

# Improving Long-Horizon Imitation through Instruction Prediction

Joey Hejna<sup>1</sup>, Pieter Abbeel<sup>2</sup>, Lerrel Pinto<sup>3</sup>

<sup>1</sup> Stanford University

<sup>2</sup> University of California, Berkeley

<sup>3</sup> New York University

jhejna@cs.stanford.edu, pabbeel@berkeley.edu, lerrel@cs.nyu.edu

## Abstract

Complex, long-horizon planning and its combinatorial nature pose steep challenges for learning-based agents. Difficulties in such settings are exacerbated in low data regimes where over-fitting stifles generalization and compounding errors hurt accuracy. In this work, we explore the use of an often unused source of auxiliary supervision: language. Inspired by recent advances in transformer-based models, we train agents with an instruction prediction loss that encourages learning temporally extended representations that operate at a high level of abstraction. Concretely, we demonstrate that instruction modeling significantly improves performance in planning environments when training with a limited number of demonstrations on the BabyAI and Crafter benchmarks. In further analysis we find that instruction modeling is most important for tasks that require complex reasoning, while understandably offering smaller gains in environments that require simple plans. More details and code can be found at <https://github.com/jhejna/instruction-prediction>.

## Introduction

Intelligent agents ought to be able to complete complex, long horizon tasks and generalize to new scenarios. Unfortunately, policies learned by modern deep-learning techniques often struggle to acquire either of these abilities. This is particularly true in planning regimes where multiple, complex, steps must be completed correctly in sequence to complete a task. Realistic constraints, such as partial observability, the underspecification of goals, or the sparse reward nature of many planning problems make learning even harder. Reinforcement learning approaches often struggle to effectively learn policies and require billions of environment interactions to produce effective solutions (Wijmans et al. 2019; Parisotto et al. 2020). Imitation learning is an alternative approach based on learning from expert data, but can still require millions of demonstrations to learn effective planners (Chevalier-Boisvert et al. 2019). Such high data constraints make learning difficult and expensive.

Unfortunately the aforementioned issues with behavior learning are only exacerbated in the low data regime. First, with limited training data agents are less likely to act perfectly at each environment step, leading to small errors

that compound overtime in the offline setting. Ultimately, this leads to sub-par performance over long horizons that can usually only be improved by carefully collecting additional expert data (Ross, Gordon, and Bagnell 2011). Second, deep-learning based policies are more likely to overfit small training datasets, making them unable to generalize to new test-time scenarios. On the other hand, humans have the remarkable ability to interpolate previous knowledge and solve unseen long-horizon tasks. After observing an environment, we might deduce plan or sequence of the steps to follow to complete our objective. However, imitation learning agents are not required to construct plans by default – they are usually trained to only output action sequences given seen observations. This begs the question: how can we make agents reason better in long-horizon tasks?

An attractive solution lies in language instructions, the same medium humans use for mental planning (Gleitman and Papafragou 2005). Several prior works directly provide agents with language instructions to follow (Anderson et al. 2018; Shridhar et al. 2020; Chen et al. 2019). Unfortunately, such approaches require the specification of exhaustive instructions at test time for systems to function. A truly intelligent agent ought to be able to devise its own plan and execute it, with only a handful of demonstrations. We propose improving policy learning in the low-data regime by having agents predict planning instructions in addition to their immediate next action. As we do not input instructions to the policy, we can plan without their specification at test time. Though prior works have used hierarchical structures that generate their own instructions to condition on (Chen, Gupta, and Marino 2021; Hu et al. 2019; Jiang et al. 2019), we surprisingly find that just predicting language instructions is in itself a powerful objective to learn good representations for planning. Teaching agents to output language instructions for completing tasks has two concrete benefits. First, it forces them to learn at a higher level of abstraction where generalization is easier. Second, by outputting multi-step instructions agents explicitly consider the future. Practically, we teach agents to output instructions by adding an auxiliary instruction prediction network to transformer-based policy networks, as in seq2seq translation (Vaswani et al. 2017). Our approach can be interpreted as translating observations or trajectories into instructions.

We test our representation learning method in limited data

settings and combinatorially complex environments. We find that in many settings higher performance can be attained by relabeling existing demonstrations with language instructions instead of collecting new ones, creating a new, scalable type of data collection for practitioners. Furthermore, our method is conceptually simple and easy to implement. This work is the first to show that direct representation learning with language can accelerate imitation learning.

To summarize, our contributions are as follows. First, we introduce a method for training transformer based planning networks on paired demonstration and instruction data via an auxiliary instruction prediction loss. Second, we test our objective in long-horizon planning based environments with limited data and find that it substantially outperforms contemporary approaches. Finally, we analyze the scenarios in which predicting instructions provides fruitful training signal, concluding that instruction modeling is a valuable objective when tasks are sufficiently complex.

## Related Work

Language in the context of policy learning has been heavily studied (Luketina et al. 2019), usually to communicate a task objective. Uniquely, we use natural language instructions to aid in learning via an auxiliary objective. Here we survey the most relevant works to our approach.

**Language Goals.** Language offers a natural medium to communicate goals to intelligent agents. As such, several prior work have focused on learning language goal conditioned policies, particularly for robotics (Nair et al. 2021; Stepputtis et al. 2020; Kanu et al. 2020; Hill et al. 2020; Akakzia et al. 2021; Goyal, Mooney, and Niekum 2021; Shridhar, Manuelli, and Fox 2021), or for games (Chevalier-Boisvert et al. 2019; Chaplot et al. 2018; Hermann et al. 2017), and sometimes even with hindsight relabeling (Chan et al. 2018; Cideron et al. 2020). Others in the area of inverse reinforcement learning use language to specify reward functions (Fu et al. 2019; Bahdanau et al. 2018; Williams et al. 2018) or shape them (Mirchandani, Karamcheti, and Sadigh 2021; Goyal, Niekum, and Mooney 2019). These works use language to give humans an easy way to specify the desired goal conditions of an environment. Unlike these works, we use language instructions that dictate the steps to reach a desired goal condition instead of just using language goals that specify the desired state. Other works, particularly in the visual navigation space, provide agents with step-by-step instructions similar to those we use, sometimes in addition to language goals. Anderson et al. (2018); Fried et al. (2018); Chen et al. (2019); Krantz et al. (2020); Chen et al. (2021a, 2019); Zang et al. (2018) condition policies on step-by-step language instructions for visual navigation, while Shridhar et al. (2020); Pashevich, Schmid, and Sun (2021); Shridhar et al. (2021) use instructions for household tasks. Many of these benchmarks ask agents to simply follow instructions, like “turn right at the end of the hallway” instead of achieving overarching goals like “go to the kitchen”. Critically unlike our method, these approaches require laboriously providing the agent with step-by-step instructions at test-time. By using instructions for representation learning instead of policy inputs, we additionally avoid needing to label entire

datasets with language instructions and can train on partially labeled datasets. Other proposed environments (Wang and Narasimhan 2021) assess understanding by prompting agents with necessary information about task dynamics, precluding the removal of text-prompting at test time.

**Language and Hierarchical Learning.** Instead of directly using instructions as policy inputs, other works use language instructions as an intermediary representations for hierarchical policies. Usually, a high-level planner outputs language instructions for a low-level executor to follow. Andreas, Klein, and Levine (2017) and Oh et al. (2017) provide agents with hand-designed high-level language instructions or policy “sketches”. Again unlike our method such approaches require instruction labels for every training task and for every new task at test-time. Jiang et al. (2019) and Shu, Xiong, and Socher (2017) provide interactive language labels to agents to train hierarchical policies with reinforcement learning. In the imitation learning setting, Hu et al. (2019) learn a hierarchical policy using behavior cloning for a strategy game. Unlike the planning problems we consider, their environment has no oracle solution and does not consider generalization to unseen tasks. Most related to our work, Chen, Gupta, and Marino (2021) use latent representations from a learned high-level instruction predictor to aid a low-level policy. However, unlike Chen, Gupta, and Marino (2021), we learn latent representations that can predict instructions, but do not explicitly condition on them at test-time. While these hierarchical approaches have shown promise, the quality of learned policies is inherently limited by the amount of language data available for training. Even with a perfect low-level policy, inaccurate languages commands will yield poor overall performance. For example, a small mis-specification in a subgoal, like changing “blue door” to “red door” would likely cause complete policy failure. This is not an issue for our loss-based approach, as our instruction prediction network can be detached from the policy. As previously mentioned, this structure lets us learn on a mix of instruction annotated and unannotated data, letting it more easily scale than hierarchical approaches particularly in data-limited scenarios. Other approaches in robotics similar to hierarchy use high-level discrete action labels alongside demonstrations to learn planning grammars (Edmonds et al. 2017, 2019). While different in flavor than our approach, such methods also share similar data limitations to the hierarchical methods previously discussed.

**Auxiliary Objectives.** The use of auxiliary objectives in policy learning has been extensively studied. Though to our knowledge no prior have used instructions, auxiliary objectives in general have been found to aid policy learning (Jaderberg et al. 2017). Laskin, Srinivas, and Abbeel (2020) and Stooke et al. (2021) demonstrated the success of contrastive auxiliary objectives in robotic reinforcement learning domains. Schwarzer et al. (2020) and Anand et al. (2019) did the same in the Atari game-playing environments. We were inspired by their effectiveness. Additionally, works like Andreas, Klein, and Levine (2018) have previously used language question and answering for representation learning in visual domains.

**Transformers.** Our approach is based on several innova-

tions involving transformer networks. Vaswani et al. (2017) previously showed state of the art results in machine translation using transformers. While the application of transformers has extended to behavior learning (Zambaldi et al. 2018; Parisotto et al. 2020; Chen et al. 2021b), prior works in the area have not leveraged the transformer decoder. Closest to our domain, Lin et al. (2021) generate captions from video. The architecture of our policy networks take inspiration from recent works adapting transformers to mediums beyond text, namely in vision (Dosovitskiy et al. 2021) and offline reinforcement learning (Chen et al. 2021b).

## Method

In this section we formally describe imitation learning with instruction prediction, then detail our implementation for both Markovian and non-Markovian environments.

### Problem Setup

The standard learning from demonstrations setup assumes access to a dataset of expert trajectory sequences containing paired observations and actions  $o_1, a_1, o_2, a_2, \dots, o_T, a_T$ . The goal of imitation learning is to learn a policy  $\pi(a_t|\cdot)$  that predicts the correct actions an agent should take. In our work we consider both Markovian and partially observed non-Markovian settings. In the non-Markovian observed case, policies are given access to previous observations in order to infer state information (Kaelbling, Littman, and Cassandra 1998), and we denote the policy as  $\pi(a_t|o_1, \dots, o_t)$ . In the Markovian setting this is unnecessary, and the policy is simply  $\pi(a_t|o_t)$ . In imitation learning it is common for policies to be goal conditioned, or even conditioned on language goals as is the case in our experiments. In goal conditioned settings an encoding of the desired task or goal  $g$  is additional input to the policy. As our approach works with or without goal conditioning we omit it from the rest of this section for brevity. A standard imitation learning technique is behavior cloning, which in discrete domains maximizes the likelihood of the actions in the dataset using a negative log likelihood objective,  $\mathcal{L}_{\text{action}} = -\sum_t \log \pi(a_t|\cdot)$ .

In this work, we assume access to oracle language instructions that tell an agent how it should complete a task to provide useful training signal. As mentioned in Section 2, for the purposes of our method we distinguish goals from instructions. Language goals describe the desired final state of the environment environment, specifying *what* to do, whereas language instructions communicate *how* an agent should reach the desired state in a step-by-step manner. Each trajectory may have several language instructions  $x^{(1)}, x^{(2)}, \dots, x^{(n)}$  corresponding to each of the  $n$  different steps to reach the desired goal configuration. For example, a language instruction like “open the door” only applies to the part of the demonstration before the agent opens the door and after it completes the last instruction. The  $i$ -th instruction  $x^{(i)}$  thus corresponds to an interval  $[T_i, T_{i+1})$  where  $T_i$  marks the time the instruction was given and  $T_{i+1}$  denotes the start of the next instruction. A depiction of an example instruction sequence can be found in Figure ???. While language instructions are an additional data require-

ment, they can be cheap to obtain, particularly in scenarios where demonstrations are expensive to collect. If demonstrations are collected in the real world, providing simultaneous instruction annotations for a demonstration is likely less labor-intensive than collecting an additional demonstration. Moreover, humans can easily re-label existing demonstrations with instructions. Video data could easily be captioned with voice-over. Similar statements can be made for simulators – if one can code an oracle policy, instructions are likely easy to generate along the way. These modifications are easy for simulators with planning stacks, as in BabyAI. Moreover, we focus on low to medium data regime, where the cost of setting up an environment and collecting more demonstrations is likely be higher than annotating an existing small set of demonstrations. Next, we describe how we train agents to predict language instructions to aid in imitation learning.

### Instruction Prediction for Imitation Learning

The central hypothesis of this work is that predicting language instructions will force agents to learn representations beneficial for long-horizon planning. In our framework, we first construct an observation encoder  $f_\theta$  that produces latent representations  $z$ . As in behavior cloning we predict actions from latents  $z$  using a policy network  $\pi_\phi$ , but we additionally use a language decoder  $g_\psi$  to predict the current instruction from  $z$ . Our general setup is shown in left half of Figure 1. We consider both non-Markovian environments, where sequences of observations must be provided to the model so it can infer the underlying state. In the non-Markovian setting, the encoder produces  $z_1, \dots, z_t = f_\theta(o_1, \dots, o_t)$ , the policy is  $\pi_\phi(a_t|z_1, \dots, z_t)$ , and the language model is  $g_\psi(x^{(i)}|z_1, \dots, z_{T_i-1})$ . For standard fully observed Markovian environments, conditioning on past observations is unnecessary and the encoder, policy, and language decoder can be written as  $z_t = f(o_t)$ ,  $\pi_\phi(a_t|z_t)$ , and  $g_\psi(x^{(i)}|z_t)$  respectively. As is common in natural language processing, we treat each language instruction  $x^{(i)}$  as a sequence of multiple text tokens  $x_1^{(i)}, x_2^{(i)}, \dots, x_{l_i}^{(i)}$  where  $l_i$  is the length of the  $i$ -th instruction. The decoder is trained using the standard language modeling loss. We construct our total imitation learning objective by maximizing the log-likelihood of both the action and instruction data. For a given trajectory in the non-Markovian case, this is written as follows

$$\mathcal{L} = -\sum_{t=1}^T \log \pi_\phi(a_t|z_1, \dots, z_t) \quad (1)$$

$$- \lambda \sum_{i=1}^n \sum_{j=1}^{l_i} \log g_\psi(x_j^{(i)}|x_1^{(i)}, \dots, x_{j-1}^{(i)}, z_1, \dots, z_{T_i-1})$$

where latent representations  $z$  are all produced by the shared encoder  $f_\theta$ . The MDP case is formulated by removing past conditioning on  $z_1, \dots, z_{t-1}$ . The first term of the loss is the standard classification loss used for behavior cloning in discrete domains. The second term of the loss corresponds to the negative log-likelihood of the language instructions. We index the language loss by instructions via the first sum. The second sum over token log

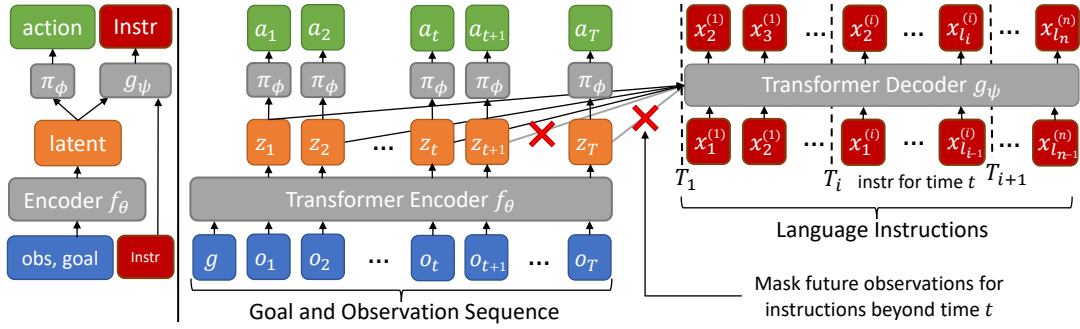


Figure 1: The left diagram depicts the general model architecture used for our approach. Notice how the policy and encoder can be completed separated from the instruction component for mixed-data training or inference. The diagram on the right depicts its implementation for the partial observed environments using a GPT-like transformer encoder. The diagram shows our masking scheme at episode step  $t$ : latent vectors from beyond time  $t$  are masked from the language decoder.

likelihoods is from the standard auto-regressive language modeling framework, where the likelihood of an instruction is the product of the conditional probabilities  $p(x^{(i)}) = \prod_{j=1}^l p(x_j^{(i)} | x_1^{(i)}, \dots, x_{j-1}^{(i)})$ . Note that when predicting the likelihood of language tokens for instruction  $i$ , the model can only condition on latents up to  $z_{T_i-1}$ . This ensures that we compute the likelihood of instruction  $i$  using only observations during or before its execution. Finally,  $\lambda$  is a weighting coefficient that trades off the importance of instruction prediction and action modeling. During training, we optimize all parameters  $\phi$ ,  $\psi$ , and  $\theta$  jointly, meaning that gradients from both behavior cloning and language prediction are propagated to the encoder weights  $\theta$ . In some of our experiments we test additional learning objectives which are also trained on top of the same latent representations  $z$  as is standard in the literature (Jaderberg et al. 2017).

Though our method is general to any network architecture, we train transformer based policies since they have been shown to be extremely effective at natural language processing tasks (Vaswani et al. 2017) and carry a good inductive bias for combinatorial planning problems (Zambaldi et al. 2018). For details on the transformer architectures we use, we defer to (Dosovitskiy et al. 2021; Chen et al. 2021b). In the following sections we describe our transformer-based models for both Markovian and non-Markovian settings.

**Non-Markovian Settings.** For environments that are non-Markovian or partially observed we use a transformer based sequence model as our policy network, similar to those employed in (Chen et al. 2021b). We operate in the entire sequence at once:  $z_1, \dots, z_T = f_\theta(o_1, \dots, o_T)$ . Causal masking similar to that in (?) ensures that at time  $t$  the representation  $z_t$  only depends on current and previous observations  $o_1, \dots, o_t$ . The same policy network  $\pi_\phi(a_t | z_t)$  is applied to each latent to produce actions for each timestep. The language decoder  $g_\psi$  is also a transformer model and employs both causal attention masks to the language inputs and cross attention masks to the latents. Causal-self attention masks on the language inputs enforce the auto-regressive modeling of the instruction tokens. Cross attention masks to the latent representations ensure that predictions for the  $i$ th instruction cannot attend to latents from timesteps after its execution as

is depicted by the red “x”s in Figure 1. This forces language prediction during training to mirror test-time as the agent cannot use future information to predict the instruction.

**Markovian settings.** For environments that are fully Markovian, we use only the most recent observation  $o_t$ . As sequence modeling is unnecessary, we use a transformer to encode individual states, leveraging their success in combinatorial environments (Zambaldi et al. 2018). Specifically, we use a Vision Transformer architecture (Dosovitskiy et al. 2021) that predict actions only for a single timestep. Observations are preprocessed into tokens and prepended with a special CLS token:  $o_t \rightarrow \text{CLS}, o_{t,1}, o_{t,2}, o_{t,3}, \dots$ . As we do not input future observations, the transformer encoder uses full unmasked self attention. At the end of the network we take the latent representation corresponding to the CLS token and use it to predict the action  $\pi_\phi(a_t | z_{t,\text{CLS}})$ . We use all latent tokens  $z_{t,\text{CLS}}, z_{t,1}, z_{t,2}, z_{t,3}, \dots$  to predict the current language instruction with  $g_\psi$ . An architecture figure can be found in Appendix C.

## Experiments

In this section we detail our experimental setup and empirical results. In particular, we investigate the benefits of instruction modeling for planning in limited data regimes. We seek to answer the following questions: How effective is instruction modeling loss? How does instruction modeling scale with both data and instruction annotations? What architecture choices are important? And finally, when is instruction modeling a fruitful objective?

### Environments

To evaluate the effectiveness of instruction prediction at enabling long-horizon planning and generalization, we test our method on BabyAI (Chevalier-Boisvert et al. 2019) and the Crafting Environment from Chen, Gupta, and Marino (2021) which both provide coarse instructions. They cover challenges in partial observability, human generated text, and more. We later examine the ALFRED environment to understand where instruction prediction is useful. Full model hyperparameters can be found in the Appendix.

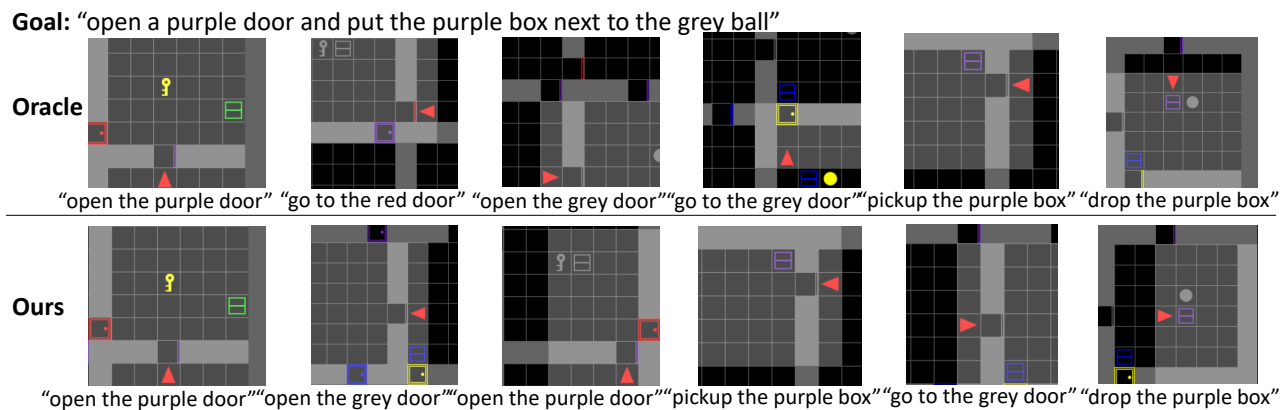


Figure 2: Snapshots of a rollouts from an oracle agent and our trained agents on the same unseen task in BabyAI. Our agent is able to predict instructions, given below each image, with high fidelity. Our learned agent employs a different strategy but still completes the task exhibiting strong generalization.

**BabyAI:** Agents must navigate partially observable grid-worlds to complete arbitrarily complex goals specified through procedurally generated language such as moving objects, opening locked doors, and more. Agents are evaluated on their ability to complete unseen missions in unseen environment configurations. We modify the BabyAI oracle agent to output language instructions based on its planning logic. We focus our experiments on the hardest environment, BossLevel, and up to 10% of the million demos in Chevalier-Boisvert et al. (2019). Because of partial observability, we employ a transformer sequence model as described in Section with the same encoder from Chevalier-Boisvert et al. (2019). The language goals from the environment are tokenized and fed as additional inputs to the policies. We evaluate on five hundred unseen tasks.

**Crafting:** This environment from Chen, Gupta, and Marino (2021) tests how well an agent can generalize to new tasks using instructions collected from humans. The original dataset contains around 5.5k trajectories with human instruction labels. Each task is specified by a specific goal item the agent should craft, encoded via language. The agent must complete from one to five independent steps to obtain the final item. As this environment is fully observed, we employ the Vision Transformer based model described in Section with the benchmark’s original state encoder.

## Baselines

We compare the effectiveness of our instruction modeling auxiliary loss to a number of baselines. The text in parenthesis indicates how we refer to the method in Tables 1, 2, 4, and 5. None of our models are pretrained, though we explore this and more additional baselines in the Appendix.

1. **Original Architecture (Orig):** The original state of the art model architectures proposed for each environment. The crafting environment uses a language-instruction hierarchy. In BabyAI, we use convolutions and FiLM layers as in Chevalier-Boisvert et al. (2019).
2. **Transformer (Xformer):** Our transformer based models

without any auxiliary objectives to determine the effectiveness of our architectures.

3. **Transformer Hierarchy (Hierarchy):** A high-level model outputs instructions for a low level executor for comparison to to hierarchical approaches.
4. **Transformer with Forward Prediction (Forward):** Instead of predicting instructions, we use the decoder to predict future actions. This baseline demonstrates the importance of using grounded information.
5. **Transformer with ATC (ATC):** Our transformer model with the active temporal contrast (ATC) self-supervised objective proposed in Stooke et al. (2021). This compares vision and instruction based representation learning.
6. **Transformer with Lang (Lang):** Our transformer based models with just instruction prediction loss.
7. **Transformer with ATC and Lang (Lang + ATC):** Our transformer based models with both instruction modeling and constrastive auxiliary losses.

## How Effective Is Instruction Prediction?

Our main experimental results can be found in Table 1, where we compare the performance of all methods on both environments with three differing dataset sizes. We find that for all environments and dataset sizes our instruction modeling objective improves or has no effect in the worst case. In BabyAI, we achieve a 70% success rate on the hardest level with fifty thousand demonstrations and instructions. For comparison, it is worth noting that the original BabyAI implementation (Chevalier-Boisvert et al. 2019) achieved a success rate of 77% with one million demonstrations. In the crafting environment, using instruction modeling boosts the success rate by about 5% or more in the 1.1k and 2.2k demonstration setting. To our knowledge our results are state of art in this environment, exceeding the reported 69% success rate on unseen tasks in Chen, Gupta, and Marino (2021) where RL is additionally used. We also find that the model is able to accurately predict instructions (Figure 2), which ap-

Env	Demos	Orig	Xformer	Hierarchy	ATC	Forward	Lang	ATC+Lang
BabyAI BossLevel	100k	38.8	48.4	41.2	62.4	43.5	<b>78.6</b>	73.6
	50k	35.3±0.1	40.2±2.2	36.8±3.5	45.8±0.6	37.0±0.6	<b>70.3±1.3</b>	64.3±0.5
	25k	32.3±2.4	39.9±0.5	37.2±3.0	37.1±1.1	38.9±0.7	<b>55.4±7.0</b>	<b>56.0±3.0</b>
	12.5k	29.9±0.9	37.3±0.1	36.4±2.6	38.4±1.4	36.0±0.4	<b>39.4±1.0</b>	38.6±0.6
Crafting	5k	9.4±1.1	75.8±3.6	63.2±10	78.8±3.1	77.9±4.3	75.9±2.4	<b>80.3±2.5</b>
	3.3k	9.3±0.4	74.5±3.3	59.9±11	<b>75.7±1.0</b>	74.5±4.9	74.5±2.8	<b>76.0±2.8</b>
	2.2k	4.9±1.0	69.4±4.9	56.5±9.9	73.9±2.1	73.6±3.2	75.2±4.4	<b>78.2±4.6</b>
	1.1k	1.7±0.8	70.1±3.8	39.4±3.8	70.1±3.7	73.0±3.8	<b>74.8±2.6</b>	71.4±2.9
ALFRED	42k	–	28.3 ± 1.0	–	–	–	28.5±1.0	–

Table 1: Success rates (in %) of all methods for varying numbers of demonstrations. The best method(s) is bolded, and the included range denotes the standard deviation (2 seeds for BabyAI and ALFRED, 4 for Crafting).




Figure 3: Data scaling with and without instructions.

pears to be correlated with performance. See the Appendix for analysis of the language outputs of the models.

Visual representation learning was not as fruitful as language based representation learning overall. The combination of ATC and instruction modeling was unfortunately not constructive in all scenarios: it performed better in some instances and worse than just language loss in others. This is consistent with results found in Chen et al. (2021c) that show that observation based auxiliary objectives often yield mixed results in the imitation learning setting. We find that our hierarchical implementations do not perform very well in comparison to plain transformer models. This is likely because with only a few demonstrations high-level language policies are likely to output incorrect instructions for unseen tasks leading low-level instruction conditioned policies to output sub-optimal actions. More analysis of the hierarchical baselines is in the Appendix. Finally, we see that the forward prediction used in the forward baseline hardly contributes to performance, indicating that grounded instructions do more than just combat compounding errors.

### How Does Instruction Prediction Scale with Data and Annotations?

Overall, we find that instruction modeling reduces the amount of data required for policies to begin to generalize well. This is particularly evident in the low to medium data regime. With too little data, agents are likely to overfit quickly and only see a minor benefit from instruction prediction. With a significant amount of data instruction modeling may become unnecessary, and the policy can learn good rep-

% w/ Instr	0%	50%	100%
50k Demos	40.2±2.2	68.6±1.4	70.3±1.3
25k Demos	39.9±0.5	50.3±1.3	55.4±7.0

Table 2: We ablate the amount of demonstrations annotated with language instructions. Values are % success rates with standard deviations.

resentations from action labels alone. However, in between these regimes we find that instruction prediction reduces the amount of data needed to generalize by forcing the model to learn more amenable representations to long-horizon planning. Figure 3 depicts how model performance changes with dataset size. In BabyAI, instruction modeling does not appear to significantly help with the smallest number of demos, however, after twelve and a half thousand demonstrations that we find that policy performance with language scales almost linearly with data before it experiences diminishing returns at two-hundred thousand demonstrations. Policies without language are unable to perform substantially better until we provide one hundred thousands demonstrations. This is not just because training with instructions helps overcome partial observability – we show similar results on a fully observed version of BabyAI in the Appendix. The Crafting environment has only fourteen training tasks versus BabyAI’s potentially infinite number, causing it to require fewer demos before performance saturates. Thus, we observe the opposite problem: instruction modeling helps when the policy is data constrained, and then is neutral when more data is introduced. In Section 4.6, we show that this saturation happens rather quickly for 5-step Crafting tasks, as there is only one in the training dataset.

A benefit of our loss-based approach is that it can easily be applied to mixed datasets that have only some instruction labels. To additionally study the scaling properties of our language prediction objective, we construct datasets in BabyAI where only half of the trajectories have paired instructions. Results can be found in Table 2. Surprisingly, *one is better off collecting 12.5k language annotations than collecting an additional 25k demonstrations* in the BabyAI environments. A similar statement can be made in the crafting environment for 1.1k demonstrations. This means that collecting instruction annotations is a feasible alternative to demonstrations.

Env	Demos	Goal	Goal + Obs
BabyAI	50K	86.4%	92.4%
Crafting	3.3K	53.8%	65.1 %
ALFRED	42K	96.9%	99.0%

Table 3: Instruction prediction accuracies for models trained with and without access to observations. When instructions cannot be predicted from goals alone, better observation representations will be learned to help predict instructions.

Level	50k Demos		25k Demos	
	Xformer	Lang	Xformer	Lang
GoTo	88.6±1.8	<b>91.0±1.4</b>	77.6±2.2	<b>82.3±3.3</b>
SynthLoc	72.5±2.3	<b>86.2±1.2</b>	60.1±0.5	<b>69.4±1.6</b>
BossLevel	40.2±2.2	<b>70.3±1.3</b>	39.9±0.5	<b>55.4±7.0</b>

Table 4: Performance, in percent, of instruction prediction when varying the BabyAI level difficulty. The included range is the standard deviation.

### What Modeling Decisions Are Important?

We ablate the use of our instruction decoder cross-attention masking in BabyAI. We find that the omission of the masking scheme leads to a 20% drop in performance, from  $70.3 \pm 1.3\%$  to  $50.1 \pm 12.1\%$ . Without masking the language decoder has an easier time predicting an instruction as it can attend to observations from after the instruction finished, creating a disparity between train and test time, ultimately leading to lower quality representations. Overall, the transformer architecture appears to be critical to high performance, likely because of its good inductive bias for reasoning about objects and their interactions. This is especially evident in the Crafting environment. As stated in Chen, Gupta, and Marino (2021), the imitation learning approaches with the original model were unable to achieve a meaningful success rate on any of the unseen tasks, whereas our baseline transformer achieves a success rate of around 70%. Our architecture choice is also extremely parameter efficient as shown in the Appendix.

Demos	Model	2 Steps	3 Steps	5 Steps
3.35k	Xformer	<b>98.1±0.4</b>	66.9±6.2	<b>22.1±3.1</b>
	Lang	96.1±3.2	<b>73.0±7.2</b>	<b>19.3±4.4</b>
	Lang+ATC	97.1±2.2	<b>75.1±10.0</b>	13.2±1.7
2.2k	Xformer	89.3±8.6	58.4±6.2	<b>20.5±1.8</b>
	Lang	<b>96.1±0.7</b>	<b>73.5±10.7</b>	17.3±1.7
	Lang+ATC	93.9±2.7	<b>78.3±13.9</b>	<b>19.8±2.7</b>
1.1k	Xformer	90.9±4.0	58.2±8.5	15.0±5.1
	Lang	<b>94.5±1.4</b>	<b>76.0±8.4</b>	13.8±6.0
	Lang+ATC	89.5±3.2	65.2±10.7	11.6±2.8

Table 5: Difficulty comparison in the crafting environment. Steps indicate the number of steps required for the agent to craft the item. Performance is given in percent success rate with standard deviations.

### When Is Instruction Prediction Useful?

We hypothesize that instruction prediction is particularly useful for combinatorially complex, long horizon tasks. Many simple tasks, like “open the door” or “grab a cup and put it in the coffee maker” communicate all required steps and consequently stand to gain little from instruction modeling. Conversely, tasks in both environments we study do not communicate all required steps to agents. Thus, as task horizon and difficulty increase one would expect instruction modeling to be more important. In BabyAI we consider two additional levels – GoTo, which only requires object localization, and SynthLoc which uses a subset of the BossLevel goals. Results in Table 4 indicate that instruction modeling is indeed more important for harder tasks. The same trend approximately holds in the Crafting environment (Table 5). All policies are able to complete two steps tasks near or above 90% success, but models with language prediction perform around 5% better with fewer than 3.3k demonstrations. The difference in performance is ever greater for the three-step tasks, where instruction prediction boosts performance from around 58% to closer to 75% in most cases. Evaluations of the five-step tasks were noisy, which we attribute to the existence of only one five-step training task with which no model was able to adequately generalize. The takeaway from these observations is that instructions offer less training signal for combinatorially simple tasks, where reaching the goal requires only a few, obvious logical steps. Thus, we expect instruction modeling to not matter when instructions can easily be predicted from goals alone.

To test this hypothesis, we train transformer models to predict instructions with and without access to observations in our primary environments and additionally in a modified version of the ALFRED benchmark (Shridhar et al. 2020). Table 3 shows the token prediction accuracy of instructions using text goals alone versus text goals and observations. While token prediction accuracies are relatively high, particularly for BabyAI, accuracy differences of 5% or more can make a large impact as instructions can share similar structure, but have a few critical tokens specifying objects. In the benchmarks where our method is impactful instructions cannot be as accurately predicted from text goals alone. However, in the ALFRED environment, which is largely based on visual understanding, instructions can easily be predicted from just the goal. This indicates that while the visual complexity of ALFRED may be high, it does not pose significant challenges in logical understanding. Further analysis of the ALFRED benchmark is provided in the Appendix. In the future as tasks become more combinatorially complex, we expect instructions to provide a more critical modeling component.

### Conclusion

We introduce an auxiliary objective that predicts language instructions for imitation learning and associated transformer based architectures. Our instruction modeling objective consistently improves generalization to unseen tasks with few demonstrations, and scales efficiently with instruction labels. We further analyze the domains where our method is successful, and make recommendations for when to apply it.

## References

- Akakzia, A.; Colas, C.; Oudeyer, P.-Y.; Chetouani, M.; and Sigaud, O. 2021. Grounding Language to Autonomously-Acquired Skills via Goal Generation. In *ICLR 2021*.
- Anand, A.; Racah, E.; Ozair, S.; Bengio, Y.; Côté, M.-A.; and Hjelm, R. D. 2019. Unsupervised state representation learning in atari. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 8769–8782.
- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and van den Hengel, A. 2018. Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Andreas, J.; Klein, D.; and Levine, S. 2017. Modular multi-task reinforcement learning with policy sketches. In *International Conference on Machine Learning*, 166–175. PMLR.
- Andreas, J.; Klein, D.; and Levine, S. 2018. Learning with Latent Language. In *NAACL-HLT*.
- Bahdanau, D.; Hill, F.; Leike, J.; Hughes, E.; Hosseini, A.; Kohli, P.; and Grefenstette, E. 2018. Learning to Understand Goal Specifications by Modelling Reward. In *International Conference on Learning Representations*.
- Chan, H.; Wu, Y.; Kiros, J.; Fidler, S.; and Ba, J. 2018. AC-TRCE: Augmenting Experience via Teacher’s Advice For Multi-Goal Reinforcement Learning. *1st Workshop on Goal Specifications for Reinforcement Learning, Workshop held jointly at ICML, IJCAI, AAMAS*.
- Chaplot, D. S.; Sathyendra, K. M.; Pasumarthi, R. K.; Rajagopal, D.; and Salakhutdinov, R. 2018. Gated-attention architectures for task-oriented language grounding. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Chen, H.; Suhr, A.; Misra, D.; Snavely, N.; and Artzi, Y. 2019. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12538–12547.
- Chen, K.; Chen, J. K.; Chuang, J.; Vázquez, M.; and Savarese, S. 2021a. Topological Planning with Transformers for Vision-and-Language Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11276–11286.
- Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021b. Decision transformer: Reinforcement learning via sequence modeling. *arXiv preprint arXiv:2106.01345*.
- Chen, V.; Gupta, A.; and Marino, K. 2021. Ask Your Humans: Using Human Instructions to Improve Generalization in Reinforcement Learning. In *International Conference on Learning Representations*.
- Chen, X.; Toyer, S.; Wild, C.; Emmons, S.; Fischer, I.; Lee, K.-H.; Alex, N.; Wang, S. H.; Luo, P.; Russell, S.; Abbeel, P.; and Shah, R. 2021c. An Empirical Investigation of Representation Learning for Imitation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Chevalier-Boisvert, M.; Bahdanau, D.; Lahlou, S.; Willems, L.; Saharia, C.; Nguyen, T. H.; and Bengio, Y. 2019. BabyAI: First Steps Towards Grounded Language Learning With a Human In the Loop. In *International Conference on Learning Representations*.
- Cideron, G.; Seurin, M.; Strub, F.; and Pietquin, O. 2020. Higher: Improving instruction following with hindsight generation for experience replay. In *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 225–232. IEEE.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Edmonds, M.; Gao, F.; Liu, H.; Xie, X.; Qi, S.; Rothrock, B.; Zhu, Y.; Wu, Y. N.; Lu, H.; and Zhu, S.-C. 2019. A tale of two explanations: Enhancing human trust by explaining robot behavior. *Science Robotics*, 4(37): eaay4663.
- Edmonds, M.; Gao, F.; Xie, X.; Liu, H.; Qi, S.; Zhu, Y.; Rothrock, B.; and Zhu, S.-C. 2017. Feeling the force: Integrating force and pose for fluent discovery through imitation learning to open medicine bottles. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3530–3537. IEEE.
- Fried, D.; Hu, R.; Cirik, V.; Rohrbach, A.; Andreas, J.; Morency, L.-P.; Berg-Kirkpatrick, T.; Saenko, K.; Klein, D.; and Darrell, T. 2018. Speaker-follower models for vision-and-language navigation. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 3318–3329.
- Fu, J.; Korattikara, A.; Levine, S.; and Guadarrama, S. 2019. From language to goals: Inverse reinforcement learning for vision-based instruction following. *arXiv preprint arXiv:1902.07742*.
- Gleitman, L.; and Papafragou, A. 2005. *Language and thought*. Cambridge University Press.
- Goyal, P.; Mooney, R. J.; and Niekum, S. 2021. Zero-shot Task Adaptation using Natural Language. *arXiv preprint arXiv:2106.02972*.
- Goyal, P.; Niekum, S.; and Mooney, R. J. 2019. Using natural language for reward shaping in reinforcement learning. *arXiv preprint arXiv:1903.02020*.
- Hermann, K. M.; Hill, F.; Green, S.; Wang, F.; Faulkner, R.; Soyer, H.; Szepesvari, D.; Czarnecki, W. M.; Jaderberg, M.; Teplyashin, D.; et al. 2017. Grounded language learning in a simulated 3d world. *arXiv preprint arXiv:1706.06551*.
- Hill, F.; Mokra, S.; Wong, N.; and Harley, T. 2020. Human instruction-following with deep reinforcement learning via transfer-learning from text. *arXiv preprint arXiv:2005.09382*.
- Hu, H.; Yarats, D.; Gong, Q.; Tian, Y.; and Lewis, M. 2019. Hierarchical Decision Making by Generating and Following Natural Language Instructions. *Advances in Neural Information Processing Systems*, 32: 10025–10034.



- Jaderberg, M.; Mnih, V.; Czarnecki, W. M.; Schaul, T.; Leibo, J. Z.; Silver, D.; and Kavukcuoglu, K. 2017. Reinforcement learning with unsupervised auxiliary tasks. *International Conference on Learning Representations*.
- Jiang, Y.; Gu, S. S.; Murphy, K. P.; and Finn, C. 2019. Language as an Abstraction for Hierarchical Deep Reinforcement Learning. *Advances in Neural Information Processing Systems*, 32: 9419–9431.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2): 99–134.
- Kanu, J.; Dessalene, E.; Lin, X.; Fermuller, C.; and Aloimonos, Y. 2020. Following instructions by imagining and reaching visual goals. *arXiv preprint arXiv:2001.09373*.
- Krantz, J.; Wijmans, E.; Majumdar, A.; Batra, D.; and Lee, S. 2020. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *European Conference on Computer Vision*, 104–120. Springer.
- Laskin, M.; Srinivas, A.; and Abbeel, P. 2020. Curl: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning*, 5639–5650. PMLR.
- Lin, X.; Bertasius, G.; Wang, J.; Chang, S.-F.; Parikh, D.; and Torresani, L. 2021. VX2TEXT: End-to-End Learning of Video-Based Text Generation From Multimodal Inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7005–7015.
- Luketina, J.; Nardelli, N.; Farquhar, G.; Foerster, J.; Andreas, J.; Grefenstette, E.; Whiteson, S.; and Rocktäschel, T. 2019. A survey of reinforcement learning informed by natural language. *arXiv preprint arXiv:1906.03926*.
- Mirchandani, S.; Karamcheti, S.; and Sadigh, D. 2021. ELLA: Exploration through Learned Language Abstraction. *arXiv preprint arXiv:2103.05825*.
- Nair, S.; Mitchell, E.; Chen, K.; Ichter, B.; Savarese, S.; and Finn, C. 2021. Learning Language-Conditioned Robot Behavior from Offline Data and Crowd-Sourced Annotation. *Conference on Robot Learning (CoRL)*.
- Oh, J.; Singh, S.; Lee, H.; and Kohli, P. 2017. Zero-shot task generalization with multi-task deep reinforcement learning. In *International Conference on Machine Learning*, 2661–2670. PMLR.
- Parisotto, E.; Song, F.; Rae, J.; Pascanu, R.; Gulcehre, C.; Jayakumar, S.; Jaderberg, M.; Kaufman, R. L.; Clark, A.; Noury, S.; et al. 2020. Stabilizing transformers for reinforcement learning. In *International Conference on Machine Learning*, 7487–7498. PMLR.
- Pashevich, A.; Schmid, C.; and Sun, C. 2021. Episodic transformer for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15942–15952.
- Ross, S.; Gordon, G.; and Bagnell, D. 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*, 627–635. JMLR Workshop and Conference Proceedings.
- Schwarzer, M.; Anand, A.; Goel, R.; Hjelm, R. D.; Courville, A.; and Bachman, P. 2020. Data-Efficient Reinforcement Learning with Self-Predictive Representations. In *International Conference on Learning Representations*.
- Shridhar, M.; Manuelli, L.; and Fox, D. 2021. CLIPort: What and Where Pathways for Robotic Manipulation. In *5th Annual Conference on Robot Learning*.
- Shridhar, M.; Thomason, J.; Gordon, D.; Bisk, Y.; Han, W.; Mottaghi, R.; Zettlemoyer, L.; and Fox, D. 2020. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Shridhar, M.; Yuan, X.; Côté, M.-A.; Bisk, Y.; Trischler, A.; and Hausknecht, M. 2021. ALFWorld: Aligning Text and Embodied Environments for Interactive Learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Shu, T.; Xiong, C.; and Socher, R. 2017. Hierarchical and interpretable skill acquisition in multi-task reinforcement learning. *arXiv preprint arXiv:1712.07294*.
- Stepputtis, S.; Campbell, J.; Phielipp, M.; Lee, S.; Baral, C.; and Ben Amor, H. 2020. Language-Conditioned Imitation Learning for Robot Manipulation Tasks. *Advances in Neural Information Processing Systems*, 33.
- Stooke, A.; Lee, K.; Abbeel, P.; and Laskin, M. 2021. Decoupling representation learning from reinforcement learning. In *International Conference on Machine Learning*, 9870–9879. PMLR.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wang, H.; and Narasimhan, K. 2021. Grounding Language to Entities and Dynamics for Generalization in Reinforcement Learning. *arXiv preprint arXiv:2101.07393*.
- Wijmans, E.; Kadian, A.; Morcos, A.; Lee, S.; Essa, I.; Parikh, D.; Savva, M.; and Batra, D. 2019. DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames. *arXiv preprint arXiv:1911.00357*.
- Williams, E. C.; Gopalan, N.; Rhee, M.; and Tellex, S. 2018. Learning to parse natural language to grounded reward functions with weak supervision. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 4430–4436. IEEE.
- Zambaldi, V.; Raposo, D.; Santoro, A.; Bapst, V.; Li, Y.; Babuschkin, I.; Tuyls, K.; Reichert, D.; Lillicrap, T.; Lockhart, E.; et al. 2018. Deep reinforcement learning with relational inductive biases. In *International Conference on Learning Representations*.
- Zang, X.; Pokle, A.; Vázquez, M.; Chen, K.; Niebles, J. C.; Soto, A.; and Savarese, S. 2018. Translating Navigation Instructions in Natural Language to a High-Level Plan for Behavioral Robot Navigation. In *EMNLP*.