

Skill Induction and Planning with Latent Language

Pratyusha Sharma Antonio Torralba Jacob Andreas

Massachusetts Institute of Technology

{pratyuss, torralba, jda}@mit.edu

Abstract

We present a framework for learning hierarchical policies from demonstrations, using sparse *natural language annotations* to guide the discovery of reusable skills for autonomous decision-making. We formulate a generative model of action sequences in which goals generate sequences of high-level subtask descriptions, and these descriptions generate sequences of low-level actions. We describe how to train this model using primarily unannotated demonstrations by *parsing* demonstrations into sequences of named high-level subtasks, using only a small number of seed annotations to ground language in action. In trained models, natural language commands index a combinatorial library of skills; agents can use these skills to *plan* by generating high-level instruction sequences tailored to novel goals. We evaluate this approach in the ALFRED household simulation environment, providing natural language annotations for only 10% of demonstrations. It achieves task completion rates comparable to state-of-the-art models (outperforming several recent methods with access to ground-truth plans during training and evaluation) while providing structured and human-readable high-level plans.¹

1 Introduction

Building autonomous agents that integrate high-level reasoning with low-level perception and control is a long-standing challenge in artificial intelligence (Fikes et al., 1972; Newell, 1973; Sacerdoti, 1973; Brockett, 1993). Fig. 1 shows an example: to accomplish a task such as *cooking an egg*, an agent must first *find the egg*, then *grasp it*, then *locate a stove or microwave*, at each step reasoning about both these subtasks and complex, unstructured sensor data. **Hierarchical planning models** (e.g. Sutton et al., 1999)—which first reason about abstract

¹Code and visualizations: <https://sites.google.com/view/skill-induction-latent-lang/>.

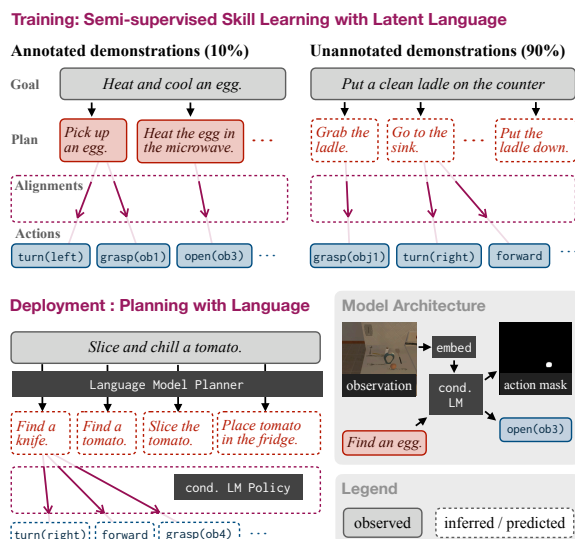


Figure 1: Hierarchical imitation learning using weak natural language supervision. During training, a small number of seed annotations are used to automatically segment and label unannotated training demonstrations with natural language descriptions of their high-level structure. When deployed on new tasks, learned policies first generate sequences of natural language subtask descriptions, then modularly translate each description to a sequence of low-level actions.

states and actions, then ground these in concrete control decisions—play a key role in most existing agent architectures. But training effective hierarchical models for general environments and goals remains difficult. Standard techniques either require detailed formal task specifications, limiting their applicability in complex and hard-to-formalize environments, or are restricted to extremely simple high-level actions, limiting their expressive power (Bacon et al., 2017; Sutton et al., 1999; Dieterich, 1999; Kaelbling and Lozano-Pérez, 2011).

Several recent papers have proposed to overcome these limitations using richer forms of supervision—especially language—as a scaffold for hierarchical policy learning. In **latent language policies** (LLPs; Andreas et al., 2018), controllers first map

from high-level goals to sequences of natural language instructions, then use instruction following models to translate those instructions into actions. But applications of language-based supervision for long-horizon policy learning have remained quite limited in scope. Current LLP training approaches treat language as a latent variable only during prediction, and require fully supervised (and often impractically large) datasets that align goal specifications with instructions and instructions with low-level actions. As a result, all existing work on language-based policy learning has focused on very short time horizons (Andreas et al., 2018), restricted language (Hu et al., 2019; Jacob et al., 2021) or synthetic training data (Shu et al., 2018; Jiang et al., 2019).

In this paper, we show that it is possible to train language-based hierarchical policies that outperform state-of-the-art baselines using only minimal natural language supervision. We introduce a procedure for *weakly* and *partially supervised* training of LLPs using ungrounded text corpora, unlabeled demonstrations, and a small set of annotations linking the two. To do so, we model *training* demonstrations as generated by latent high-level plans: we describe a deep, structured latent variable model in which goals generate subtask descriptions and subtask descriptions generate actions. We show how to learn in this model by performing inference in the infinite, combinatorial space of latent plans while using a comparatively small set of annotated demonstrations to seed the learning process.

Using an extremely reduced version of the ALFRED household robotics dataset (Shridhar et al., 2020)—with 10% of labeled training instructions, no alignments during training, and no instructions at all during evaluation—our approach performs comparably a state-of-the-art model that makes much stronger dataset-specific assumptions (Blukis et al., 2021), while outperforming several models (Zhang and Chai, 2021; Suglia et al., 2021; Kim et al., 2021) that use more information during both training and evaluation. Our method correctly segments and labels subtasks in unlabeled demonstrations, including subtasks that involve novel compositions of actions and objects. Additional experiments show that pretraining on large (ungrounded) text corpora (Raffel et al., 2020) contributes to this success, demonstrating one mechanism by which background knowledge encoded in language can benefit tasks that do not involve language as an

input or an output.

Indeed, our results show that relatively little information about language grounding is needed for effective learning of language-based policies—a rich model of natural language text, a large number of demonstrations, and a small number of annotations suffice for learning compositional libraries of skills and effective policies for deploying them.

2 Preliminaries

We consider learning problems in which agents must perform multi-step tasks (like *cooking an egg*; Fig. 1) in interactive environments. We formalize these problems as undiscounted, episodic, partially observed Markov decision processes (POMDPs) defined by a tuple $(\mathcal{S}, \mathcal{A}, T, \Omega, O)$, where \mathcal{S} is a set of **states**, \mathcal{A} is a set of **actions**, $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is an (unknown) **state transition function**, Ω is a set of **observations**, and $O : \mathcal{S} \rightarrow \Omega$ is an (unknown) **observation function**.² We assume that observations include a distinguished **goal specification** g that remains constant throughout an episode; given a dataset \mathcal{D} of consisting of **goals** g and **demonstrations** \mathbf{d} (i.e. $\mathcal{D} = \{(\mathbf{d}_1, g_1), (\mathbf{d}_2, g_2), \dots\}$; $\mathbf{d} = [(o_1, a_1), (o_2, a_2), \dots]$; $o \in \Omega, a \in \mathcal{A}$), we aim to learn a goal-conditional policy $\pi(a_t \mid \mathbf{a}_{:t-1}, \mathbf{o}_{:t}, g) = \pi(a_t \mid a_1, \dots, a_{t-1}, o_1, \dots, o_t, g)$ that generalizes demonstrated behaviors to novel goals and states.

For tasks like the ones depicted in Fig. 1, this learning problem requires agents to accomplish multiple subgoals (like *finding an egg* or *operating an appliance*) in a feasible sequence. As in past work, we address this challenge by focusing on **hierarchical** policy representations that plan over temporal abstractions of low-level action sequences. We consider a generic class of hierarchical policies that first predict a sequence of **subtask specifications** τ from a distribution $\pi^C(\tau_i \mid \tau_{:i-1}, g)$ (the *controller*), then from each τ generate a sequence of actions $a_1 \dots a_n$ from a distribution $\pi^E(a_i \mid \mathbf{a}_{:i-1}, \mathbf{o}_{:i}, \tau)$ (the *executor*).³ At each timestep, π^E may either generate an action from \mathcal{A} ; or a special termination signal STOP; after STOP is selected, control is returned to π^C and a new τ is generated. This process is visualized

²For notational convenience, we assume without loss of generality that T and O are deterministic.

³In past work, π^E often conditions on the current observation as well as goal and history of past subtask specifications; we found that this extra information was not needed for the tasks studied here.

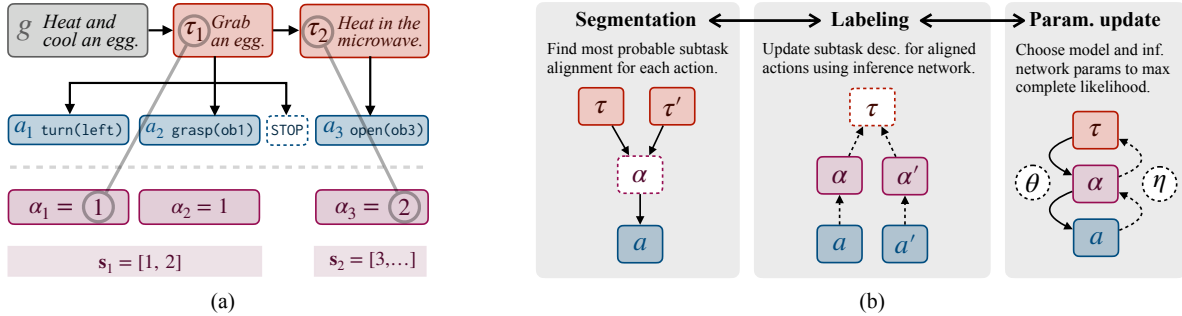


Figure 2: (a) When a hierarchical policy is deployed, π^C generates a sequence of subtask specifications, and π^E translates each of these to a low-level action sequence ending in STOP. At training time, this hierarchical structure is not available, and must be inferred to train our model. To do so, we assign each action a_i an auxiliary **alignment** variable α_i identifying the subtask that produced it. Alignments divide an action sequence into a sequence of **segments** s containing actions aligned to the same subtask. Automatically segmenting training demonstrations makes it possible to learn modular, reusable policies for individual subtasks without direct supervision. (b) Overview of the proposed learning algorithm (SL)³, which alternates between segmenting (by aligning) actions to fixed subtask specifications; labeling segments given fixed alignments, and updating model parameters.

in Fig. 2(a). Trajectories generated by hierarchical policies themselves have hierarchical structure: each subtask specification τ generates a **segment** of a trajectory (delimited by a STOP action) that accomplishes a specific subgoal.

Training a hierarchical policy requires first defining a space of subtask specifications τ , then parameterizing controller and executor policies that can generate these specifications appropriately. Most past research has either pre-defined an inventory of target skills and independently supervised π^C and π^E (Sutton et al., 1999; Kulkarni et al., 2016; Dayan and Hinton, 1992); or performed unsupervised discovery of a finite skill inventory using clustering techniques (Dietterich, 1999; Fox et al., 2017).

Both methods have limitations, and recent work has explored methods for using richer supervision to guide discovery of skills that are more robust than human-specified ones and more generalizable than automatically discovered ones. One frequently proposed source of supervision is language: in **latent language policies**, π^C is trained to generate goal-relevant instructions in natural language, π^E is trained to follow instructions, and the space of abstract actions available for planning is in principle as structured and expressive as language itself. But current approaches to LLP training remain impractical, requiring large datasets of independent, fine-grained supervision for π^C and π^E . Below, we describe how to overcome this limitation, and instead learn from large collections of unlabeled demonstrations augmented with only a small amount of natural language supervision.

3 Approach

Overview We train hierarchical policies on unannotated action sequences by inferring latent natural language descriptions of the subtasks they accomplish (Fig. 2(b)). We present a learning algorithm that jointly partitions these action sequences into smaller segments exhibiting reusable, task-general skills, labels each segment with a description, trains π^C to generate subtask descriptions from goals, and π^E to generate actions from subtask descriptions.

Formally, we assume access to two kinds of training data: a large collection of **unannotated demonstrations** $\mathcal{D} = \{(\mathbf{d}_1, g_1), (\mathbf{d}_2, g_2), \dots\}$ and a smaller collection of **annotated demonstrations** $\mathcal{D}^{\text{ann}} = \{(\mathbf{d}_1, g_1, \tau_1), (\mathbf{d}_2, g_2, \tau_2), \dots\}$ where each τ consists of a sequence of natural language **instructions** $[\tau_1, \tau_2, \dots]$ corresponding to the subtask sequence that should be generated by π^C . We assume that even annotated trajectories leave much of the structure depicted in Fig. 2(a) unspecified, containing no explicit segmentations or STOP markers. (The number of instructions $|\tau|$ will in general be smaller than the number of actions $|\mathbf{d}|$.) Training π^E requires inferring the correspondence between actions and annotations on \mathcal{D}^{ann} while inferring annotations themselves on \mathcal{D} .

Training objective To begin, it will be convenient to have an explicit expression for the probability of a demonstration given a policy (π^C, π^E) . To do so, we first observe that the hierarchical generation procedure depicted in Fig. 2(a) produces a latent *alignment* between each action and the subtask

τ that generated it. We denote these alignments α , writing $\alpha_i = j$ to indicate that a_i was generated from τ_j . Because π^C executes subtasks in sequence, alignments are *monotonic*, satisfying $\alpha_i = \alpha_{i-1}$ or $\alpha_i = \alpha_{i-1} + 1$. Let $\text{seg}(\alpha)$ denote the **segmentation** associated with α , the sequence of sequences of action indices $[[i : \alpha_i = 1], [i : \alpha_i = 2], \dots]$ aligned to the same instruction (see Fig. 2(a)). Then, for a fixed policy and POMDP, we may write the joint probability of a demonstration, goal, annotation, and alignment as:

$$p(\mathbf{d}, g, \tau, \alpha) \propto \prod_{\mathbf{s} \in \text{seg}(\alpha)} \left[\pi^C(\tau_{\mathbf{s}} \mid \tau_{<\mathbf{s}}, g) \times \left(\prod_{i \in 1..|\mathbf{s}|} \pi^E(a_i \mid \mathbf{a}_{\mathbf{s}_{:i-1}}, \mathbf{o}_{\mathbf{s}_{:i}}, \tau_{\alpha_i}) \right) \times \pi^E(\text{STOP} \mid \mathbf{a}_{\mathbf{s}}, \mathbf{o}_{\mathbf{s}}) \right]. \quad (1)$$

Here $<\mathbf{s}$ (in a slight abuse of notation) denotes all segments preceding \mathbf{s} , and \mathbf{s}_i is the index of the i th action in \mathbf{s} . The constant of proportionality in Eq. (1) depends only on terms involving $T(s' \mid s, a)$, $O(o \mid s)$ and $p(g)$, all independent of π^C or π^E ; Eq. (1) thus describes the component of the data likelihood under the agent’s control (Ziebart et al., 2013).

With this definition, and given \mathcal{D} and \mathcal{D}^{ann} as defined above, we may train a latent language policy using partial natural language annotations via ordinary maximum likelihood estimation, imputing the missing segmentations and labels in the training set jointly with the parameters of π^C and π^E (which we denote θ) in the combined annotated and unannotated likelihoods:

$$\arg \max_{\hat{\tau}, \hat{\alpha}, \hat{\theta}} \mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta}) + \mathcal{L}^{\text{ann}}(\hat{\alpha}, \hat{\theta}) \quad (2)$$

where

$$\mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta}) = \sum_{(\mathbf{d}, g) \in \mathcal{D}} \log p(\mathbf{d}, g, \hat{\tau}, \hat{\alpha}) \quad (3)$$

$$\mathcal{L}^{\text{ann}}(\hat{\alpha}, \hat{\theta}) = \sum_{(\mathbf{d}, g, \tau) \in \mathcal{D}^{\text{ann}}} \log p(\mathbf{d}, g, \tau, \hat{\alpha}) \quad (4)$$

and where we have suppressed the dependence of $p(\mathbf{d}, g, \tau, \alpha)$ on $\hat{\theta}$ for clarity. This objective involves continuous parameters $\hat{\theta}$, discrete alignments $\hat{\alpha}$, and discrete labelings $\hat{\tau}$. We optimize it via block coordinate ascent on each of these components in turn: alternating between **re-segmenting**

demonstrations, **re-labeling** those without ground-truth labels, and **updating parameters**. The full learning algorithm, which we refer to as **(SL)³** (*semi-supervised skill learning with latent language*), is shown in Algorithm 1, with each step of the optimization procedure described in more detail below.

Segmentation: $\arg \max_{\hat{\alpha}} \mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta}) + \mathcal{L}^{\text{ann}}(\hat{\alpha}, \hat{\theta})$

The segmentation step associates each low-level action with a high-level subtask by finding the highest scoring alignment sequence α for each demonstration in \mathcal{D} and \mathcal{D}^{ann} . While the number of possible alignments for a single demonstration is exponential in demonstration length, the assumption that π^E depends only on the current subtask implies the following recurrence relation:

$$\begin{aligned} & \max_{\alpha_{1:n}} p(\mathbf{d}_{1:n}, g, \tau_{1:m}, \alpha_{1:n}) \\ &= \max_i \left(\max_{\alpha_{1:i}} p(\mathbf{d}_{1:i}, g, \tau_{1:m-1}, \alpha_{1:i}) \right. \\ & \quad \left. \times p(\mathbf{d}_{i+1:n}, g, \tau_m, \alpha_{i+1:n} = m) \right) \quad (5) \end{aligned}$$

This means that the highest-scoring segmentation can be computed by an algorithm that recursively identifies the highest-scoring alignment to each prefix of the instruction sequence at each action (Algorithm 2), a process requiring $O(|\mathbf{d}||\tau|)$ space and $O(|\mathbf{d}|^2|\tau|)$ time. The structure of this dynamic program is identical to the forward algorithm for hidden semi-Markov models (HSMMs), which are widely used in NLP for tasks like language generation and word alignment (Wiseman et al., 2018). Indeed, Algorithm 2 can be derived immediately from Eq. (1) by interpreting $p(\mathbf{d}, g, \tau, \alpha)$ as the output distribution for an HSMM in which emissions are actions, hidden states are alignments, the emission distribution is π^E and the transition distribution is the deterministic distribution with $p(\alpha + 1 \mid \alpha) = 1$.

This segmentation procedure does not produce meaningful subtask boundaries until an initial executor policy has been trained. Thus, during the first iteration of training, we estimate a segmentation by fitting a 3-state hidden Markov model to training action sequences using the Baum–Welch algorithm (Baum et al., 1970), and mark state transitions as segment boundaries. Details about the initialization step may be found in Appendix B.

Algorithm 1: (SL)³: Semi-Supervised Skill Learning with Latent Language

Input: Unannotated demonstrations

$\mathcal{D} = \{(\mathbf{d}_1, g_1), (\mathbf{d}_2, g_2), \dots\}$;

Annotated demonstrations

$\mathcal{D}^{\text{ann}} = \{(\mathbf{d}_1, g_1, \tau_1), (\mathbf{d}_2, g_2, \tau_2), \dots\}$

Output: Inferred alignments $\hat{\alpha}$, labels $\hat{\tau}$, and parameters θ for π^{C} and π^{E} .

// Initialization

Initialize policy parameters θ using a pretrained language model (Raffel et al., 2020).

Initialize inference network parameters

$\eta \leftarrow \arg \max_{\eta} \sum_{\mathbf{d} \in \mathcal{D}^{\text{ann}}} \sum_{\mathbf{s}, \tau} \log q_{\eta}(\tau | \mathbf{a}_{\mathbf{s}}, \mathbf{o}_{\mathbf{s}})$.

for iteration $t \leftarrow 1 \dots T$ do

 // Segmentation

 // Infer alignments between actions and subtasks.

 if $t = 1$ then

 Initialize $\hat{\alpha}$ using the Baum–Welch algorithm (Baum et al., 1970)

 else

$\hat{\alpha} \leftarrow \arg \max_{\alpha} \mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta}) + \mathcal{L}^{\text{ann}}(\hat{\alpha}, \hat{\theta})$
 [Algorithm 2].

 end

 // Labeling

 // Infer subtask labels for unannotated demos \mathcal{D} .

$\hat{\tau} \leftarrow \arg \max_{\tau} \mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta})$

 // Parameter Update

 // Fit policy and proposal model parameters.

$\hat{\theta} \leftarrow \arg \max_{\theta} \mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta}) + \mathcal{L}^{\text{ann}}(\hat{\alpha}, \hat{\theta})$

$\hat{\eta} \leftarrow \arg \max_{\eta} \sum_{\mathbf{d}} \sum_{\mathbf{s}, \tau} \log q_{\eta}(\hat{\tau} | \mathbf{a}_{\mathbf{s}}, \mathbf{o}_{\mathbf{s}})$

end

Algorithm 2: Dynamic program for segmentation

Input: Demonstration $\mathbf{d} = [(o_1, a_1), \dots, (o_n, a_n)]$;

Task specifications $\tau = [\tau_1, \dots, \tau_m]$.

Executor $\pi^{\text{E}}(\mathbf{a} | \mathbf{o}, \tau) = \prod_i \pi^{\text{E}}(a_i | \mathbf{a}_{:i-1}, \mathbf{o}_{:i}, \tau)$

Output: Maximum *a posteriori* alignments α .

scores \leftarrow an $n \times m$ matrix of zeros

// scores $[i, j]$ holds the log-probability of the

// highest-scoring sequence whose final action i is

// aligned to subtask j .

for $i \leftarrow 1 \dots n$ do

 for $j \leftarrow 1 \dots |\tau|$ do

 scores $[i, j] \leftarrow -\infty$

 for $k \leftarrow 1 \dots i - 1$ do

 scores $[i, j] \leftarrow \max$ (

 scores $[i, j]$,

 scores $[k, j - 1]$

 + $\log \pi^{\text{E}}(\mathbf{a}_{k+1:i} | \mathbf{o}_{k+1:i}, \tau_j)$)

 end

 end

end

The optimal alignment sequence may be obtained from scores via back-tracing (Rabiner, 1989).

Labeling: $\arg \max_{\hat{\tau}} \mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta})$

Inference of latent, language-based plan descriptions in unannotated demonstrations involves an intractable search over string-valued τ . To approximate this search tractably, we used a **learned, amortized inference procedure** (Wainwright and Jordan, 2008; Hoffman et al., 2013; Kingma and Welling, 2014) to impute descriptions given fixed segmentations. During each parameter update step (described below), we train an inference model $q_{\eta}(\tau | \mathbf{a}_{\mathbf{s}^{(i)}}, \mathbf{a}_{\mathbf{s}^{(i+1)}}, g)$ to approximate the posterior distribution over descriptions for a given segment given a goal, the segment’s actions, and the actions from the subsequent segment.⁴ Then, during the labeling step, we label complete demonstrations by choosing the highest-scoring instruction for each trajectory independently:

$$\arg \max_{\tau} \log p(\mathbf{d}, g, \tau, \alpha) \approx$$

$$\left[\arg \max_{\tau} q(\tau | \mathbf{a}_{\mathbf{s}^{(i)}}, \mathbf{a}_{\mathbf{s}^{(i+1)}}, g) \mid \mathbf{s}^{(i)} \in \text{seg}(\alpha) \right] \quad (6)$$

Labeling is performed only for demonstrations in \mathcal{D} , leaving the labels for \mathcal{D}^{ann} fixed during training.

Param update: $\arg \max_{\hat{\theta}} \mathcal{L}(\hat{\tau}, \hat{\alpha}, \hat{\theta}) + \mathcal{L}^{\text{ann}}(\hat{\alpha}, \hat{\theta})$

This is the simplest of the three update steps: given fixed instructions and alignments, and π^{E} , π^{C} parameterized as neural networks, this objective is differentiable end-to-end. In each iteration, we train these to convergence (optimization details are described in Section 4 and Appendix C). During the parameter update step, we also fit parameters η of the proposal model to maximize the likelihood $\sum_{\mathbf{d}} \sum_{\mathbf{s}, \tau} \log q_{\eta}(\hat{\tau} | \mathbf{a}_{\mathbf{s}}, \mathbf{o}_{\mathbf{s}})$ with respect to the current segmentations $\hat{\mathbf{s}}$ and labels $\hat{\tau}$.

As goals, subtask indicators, and actions may all be encoded as natural language strings, π^{C} and π^{E} may be implemented as conditional language models. As described below, we initialize both policies with models *pretrained* on a large text corpora.

4 Experimental Setup

Our experiments aim to answer two questions. First, does the latent-language policy representation described in Section 3 improve downstream performance on complex tasks? Second, how many natural language annotations are needed to train

⁴In our experiments, conditioning on observations or longer context did not improve the accuracy of this model.

an effective latent language policy given an initial dataset of unannotated demonstrations?

Environment We investigate these questions in the ALFRED environment of Shridhar et al. (2020). ALFRED consists of a set of interactive simulated households containing a total of 120 rooms, accompanied by a dataset of 8,055 expert task demonstrations for an embodied agent annotated with 25,743 English-language instructions. Observations o are bitmap images from a forward-facing camera, and actions a are drawn from a set of 12 low-level navigation and manipulation primitives. Manipulation actions (7 of the 12) additionally require predicting a mask over the visual input to select an object for interaction. See Shridhar et al. (2020) for details.

While the ALFRED environment is typically used to evaluate *instruction following* models, which map from detailed, step-by-step natural language descriptions to action sequences (Shridhar et al., 2020; Singh et al., 2020; Corona et al., 2021), our experiments focus on an *goal-only* evaluation in which agents are given goals (but not fine-grained instructions) at test time. Several previous studies have also considered goal-only evaluation for ALFRED, but most use extremely fine-grained supervision at *training time*, including full supervision of symbolic plan representations and their alignments to demonstrations (Min et al., 2021; Zhang and Chai, 2021), or derived sub-task segmentations using ALFRED-specific rules (Blukis et al., 2021). In contrast, our approach supports learning from partial, language-based annotations without segmentations or alignments, and this data condition is the main focus of our evaluation.

Modeling details π^C and π^E are implemented as sequence-to-sequence transformer networks (Vaswani et al., 2017). π^C , which maps from text-based goal specifications to text-based instruction sequences, is initialized with a pre-trained T5-small language model (Raffel et al., 2020). π^E , which maps from (textual) instructions and (image-based) observations to (textual) actions and (image-based) object selection masks is also initialized with T5-small; to incorporate visual input, this model first embeds observations using a pretrained ResNet18 model (He et al., 2016) and transforms these linearly to the same dimensionality as the word embedding layer. Details about the architecture of π^C and π^E may be found in Appendix C.

Model variants for exploration In ALFRED, navigation in the goal-only condition requires exploration of the environment, but no exploration is demonstrated in training data, and techniques other than imitation learning are required for this specific skill. To reflect this, we replace all annotations containing detailed navigation instructions *go to the glass on the table to your left* with generic ones *find a glass*. Examples and details of how navigation instructions are modified can be found in Appendix E and Fig. 7. The ordinary (SL)³ model described above is trained on these abstracted instructions.

A key advantage of (SL)³ is modularity: individual skills may be independently supervised or re-implemented. To further improve (SL)³’s navigation capabilities, we introduce two model variants in which sub-task specifications beginning *Find...* are executed by either a planner with ground-truth environment information or a specialized navigation module from the HLSM model (Blukis et al., 2021) rather than π^E . Outside of navigation, these models preserve the architecture and training procedure of (SL)³, and are labeled (SL)³+**planner** and (SL)³+**HLSM** in experiments below.

Baselines and comparisons We compare the performance of (SL)³ to several baselines:

seq2seq: A standard (non-hierarchical) goal-conditioned policy, trained on the (g, d) pairs in $\mathcal{D} \cup \mathcal{D}^{\text{ann}}$ to maximize $\sum_{\mathbf{a}, \mathbf{o}, g} \log \pi(\mathbf{a} \mid \mathbf{o}, g)$, with π parameterized similar to π^E .

seq2seq2seq: A supervised hierarchical policy with the same architectures for π^C and π^E as in (SL)³, but with π^C trained to generate subtask sequences by maximizing $\sum_{\tau, g} \log \pi^C(\tau \mid g)$ and π^E trained to maximize $\sum_{\mathbf{a}, \mathbf{o}, \tau, g} \log \pi^E(\mathbf{a} \mid \mathbf{o}, \tau, g)$ using only \mathcal{D}^{ann} . Because π^E maps from complete task sequences to complete low-level action sequences, training of this model involves no explicit alignment or segmentation steps.

no-pretrain, no-latent: Ablations of the full (SL)³ model in which π^C and π^E are, respectively, randomly initialized or updated only on $\mathcal{L}^{\text{ann}}(\hat{\alpha}, \hat{\theta})$ during the parameter update phase.

We additionally contextualize our approach by comparing it to several state-of-the-art models for the *instruction following* task in ALFRED: **S+** (Shridhar et al., 2020), **MOCA** (Singh et al., 2020), **Modular** (Corona et al., 2021), **HiTUT** (Zhang and Chai, 2021), **ABP** (Kim et al., 2021), **ET** (Pashevich et al., 2021), **EmBERT** (Suglia et al., 2021), and **FILM** (Min et al., 2021). Like seq2seq, these

are neural sequence-to-sequence models trained to map instructions to actions; they incorporate several standard modeling improvements from the instruction following literature, including progress monitoring (Ma et al., 2019) and pretrained object recognizers (Singh et al., 2020). Many of these models are trained with stronger supervision than (SL)³, including instructions and alignments during training, and ground truth instructions *during evaluation*; see Table 3 for details.

Evaluation Following Shridhar et al. (2020), Table 1(a) computes the *online, subtask-level* accuracy of each policy, and Table 1(b) computes the end-to-end *success rate* of each policy. See the ALFRED paper for details of these evaluations. For data-efficiency experiments involving a large number of policy variants (Table 2, Fig. 4), we instead use an *offline evaluation* in which we measure the fraction of subtasks in which a policy’s predicted actions (ignoring object selection masks) *exactly match* the ground truth action sequence.

5 Results

Table 1 compares (SL)³ with flat and hierarchical imitation learning baselines. The table includes two versions of the model: a 100% model trained with full instruction supervision ($|\mathcal{D}|=0$, $|\mathcal{D}^{\text{ann}}|=21000$) and a 10% model trained with only a small fraction of labeled demonstrations ($|\mathcal{D}|=19000$, $|\mathcal{D}^{\text{ann}}|=2000$). seq2seq and seq2seq2seq models are *always* trained with 100% of natural language annotations. Results are shown in Table 1. We find:

(SL)³ improves on flat policies: In both the 10% and 100% conditions, it improves over the *subtask* completion rate of the seq2seq (goals-to-actions) model by 25%. When either planner- or mapping-based navigation is used in conjunction with (SL)³, it achieves *end-to-end* performance comparable to the HLSM method, which relies on similar supervision. Strikingly, it outperforms several recent methods with access to even more detailed information at training or evaluation time.

Language-based policies can be trained with sparse natural language annotations: Performance of (SL)³ trained with 10% and 100% natural language annotations is similar (and in both cases superior to seq2seq and seq2seq2seq trained on 100% of data). Appendix Fig. 4 shows more detailed supervision curves. Ablation experiments in Table 2 show that inference of latent training plans is important for this result: with no inference of

(a) **Online** subtask success rate for (SL)³ and baselines

Model	Avg	Clean	Cool	Heat	Pick	Put	Slice	Toggle	GoTo
(SL) ³ (10%)	50	56	75	74	50	48	54	32	13
(SL) ³ (100%)	53	68	82	75	50	45	55	32	15
seq2seq	25	16	33	64	20	15	25	13	14
seq2seq2seq	39	15	69	58	29	42	50	32	15

(b) **End-to-end** task success rates for (SL)³ and other models.

Goal + partial plan sup.		Extra information	
Model	SR	Model	SR
(SL) ³ (10%)	0.0	FILM (Min+21)	20.1
(SL) ³ +HLSM (10%)	15.5	(SL) ³ +planner (10%)	40.4
HLSM (Blukis+21)*	17.2	HiTUT (Zhang+21)	11.1
seq2seq	0.0	EmBERT (Suglia+21)	5.7
seq2seq2seq	0.0	ET (Pashkevich+21)	7.3
		ABP (Kim+21)	12.6
		S+ (Shridhar+20)	0.1
		MOCA (Singh+21)	5.4

Table 1: (a) Evaluation of (SL)³ and baselines using the *subtask* evaluation from Shridhar et al. (2020). All models in this section were trained with both goals g and annotated subtask descriptions τ , but observed only goals during evaluation. (b) Evaluation of (SL)³ and concurrent work using the *success rate* evaluation from Shridhar et al. (2020). Models in the left column use only goals and partial subtask descriptions at training time, and only goals at test time. (The HLSM model also uses a rule-based, ALFRED-specific procedure for converting action sequences to high-level plan specifications.) Models on the right use extra information, including ground-truth training segmentations and alignments, and ground-truth test-time plans. *Result of our HLSM reproduction using public code and trained models.

latent instructions (i.e. training only on annotated demonstrations), performance drops from 56% to 52%. Fig. 3 shows an example of the structure inferred for an unannotated trajectory: the model inserts reasonable segment boundaries and accurately labels each step.

Language model pretraining improves automated decision-making. Ablation experiments in Table 2 provide details. Language model pretraining of π^C and π^E (on *ungrounded* text) is crucial for good performance in the low-data regime: with 10% of annotations, models trained from scratch complete 49% of tasks (vs 56% for pretrained models). We attribute this result in part to the fact that pretrained language models encode information about the common-sense structure of plans, e.g. the fact that *slicing a tomato* first requires *finding a knife*. Such models are well-positioned to adapt to “planning” problems that require modeling relations between natural language strings. These

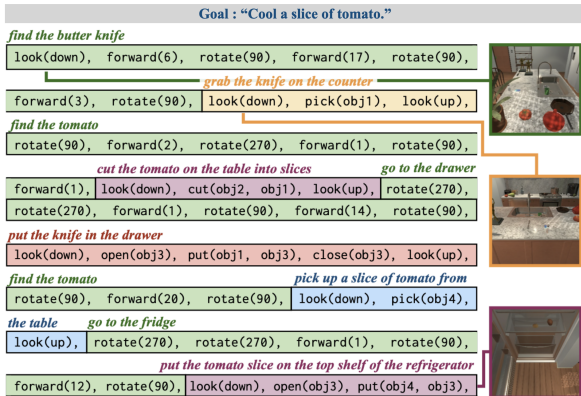


Figure 3: Example of an inferred segmentation and labeling for an unannotated trajectory. The trajectory is parsed into a sequence of 10 segments and q_η assigns high scoring natural-language labels to the segmented actions. These are consistent with the objects, receptacles and sub-tasks. The overall sequence of latent-language skills is a good plan for the high-level goal.

Model	Average
(SL) ³ (10%)	56
(SL) ³ (100%)	58
(SL) ³ (ground-truth α)	65
no-pretrain	49
no-latent	52

Table 2: Ablation experiments. Providing ground-truth alignments at training time improves task completion rates, suggesting potential benefits from an improved alignment procedure. Pretraining and inference of latent task representations contribute 7% and 4% respectively to task completion rate with 10% of annotations.

experiments point to a potentially broad role for pretrained language models in tasks that do not involve language as an input or an output.

One especially interesting consequence of the use of language-based skills is our model’s ability to produce high-level plans for **out-of-distribution goals**, featuring objects or actions that are not part of the ALFRED dataset at all. Examples are provided in Fig. 5 and discussed in Appendix A. While additional modeling work is needed to generate low-level actions for these high-level plans, they point to *generalization* as a key differentiator between latent language policies and ordinary hierarchical ones.

6 Related Work

Our approach draws on a large body of research at the intersection of natural language processing, representation learning, and autonomous control.




Figure 4: Offline subtask success rate as a function of the fraction of annotated examples. Only a small fraction of annotations (5–10%) are needed for good performance; inference of latent instructions is beneficial in the low-data regime.

Language-based supervision and representation The use of natural language annotations to scaffold learning, especially in computer vision and program synthesis applications, has been the subject of a number of previous studies (Branan et al., 2009; Frome et al., 2013; Andreas et al., 2018; Wong et al., 2021). Here, we use language to support policy learning, specifically by using natural language instructions to discover compositional subtask abstractions that can support autonomous control. Our approach is closely related to previous work on learning skill libraries from *policy sketches* (Andreas et al., 2017; Shiarlis et al., 2018); instead of the fixed skill inventory used by policy sketches, (SL)³ learns an open-ended, compositional library of behaviors indexed by natural language strings.

Hierarchical policies Hierarchical policy learning and temporal abstraction have been major areas of focus since the earliest research on reinforcement learning and imitation learning (McGovern and Barto, 2001; Konidaris et al., 2012; Daniel et al., 2012). Past work typically relies on direct supervision or manual specification of the space of high-level skills (Sutton et al., 1999; Kulkarni et al., 2016) or fully unsupervised skill discovery (Dietterich, 1999; Bacon et al., 2017). Our approach uses policy architectures from this literature, but aims to provide a mechanism for supervision that allows fine-grained control over the space of learned skills (as in fully supervised approaches) while requiring only small amounts of easy-to-gather human supervision.

Language and interaction Outside of language-based supervision, problems at the

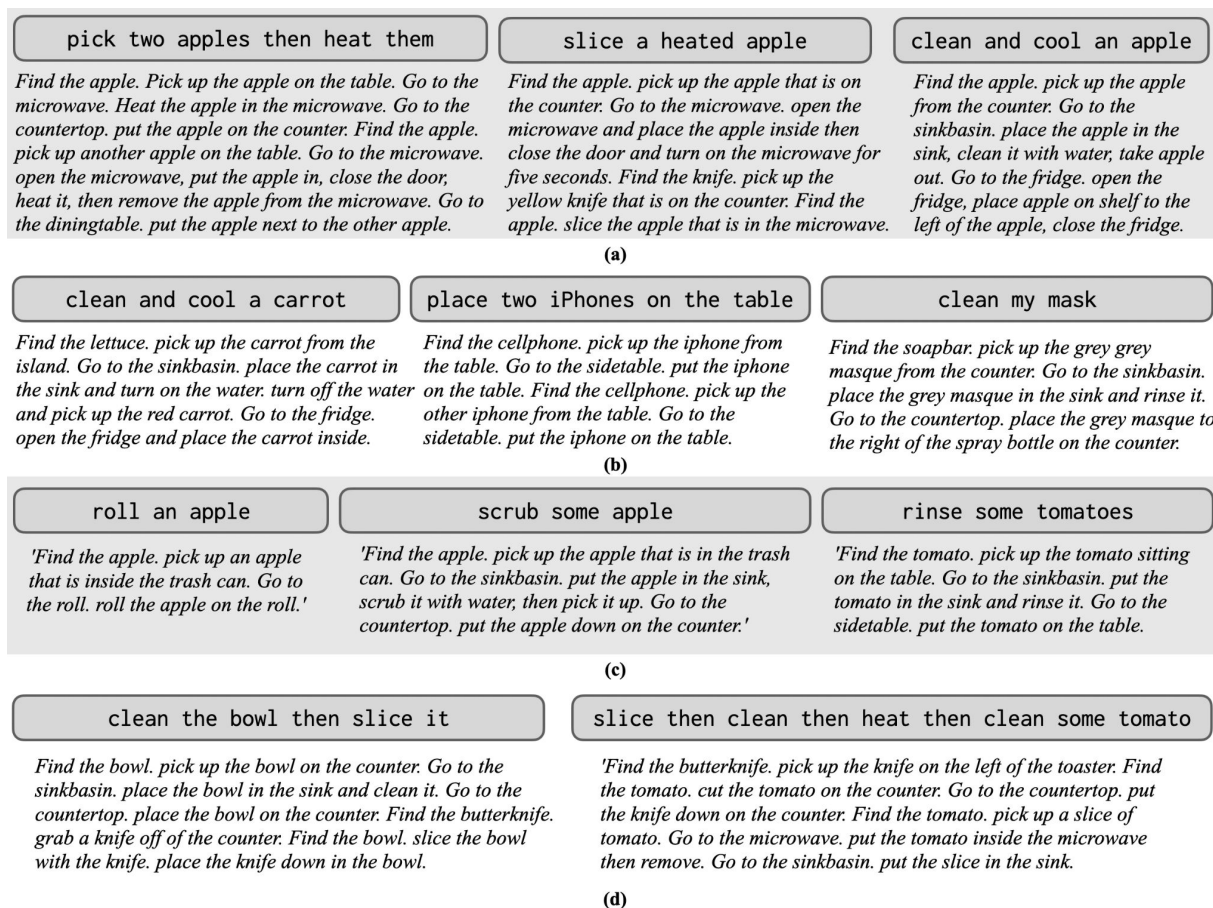


Figure 5: Successes and failures of π^C in out-of-distribution (OOD) settings including novel (a) sub-task orders (b) objects (c) verbs. The use of a pretrained LM as the backbone of the planning model means that models produce correct or plausible plans for many of these out-of-distribution goals. (d) Other failure modes: The model fails to predict actions based on the true affordances of objects and cannot generate arbitrarily long plans.

intersection of language and control include *instruction following* (Chen and Mooney, 2011; Branavan et al., 2009; Tellex et al., 2011; Anderson et al., 2018; Misra et al., 2017), *embodied question answering* (Das et al., 2018; Gordon et al., 2018) and dialog tasks (Tellex et al., 2020). As in our work, representations of language learned from large text corpora facilitate grounded language learning (Shridhar et al., 2021), and interaction with the environment can in turn improve the accuracy of language generation (Zellers et al., 2021); future work might extend our framework for semi-supervised inference of plan descriptions to these settings as well.

7 Conclusion

We have presented (SL)³, a framework for learning hierarchical policies from demonstrations sparsely annotated with natural language descriptions. Using these annotations, (SL)³ infers the latent structure of unannotated demonstrations, automatically

segmenting them into subtasks and labeling each subtask with a compositional description. Learning yields a hierarchical policy in which natural language serves as an abstract representation of subgoals and plans: a *controller* sub-policy maps from goals to natural language plan specifications, and a modular *executor* that maps each component of the plan to a sequence of low-level actions. In simulated household environments, this model can complete abstract goals (like *slice a tomato*) with accuracy comparable to state-of-the-art models trained and evaluated with fine-grained plans (*find a knife, carry the knife to the tomato, ...*).

While our evaluation has focused on household robotics tasks, the hierarchical structure inferred by (SL)³ is present in a variety of learning problems, including image understanding, program synthesis, and language generation. In all those domains, generalized versions of (SL)³ might offer a framework for building high-quality models using only a small amount of rich natural language supervision.

Acknowledgements

We would like to thank Valts Blukis and Shikhar Murty for helpful discussions. Also thanks to Joe O’ Connor, Gabe Grand and the anonymous reviewers for their feedback on an early draft of the paper.

References

- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, I. Reid, Stephen Gould, and A. V. Hengel. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3674–3683.
- Jacob Andreas, D. Klein, and Sergey Levine. 2017. Modular multitask reinforcement learning with policy sketches. *International Conference of Machine Learning*.
- Jacob Andreas, Dan Klein, and Sergey Levine. 2018. Learning with latent language. New Orleans, Louisiana. Association for Computational Linguistics.
- P. Bacon, Jean Harb, and Doina Precup. 2017. The option-critic architecture. In *AAAI*.
- L. Baum, T. Petrie, George W. Soules, and Norman Weiss. 1970. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *Annals of Mathematical Statistics*, 41:164–171.
- Valts Blukis, Chris Paxton, D. Fox, Animesh Garg, and Yoav Artzi. 2021. A persistent spatial semantic representation for high-level natural language instruction execution. *ArXiv*, abs/2107.05612.
- S. Branavan, Harr Chen, Luke Zettlemoyer, and R. Barzilay. 2009. Reinforcement learning for mapping instructions to actions. In *ACL*.
- R. Brockett. 1993. Hybrid models for motion control systems.
- David L. Chen and R. Mooney. 2011. Learning to interpret natural language navigation instructions from observations. In *AAAI 2011*.
- Rodolfo Corona, Daniel Fried, Coline Devin, D. Klein, and Trevor Darrell. 2021. Modular networks for compositional instruction following. In *NAACL*.
- Christian Daniel, G. Neumann, and Jan Peters. 2012. Hierarchical relative entropy policy search. *J. Mach. Learn. Res.*, 17:93:1–93:50.
- Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. 2018. Embodied question answering. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2135–213509.
- P. Dayan and Geoffrey E. Hinton. 1992. Feudal reinforcement learning. In *NIPS*.
- Thomas G Dietterich. 1999. Hierarchical reinforcement learning with the MAXQ value function decomposition.
- R. Fikes, P. Hart, and N. Nilsson. 1972. Learning and executing generalized robot plans. *Artif. Intell.*, 3:251–288.
- Roy Fox, S. Krishnan, I. Stoica, and Ken Goldberg. 2017. Multi-level discovery of deep options. *ArXiv*, abs/1703.08294.
- Andrea Frome, G. Corrado, Jonathon Shlens, Samy Bengio, J. Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. 2013. Devise: A deep visual-semantic embedding model. In *NIPS*.
- Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, D. Fox, and Ali Farhadi. 2018. Iqa: Visual question answering in interactive environments. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4089–4098.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- M. Hoffman, David M. Blei, Chong Wang, and J. Paisley. 2013. Stochastic variational inference. *ArXiv*, abs/1206.7051.
- Hengyuan Hu, Denis Yarats, Qucheng Gong, Yuandong Tian, and M. Lewis. 2019. Hierarchical decision making by generating and following natural language instructions. In *NeurIPS*.
- Athul Paul Jacob, M. Lewis, and Jacob Andreas. 2021. Multitasking inhibits semantic drift. *ArXiv*, abs/2104.07219.
- Yiding Jiang, S. Gu, K. Murphy, and Chelsea Finn. 2019. Language as an abstraction for hierarchical deep reinforcement learning. In *NeurIPS*.
- L P Kaelbling and T Lozano-Pérez. 2011. Hierarchical task and motion planning in the now. *2011 IEEE International*.
- Byeonghwi Kim, Suvaansh Bhambri, Kunal Pratap Singh, Roozbeh Mottaghi, and Jonghyun Choi. 2021. Agent with the big picture: Perceiving surroundings for interactive instruction following. In *Embodied AI Workshop CVPR*.
- Diederik P. Kingma and Max Welling. 2014. Auto-encoding variational bayes. *CoRR*, abs/1312.6114.
- G. Konidaris, S. Kuindersma, R. Grupen, and A. Barto. 2012. Robot learning from demonstration by constructing skill trees. *The International Journal of Robotics Research*, 31:360 – 375.

- Tejas D. Kulkarni, Karthik Narasimhan, A. Saeedi, and J. Tenenbaum. 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *NIPS*.
- I. Loshchilov and F. Hutter. 2019. Decoupled weight decay regularization. In *ICLR*.
- Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, G. Al-Regib, Z. Kira, R. Socher, and Caiming Xiong. 2019. Self-monitoring navigation agent via auxiliary progress estimation. *ArXiv*, abs/1901.03035.
- A. McGovern and A. Barto. 2001. Automatic discovery of subgoals in reinforcement learning using diverse density. In *ICML*.
- So Yeon Min, Devendra Singh Chaplot, Pradeep Ravikumar, Yonatan Bisk, and Ruslan Salakhutdinov. 2021. [FILM: following instructions in language with modular methods](#). *CoRR*, abs/2110.07342.
- Dipendra Kumar Misra, J. Langford, and Yoav Artzi. 2017. Mapping instructions and visual observations to actions with reinforcement learning. In *EMNLP*.
- A. Newell. 1973. Human problem solving.
- Alexander Pashevich, Cordelia Schmid, and Chen Sun. 2021. [Episodic transformer for vision-and-language navigation](#). *CoRR*, abs/2105.06453.
- Lawrence R. Rabiner. 1989. A tutorial on hidden markov models and selected applications. *Proceedings of the IEEE*.
- Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, W. Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *ArXiv*, abs/1910.10683.
- E. Sacerdoti. 1973. Planning in a hierarchy of abstraction spaces. *Artif. Intell.*, 5:115–135.
- K. Shiarlis, Markus Wulfmeier, S. Salter, S. Whiteson, and I. Posner. 2018. Taco: Learning task decomposition via temporal alignment for control. In *ICML*.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. [ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks](#). In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and M. Hausknecht. 2021. [Alfworld: Aligning text and embodied environments for interactive learning](#). *ArXiv*, abs/2010.03768.
- Tianmin Shu, Caiming Xiong, and R. Socher. 2018. Hierarchical and interpretable skill acquisition in multi-task reinforcement learning. *ArXiv*, abs/1712.07294.
- Kunal Pratap Singh, Suvaansh Bhambri, Byeonghwi Kim, Roozbeh Mottaghi, and Jonghyun Choi. 2020. Moca: A modular object-centric approach for interactive instruction following. *arXiv preprint arXiv:2012.03208*.
- Alessandro Suglia, Qiaozi Gao, Jesse Thomason, Govind Thattai, and Gaurav Sukhatme. 2021. [Embodied BERT: A transformer model for embodied, language-guided visual task completion](#). *arXiv*.
- R S Sutton, D Precup, and S Singh. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artif. Intell.*
- Stefanie Tellex, N. Gopalan, H. Kress-Gazit, and Cynthia Matuszek. 2020. Robots that use language.
- Stefanie Tellex, T. Kollar, Steven Dickerson, Matthew R. Walter, A. Banerjee, S. Teller, and N. Roy. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *AAAI*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).
- Martin J. Wainwright and M.I. Jordan. 2008. Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.*, 1:1–305.
- Sam Wiseman, S. Shieber, and Alexander M. Rush. 2018. Learning neural templates for text generation. *ArXiv*, abs/1808.10122.
- Catherine Wong, Kevin Ellis, J. Tenenbaum, and Jacob Andreas. 2021. Leveraging language to learn program abstractions and search heuristics. In *ICML*.
- Rowan Zellers, Ari Holtzman, Matthew E. Peters, R. Mottaghi, Aniruddha Kembhavi, Ali Farhadi, and Yejin Choi. 2021. [Piglet: Language grounding through neuro-symbolic interaction in a 3d world](#). In *ACL/IJCNLP*.
- Yichi Zhang and Joyce Chai. 2021. [Hierarchical task learning from language instructions with unified transformers and self-monitoring](#). *CoRR*, abs/2106.03427.
- Brian D Ziebart, J Andrew Bagnell, and Anind K Dey. 2013. The principle of maximum causal entropy for estimating interacting processes. *IEEE Transactions on Information Theory*, 59(4):1966–1980.

A Out-of-distribution Generalization

One of the advantages of language-based skill representations over categorical representations is open-endedness: (SL)³ does not require pre-specification of a fixed inventory of goals or actions. As a simple demonstration of this potential for extensibility, we design goal prompts consisting of novel object names, verbs and skill combinations not seen at training time, and test the model’s ability to generalize to out-of-distribution samples across the three categories. Some roll-outs can be seen in Fig. 5. We observe the following:

Novel sub-task combinations We qualitatively evaluate the ability of the model to generalize systematically to novel subtask combinations and subtask ordering not encountered at training time. Examples are shown in Fig. 5. For example, we present the model with the goal *slice a heated apple*; in the training corpus, objects are only heated *after* being sliced. It can be seen in Fig. 5 that the model able correctly orders the two subtasks. The model additionally generalizes to new combinations of tasks such as *clean and cool an apple*.

Novel objects and verbs The trained model also exhibits some success at generalizing novel object categories such as *carrot* and *mask*. In the carrot example, an incorrect *Find the lettuce* example is generated at the first step, but subsequent subtasks refer to a carrot (and apply the correct actions to it). The model also generalizes to new but related verbs such as *scrub* but fails at ones like *squash* that are unrelated to training goals.

Limitations One shortcoming of this approach is that affordances and constraints are incompletely modeled. Given a (physically unrealizable) goal *clean the bowl and then slice it*, the model cannot detect the impossible goal and instead generates a plan involving slicing the bowl. Another shortcoming of the model is the ability to generalize to goals that may involve considerably larger number of subgoals than goals seen at training time. For plans that involve very long sequences of skills (*slice then clean then heat...*) the generated plan skips some subtasks Fig. 5.

B Initialization: Segmentation Step

The training data contains no STOP actions, so π^E cannot be initialized by training on \mathcal{D}^{ann} . Using a randomly initialized π^E during the segmentation

step results in extremely low-quality segmentations. Instead, we obtain an initial set of segmentations via *unsupervised* learning on low-level action sequences.

In particular, we obtain initial segmentations using the Baum–Welch algorithm for unsupervised estimation of hidden Markov models (Baum et al., 1970). We replace string-valued latent variables produced by π^C with a discrete set of hidden states (in our experiments, we found that three hidden states sufficed). Transition and emission distributions, along with maximum *a posteriori* sequence labels, are obtained by running the expectation–maximization algorithm on state sequences. We then insert segment boundaries (and an implicit STOP action) at every transition between two distinct hidden states. Evaluated against ground-truth segmentations from the ALFRED training set, this produces an action-level accuracy of 87.9%. The detailed algorithm can be found in Baum et al. (1970).

C Model Architecture: Details

The controller policy π^C is a fine-tuned T5-small model. The executor policy π^E decodes the low-level sequence of actions conditioned on the first-person visual observations of the agent. We use the same architecture across the remaining baselines too. Fig. 6 depicts the architecture of the image-conditioned T5 model. In addition to task specifications, we convert low-level actions to templated commands: for example, `put(cup, table)` becomes *put the cup on the table*. These are parsed to select actions to send to the ALFRED simulator. During training, both models are optimized using the AdamW algorithm (Loshchilov and Hutter, 2019) with a learning rate of 1e-4, weight decay of 0.01, and $\epsilon = 1e-8$. We use a MaskRCNN model to generate action masks, selecting the predicted mask labeled with the class of the object name generated by the action decoder. The same model architecture is used across all baselines.

D Role of trajectory length

We conduct an additional set of ablation experiments aimed at clarifying what aspects of the demonstrated trajectories (SL)³ is better able to model than baselines. We begin by observing that most actions in our data are associated with navigation, with sequences of object manipulation actions (like those depicted in Fig. 3) constitut-

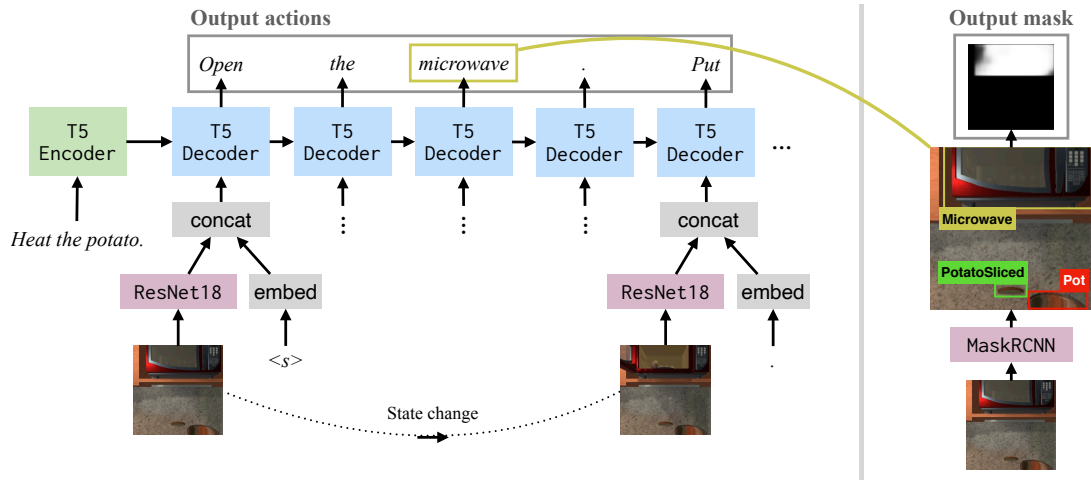


Figure 6: Model architecture for π^E , seq2seq and seq2seq2seq: Language parametrized sub-task/goal is input to the encoder and actions templated in natural language are generated sequentially token-wise. The predictions made are conditioned on the visual field of view of the agent at every time step along with the token generated the previous time step. At the end of every low-level action (when '.' is generated) the action the executed. For manipulation actions, the mask corresponding to the the object predicted is selected from the predictions of a MaskRCNN model on the visual state. Navigation actions do not operate over objects. Once the action is taken, the environment returns the updated visual state and the policy continues to be unrolled until termination (STOP).

ing only about 20% of each trajectory. We construct an alternative version of the dataset in which all navigation subtasks are replaced with a single TeleportTo action. This modification reduces average trajectory length from 50 actions to 9. In this case (SL)³ and seq2seq2seq perform comparably well (55.6% success rate and 56.7% success rate respectively), and only slightly better than seq2seq (53.6% success rate). Thus, while (SL)³ (and all baselines) perform quite poorly at navigation skills, *identifying* these skills and modeling their conditional independence from other trajectory components seems to be crucial for effective learning of other skills in the long-horizon setting. Hierarchical policies are still useful for modeling these shorter plans, but by a smaller margin than for long demonstrations.

E Navigation Instructions

The original ALFRED dataset contains detailed instructions for navigation collected post-hoc after the demonstrations are generated. For example, the sub-task specification associated with finding an apple might be given as *Go straight and turn to the right of the fridge and take a few steps ahead and look down*. Such instructions cannot be used for high-level planning, as they can only be generated with advance knowledge of the environment layout; successful behavior in novel environments requires exploration or explicit access to the environment's

'turn to the right twice and to to the end of the counter top and turn to the left and go to the end of the counter top'	→	"Find the knife."
'turn to the left and go to the front of the refrigerator and turn to the left and go to the refrigerator'	→	"Go to the refrigerator"
'turn to the right twice and take a few steps and turn to the left and go to the microwave'	→	"Go to the microwave"
'Go to the counter across the room from the stove.' Pick up the butter knife on the counter.'	→	"Find the butterknife"
'Turn right, turn right, walk straight to the oven'	→	"Go to the microwave"

Figure 7: Modified navigation annotations. Navigation instructions are converted to simpler object/location-oriented navigation goals using by creating templated plans from ALFRED dataset metadata.

map.

To address the mismatch between the agent's knowledge and the information needed to generate detailed navigation instructions, we navigation instructions in the ALFRED dataset with templated instructions of the form *Go to the* [object] (for appliances and containers) and *Find the* [object] (for movable objects). Because the ALFRED dataset provides PDDL plans for each demonstration, we can obtain the name of the target [object] directly from these plans. Examples are shown in Fig. 7.

Method	Training time					Inference time	
	Goal	Instructions	Program	Alignments	Depth	Goal	Instructions
(SL) ³	✓	10%	✗	✗	✗	✓	✗
seq2seq	✓	✗	✗	✗	✗	✓	✗
seq2seq2seq	✓	✓	✗	✗	✗	✓	✗
S+(Shridhar et al., 2020)	✓	✓	✗	✓	✗	✓	✓
MOCA(Singh et al., 2020)	✓	✓	✗	✓	✗	✓	✓
Modular (Corona et al., 2021)	✓	✓	✓	✓	✗	✓	✓
ABP (Kim et al., 2021)	✓	✓	✗	✓	✗	✓	✓
EmBERT (Suglia et al., 2021)	✓	✓	✗	✓	✗	✓	✓
ET (Pashevich et al., 2021)	✓	✓	✗	✓	✗	✓	✓
HLSM(Blukis et al., 2021)	✓	✗	*	*	✓	✓	✗
HiTUT (Zhang and Chai, 2021)	✓	✗	✓	✓	✓	✓	✗
FILM (Min et al., 2021)	✓	✗	✓	✓	✓	✓	✗

Table 3: Detailed comparison of information available to models and baselines at training time and inference.

*Re-derived using a rule-based segmentation procedure