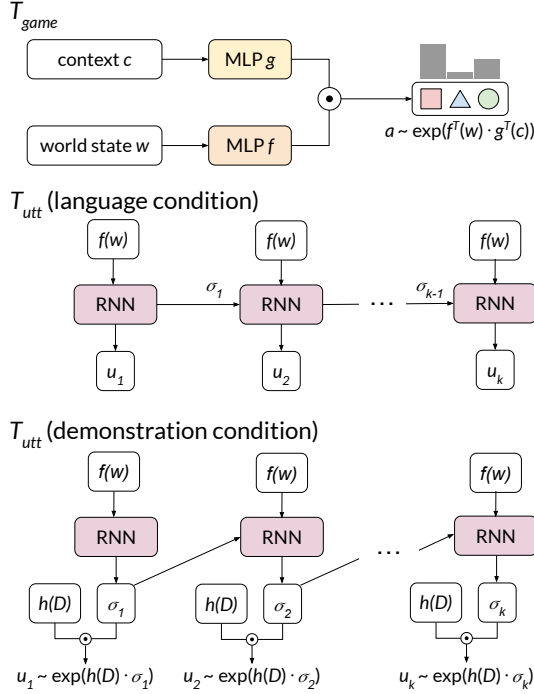


Figure 2: Overview of teacher architecture, including  $T_{game}$  (game-playing component) and  $T_{utt}$  (utterance-generating component, for both language and demonstrations).

Here,  $J_T$  is the expected reward of the game-playing part of the teacher;  $J_{S,T}$  is the expected reward of the student acting on utterances sent from the teacher; and  $J(S, T)$  is the combined objective.

To maintain backwards differentiability, we optimize  $J(S, T)$  by differentiating through messages sampled from the teacher via the straight-through Gumbel-Softmax trick (Jang et al., 2017; Maddison et al., 2017).

### Message Conditions

The differences between agents in our experiments are determined by the space of utterances for communication  $\mathcal{U}$ , which we describe below.

#### Language

In the **language** condition, agents learn a generic discrete communication protocol, as in the emergent communication literature (Lazaridou & Baroni, 2020). Here, an utterance  $u$  is a sequence of up to  $K$  discrete tokens:  $u = (u_1, \dots, u_k)$ ,  $u_i \in \{1, \dots, V\}$ , where  $k \leq K$  and  $V$  is the size of the discrete vocabulary. In this condition, the RNNs used by the teacher and student are standard GRUs as commonly used in natural language processing, with word embedding matrices corresponding to (arbitrary) language tokens, and a language modeling head that projects from the GRU’s hidden state at time  $i$  to a distribution over tokens for time  $i + 1$ .

#### Demonstrations

In the **demonstration** condition, we instead ask the teacher to (1) use  $T_{game}$  to generate a dataset of demonstrations of optimal choices for the given world state, and (2) send a utterance that takes the form of 1 or more demonstrations.

Specifically, for each world state  $w$ , we construct a dataset of demonstrations  $D_w = \{(c^i, \hat{a}^i)\}_{i=1}^{|\mathcal{C}|}$ . For all possible contexts  $c^i \in \mathcal{C}$ , we retrieve the teacher’s prediction for this context:  $\hat{a}^i = \arg \max_a T_{game}(a | c^i, w)$ . Then, the teacher sends an utterance to the student where each utterance “token” is actually a demonstration:  $u = (d_1, \dots, d_k)$ , where  $d_i \in D_w$ .

This means that for demonstration agents, at each time  $i$  the teacher RNN selects from a set of demonstrations with the following equation:

$$p_{RNN}(d_i | d_{<i}; f(w)) \propto \exp(h(d_i)^\top o_i) \quad (8)$$

i.e. a softmax over possible demonstrations in  $D_w$ , where  $o_i$  is the hidden state of the RNN at time  $i$  and  $h^T(d_i)$  is a learned demonstration encoding (again, a small MLP). Similarly, the student’s RNN encoder does not use word embeddings to encode signal tokens  $s_i$ , but rather, a separately learned demonstration MLP  $h^S(d_i)$ . The demonstration encoder  $h^S(d_i)$  reuses the context encoder  $g^S(c_a)$  that is used for the agents’ game choice (Equation 4).

### Experiment Details

#### Data

Agents perform the signaling bandits task given a world state with 2 features: color and shape (Figure 1). Both of these two feature can take on  $n$  possible values. Each of those  $n$  values has an associated reward: the rewards are allocated such that maximum for the a color value is 6 and the minimum for a color value is -6, while the maximum for a shape value is 3 and the minimum is -3, with intermediate values associated with intermediate rewards. For example, a world state with  $n = 3$  values for each feature (as shown in Figure 1A) means that each one of the 3 colors is associated with one of the values in  $\{-6, 0, 6\}$ , and each one of the 3 shapes is associated with one of the values in  $\{-3, 0, 3\}$ . For all games, the referring context  $c$  over which teachers and students disambiguate consists of 3 objects. Thus for a task setting with  $n$  unique values for each of the color and shape features, there are  $n! \cdot n!$  unique world states and  $\binom{n^2}{3}$  unique possible referring contexts that can be passed to a student as part of a demonstration.

During training time, teacher agents see world states sampled from 80% of the possible world states, and are evaluated on world states sampled from the remaining 20%.

#### Model and training details

For the teacher model (Figure 2), the game-playing portion of the agent is implemented as follows. As aforementioned,  $f^T$  and  $g^T$  are MLPs, both of which use ReLU activation functions and produce outputs of 64-d. The utterance-generating portion of the agent differs based on the type of utterance produced.

For the language-based teacher, the RNN that defines  $p_{RNN}$  produces a 100-d hidden state which is then projected down via a linear layer, over which we sample a vocabulary token using Gumbel-Softmax with  $\tau = 1$ . For the demonstration-based teacher,  $h$  is a multi-layer perceptron that produces a 64-d representation of all possible demonstrations, while the RNN produces a 64-d hidden state representation; we sample a demonstration by taking the dot product of the demonstration representation and the hidden state via Gumbel-Softmax with  $\tau = 1$ .

The student model is implemented similarly to the game-playing portion of the teacher:  $f^S$  is an MLP with output size 64-d, and  $g^S$  is also an MLP with output size 64-d.

For all simulations, each teacher-student pair is trained with a batch size of 32 games, with 4000 updates for a total of  $32 \times 4000 = 128000$  games. The parameters of the teacher and student are jointly updated after each training step, using the Adam optimizer (Kingma & Ba, 2015). For the language experiments, we ran a small search and found empirically that a vocabulary of size  $V = 80$  yielded good results, so all language experiments use that value.

We measure communicative success by computing the normalized reward of the student agent. A reward of 0 indicates random chance performance while a reward of 1 indicates optimal performance (i.e. the reward achieved if the student always selected the highest utility object).<sup>1</sup>

## Results

We first present results showing that the language and demonstration agents learn successful communication protocols, in line with expected qualitative patterns. Then, we investigate the tradeoffs between language and demonstration as task difficulty increases and agent competence decreases.

### Experiment 1: Increasing capacity improves agent performance across modalities

In our first simulation, we vary the size of the communication channel for language and demonstration agents. For the language agents, this is operationalized through changing the length of the message. For the demonstration agents, this is operationalized through changing the number of demonstrations produced by the teacher and shown to the student.

As a baseline we use **random demonstration** agents, for which the teacher generates demonstration data, but the demonstrations are randomly selected, so it may not be ideal at conveying key features of the world state. For clarity, we refer to the demonstration agent that selects examples as formulated above as the **pedagogical demonstration** agent.

Figure 3 shows that as the channel capacity increases for the language and the 2 demonstration modalities, student performance improves. The size of the world state is fixed at  $n = 4$ . For the language student, increasing the length of message




Figure 3: Normalized reward of the student as a function of channel capacity (number of messages for language agents and number of *pedagogical* or *random* demonstrations for demonstration agents) on the validation set. Each point represents the mean of 5 runs. Error bars show standard error of the mean.

leads to gains until message length 10, at which the performance levels off. For the pedagogical demonstration student, performance saturates around  $k = 2$  demos. For the random demonstration student, performance improves monotonically with increased capacity.

We also observe that the pedagogical demonstration agent outperforms the random demonstration agent. This is consistent with observed behavior in humans, namely that teachers reason about students’ knowledge to select examples, rather than choosing at random (Ho et al., 2016; Shafto et al., 2014).

Finally, Figure 2 shows that the language agent and the pedagogical demonstration agent achieve similar performance with low channel capacity, but as capacity increases, the language agent outperforms the demo agent. This difference in performance cannot be explained by the number of possible signals that can be sent by the language agent vs. the demonstration agent: with a message length of 5, the agent can produce  $\sum_{i=1}^{k=5} 80^i = O(10^9)$  unique messages, while a demo agent with  $k = 5$  can send  $\binom{560}{5} = O(10^{11})$  unique demo combinations,<sup>2</sup> yet that language agent outperforms the demo agent.

These results confirm that the agents exhibit the same qualitative patterns that we would expect in human participants, namely that increasing the size of the communication channel and selecting teaching examples conditioned on a student both lead to improved performance. Having established this, we can now examine how manipulating task difficulty affects the effectiveness of different teaching modalities.

### Experiment 2: Effectiveness of teaching by language and demonstration is modulated by task difficulty

Next, we show how changing task difficulty, by making the task inherently more challenging, affects performance of the language and pedagogical demonstration agents. To operationalize task difficulty, we vary the number of possible values that each feature (color or shape) can take on

<sup>1</sup>Code for experiments and visualizations is available at <https://github.com/dharakyu/language-and-demos>.

<sup>2</sup>This is a conservative lower bound on the number of demos, assuming that agents do not send duplicate demos and that the order in which they are selected doesn’t matter.

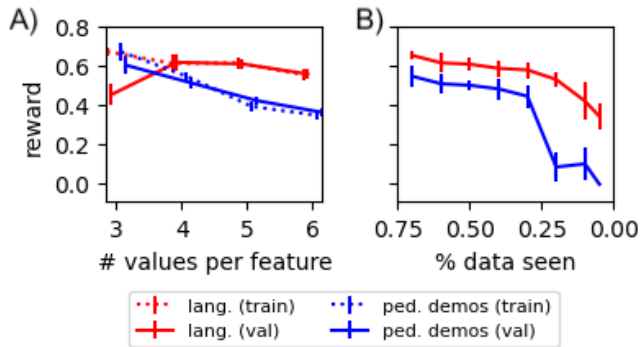


Figure 4: A) Normalized reward on train (dotted line) and validation (solid line) sets as a function of task difficulty.  $n$  possible values per feature means that there are  $n^2$  unique objects,  $n! \cdot n!$  unique reward assignments to those objects, and  $\binom{n^2}{3}$  unique referring contexts. B) Normalized reward on validation set as a function of the amount of training data. The y-axis is the percent of possible world states seen during training, ranging from 70% to 5%. Agents are evaluated on world states sampled from the remaining data.

as defined by the world state, from  $n = 3$  to  $n = 6$ . For example, a world state with  $n = 4$  has 4 unique color values (red, blue, green, purple) and 4 unique shape values (circle, square, triangle, pentagon) which each have an associated reward. Intuitively, increasing the number of possible values increases difficulty because the teacher must learn to convey more granular information to maintain the same level of performance. We otherwise keep the size of the communication channel constant: the language-based teacher has a fixed message length of 10, while the pedagogical demonstration-based teacher can send 2 demonstrations.<sup>3</sup>

Figure 4A shows that demonstration enables more effective knowledge transmission in the simplest setting ( $n = 3$ ), but language leads to better performance as task difficulty increases. To show agents’ generalization ability, we show the reward earned on both the train and validation sets. In the simplest setting, with a world state containing 3 colors and 3 shapes, the language and pedagogical demonstration agents achieve similar performance on the train set, but the demonstration agent outperforms the language agent on the unseen validation set. The disparity between train and validation for the language agent suggests that the learned communication protocol overfits to the train set with relatively limited data and thus cannot generalize to unseen world states.

Yet as the number of object property values increases, the learning dynamics change. Notably, the language agent no

<sup>3</sup>These values were chosen based on the saturation points for performance from Experiment 1. We confirmed that the key qualitative patterns hold (namely, that language performance stays relatively flat while demonstration performance declines) even as channel capacity changes by running additional experiments with varying language message/demonstration lengths. The code to run these experiments is available in the publicly-released repository.

longer exhibits overfitting: train and validation performance is similar. Additionally, the language student achieves similar performance in the  $n = 4, n = 5$  conditions, with a small dropoff for  $n = 6$ . The demonstration student, on the other hand, performs worse as difficulty scales. The reason for this may be that as the number of possible demonstrations becomes combinatorially large as the task is scaled up, whereas the quantity of abstractions that might be conveyed (e.g. “red is good”) increases at a slower rate.

These results are consistent with empirical studies that show teaching by language is more effective for when the task being taught to the student is difficult (Sumers et al., 2023). In addition to matching empirical observations, our simulations make the novel qualitative prediction that teaching by demonstration may be preferred over language for highly simple tasks, in the absence of established linguistic conventions.

### Experiment 3: Language enables generalization even as teacher competence declines

In this section, we explore how the competence of the teacher affects student performance across modalities. Teaching is often studied in settings assuming a fully knowledgeable teacher, yet in many naturalistic contexts, teachers do not have oracle knowledge but rather partial, noisy information (Vélez & Gweon, 2019). In our simulations, we operationalize competence by varying the *amount* of data shown to the teacher during training. This is theoretically motivated by the cultural evolution literature, and in particular, the fact that humans see limited data over the course of a lifespan but nonetheless can aggregate this information over generations to develop increasingly complex cultural technologies (Tomasello, 1999).

In previous experiments, we fixed the train set to consist of examples sampled from 80% of all possible world states, while the validation set came from sampling from the remaining 20%. Here we vary the percentage of world states from which the train set is sampled, to emulate how human teachers have limited access to the structure of a task. Again, we fix the message length for the language agent at 10 and the number of demos for the pedagogical demonstration agent at 2.

As seen in Figure 4B, as the pedagogical demonstration-based teacher sees increasingly limited training data, student performance declines, with a substantial drop occurring at around 0.2. In the lowest-data setting, the student fails to learn at all, exhibiting random chance performance. On the other hand, while the language student’s performance declines, it still is able to achieve well above chance performance with highly limited data. These results suggest that language more effectively enables communication about the underlying abstractions that characterize a world state.

### Analysis of agent strategies: Language conveys generalizations

What properties of language enable information transfer even with limited training data? We analyze the messages and demonstrations sent by low-competence teachers that see just 5% of all possible world states during training, for one run




Figure 5: Number of unique utterances (messages or demonstrations) produced across the language and demonstration conditions, on the train and validation sets. Novel utterances are utterances that were produced only on the validation set.

	color feature				shape feature			
	red	blue	green	purple	circle	square	triangle	pentagon
overall	<b>0.49</b>	-0.11	<b>0.03</b>	-0.40	-0.14	<b>0.41</b>	<b>-0.07</b>	-0.20
message 1	<b>0.14</b>	<b>0.20</b>	0.03	0.03	0.03	<b>0.61</b>	-0.23	-0.26
message 2	-0.13	-0.13	<b>0.50</b>	-0.25	-0.19	<b>0.50</b>	-1.13	<b>0.81</b>

Figure 6: Mean reward of each possible value for each of the color and shape features associated with random message 1, random message 2, and all messages. Within each feature, the two highest rewards are indicated in bold.

in each condition. Figure 5 shows the **number of unique utterances**, which are messages for the language agents and demonstrations for the demonstration agents, seen on the train and validation sets, and the **number of unique novel utterances** (pink portion of validation set bars) produced by the agent on the validation set that were not produced during training. Teachers of both modalities produce far fewer unique signals than unique world states; language-based teachers produce just 2 unique messages in total and no novel ones during test time, while demonstration-based teachers produce substantially more, including several novel ones at test time.

The relatively strong performance of the language agents with so few unique messages (far fewer than the number of world states) necessarily means that the messages themselves contain *abstractions* about the given world state, as opposed to one-to-one mappings between messages and world states. Concretely, the semantics of a message may be closer to “red and blue are good, green and purple are bad” as opposed to something like “red has value 6, blue has value 2, green has value -2, purple has value -6”. To test this hypothesis, we selected two messages sent by the language-based teacher that is trained on 80% of the data. Figure 6 shows the mean reward associated with each of the values for the color (red, blue, etc.) and shape (circle, square, etc.) features, for message 1, message 2, and the full set of messages, produced on the previously-unseen validation set. For example, the value of 0.14 for red for message 1 indicates that for all reward assignments associated with message 1, the mean value of red was

0.14. We see that while the overall set of validation games is biased toward red, green, square and triangle having relatively high values, message 1 is associated with high rewards for red, blue, square and circle, and message 2 is associated with high rewards for green, square and pentagon. This provides evidence that the teacher is conveying information about the *relative* rewards associated with features, rather than exact quantities, which in of itself is a form of abstraction. This result aligns with findings from behavioral studies that humans use language to communicate abstractions in a teaching setting (Chopra et al., 2019; Summers et al., 2023).

## Discussion

We studied the tradeoffs between language- and demonstration-mediated teaching using neural network models of multi-agent interaction. Our results suggest that in the absence of strong communicative conventions, demonstration is more effective for knowledge transmission when the task is simple, while language confers an advantage as task difficulty scales. Furthermore, language maintains the ability to convey generalized knowledge about a world state even when agents have limited access to novel data. A qualitative analysis reveals that language-based agents may accomplish this through communicating higher-level abstractions about the specific world state, such as the relative utilities associated with each feature value.

In addition to providing a formal account that supports findings from empirical studies (Sumers et al., 2023), our models also make novel predictions about the communicative success of teaching with language vs. demonstrations in humans. In particular, the modeling predicts that demonstration is more effective and thus may be preferred as a teaching modality for simple tasks. It also predicts that language maintains some degree of effectiveness for teaching on unseen world states even as agents see limited training data, while demonstrations cannot enable generalization to new examples.

More broadly, these simulations show that under rational pressures, another class of learning agents—neural networks—which are likely optimized in a very different manner compared to humans and in fact have been shown to develop communication systems distinctly unlike human language (Chaabouni et al., 2020; Chaabouni et al., 2019), nonetheless exhibit similar tradeoffs in teaching by language vs. demonstration. The existence of these similarities is a form of complementary evidence for empirical results in human studies and bolsters our confidence in using these models as a useful source of predictions. In particular, these results lend credence to the idea that language may have *become* ubiquitous as a teaching modality because of the ease in which it can emerge to convey complex information (Laland, 2017).

In conclusion, this work provides a formal framework and a methodological toolkit for studying the emergence of teaching strategies. Iteratively refining these models and testing their predictions against human behavioral data is an exciting avenue for future work.

## Acknowledgements

We thank Ted Sumers and the anonymous reviewers for useful feedback. JM is supported by the Open Philanthropy AI Fellowship.

## References

- Chaabouni, R., Kharitonov, E., Bouchacourt, D., Dupoux, E., & Baroni, M. (2020). Compositionality and generalization in emergent languages. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4427–4442.
- Chaabouni, R., Kharitonov, E., Dupoux, E., & Baroni, M. (2019). Anti-efficient encoding in emergent communication.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Chopra, S., Tessler, M. H., & Goodman, N. D. (2019). The first crank of the cultural ratchet: Learning and transmitting concepts through language. *Proceedings of the Cognitive Science Society*.
- Clark, H. H. (1996). *Using language*. Cambridge university press.
- Ellis, S., & Rogoff, B. (1982). The strategies and efficacy of child versus adult teachers. *Child development*, 730–735.
- Hawkins, R. D., Frank, M. C., & Goodman, N. D. (2020). Characterizing the dynamics of learning in repeated reference games. *Cognitive Science*, 44(6).
- Ho, M. K., Littman, M., MacGlashan, J., Cushman, F., & Austerweil, J. L. (2016). Showing versus doing: Teaching by demonstration. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems*. Curran Associates, Inc.
- Jang, E., Gu, S. S., & Poole, B. (2017). Categorical reparameterization with gumbel-softmax. *5th International Conference on Learning Representations (ICLR)*.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Y. Bengio & Y. LeCun (Eds.), *3rd international conference on learning representations, ICLR 2015, san diego, ca, usa, may 7-9, 2015, conference track proceedings*.
- Laland, K. N. (2017). The origins of language in teaching. *Psychonomic Bulletin & Review*, 24, 225–231.
- Lazaridou, A., & Baroni, M. (2020). Emergent multi-agent communication in the deep learning era.
- Lewis, D. K. (1969). *Convention: A philosophical study*. Cambridge, MA, USA: Wiley-Blackwell.
- Maddison, C. J., Mnih, A., & Teh, Y. W. (2017). The concrete distribution: A continuous relaxation of discrete random variables. *5th International Conference on Learning Representations, ICLR 2017*.
- Mu, J., & Goodman, N. (2021). Emergent communication of generalizations. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55–89.
- Strauss, S., Ziv, M., & Stein, A. (2002). Teaching as a natural cognition and its relations to preschoolers’ developing theory of mind. *Cognitive development*, 17(3-4), 1473–1487.
- Sumers, T. R., Hawkins, R. X. D., Ho, M. K., & Griffiths, T. L. (2021). Extending rational models of communication from beliefs to actions. *Proceedings of the Cognitive Science Society*.
- Sumers, T. R., Ho, M. K., Hawkins, R. D., & Griffiths, T. L. (2023). Show or tell? exploring when (and why) teaching with language outperforms demonstration. *Cognition*, 232, 105326.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*.
- Tessler, M. H., Bridgers, S., & Tenenbaum, J. B. (2020). How many observations is one generic worth? *Proceedings of the Annual Conference of the Cognitive Science Society*.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Harvard University Press.
- Vélez, N., & Gweon, H. (2019). Integrating incomplete information with imperfect advice. *Topics in Cognitive Science*, 11(2), 299–315.