Tell Me Dave: Context-Sensitive Grounding of Natural Language to Manipulation Instructions

Dipendra K Misra, Jaeyong Sung, Kevin Lee and Ashutosh Saxena.
Computer Science Department, Cornell University (USA).
Email contact: {dkm,jysung,kkl53,asaxena}@cs.cornell.edu

***Abstract*—It is important for a robot to be able to interpret natural language commands given by a human. In this paper, we consider performing a sequence of mobile manipulation tasks with instructions described in natural language (NL). Given a new environment, even a simple task such as of boiling water would be performed quite differently depending on the presence, location and state of the objects. We start by collecting a dataset of task descriptions in free-form natural language and the corresponding grounded task-logs of the tasks performed in an online robot simulator. We then build a library of verb-environment-instructions that represents the possible instructions for each verb in that environment—these may or may not be valid for a different environment and task context.**

We present a model that takes into account the variations in natural language and ambiguities in grounding them to robotic instructions with appropriate environment context and task constraints. Our model also handles incomplete or noisy NL instructions. Our model is based on an energy function that encodes such properties in a form isomorphic to a conditional random field. We evaluate our model on tasks given in a robotic simulator and show that it successfully outperforms the state-of-the-art with 61.8% accuracy. We also demonstrate a grounded robotic instruction sequence on a PR2 robot through Learning from Demonstration approach. [1]

Fig. 1: **Personal household robot setting** where user describes a task in natural language and the robot infers an action plan based on the environment that accomplishes the task.

I. Introduction

For robots working in the presence of humans, it is important for them to be able to ground a task described in natural language for the given environment. In this paper, we present a learning algorithm that maps natural language instruction to a mobile manipulation instruction sequence depending upon the environment context.

We consider a personal household robot setting as shown in Figure 1 where the robot interacts with non-expert users only using natural language and is asked to do different set of tasks. For a given task described in a natural language instruction, the robot must come up with a valid sequence of instructions that accomplishes the task. This is challenging for several reasons. Consider the task of boiling water shown in Figure 2—it consists of a series of steps to be performed that a user describes in natural language (NL). Each step is challenging because there are variations in the natural language, and because the environment context (i.e., the objects and their state) determines how to perform it. For example, to accomplish the instruction *'heating the water'*, one can either use a stove for heating or a microwave (if it is available). For each option, several steps of grasping, placing, turning the stove, etc. would need to be performed. In another example, the NL instruction *'fill the cup with water'* may require the robot to pick up a cup and fill it with water either from the tap (if there is one) or fill it from a fridge water-dispenser or whatever other means are available in the environment. We also note that if the cup already has water then we may need to do nothing. This mapping (or grounding) of the NL instruction into a sequence of mobile manipulation instructions thus varies significantly with the *task constraints* and the *environment context*. In this paper, our goal is to develop an algorithm for learning this grounding.

Different aspects of the language grounding problem have been explored by recent works (Beetz et al., 2011; Bollini et al., 2012; Guadarrama et al., 2013; Marco et al., 2012; Tellex et al., 2011). In particular, Guadarrama et al. (2013) focused on using spatial relations to ground noun-phrases to objects in the environment. They used an injective mapping from verbs to controller instructions based on pre-defined templates. As we show in our experiments, pre-defined templates do not work well with the variations in NL and with changing environment and task context (see Figure 3 for an example). Beetz et al. (2011) and Marco et al. (2012) consider translating web recipe into a robot making pancakes and focus on translating the knowledge into a knowledge reasoning system. However, our problem requires data-driven retrieval of

[1]The first version of this work was presented at Robotics Science and Systems (RSS) conference in July 2014 (Berkeley) (Misra et al., 2014).

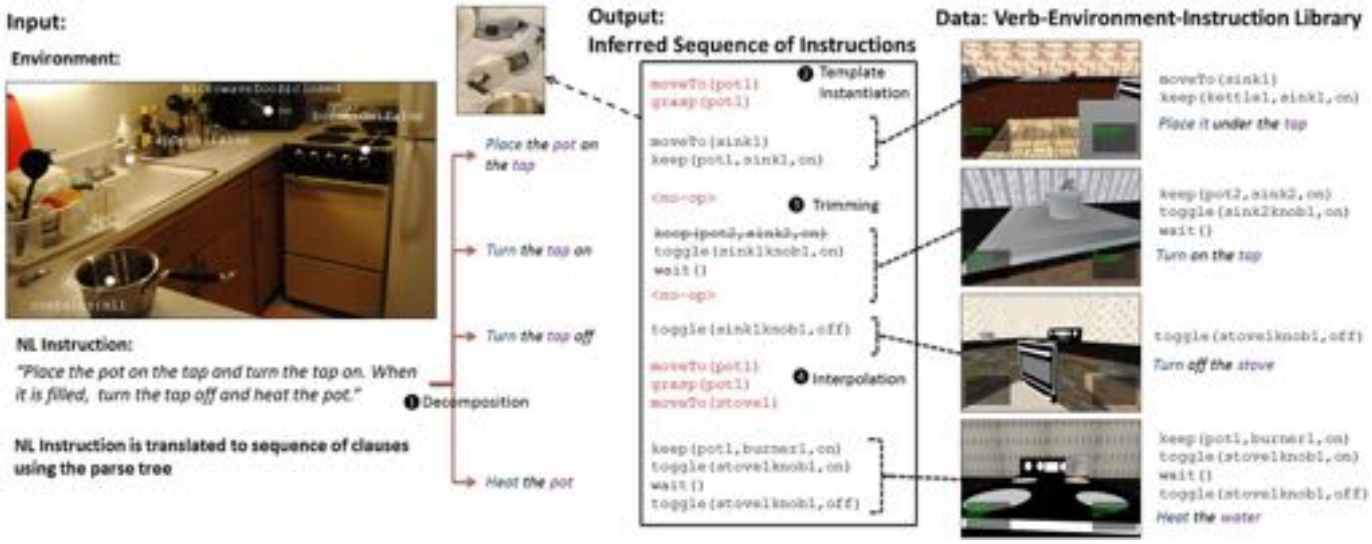

Fig. 2: **Natural Language Instructions to a sequence of controller instructions** for a given new environment. Our approach takes a description in natural language and sequences together robotic instructions that are appropriate for a given environment and task. Note that the NL instructions are often ambiguous or incomplete.

relevant pieces of instructions that are contextually-relevant for that sub-task. Therefore, our work focuses on considering large variations in the NL instructions for generalizing to different tasks in changing environments. Bollini et al. (2012) showed that mapping from natural language to recipe is possible by designing a probabilistic model for mapping NL instructions to robotic instructions, and by designing an appropriate state-action space. They then perform a tree search in the action space to come up with a feasible plan. Since in our problem the search space is very large, their tree search becomes infeasible.

Natural language instructions given in a task and environment context are often incomplete since the remaining information can be inferred from the context. Malmaud et al. (2014) noticed many ambiguous and incomplete instructions in the context of interpreting cooking recipes. This pattern was found in our user study as well. Table I, shows some natural language instructions from our dataset along with the different challenges that exist in grounding them.

One key property of our model is that we can handle missing or *incomplete NL instructions*, for which the robotic instructions have to be inferred from context. For example, the NL instruction *'heat the pot'* does not explicitly say that the pot must be placed on the stove first, and it has to be inferred from the task constraints and the environment context. Furthermore, sometimes one should not follow the NL instructions precisely and instead come up with alternatives that are suitable for the robot to perform in that particular situation.

In this work, we focus on developing a method that models the variations in NL and the ambiguities in grounding them to robotic instructions given an environment context and task constraints. Our model considers the trade-off in following NL instructions as closely as possible while relying on previously-seen contextually-relevant instructions in the training dataset. In detail, we take a data-driven approach where we first collect a database of NL instructions and robotic instructions sequences performed for different tasks in an online simulated game. Using this data, we build a verb-environment-instruction library (VEIL). We then present a machine learning approach that models the relations between the language, environment states and robotic instructions. Our model is isomorphic to a conditional random field (CRF), which encodes various desired properties in the form of potential functions on the edges. With a sampling based inference algorithm, we show that our approach produces valid and meaningful instructions, even when the environment is new or the NL instructions are incomplete and not precisely valid for that environment.

We evaluate our approach on our VEIL dataset for six different tasks. Each task has 5 different environments with free-form natural language instructions and robotic instruction logs, collected from several users. The tasks comprise performing several steps in sequence and there are often different ways of performing the task in different environments. We compare our method against our implementation of Guadarrama et al. (2013) and Bollini et al. (2012), and show significant improvements. More importantly, we find that our method handles generalization to new environments and variations in language well, and is also able to handle incomplete NL instructions in many cases. Finally, we train our PR2 robot using Learning from Demonstration approach for several robotic instructions for several objects. We then test a full predicted sequence on the PR2 robot for making a dessert, following NL instructions given by a user.

In summary, the key contributions of this paper are:

- We encode the environment and task context into an energy function over a CRF which allows grounding of

TABLE I: **Dataset Examples** of natural language descriptions from VEIL-300 along with the difficulty in grounding them. Our dataset shows that users may often provide minimal information in their instructions if the context of the problem is clear.

Task	Natural Language Description	Challenges
Cooking Ramen	*Heat up the water, then cook the ramen*	Ambiguous (How to cook?)
Serving Affogato	*Drop a scoop of ice cream, drizzle syrup on top, then add coffee to the bowl.*	Missing object (Add drizzle on top of what?)
Making Coffee	*Add milk, sugar and coffee to a cup. Next ignite the stove and place the cup on the stove and wait. After 5 minutes close the stove and take out the cup.*	Verb with many arguments
Boiling Water	*Microwave for 20 minutes and place the cup on the desk.*	Missing object with resolution in future. High level verb (microwave)
Serving Affogato	*Microwave the coffee, scoop some ice cream, and drizzle syrup on top.*	Incomplete (adding scoops to coffee not specified)
Cooking Ramen	*Fill the pot with ramen and water. Microwave it.*	Ambiguous (How to fill?). Anaphoric reference (Microwave what?).

the NL instructions into environment for tasks.

- Our model is able to handle missing NL instructions and free-form variations in the language.
- Our method can handle mobile manipulation tasks with long instruction sequence. Our setting has a large state space of the objects, and a large robotic action space.
- We contribute an online data collecting method, and the resulting VEIL dataset comprising free-form natural language instructions and corresponding robot instruction logs. Our experiments show good results on the dataset and our model outperforms the related work.

The paper is organized as follows. We give a summary of related works and their limitations in the context of this problem in Section II. We give an overview of our algorithm and the representation of input - environment and language and the output - controller instruction sequence in Section III. We describe our model in Section IV and the VEIL dataset format that we use for solving the model in Section V. We provide details of the features used by our model and the inference and learning procedure in Section VI. Our crowd-sourcing system that we developed for collecting the VEIL dataset is described in section VII. We describe our experiments and results and give details of our robot experiment in Section VIII and IX respectively. We describe future work and discuss the limitations of our approach in Section IX. We give the concluding remarks in Section X.

II. Related Work

In this section, we first describe related works in the field of semantic parsing, environment representation and mobile manipulation instructions. In this paper, we use several ideas and results from these fields. We then describe related work for the problem of grounding natural language.

Semantic Parsing. Semantic parsing is the problem of representing natural language by a formal representation that preserves the meaning. In spite of success in the area of syntactic parsing, semantic parsing remains an open challenging task. For application in robotics, this problem is further complicated by the presence of an environment context. This often results in incomplete and ambiguous natural language instructions which implicitly use context—environment and task objective, for resolution. As an example, consider the following sentence from our dataset: *"Microwave the coffee, scoop some ice cream, drizzle syrup on top"*. This sentence is incomplete since it does not specify that ice-cream should be added to the coffee after scooping it and that syrup needs to be drizzled on top of the coffee cup. However, these missing instructions are obvious to humans since drizzling syrup on the floor will make little sense.

In order to map NL instructions into their meaning representations, Guadarrama et al. (2013) parse the NL instructions into one of the many manually created templates.This approach becomes unscalable with the increasing complexity of the task.

Recently, learning methods based on *combinatory categorial grammar(CCG)* (Steedman, 1996, 2000) have been used with success for different applications (Zettlemoyer and Collins, 2007). Matuszek et al. (2012a) use probabilistic CCG parsers to convert natural language commands to a Robot Control Language (subset of typed lambda calculi) such as given below:

"exit the room and go left"

(do-sequentially (take-unique-exit) (turn-left))

Their approach does not handle the ambiguities of natural language that can be only be resolved by incorporating environment in the parsing step, and neither can it handle missing actions, incorrect or missing object references. The probabilistic CCG parsers such as UBL (Kwiatkowski et al., 2010) take a manually defined seed lexicon as input. This seed lexicon contains lexical entries which represent mapping between words and formal expressions, which is instruction sequence in our case. The accuracy of these semantic parsers is sensitive to this seed lexicon and providing it requires domain expertise. Yet another challenge is that annotating the NL instructions with their formal expressions in typed-lambda calculi is tedious and requires expert knowledge. In contrast, we show that collecting instruction sequence can be easily crowd-sourced.

In another work, Matuszek et al. (2012b) jointly model language and perception for the task of selecting a subset of objects that are described by natural language. The joint modeling gives significant improvement over separate models for language and perception. However, they do not handle the

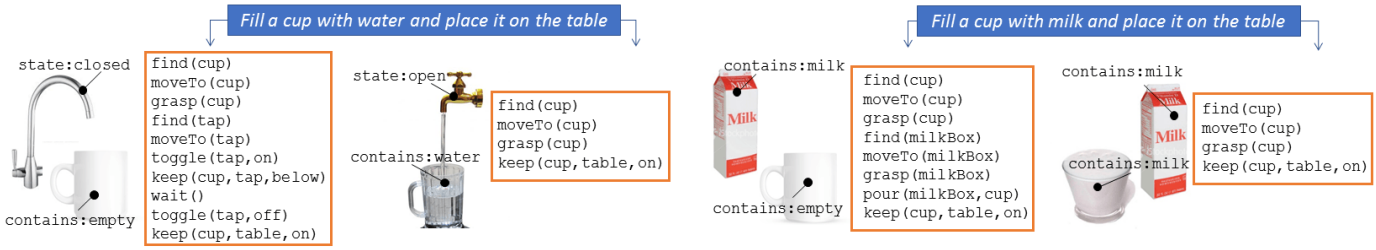

Fig. 3: **Many-many correspondence between language and controller instructions depending on the environment.** The two sentences differ only in one object(water/milk) and they ground to two different or same instruction-sequences depending upon the environment. This shows that the mapping is dependent upon both the objects in the given sentence and the environment. This means that we cannot assume a fixed mapping for sentence template such as *"fill x with y"*.

complex sequence of manipulation tasks that we address in this paper.

Some research have taken the direction to parse natural language into intermediate representations instead of directly parsing into formal semantic expressions. This is also the direction we take in this paper. Tellex et al. (2011) use *tree of Spatial Description Clause (SDC)* for this purpose. Others use *Linear Temporal Logic* to represent the language task which generates robot controllers which can be proven to be correct (Finucane et al., 2010; Kress-Gazit et al., 2007; Wongpiromsarn et al., 2010). However, these works focus on creating formal descriptions and creating controllers, and not on handling ambiguous NL variations or data-driven grounding into the environment and task context.

Environment Representation. There has been a lot of work done in the computer vision community on representing environment. Previous works represent environment as a graph whose nodes represent objects such as cup, microwave, book, television etc. and whose edges represent object-object relationships which can be spatial relations (e.g., *near*, *far*, *left*), *is-part of*, or *is-type of* relationship etc.

Wu et al. (2014) produce a hierarchical labeling of a RGB-D scene using *is-part of* and *is-type of* relationships between objects. Aydemir et al. (2011) use spatial relations between objects to search for a specified object. These works form the ideas behind our environment representation. We use spatial relations between solid objects such as microwave, fridge and is-part of relationship between components of a given object such as fridge-door, fridge-platform, main-fridge-body etc.

Mobile Manipulation Tasks. Previous decade has seen significant work on different manipulation and navigational skills such as grasping (Kroemer et al., 2010; Lenz et al., 2013), mixing (Bollini et al., 2012), pushing (Srinivasa et al., 2010), placing (Barry et al., 2013; Jiang et al., 2012), constructing semantic maps (Walter et al., 2013), and high degree of freedom arm planners (e.g., (Alterovitz et al., 2011; Ratliff et al., 2009)). These works form the building blocks for executing the output instructions for our model. On the other hand, rather than building specific manipulation or navigation primitives, Learning from Demonstration (LfD) approach allows even non-experts to train the robot by using the arms of the robot (Argall et al., 2009). For actual robotic experiment on a PR2 robot, since our goal is to verify the applicability of grounded robotic instruction sequences on a real robot, we take LfD-based approach where we train the robot via our teleoperation method among many other LfD alternatives.

Traditionally, symbolic planners (Rintanen, 2012) have been used to accomplish sequencing complicated controller instructions. Since real environments have uncertainty and non-determinism, Kaelbling and Lozano-Pérez (2011) start with an abstract plan and recursively generate plans as needed. Or, the tasks are defined through expert designed state machines (Nguyen et al., 2013), which does not generalize well when the environment or the task changes. Rather than relying on symbolic representation of the environment, Sung et al. (2014) rely on a set of visual attributes to represent each object in the environment and dynamically choose the controller sequence from a list of possible sequences that minimizes the score function based on the current environment and the potential candidate for the next instruction. Others use demonstrations for learning different behaviors (Niekum et al., 2013). These approaches solve only parts of the problem that we address in this work—of creating valid plans and using a score function for data-driven retrieval of instruction sequence. Our work addresses not only the validity of instruction sequences and data-driven retrieval of low-level instructions, but it also models the ambiguity and grounding of natural language instructions in the environment context. Furthermore, the tasks considered by our work are complex manipulation tasks requiring long instruction sequences.

Grounding Natural Language. Several recent works in robotics have looked at the problem of grounding natural language. This has been driven by advances in HRI, better vision algorithms and more reliable manipulation instructions. Other than the works discussed in the introduction (Beetz et al., 2011; Bollini et al., 2012; Guadarrama et al., 2013; Marco et al., 2012), the problem of navigation has been addressed by using learned models for verbs such as *follow*, *meet*, *go*, and dynamic spatial relations such as *walk close to the wall* (Fasola and Matarić, 2013; Kollar et al., 2010). To detail, Kollar et al. (2010) use a maximum-likelihood approach to infer the path taken by the robot. Translation of such weakly

specified actions into robotic behaviors is very important and form the ideas for our robotic instruction set in Table II. Artzi and Zettlemoyer (2013) focus on mapping natural language instructions representing navigational tasks, to a sequence of instructions for a 2D world scenario. In contrast to these works, the chief linguistic object of our investigation are high level verbs such as *serve, cook, throw* which have more complex representation than simple navigational verbs. In contrast to Artzi and Zettlemoyer (2013), we also consider more complex 3D scenarios with a larger set of manipulation instructions.

Several works (Fasola and Matarić, 2013; Guadarrama et al., 2013) have looked at the problem of grounding intricate noun-phrases in the language to the objects in the environment. Guadarrama et al. (2014) ground open vocabulary descriptions of objects to the described objects in the environment. There have also been success in grounding concepts (Chao et al., 2011) and objects (Lemaignan et al., 2012; Ros et al., 2010) through human-robot interaction. Works like Chu et al. (2013) look at mapping from text to haptic signals. Kulick et al. (2013) consider active learning for teaching robot to ground relational symbols. The inverse problem of generating natural language queries/commentaries has also seen increasing interest in robotics (Chen et al., 2010; Tellex et al., 2013). In a recent work, Duvallet et al. (2014) explored the direction of using natural language as a sensor to come up with prior distribution over unseen regions in the environment.

In the area of computer vision, some works have considered relating phrases and attributes to images and videos (Farhadi et al., 2010a; Jiang et al., 2013; Koppula and Saxena, 2013; Koppula et al., 2011; Ramanathan et al., 2013; Wu et al., 2014). These works focus primarily on labeling the image/video by modeling the rich perceptual data rather than modeling the relations in the language and the entities in the environment. Thus, our work complements these works.

In NLP community a lot of literature exists on parsing natural language sentences (e.g., (Klein and Manning, 2003)) and grounding text in different domains such as linking events in a news archive and mapping language to database queries (Berant et al., 2013; Nothman et al., 2012; Poon, 2013; Yu and Siskind, 2013). These techniques form the basis of ours in syntactic parsing and representation. However most of these works use only text data, and do not address grounding the physical task or the environment. In another work, Branavan et al. (2010) consider mapping natural language queries to sequence of GUI action commands for the windows operating system. They also consider incomplete instructions and learn environment model but they do not consider the real world 3D environments that robots encounter nor the complex controller instructions that a robot can perform.

III. Overview

In this section, we give an overview of our approach and the representation we use for natural language, environment, mobile manipulation instructions and the domain knowledge.

Given an environment E containing objects and the robot, a human gives the instructions for performing a task in natural language (see Figure 2). The instructions in natural language L consist of a set of sentences, and our goal is to output a sequence of controller instructions $\mathcal{I}$ that the robot can execute. Each of these low-level robot instructions often have arguments, e.g., $\texttt{grasp}(\texttt{object}_\texttt{i})$.

$$E, L \rightarrow \mathcal{I}$$

This mapping is hard to learn for two reasons: (a) the output space $\mathcal{I}$ is extremely large, and (b) the mapping changes significantly depending on the task context and the environment.

For a given E and L, certain instructions $\mathcal{I}$ are more likely than others, and we capture this likelihood by using an energy function $\Gamma(\mathcal{I}|E, L)$. We will use this energy function to encode desired properties and constraints in our problem. Once having learned this energy function, for a given new language and environment, we can simply minimize this energy function to obtain the optimal sequence of instructions:

$$\mathcal{I}^* = \text{argmin}_{\mathcal{I}}\ \Gamma(\mathcal{I}|E, L)$$

There are several steps that we need to take in order to define this energy function: we need to convert the language L into a set of verb clauses $\mathcal{C}$, we need to represent the environment E with a usable representation that contains information about the objects, we need to describe what are the low-level instructions $\mathcal{I}$ and how do they connect to the actual execution on the robot, and we need to figure out how to represent and store the training data for their use in inference.

A. Language Representations by a Set of Verb Clauses

We follow previous work (Tellex et al., 2011) in parsing natural language commands into a structured intermediate representation. These structured representations only contain information relevant to the manipulation tasks and are easier to work with. We use a sequence of *verb clauses* as the representation of natural language which are ordered temporally. Each verb clause informally represents an atomic natural language command dictated by the main verb; more formally, a verb clause $\mathcal{C}$ is a tuple:

$$\mathcal{C} = (\nu, [obj], \rho)$$

containing the verb ν, the set of *language-objects* $[obj]$ on which it acts and a relationship matrix $\rho : obj \times obj \rightarrow Rel$ where Rel is the space of relationship (e.g., 'with', 'from'). For example, the following natural language command comprises of four verb clauses:

$$\underbrace{\textit{Take the cup with water and}}_{(\texttt{take},[\texttt{cup},\texttt{water}],\texttt{with:cup}\rightarrow\texttt{water})}\ \underbrace{\textit{then ignite the stove.}}_{(\texttt{ignite},[\texttt{stove}],\emptyset)}$$
$$\underbrace{\textit{Now place the cup on the stove}}_{(\texttt{place},[\texttt{cup},\texttt{stove}],\texttt{on:cup}\rightarrow\texttt{stove})}\ \underbrace{\textit{and wait.}}_{(\texttt{wait},\emptyset,\emptyset)}$$

In order to convert unconstrained natural language L to a sequence of verb-clauses $\{\mathcal{C}\}$ we use the following steps:

TABLE II: **List of Low-level Instructions** that could be executed by the robot. Each instruction is parametrized by the required objects. We implement a subset of these instructions on a PR2 robot (see Section IX).

Instruction	Description
find(obj)	Find obj in the environment (Anand et al., 2012).
keep(obj1, obj2, R)	Keeps obj1 with respect to obj2 such that relation R holds (Jiang et al., 2012).
grasp(obj)	Grasp obj (Lenz et al., 2013).
release(obj)	Releases obj by opening gripper.
moveTo(obj)	Move to obj by using motion planner (Ratliff et al., 2009).
press(obj)	Presses obj using end effector force controller (Bollini et al., 2011).
turn(obj, angle)	Turns the obj by a certain angle.
open(obj)	Opens the obj (Endres et al., 2013).
close(obj)	Closes the obj (Endres et al., 2013).
scoopFrom(obj1, obj2)	Takes scoop from obj2 into obj1.
scoopTo(obj1, obj2)	Puts the scoop from obj1 into obj2.
squeeze(obj1, obj2, rel)	Squeezes obj1 satisfying relation rel with respect to obj2.
wait()	Wait for some particular time.

- First, the constituency parse tree of L is found using the Stanford Parser (Klein and Manning, 2003).
- A verb clause $\mathcal{C}$ is created for each corresponding node of verb type in the parse tree. The clauses are ordered by the order of their respective verb in L. The verbs are also stemmed before being initialized as ν.
- Next, we compute the set of all *maximal nouns*, which are noun-type nodes in the parse tree whose parents are not noun-type nodes e.g., *"get the red cup"* has the node [NP:*the red cup*] as the only maximal noun.
- These *maximal nouns* are attached to the clause whose verb node is closest in the parse tree. In the ambiguous case, the nodes get attached to the nearest immediate left clause.
- Finally, all relationship-type nodes between *maximal nouns* of the same clause are inserted in the corresponding relationship matrix ρ.

Note that *language-objects* are not same as physical objects in an environment and not all *language-objects* (e.g., water) represent an actual physical object.

B. Object and Environment Representation Using Graphs

For performing a task in the environment, a robot would need to know not only the physical location of the objects (e.g., their 3D coordinates and orientation), but also their functions. These functions are often represented as symbolic attributes (Anand et al., 2012; Farhadi et al., 2010b) or using a more functional representation (Cakmak et al., 2007; Höfer et al., 2014; Tenorth et al., 2010). For example, a microwave consists of four parts: main body, door, buttons and display screen. Each part has its own states (e.g., a door could be open/closed), and sometimes its state is affected by another part, e.g., a button when pressed could unlatch the microwave door. We represent each object as a directed-graph G where the vertices are parts of the object (also storing their states), and edges represent the functional dependency between two object parts. Now for a given environment, we define E to store aforementioned representation of every object in the environment along with spatial relations between objects. Here we consider five basic spatial relations: Grasping, Inside, On, Below and Near.

TABLE III: **List of types of STRIPS predicate** used to represent precondition and effect of controller instruction primitives which are listed in Table II.

STRIPS Predicate	Truth Condition
state obj st	Boolean state st of object obj has value $True$
affordance_type obj	object obj has the given affordance
relationship_type obj1 obj2	Given relationship exists between object $obj1$ and $obj2$

C. Representing Robotic Instructions

We have defined a set of low-level instructions that could be executed by the robot in Table II. Each controller instruction is specified by its name and its parameters (if any). For example, moveTo(obj) is an instruction which tells the robot to move close to the specified object obj. We ground the parameters in objects instead of operation-space constants such as '2m North' because it is objects that are relevant to the task (Cakmak et al., 2007).

In order to completely execute a sub-task such as *'keeping a cup on stove'*, we need a sequence of instructions $\mathcal{I}$. An example of such a sequence for the sub-task might look as:

find(cup1); moveTo(cup1); grasp(cup1); find(stove1); moveTo(stove1); keep(cup1, stoveBurner3, on)

Note that the space of possible sequence of instructions $\mathcal{I}$ is extremely large, because of the number of possible permutations of the instructions and the arguments.

D. Representing Domain Knowledge

In this paper, the domain knowledge for a given set of environment states and a set of controller instruction primitives refer to the set of preconditions under which an instruction can be executed in an environment and its effect on the environment. In a more general setting, it can also include post-conditions and transition probabilities between two environments. The domain knowledge is provided as input to our algorithm.

We use the STRIPS formal language (Fikes and Nilsson, 1972) to write the preconditions and effect of each instruction listed in Table II. To make this representation scalable for a large set of objects, we consider three types of predicates as listed in Table III. We consider common affordances such as- *graspable, turnable, pourable etc.* and five basic spatial relationships as described earlier. Our set of object affordance is inspired by previous work (Kjellström et al., 2011; Koppula et al., 2013; Montesano et al., 2008).

This domain knowledge encodes high-level symbolic information about the world such as—a cup is filled with

water if kept below an open tap. For example, this information can be encoded by the following STRIPS function:

*(**:action** keep **:parameters** (obj1 obj2 R)*
***:precondition** (and (grasping Robot obj1) (near Robot obj2))*
***:effect** (when (and (= obj2 sink) (= R on) (containable obj1))*
(and (when (state sink TapIsOn) (state obj1 Water))
(on obj1 sink)))

This function describes an instruction primitive *keep* which accepts three arguments, object $obj1, obj2$ and relationship R. The second line defines the precondition for this instruction, which evaluates to true only when the robot is grasping the object $obj1$ and is near the object $obj2$. The third line defines the effect of this instruction, which says that when the object $obj2$ is *sink*, the relationship R is *on* and the object $obj1$ has the *containable* affordance then fill the object $obj1$ with water if the sink is open and also place the object $obj1$ on the sink.

Encoding this information for an environment is very tedious and there is a large body of literature on learning them from observations (Mourao et al., 2012). While we use this domain knowledge in our main model we also show that in the absence of this knowledge, the accuracy of our algorithm does not reduce significantly.

IV. Model

In this section we describe our energy function $\Gamma(\mathcal{I}|L, E)$, which is isomorphic to a conditional random field (Fig. 4) and is comprised of several nodes and factors (ψ).

It has the following observed nodes: natural language instruction L which is deterministically decomposed into a sequence of verb clauses $\mathcal{C}$ and the initial environment E'_0. As the instructions are executed, the state of the environment changes, and we use E_i, E'_i to represent the environment-state at step i. Our goal is to infer the sequence of controller instructions $\mathcal{I} = \{\mathcal{I}_{i\ell}, \mathcal{I}_i\}_{i=1}^k$.

Note that for each clause $\mathcal{C}_i$, we have two instruction sequences $\mathcal{I}_i$ and $\mathcal{I}_{i\ell}$—this is because natural language instructions are often incomplete and $\mathcal{I}_{i\ell}$ represents this missing instruction sequence. For example, if the natural language instruction is *place cup in microwave*, then one may need to open the microwave door first ($\mathcal{I}_{i\ell}$) and then place the cup within ($\mathcal{I}_i$).

Following the independence assumptions encoded in the graphical model in Figure 4, the energy function is written as a sum of the factor functions:

$$\Gamma(\mathcal{I}|E, \mathcal{C}) = \sum_{i=1}^{k} \psi_i(\mathcal{I}_i, \mathcal{C}_i, E_i) + \psi'_i(\mathcal{I}_i, \mathcal{C}_i, E'_i) + \psi_{i\ell}(\mathcal{I}_{i\ell}, E'_{i-1}) + \psi'_{i\ell}(\mathcal{I}_{i\ell}, E_i)$$

We encode several desired properties into these factor functions (we describe these in detail in Section VI):

- The sequence of instructions should be *valid*. Intuitively $\psi_i(\cdot)$ represents the pre-condition satisfaction score and $\psi'_i(\cdot)$ represents the post-condition satisfaction score.

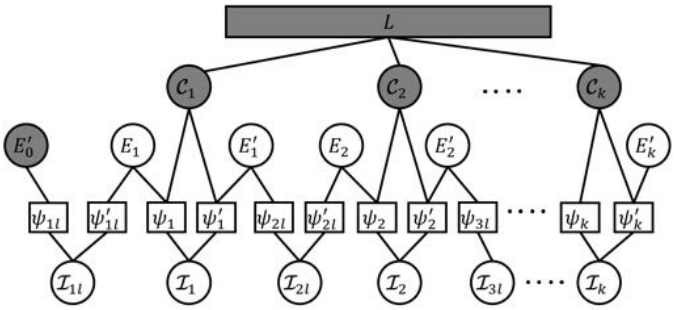

Fig. 4: **Graphical model representation of our energy function.** The clauses $\mathcal{C}$ and initial environment E'_0 is given, and we have to infer the instructions $\mathcal{I}_i, \mathcal{I}_{i\ell}$ and the environment at other time steps. The ψ's are the factors in the energy function.

- The output instructions should follow the natural language instructions as closely as possible. Thus, $\psi_i(\cdot)$ and $\psi'_i(\cdot)$ depend on the clause $\mathcal{C}_i$.
- Length of the instruction set. Shorter instruction sequences are preferred for doing the same task.
- Prior probabilities. An instruction sequence with too many unlikely instructions is undesirable.
- Ability to handle missing natural language instructions. The $\psi_{i\ell}(\cdot)$ and $\psi'_{i\ell}(\cdot)$ represent the potential function for the missing instruction set, and they do not depend on any clause.

For each verb clause $\mathcal{C}_i$, the first step is to sample a few sequence of instructions $\mathcal{I}_i^{(s)}$ from the training data. With the sample of instructions for each clause obtained, we will then run our inference procedure to minimize the energy function to give a complete set of instructions satisfying the aforementioned properties. We first describe how we obtain these samples by building a verb-environment-instruction library (VEIL) from the training data. This library plays the role of a lexicon in our grounding approach.

V. Verb-Environment-Instruction Library (VEIL)

Space of all possible instruction sequence is countably infinite, and even when the length of sequence is limited and invalid instruction sequences are pruned, the sample space still remains very large. Moreover the instruction sequence structure is often shared between different tasks such as for example between the NL commands *"fill a cup with coke"* and *"fill a mug with beer"*. This motivates the sampling based approach that we take.

For a given verb clause $\mathcal{C}_i$, we use the training data $\mathcal{D}$ to come up with a sample of the instruction-templates $\{\mathcal{I}_i^{(s)}\}$. During training we collect a large set of verb clauses $\mathcal{C}^{(j)}$, the corresponding instruction-sequence $\mathcal{I}^{(j)}$ and the environment $E^{(j)}$. The parameters of the instructions in the training dataset are specific values, such as `cup01`, and that particular object may not be available in the new test environment. Therefore, we replace the parameters in the training instructions with generalized variables Z. For example, {`moveTo(cup01)`; `grasp(bottle01)`} would be stored as

generalized instruction sequence {moveTo(z_1); grasp(z_2)}. The original mapping $\{z_1 \rightarrow \texttt{cup01}, z_2 \rightarrow \texttt{bottle01}\}$ is stored in $\xi^{(j)}$.

In order to come up with proposal templates for a given clause $\mathcal{C}_i$, we return all entries containing the corresponding verb clause, environment, instruction and the grounding: $\mathcal{D}^{(i)} = \{(\mathcal{C}^{(j)}, E^{(j)}, \mathcal{I}^{(j)}, \xi^{(j)}) \,|\, \nu(\mathcal{C}^{(j)}) = \nu(\mathcal{C}_i)\}$ for $j = 1 \ldots |\mathcal{D}|$ in the dataset. The information in the particular s^{th} sample is represented as $D_s = (\mathcal{C}_s, E_s, \mathcal{I}_s, \xi_s)$. Note that we capture both the language context C and the environment context E in the structure of these samples.

A. Instantiation Algorithm

Since we only store generalized instructions (with actual objects replaced by general variables Z), we need a way to instantiate the generalized instructions for a new environment and language context.

Given the sth instruction-template $D_s = (\mathcal{C}_s, E_s, I_s, \xi_s)$, a clause $\mathcal{C}_j$ and a new environment E_j, the aim of the instantiation algorithm is to return the instantiated instruction template $\mathcal{I}_j$. In order to accomplish this task, we should find a grounding of the generalized variables $Z(s)$ to specific objects in the environment E_j. In order to do so, we first define the function $Ground(a, E)$ which takes a language-object a such as '*mug*','*stove*' etc. and an environment E and returns the object in E which is being described by the language-object. We use a simple syntax based similarity between language-object a and categories of the environment objects to perform this grounding e.g., the language object *'milk'* has high syntactical similarity with the object $milkBox$ but not with mug. If no suitable match is found, such as for non-physical language-objects like water or objects which are mentioned in the language but not actually present, the function returns $Null$.

We then take the following rule based approach:

1) We first map to the objects that have a relation. For example, if we have `cup on table` in the clause and the template contains $\texttt{z}_1$ `on` $\texttt{z}_2$, then we map z_1, z_2 to the grounding of `cup` and `table` respectively. More formally, $\forall z_1, z_2 \in obj(\mathcal{C}_s)$, $\forall a, b \in obj(\mathcal{C}_j)$ such that $\rho(\mathcal{C}_s)[z_1, z_2] = \rho(\mathcal{C}_j)[a, b]$ we map variable $z_1 \mapsto Ground(a, E_j)$ and $z_2 \mapsto Ground(b, E_j)$.
2) For an unmapped variable z, we map it in the same way as it was mapped in the sample s. More formally, $\forall z_1 \in obj(\mathcal{C}_s), a \in obj(\mathcal{C}_j)$ such that the object $\xi_s(z_1)$ has the same category as $Ground(a, E_j)$, we map $z_1 \mapsto Ground(a, E_j)$.
3) For every remaining unmapped variable z in $Z(s)$, we map it to the object in E_j that has the most common state-value pairs with the object $\xi_s(z)$.

Each rule only maps the variables which are not already mapped to an object. Also the mapping $z \mapsto Ground(a, E)$ is performed only when $Ground(a, E)$ is not $Null$. Note that the last rule will always map every remaining unmapped variable.

The new mapping is stored as ξ and the instructions returned after replacing the variables using ξ is given by $\mathcal{I}_j$. We further define the predicate $replace$ such that $replace(\mathcal{I}_s, Z(s), \xi)$ will return the instantiated instruction sequence $\mathcal{I}_j$ after replacing all variables $z \in Z(s)$ with the object $\xi(z)$.

VI. Energy Function and Inference

Now for each clause $\mathcal{C}_i$, we have obtained a sample instruction set $\mathcal{I}_i^{(s)}$ (together with the clause and environment data in $\mathcal{D}_s$). We now need to come up with a full sequence $\{\mathcal{I}_{i\ell}^*, \mathcal{I}_i^*\}$ based on these initial samples. Note that we only have samples for $\mathcal{I}_i$ and not for $\mathcal{I}_{i\ell}$, which are for handling the missing natural language instructions. Furthermore, the sample instruction set is not consistent and valid, and also does not apply directly to the current environment. Based on these samples, we need to optimize and infer a sequence that minimizes the energy function.

In the following subsections, we now describe the different terms in the factor functions that encode the desired properties aforementioned.

A. Term $\psi_i(\mathcal{I}_i, \mathcal{C}_i, E_i; w)$

This term consists of features that encapsulate properties such as pre-condition score, instruction length, prior probabilities, etc. For a given sample $D_s = (\mathcal{C}_s, E_s, \mathcal{I}_s, \xi_s)$, we define the cost of setting $\mathcal{I}_i := replace(\mathcal{I}_s, Z, \xi)$ as:

$$\psi_i(\mathcal{I}_i, \mathcal{C}_i, E_i | D_s, \xi; w) = w^T \begin{bmatrix} \Delta_{env}(D_s, \xi); & \Delta_{nl}(\mathcal{C}_s, \mathcal{C}_i); & \Delta_{sim}(Z(s), \xi_s, \xi); \\ \Delta_{pcc}(\mathcal{C}_i, \mathcal{C}_s); & \Delta_{jmp}(\mathcal{I}_i, E_i); & \Delta_{dscpl}(\mathcal{I}_i); \\ \Delta_{inpr}(\mathcal{I}_i); & \Delta_{para}(\mathcal{I}_i); & \Delta_{trim}(\mathcal{I}_i, \mathcal{I}_s) \end{bmatrix}$$

We describe each term in the following:

1) Environmental Distance $\Delta_{env}(D_s, \xi)$: It is more likely to have instructions that were made in similar environments in the training data as compared to the test environment. Hence, if a variable $z \in Z(s)$ is mapped to a `cup` in the new environment and was mapped to a `pot` in the original environment, then we prefer the template D_s if `cup` and `pot` have similar state-value pairs (e.g., both are empty). We encode it as the average difference between states of objects $\xi_s(z)$ and $\xi(z)$ and for all $z \in Z(s)$. We represent the union of the states of $\xi_s(z)$ and $\xi(z)$ by $T(z)$.

$$\Delta_{env}(D_s, \xi) = \frac{1}{|Z(s)|} \sum_{z \in Z(s)} \frac{1}{|T(z)|} \sum_{t \in T(z)} \mathbf{1}(\xi(z)[t] \neq \xi_s(z)[t])$$

2) Natural Language Similarity $\Delta_{nl}(\mathcal{C}_s, \mathcal{C}_i)$: We prefer to use those instruction templates whose verb clause was similar to the test case. Therefore, we measure the unordered similarity between the language-objects of $\mathcal{C}_s$ and $\mathcal{C}_i$ by computing their Jaccard index:

$$\Delta_{nl}(\mathcal{C}_s, \mathcal{C}_i) = \frac{|obj(\mathcal{C}_s) \cap obj(\mathcal{C}_i)|}{|obj(\mathcal{C}_s) \cup obj(\mathcal{C}_i)|}$$

Algorithm 1: We use dynamic programming to minimize the energy function

global $\mathcal{D}, \mathcal{C}, \alpha$
function **Forward-Inference**(j)
for *each* E'_{j-1} *such that* $\alpha_{j-1}[E'_{j-1}]$ *is defined* **do**
  for $D_s \in \mathcal{D}^{(j)}$ **do** *// iterate over all VEIL templates with the same verb as* $\nu(\mathcal{C})$
    $\mathcal{I} \leftarrow instantiate(\mathcal{I}_s, \mathcal{C}_s, \mathcal{C}_j, E'_{j-1})$ *// instantiate the VEIL template using the test clause and the environment*
    for $t \in [0..|\mathcal{I}|]$ **do**
      $\mathcal{I}_j = \mathcal{I}[t \cdots]$ *// trim the instantiated sequence as a check for noise*
      $cstr = getContraints(\mathcal{I}_j, E'_{j-1})$ *// find constraints needed to execute* $\mathcal{I}_t$ *by looking at STRIPS preconditions*
      $\mathcal{I}_{j\ell} = callSymbolicPlanner(E'_{j-1}, cstr)$ */* Call the planner to find an instruction sequence such that the resultant environment* E_j *satisfies the constraints */*
      $E_j = \Phi(E'_{j-1}, \mathcal{I}_{j\ell})$
      $E'_j = \Phi(E_j, \mathcal{I}_j)$
      $\alpha_j[E'_j] = \min\{\alpha_j[E'_j], \alpha_{j-1}[E'_{j-1}] + \psi_j(\mathcal{I}_j, \mathcal{C}_j, E_j) + \psi'_j(\mathcal{I}_j, \mathcal{C}_j, E'_j) + \psi_{j\ell}(\mathcal{I}_{j\ell}, E'_{j-1}) + \psi'_{j\ell}(\mathcal{I}_{j\ell}, E_j)\}$

3) Parameter Similarity Cost $\Delta_{sim}(Z(s), \xi_s, \xi)$*:* We want the objects in the instantiated sequence to be similar (i.e. have same category) to the one in the training set. Therefore, we define:

$$\Delta_{sim}(Z(s), \xi_s, \xi) = \frac{1}{|Z(s)|} \sum_{z \in Z(s)} \mathbf{1}\{\xi_s(z) \neq \xi(z)\}$$

4) Parameter Cardinality Cost $\Delta_{cvr}(\mathcal{C}_i, \mathcal{C}_s)$*:* Different clauses with the same verb can have different number of *language-objects*. For example, the sentences *'add ramen to the crockpot'* and *'add milk and sugar to the mug'* have different number of language objects (2 and 3 resp.). We thus, prefer to use the template which has the same number of language objects as the given clause.

$$\Delta_{pcc}(\mathcal{C}_i, \mathcal{C}_s) = \mathbf{1}(|obj(C_i)| \neq |obj(C_s)|)$$

5) Jump Distance $\Delta_{jmp}(\mathcal{I}_i, E_i)$*:* The jump distance is a boolean feature which is 0 if program $\mathcal{I}_i$ can be executed by the robot given the starting environment E_i and 1 otherwise.

6) Description Length $\Delta_{dscpl}(\mathcal{I}_i)$*:* Other things remaining constant, we believe a smaller sequence is preferred. Therefore, we compute the sum of norms of each instruction in the sequence $\mathcal{I}_i$, where we define norm of an instruction I as the number of parameters that I takes plus 1.

7) Instruction Prior $\Delta_{inpr}(\mathcal{I}_i)$*:* We want our algorithm to give preference to instruction sequence which are more likely to occur. This is particularly useful while dealing with ambiguous or incomplete sentences. For example, the instruction `keep`(`book`, `shelf`, `on`) has higher prior probability than `keep`(`book`, `fridge`, `inside`), although the later is not restricted. Thus, if we have an incomplete sentence *"keep the book"* then the grounded instruction sequence will be preferred if its keeping the book on a shelf rather than inside a fridge. We therefore add the average of prior probability $prior(I)$ of every instruction I in the sequence. We compute it by counting the number of times it appears in the training set:

$$\Delta_{inpr}(\mathcal{I}_i) = \frac{1}{||\mathcal{I}_i||} \sum_{I \in \mathcal{I}_i} prior(I)$$

8) Parameter Instruction Prior $\Delta_{para}(\mathcal{I}_i)$*:* Here we add the prior probability of how often a parameter (e.g., `fridge`) is used for a particular instruction (e.g., `grasp`). We compute it from the training data. Verb-Correlation score is the average taken over all instructions in $\mathcal{I}_i$ of the probabilities that the grounded parameter appears in the instruction at the specified position.

$$\Delta_{para}(\mathcal{I}_i) = \frac{1}{|\mathcal{I}_i|} \sum_{I \in \mathcal{I}_i} \frac{1}{|d_I|} \sum_{j \in [1, |d_I|]} P(n_I, d_I[j], j)$$

Here $P(n_I, d_I[j], j)$ is the prior probability that the parameter $d_I[j]$ appears at position j for the instruction primitive n_I.

9) Trimming Cost $\Delta_{trim}(\mathcal{I}_i, \mathcal{I}_s)$*:* Often we do not use the full sequence of instructions from the set D_s but trim them a little bit. This removes irrelevant instruction belonging to neighbouring clauses in the training corpus. We define this trimming cost: $\Delta_{trim} = (|\mathcal{I}_s| - |\mathcal{I}_i|)^2$.

B. Term $\psi'_i(\mathcal{I}_i, \mathcal{C}_i, E'_i; w)$

This term consists of a single consistency term, given as:

$$\psi'_i(\mathcal{I}_i, \mathcal{C}_i, E'_i; w) = w_{cons} \Delta_{cons}(E'_i, \mathcal{C}_i)$$

The purpose of this consistency score $\Delta_{cons}(E'_i, \mathcal{C}_i)$ is to capture the fact that at the end of execution of $\mathcal{I}_i$, the resultant environment E'_i should have fulfilled the semantics of the clause $\mathcal{C}_i$. Thus if the clause intends to *ignite the stove* then the stove should be in *on* state in the environment E'_i. For this we compute the probability table $P(obj, s, v)$ from training data using only those datapoints that have the same verb as $\mathcal{C}_i$, which gives the probability that obj in clause $\mathcal{C}_i$ can have state s with value v. We use this table to find the average probability that objects of clause C'_i have the given end state values in E'_i.

C. Term $\psi_{i\ell}(\mathcal{I}_{i\ell}, E'_{i-1}; w)$

This term is for the instruction sequence $\mathcal{I}_\ell$ that does not correspond to a NL instruction—i.e., its purpose is to handle

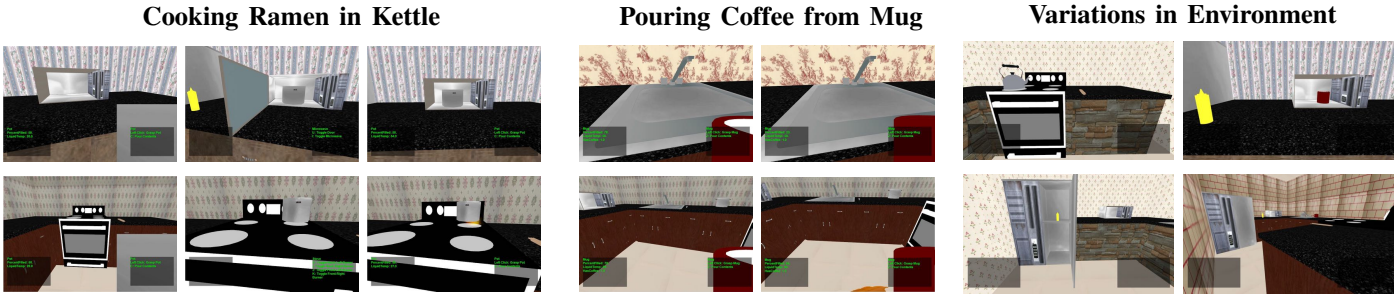

Fig. 5: **[Left] Different ways of performing the same task:** Users have multiple options to do a task, example ramen can be cooked using a microwave or a stove. Thus, algorithm should closely follow the sentence specification and not simply work with high level objective of making ramen. **[Center] Unrestrictive actions:** Actions are unrestricted and can be done even when they have no semantic meaning. Example, coffee can be poured in a sink or on floor. **[Right] Environment Variation:** Shows variations in an environment created by modifying object set, state values of objects and object-object relationship.

.

missing NL instructions. Therefore, it consists of a subset of features used in ψ_i:

$$\psi_{i\ell}(\mathcal{I}_{i\ell}, E'_{i-1}; w) = w^T \big[\Delta_{jmp}(\mathcal{I}_{i\ell}, E'_i); \quad \Delta_{dscpl}(\mathcal{I}_{i\ell}); \quad \Delta_{inpr}(\mathcal{I}_{i\ell}); \quad \Delta_{para}(\mathcal{I}_{i\ell})\big]$$

D. Term $\psi'_{il}(\mathcal{I}_{il}, E_i; w)$

This consists of a single consistency term:

$$\psi'_i(\mathcal{I}_{il}, E_i; w) = w_{cons,l}\Delta_{cons,l}(E'_i)$$

This consistency term is defined similarly to Δ_{cons} however since we do not have a given clause, we therefore build the table $P(obj, s, v)$ using all the datapoints in the training dataset. This term prevents the robot from performing an action which takes it into an unlikely environment.

E. Inference Procedure

Given the training dataset $\mathcal{D}$, a sequence of clauses $\{C\}$ and the starting environment E'_0, the goal is to find the instruction sequence that minimizes the energy function. Since the structure of our model is a chain, we use an approach similar to the forward-backward inference to obtain the $\{\mathcal{I}^*_{i\ell}, \mathcal{I}^*_i\}_{i=1}^k$.

The inference procedures computes the forward variable $\alpha_j[E'_j]$ which stores the cost of the minimum-cost assignment to the node $\{\mathcal{I}_{il}, \mathcal{I}_i\}_{i\leq j}$ such that the environment at chain depth j is E'_j. As a base case we have $\alpha_0[E'_0] = 0$.

We also define the environment simulator Φ as taking an environment E and an instruction sequence $\mathcal{I}$ and outputting the final environment $\Phi(E, \mathcal{I})$ after simulating the execution of the sequence in the given environment.

Our algorithm calls the function **Forward-Inference**(j) (Algo. 1) to compute α_j given $\alpha_i \forall_{i<j}$. To do so, the algorithm iterates over all samples $\mathcal{D}^{(j)}$ for clause C_j which are created as described in Sec. V. For the case when the verb $\nu(C_j)$ was unseen in the training data, we define $\alpha_j = \alpha_{j-1}$.

Each sample D_s is instantiated as described in Sec. V-A which gives us the instruction sequence $\mathcal{I}$. Instantiation is followed by all possible trimming of $\mathcal{I}$ giving us $\mathcal{I}_j = \mathcal{I}[t \cdots]$ for all t. We note here that the *no-op* is also considered when $\mathcal{I}$ is totally trimmed.

The trimmed sequence $\mathcal{I}_j$ may not be executable due to some missing instructions. To handle this, the sub-routine *getConstraints* looks at the pre-condition of every instruction in $\mathcal{I}_j$ to find the hard-constraints required for executing $\mathcal{I}_j$. These constraints along with the environment are passed onto a symbolic-planner (Rintanen, 2012) which outputs the missing instructions $\mathcal{I}_{j\ell}$. The cost of the new assignment $\mathcal{I}_{j\ell}, \mathcal{I}_j$ is used to update the value of α_j as shown in line 13.

Once the $\alpha_j[E]$ has been computed $\forall_j$ and all reachable E, the optimum assignment is computed by backward-traversal.

F. Learning Method

For training in our experiments, we divide the dataset into k-folds for cross-validation and use stochastic gradient descent to learn the weights. While doing experiments, we manually fixed the weight $w_{jmp} = \infty$ since we do not want the inferred sequence to be unexecutable. The remaining weights were initialized to 0.

VII. Robot Simulator for Data Collection.

To train our algorithm, we need a dataset in the VEIL format as described in Section V. This is challenging for two reasons: firstly, it is difficult to find natural language commands paired with environment descriptions using approaches such as web-mining; and secondly, it is difficult to elicit controller instruction sequences from non-expert users.

We therefore designed a crowd-sourcing system to solve these challenges. Our system consists of a virtual simulator which presents a 3D environment to users. Figure 5 shows sample environments that are presented to the users. Each environment has around 30 objects of which around 50% are interactive. The interactive objects have 2-5 states in addition to location and orientation. The value of these states can be boolean, real number or strings from a finite set, and can be modified by the interaction of a user with the environment. For example, the object of category *mug* has the following states: *percentage-filled, is-grasped, has-coffee, has-water, has-ice_cream, has-syrup*. Objects also have affordances such

Fig. 6: **Robotic Simulator.** User points on the stove to see the possible actions. User can also see different states of the stove and the grasped kettle thus taking appropriate actions. The natural language segment on which the user is working, is shown above the simulator while the recorded instructions are shown in the right panel.

as *graspable, pourable, placeable,* etc. There could be more than one object of the same category and appearance.

Users of our system are presented with two type of tasks—*language task* and *game task*. In the language task, they write natural language commands for accomplishing a given objective after exploring the environment. In the game task, users are presented with natural language commands written by other users and have to control the virtual robot to accomplish the given command. As shown in Figure 6, users can point on objects to view their states and can modify them by taking different actions. Whenever a user takes an action, it is stored in the database in a symbolic form. For instance, if a user drops a cup on a table then the action is recorded as $\texttt{keep}(\texttt{cup}, \texttt{table}, \texttt{on})$. Since that may not have been the intention of the user, therefore we also store the discretized trajectory of the action.

In the following section, we describe how we use this system to collect a VEIL dataset for testing our model.

VIII. Experiments and Results

Dataset Details. For evaluating our approach, we have collected a dataset called *VEIL-300*. Each data-point consists of a natural language command, the starting environment, ground-truth instruction sequence and the mapping between segments of the command and the instruction sub-sequence.

The natural language command describes the high level task but does not necessarily have the exact knowledge of the environment (e.g., it may refer to a cup while there is no actual cup in the environment). Furthermore, it may miss several steps in the middle, and use ambiguous words for describing a task. As we can see from the examples in Table I, the context plays a very important part in grounding these sentences. For example, in sentence 5 from Table I, scooping ice-cream using a spoon and keeping it on a table would make little sense than adding it to the coffee that has been just prepared.

The dataset is then processed to be in the VEIL format (see Section V).

We considered six tasks: *boiling water*, *making coffee*, *making ramen*, *making affogato*, *prepare the room for party* and *clean the room*. Each task was performed in 5 different environments and each environment had 10 natural language commands (thus giving us a dataset of size 6x5x10=300).

For the first two tasks, we only gave 10 "general" natural language instructions ($5 \times 10 = 50$). This allowed us to evaluate whether our algorithm can ground natural language instructions in different environments.

For the last four tasks, there was a different natural language instruction for each of the 50 datapoints. During training, we use 10-fold cross validation, with the six tasks trained in 3 pairs of 2 – {(task1, task2), (task3, task4), (task5, task6)}. While testing, the algorithm always tests on a new environment.

Evaluation Metric. We consider two different type of evaluation metrics:

1) Instruction Edit Distance (IED). We use a string edit distance namely the *Levenshtein distance*, for measuring the edit distance between the two instruction sequence. This gives some idea on how similar our model's output $\hat{\mathcal{I}}$ is to the ground-truth sequence $\mathcal{I}^g$. However, it is limited in that it does not handle the case where one wrong instruction in the middle can completely change the resulting state. For instance, it is not acceptable to forget filling a pot with water while making ramen, even if the rest of the sequence is correct.

2) Environment Edit Distance (EED). This metric relies on that *a successful execution of a task given the NL instruction is the attainment of a sequence of states for concerned objects in a possibly time-ordered way*. For this, the metric computes the edit distance between the ground-truth sequence of environments $(E^g_k)^m_{k=0}$ and the predicted sequence of environments $(\hat{E}_k)^n_{k=0}$. However, there are two subtle issues: finding the set of concerned objects ∂, and finding the correct representation of difference, $\text{dist}_{env}(E^g_i, E^a_j)$, between two environments. This is because not all the object states are important in an environment, for a given task (e.g., closing the microwave door after its use is irrelevant to the task of heating water).

We use the objects in the ground-truth sequence $\mathcal{I}^g$ as the set of concerned objects ∂. Further we define dist_{env} as a 0-1 loss function on whether the two environments completely agree on the states of ∂. The recursive equation for computing EED is given below, where $EED(.,.,i,j)$ represents the distance between $(E^g_k)^m_{k=i}$,$(\hat{E}_k)^n_{k=j}$

$$\text{EED}((E^g_k)^m_{k=0}, (\hat{E}_k)^n_{k=0}, i=0, j=0) =$$
$$\min\{\text{EED}(\cdot,\cdot,i,j+1), \text{EED}(\cdot,\cdot,i+1,j) + \text{dist}_{env}(E^g_i, \hat{E}_j, \partial)\}$$
$$\text{where EED}(\cdot,\cdot,i,j) = m-i \text{ if } i=m \text{ or } j=n$$

We normalize these two metrics and report them as percentages in Table IV, where higher numbers are better.

Baselines. We compare our model against the following:

- *Chance*: Randomly chooses an instruction sequence of fixed length. This shows how large the output space is.

TABLE IV: **Quantitative results** on six different tasks from our dataset. Results of baselines, different variants of our method are reported on two different metrics— *IED* and *EED*, which are normalized to 100 with larger numbers being better.

	making coffee		*boiling water*		*making ramen*		*making affogato*		*party night*		*clean room*		Average	
	IED	EED	IED	EED	IED	EED	IED	EED	IED	EED	IED	EED	IED	EED
Chance	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2
Predefined Templates (Guadarrama et al., 2013)	14.3	22.8	21.5	38.8	32.8	28.7	14.0	17.0	7.1	10.0	16.7	13.7	17.8	21.8
Instruction-Tree (Bollini et al., 2012)	14.4	25.7	11.2	38.8	26.3	25.6	12.5	15.9	22.7	49.7	20.5	14.3	17.9	28.3
No NL Instructions	39.2	43.9	21.8	43.3	1.2	15.7	12.8	11.5	15.0	17.9	10.8	9.4	16.8	23.6
Nearest Environment Instantiation	63.4	47.3	43.5	50.1	52.1	51.4	45.4	49.2	47.5	58.3	61.7	41.7	52.3	49.7
Our model, No Domain Knowledge	78.6	71.7	55.7	55.2	66.1	49.5	51.2	51.3	51.8	56.4	**69.2**	47.1	62.1	55.2
Our model, No latent nodes.	**78.9**	73.8	**57.0**	58.4	66.3	**53.9**	51.5	51.2	**53.5**	**62.6**	66.2	**51.5**	**62.2**	58.6
Our full model.	76.5	**78.3**	54.5	**61.3**	**67.3**	49.4	**53.5**	**56.2**	53.2	**62.6**	66.0	**51.5**	61.8	**59.9**

- *Predefined Templates, based on Guadarrama et al. (2013)* We developed a method similar to Guadarrama et al. (2013), in which we manually created a library of proposal-templates for every verb. For a given situation, we disambiguate the language-objects and accordingly instantiate the templates. We further extend their approach by also considering many-many mappings.
- *Instruction-Tree Search, based on Bollini et al. (2012)*: We define a log-linear reward function that searches over the sequence of instructions for every clause and chooses the one which maximizes the reward. The reward function contains bag-of-word features (correlation between words in the clause and objects in the instruction sequence). Invalid instruction sequences are pruned and gradient descent is used for training the parameters.
- *Nearest Environment Instantiation*: This method looks up for the nearest template by minimizing the distance between the given environment and that of the template, from the VEIL library. Template is instantiated according to the new environment, following Section V-A.
- *Our model - No NLP.* Our model in which we do not give the natural language command, i.e., the given clauses $\mathcal{C}$ are missing in the model. However, all the other terms are present.
- *Our model - No Domain Knowledge.* Our model in which the robot does not have the knowledge of the results of its action on the environment, and instead relies only on the training data. (This is to compare against the symbolic planner based approaches.) For this we disable the latent, jump features and post-condition features. We also simplify the inference procedure 1, which uses the simulator Φ and hence the domain knowledge, by doing a greedy assignment based on the initial environment.
- *Our model - No latent nodes.* Our model in which the latent instruction and environment nodes are missing. This model cannot handle the missing natural language instructions well.
- *Our full model.* This is our full model.

Results. Table IV shows our results. We note that the chance performance is very low because the output space (of instruction sequence) is very big, therefore the chances of still getting a correct sequence are very low.

We compare our results to our implementation of two recent notable works in the area. Table IV shows that the method *Predefined Templates (Guadarrama et al., 2013)* focused on disambiguating spatial relations but was extremely brittle to ambiguity in grounding, therefore giving low performance.

Method *Instruction-Tree (Bollini et al., 2012)* was able to give reasonable results for some datapoints. However this approach has problem working with large search trees. Furthermore, the bag-of-word features do not take into account the environment context; for instance, the natural language instruction might say *"keep the cup in microwave"* but the cup might already be inside the microwave (unless such constraints are hard-coded). This approach thus fails when the language is vague, for example, for the following sentence, *"heat the water and add ramen"*. However, our approach takes this vague sentence and grounds it in the environment using our model. Our energy function incorporates several features and thus is able to often give reasonable output for such natural language instructions.

Following are some natural language instructions that our algorithm was able to ground successfully:

- *"serve the chips and then bring the pillows to the couches and turn on the tv"*
- *"clean up the trash and replace the pillows and turn off the tv"*
- *"fill the pot with ramen and water and then cook it"*

and following are some natural language instructions that our algorithm was not able to ground:

- *"throw away the chips and throw away the remaining bag"*
- *"Microwave for 12 minutes and place it on the table"*
- *"heat up water in the pot and place ramen inside pot"*

Common reasons for errors include the required VEIL template not being present, incorrect object grounding and other linguistic complications that we are not handling in this paper such as conditionals, quantifiers etc.

We analyze the results in light of the following questions:

Is Language important? If we enforce all the constraints of the task and provide the end-state of the environment, one may argue that just using a symbolic planner may give reasonable programs. However, the success of a task depends on the way how things are done. Natural language gives an approximate guide that our model tries to follow. For example, though we may provide that the cup has water in the goal environment. But if the natural language command is *"fill a cup with water*

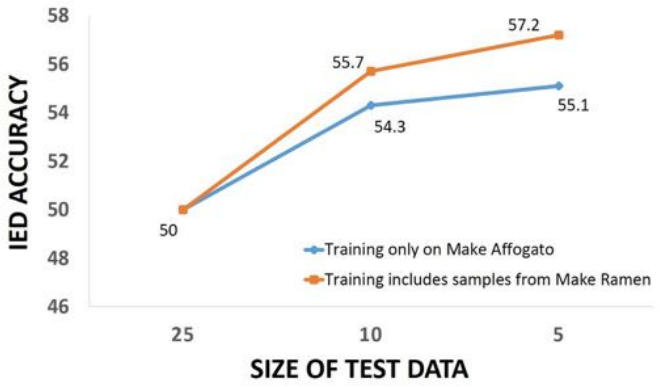

Fig. 7: Shows IED accuracy on Making Affogato dataset versus size of test data. Cross-validation is used with test-data of the given size. In another experiment, we add Making Ramen samples to the training and we see an improvement of 1-2% for test data of size 5 and 10. This shows that samples from other task were useful.

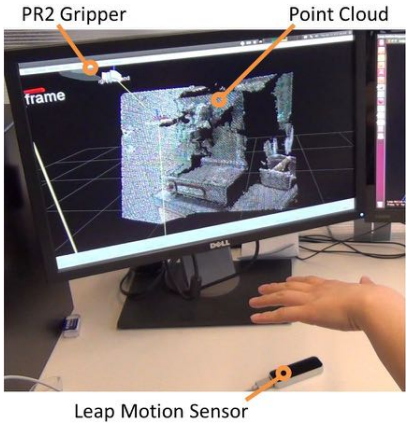

Fig. 8: The demonstration system using Leap motion sensor. The expert demonstrates the instruction for object shown on the screen by hovering over the sensor and moving the gripper of the PR2 on the screen.

from the sink" then an instruction sequence that uses fridge to fill the cup with water will result in failure.

We see that *Our Model - No NLP* gives 16.8% on average on the IED metric as compared to 61.8% for our full model.

In fact, we see evidence of such behavior in our results also. While our model can handle ambiguous and incomplete NL instructions, e.g., *'heat up the water and then cook the ramen'* that resulted in success, in some of the test cases the NL instructions were quite ambiguous, e.g., *'Microwave for 12 minutes and place it on the table'* on which our model failed.

How important is the latent node? Overall, Table IV shows that the results improve by about 1.3% on the EED metric after adding the latent node. We found that it was especially helpful in scenarios where instructions were partially missing.

Fig. 10: **Snapshots of PR2** using microwave for making coffee when only microwave is available for heating the water and preparing ramen.

For example, for the instruction in Fig. 2 -

"place the pot on the tap and turn the tap on... Turn the tap off and heat the pot."

there is no template that can fit in for the first clause ($\texttt{place}, [\texttt{pot}, \texttt{tap}], \texttt{on} : \texttt{pot} \rightarrow \texttt{tap}$). One such template after initialization has the form-

$$\texttt{moveTo}(\texttt{sink});\ \texttt{keep}(\texttt{pot}, \texttt{sink}, \texttt{on})$$

However the robot cannot execute this sequence since it is not already grasping the `pot`. In such cases, the latent nodes help us with modeling these missing instructions and gives us the following executable output-

$$\texttt{moveTo}(\texttt{pot});\ \texttt{grasp}(\texttt{pot});\ \texttt{moveTo}(\texttt{sink});\ \texttt{keep}(\texttt{pot}, \texttt{sink}, \texttt{on})$$

How well does our model generalize to new environments and tasks? In this test, we wanted to examine how well our model can make use of examples from different environment and tasks. It is not obvious whether the templates learned for one task, such as *making ramen* will be useful for another task such as *making affogato*. For studying this, we performed another experiment in which we trained and tested the model on *making affogato* task only and then we added samples from *making ramen* to the training. We found that because the VEIL library from the *making ramen* task was not available for training, the performance dropped by 1.5-2% on the *IED* metric. The results are shown in Figure 7. This suggests that examples from other tasks are helpful.

How does the algorithm handle noisy and incomplete sentences? The algorithm handles noisy object references as part of the instantiation procedure. For example, consider that the robot is asked to get a cup of water and there is no object of category cup in the given environment. In this case, the function $Ground(\text{"}cup\text{"}, E)$ will return $Null$ and the instantiation procedure V-A will then find an object in the given environment which shares the most state-value pairs with the object that it saw in the training environment. This object is then taken as the intended object.

Our algorithm also handles several incomplete sentences such as *"microwave a cup of water"*, which does not specify filling the cup with water in case it is empty, or opening the door of the microwave if it is closed etc. The algorithm handles this in two ways– either by using a similar sample, that was present in the VEIL template which contains the missing steps; or if the missing steps constitute a hard constraint (such as

Fig. 9: **Robot Experiment.** Given the natural language instruction for making the 'Affogato' dessert: *"Take some coffee in a cup. Add ice-cream of your choice. Finally, add raspberry syrup to the mixture."*, our algorithm outputs a controller instruction sequence that the PR2 uses to make the dessert. (Please see the video.)

opening the door of a microwave before trying to put a cup in) then they are handled by the latent node $\mathcal{I}_\ell$ using a symbolic planner.

What if the robot does not know the result of its action? The algorithm implicitly assumes that the robot knows the result of its interaction with the environment (It is being used to compute certain features, doing the planning and in inference). In order to test how crucial it is, we ran the baseline *Our Model - No Domain Knowledge* and as the results in Table IV show, the accuracy falls by only 4.7 % on the EED metric. However, without the domain knowledge the guarantee that the output sequence is executable is lost. The robot will then have to recover at run-time if the lower level controllers fail.

IX. Robot Experiment

We use our grounding algorithm in an end to end robotic system, which can take NL instructions from users and manipulate real world objects to achieve the task. In this section we describe our robotic system for the task of making affogato dessert. We take the following NL instruction, given by a user, as our working example: *"Take some coffee in a cup. Add ice cream of your choice. Finally, add raspberry syrup to the mixture."*

Grounded instruction sequence given by our grounding algorithm consists of instruction listed in Table II, with appropriate parameters which are chiefly objects. Although there are many related works for each instruction (Anand et al., 2012; Bollini et al., 2011; Endres et al., 2013; Jiang et al., 2012; Lenz et al., 2013; Ratliff et al., 2009), it is still a very active area of research and it is quite challenging to execute these instructions reliably in a sequence. Thus, we take a Learning from Demonstration approach (Argall et al., 2009) for the set of instructions relevant to the task, in order to test the validity of the grounded sequence given by our model.

First, the robot observes the scene using RGB-D sensor (Microsoft Kinect). Given the point cloud from the sensor, the scene is segmented into smaller patches representing object parts (Anand et al., 2012). In order to reuse demonstrations regardless of orientation and location of the object, the object frame using the segmented object part is found for the use of demonstration and execution. Reliable frame of the object can be established by aligning axis with the principal axis of the point cloud computed using PCA (Hsiao et al., 2010). The segmentation and alignment of the axis with the segmented point-cloud allows the manipulation to be location and orientation invariant when the demonstration is trained with respect to this frame.

Among many approaches of Learning from Demonstration such as kinesthetic teaching, teleoperation, and so forth (Argall et al., 2009), we use teleoperation-based approach for demonstrating the task to the robot. Figure 8, shows our system built using the Leap motion sensor where a user can teach the robot how to execute the instruction with certain parameters (objects). By observing the object shown on screen, the user controls the gripper of the robot by hovering over the Leap motion sensor. The center of the object frame is virtually placed few centimeters above the Leap Motion sensor, which allows an expert to easily demonstrate the sequence by controlling the gripper with his palm. Rather than recording the full movement of the demonstrator which could be not smooth, the recorded sequence for the robot to execute is based on a sequence of keyframe similar to (Akgun et al., 2012). Each keyframe is recorded by pressing the key as the user demonstrates the instruction. Also, rather than recording the full joint configuration, each keyframe records only the location and orientation of the end-effector so that it can used regardless of the location of the object relative to the robot.

For the robotic experiment, demonstration given by the user can be executed in a new setting by segmenting the point cloud, finding the object frame and then executing the demonstration using the trajectory controller. We utilize the impedance controller (`ee_cart_imped`) (Bollini et al., 2011) for our PR2 robot to follow the end-effectory trajectory of learned demonstration. Once the object frame is centered correctly, the robot can successfully execute the instruction, as long as there is no collision and the object is in the workspace of the robot. This is because each control sequence was trained with respect to the object frame.

Figure 9 and Figure 10 shows several snapshots of PR2 making Affogato, making coffee, and preparing ramen and a full video of PR2 making affogato dessert is available at: `http://tellmedave.com`

X. Future Work

There are many directions in which we can extend the approach presented in this paper. In brief, following are some

TABLE V: Neo-Davidsonian semantics for representing verbs with varying number of arguments, modifiers and relative clauses. Variable e is the Neo-Davidsonian event.

	Sentence	Neo-Davidsonian Semantics
1.	*"move to the cup"*	$\exists$e.moveto(e) $\wedge$ to(e, cup$_{02}$)
2.	*"slowly move to the cup"*	$\exists$e.moveto(e) $\wedge$ to(e, cup$_{02}$) $\wedge$speed(e, slow)
3.	*"slowly move to the cup by staying close to the wall"*	$\exists$e.moveto(e) $\wedge$ to(e, cup$_{02}$)$\wedge$ speed(e, slow) $\wedge$ close(e, wall)

such directions:

Improving the Language Model: In this paper, our main focus has been on grounding high level verbs which posses the most complex semantic structure. For grounding noun-phrases such as *"the white cup", "red mug by the door"*; we used a simple syntactical matching based on object categories. While this worked relatively well for our dataset, this needs to be improved to handle complex object descriptions. These complex object descriptions can contain determiners (e.g., *"bring me the cup"*), anaphoric reference (e.g., *"keep it on the table"*), recursive object descriptions (e.g., *"microwave the cup which is on the table, close to fridge"*) etc.

As mentioned before, there are several works (Guadarrama et al., 2013, 2014) which aim to solve this problem. These algorithm can be easily integrated into our model by replacing the *Ground* function in Section V-A with them.

Besides handling complex object descriptions, the language model needs to handle conditionals (e.g., *"if there is a cup"*), quantifiers as well as domain specific expressions such as time (e.g., *"microwave the cup for 12 minutes"*) and quantity (e.g., *"take 2 ounces of potatos"*). Each of these challenge is an interesting problem in the field of natural language understanding. We refer the interested readers to Artzi and Zettlemoyer (2013) for parsing quantifiers and spatial relations and Lee et al. (2014) for parsing time expressions.

Richer Representations: In this paper, we ground natural language commands to a sequence of instructions, each having a single predicate– `controller-name(arguments)`. These instructions are then mapped onto trajectories (see Section IX). One advantage of this is that instructions with one predicate can be easily recorded in our crowd-sourcing system from user interaction. However, single predicate instructions cannot handle variable number of arguments which may be present in the natural language command. For example, the verb *move* can accept varied number of arguments and modifiers as shown in Table V. A single predicate cannot handle all of these cases. Whereas having a predicate for every situation leads to combinatorial explosion. In the field of semantics, a solution is to use Neo-Davidsonian semantics which defines a common event shared between multiple atoms (instantiated predicate). Table V shows examples of Neo-Davidsonian semantics.

In future, we want to ground natural language to a sequence of instruction where each instruction is a conjunction of atoms coupled by a Neo-Davidsonian event.

Learning Environment Model: As mentioned before, in this paper we assume access to domain knowledge of the environment. We also assumed that the world is deterministic. In future, we want to learn the environment model along with learning to ground natural language. One direction could be to use the reinforcement learning setting of Branavan et al. (2010).

Real-time Inference: The inference algorithm presented in this paper uses all the samples corresponding to a verb, which makes the algorithm impractical for a very large dataset. In future, we plan to use better approximate inference techniques to make the inference real-time while maintaining accuracy.

In future, we also want to test our algorithm on much larger datasets as well as improving our crowd-sourcing system by integrating it with the RoboBrain platform (Saxena et al., 2014).

XI. Conclusion

In this paper, we presented a novel approach for grounding free-form natural language instructions into a controller instruction sequence for a given task and environment, that can be executed by a robot to perform the task. This is challenging since the grounding is highly dependent on the environment. Another challenge is that the natural language instructions can be incomplete and ambiguous as shown in Table I. To solve these issues, we represented this context in a VEIL dataset format which was collected using crowd-sourcing. We presented a new learning model that encodes certain desired properties into an energy function—expressed as a model isomorphic to a conditional random field with edges representing relations between natural language, environment and controller instructions. We showed that our model handles incomplete natural language instructions, variations in natural language, as well as ambiguity in grounding. We also showed that we outperform related work in this area on both syntax-based and semantic-based metrics.

Acknowledgement

We thank Claire Cardie for useful discussions and Kejia Tao and Aditya Jami for help with the crowd-sourcing system. This work was supported in part by ONR Grant N00014-14-1-0156, and Microsoft Faculty Fellowship and NSF Career award to Saxena.

References

B. Akgun, M. Cakmak, K. Jiang, and A. Thomaz, "Keyframe-based learning from demonstration," *International Journal of Social Robotics*, vol. 4, no. 4, pp. 343–355, 2012.

R. Alterovitz, S. Patil, and A. Derbakova, "Rapidly-exploring roadmaps: Weighing exploration vs. refinement in optimal motion planning," in *ICRA*, 2011.

A. Anand, H. Koppula, T. Joachims, and A. Saxena, "Contextually guided semantic labeling and search for 3d point clouds," *IJRR*, 2012.

B. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *RAS*, 2009.

Y. Artzi and L. Zettlemoyer, "Weakly supervised learning of semantic parsers for mapping instructions to actions," *ACL*, 2013.

A. Aydemir, K. Sjoo, J. Folkesson, A. Pronobis, and P. Jensfelt, "Search in the real world: Active visual object search based on spatial relations," in *ICRA*. IEEE, 2011, pp. 2818–2824.

J. Barry, K. Hsiao, L. P. Kaelbling, and T. Lozano-Pérez, "Manipulation with multiple action types," in *Exp Robo.*, 2013, pp. 531–545.

M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Mosenlechner, D. Pangercic, T. Ruhr, and M. Tenorth, "Robotic roommates making pancakes," in *Humanoids*, 2011.

J. Berant, A. Chou, R. Frostig, and P. Liang, "Semantic parsing on freebase from question-answer pairs," in *Proceedings of EMNLP*, 2013.

M. Bollini, J. Barry, and D. Rus, "Bakebot: Baking cookies with the pr2," in *The PR2 Workshop, IROS*, 2011.

M. Bollini, S. Tellex, T. Thompson, N. Roy, and D. Rus, "Interpreting and executing recipes with a cooking robot," in *ISER*, 2012.

S. Branavan, L. Zettlemoyer, and R. Barzilay, "Reading between the lines: Learning to map high-level instructions to commands," in *ACL*. ACL, 2010, pp. 1268–1277.

M. Cakmak, M. Dogar, E. Ugur, and E. Sahin, "Affordances as a framework for robot control," in *EpiRob*, 2007.

C. Chao, M. Cakmak, and A. Thomaz, "Towards grounding concepts for transfer in goal learning from demonstration," in *ICDL*, 2011.

D. L. Chen, J. Kim, and R. J. Mooney, "Training a multilingual sportscaster: Using perceptual context to learn language," *Journal of Artificial Intelligence Research*, vol. 37, no. 1, pp. 397–436, 2010.

V. Chu, I. McMahon, L. Riano, C. McDonald, Q. He, J. Perez-Tejada, M. Arrigo, N. Fitter, J. Nappo, T. Darrell *et al.*, "Using robotic exploratory procedures to learn the meaning of haptic adjectives," in *IROS*, 2013.

F. Duvallet, M. R. Walter, T. Howard, S. Hemachandra, J. Oh, S. Teller, N. Roy, and A. Stentz, "Inferring maps and behaviors from natural language instructions," in *Proceedings of the International Symposium on Experimental Robotics (ISER)*, 2014.

F. Endres, J. Trinkle, and W. Burgard, "Learning the dynamics of doors for robotic manipulation," in *IROS*, 2013.

A. Farhadi, I. Endres, and D. Hoiem, "Attribute-centric recognition for cross-category generalization," in *CVPR*, 2010.

A. Farhadi, M. Hejrati, M. A. Sadeghi, P. Young, C. Rashtchian, J. Hockenmaier, and D. Forsyth, "Every picture tells a story: Generating sentences from images," in *Computer Vision–ECCV 2010*. Springer, 2010, pp. 15–29.

J. Fasola and M. Matarić, "Using semantic fields to model dynamic spatial relations in a robot architecture for natural language instruction of service robots," in *IROS*, 2013.

R. E. Fikes and N. J. Nilsson, "Strips: A new approach to the application of theorem proving to problem solving," *Artificial intelligence*, vol. 2, no. 3, pp. 189–208, 1972.

C. Finucane, G. Jing, and H. Kress-Gazit, "Ltlmop: Experimenting with language, temporal logic and robot control," in *IROS*, 2010.

S. Guadarrama, L. Riano, D. Golland, D. Gouhring, Y. Jia, D. Klein, P. Abbeel, and T. Darrell, "Grounding spatial relations for human-robot interaction," in *IROS*, 2013.

S. Guadarrama, E. Rodner, K. Saenko, N. Zhang, R. Farrell, J. Donahue, and T. Darrell, "Open-vocabulary object retrieval," in *RSS*, 2014.

S. Höfer, T. Lang, and O. Brock, "Extracting kinematic background knowledge from interactions using task-sensitive relational learning," in *ICRA*, 2014.

K. Hsiao, S. Chitta, M. Ciocarlie, and E. Jones, "Contact-reactive grasping of objects with partial shape information," in *IROS*, 2010.

Y. Jiang, M. Lim, C. Zheng, and A. Saxena, "Learning to place new objects in a scene," *IJRR*, vol. 31, no. 9, 2012.

Y. Jiang, H. Koppula, and A. Saxena, "Hallucinated humans as the hidden context for labeling 3d scenes," in *CVPR*, 2013.

L. P. Kaelbling and T. Lozano-Pérez, "Hierarchical task and motion planning in the now," in *ICRA*, 2011.

H. Kjellström, J. Romero, and D. Kragić, "Visual object-action recognition: Inferring object affordances from human demonstration," *Computer Vision and Image Understanding*, vol. 115, no. 1, pp. 81–90, 2011.

D. Klein and C. Manning, "Accurate unlexicalized parsing," in *ACL*, 2003.

T. Kollar, S. Tellex, D. Roy, and N. Roy, "Grounding verbs of motion in natural language commands to robots," in *ISER*, 2010.

H. Koppula and A. Saxena, "Anticipating human activities using object affordances for reactive robotic response," in *RSS*, 2013.

H. S. Koppula, R. Gupta, and A. Saxena, "Learning human activities and object affordances from rgb-d videos," *IJRR*, vol. 32, no. 8, pp. 951–970, 2013.

H. Koppula, A. Anand, T. Joachims, and A. Saxena, "Semantic labeling of 3d point clouds for indoor scenes," in *NIPS*, 2011.

H. Kress-Gazit, G. Fainekos, and G. Pappas, "From structured english to robot motion," in *IROS*, 2007.

O. Kroemer, R. Detry, J. Piater, and J. Peters, "Combining active learning and reactive control for robot grasping," *RAS*, vol. 58, no. 9, pp. 1105–1116, 2010.

J. Kulick, M. Toussaint, T. Lang, and M. Lopes, "Active learning for teaching a robot grounded relational symbols," in *IJCAI*, 2013.

T. Kwiatkowski, L. Zettlemoyer, S. Goldwater, and M. Steedman, "Inducing probabilistic ccg grammars from logical form with higher-order unification," in *EMNLP*. ACL, 2010, pp. 1223–1233.

K. Lee, Y. Artzi, J. Dodge, and L. Zettlemoyer, "Context-dependent semantic parsing for time expressions," in *ACL*, 2014.

S. Lemaignan, R. Ros, E. A. Sisbot, R. Alami, and M. Beetz, "Grounding the interaction: Anchoring situated discourse in everyday human-robot interaction," *IJSR*, vol. 4, no. 2, pp. 181–199, 2012.

I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," in *RSS*, 2013.

J. Malmaud, E. Wagner, N. Chang, and K. Murphy, *Cooking with Semantics*. ACL 2014 Workshop on Semantic Parsing, 2014, pp. 33–38.

D. Marco, M. Tenorth, K. Häussermann, O. Zweigle, and P. Levi, "Roboearth action recipe execution," in *IAS*, 2012.

C. Matuszek, N. Fitzgerald, L. Zettlemoyer, L. Bo, and D. Fox, "A joint model of language and perception for grounded attribute learning," in *ICML*, 2012, pp. 1671–1678.

C. Matuszek, E. Herbst, L. Zettlemoyer, and D. Fox, "Learning to parse natural language commands to a robot control system," in *ISER*, 2012.

D. Misra, J. Sung, K. Lee, and A. Saxena, "Tell me dave: Context-sensitive grounding of natural language to mobile manipulation instructions," in *RSS*, 2014.

L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor, "Learning object affordances: From sensory–motor coordination to imitation," *Robotics, IEEE Transactions on*, vol. 24, no. 1, pp. 15–26, 2008.

K. Mourao, L. Zettlemoyer, R. Petrick, and M. Steedman, "Learning strips operators from noisy and incomplete observations," *UAI*, 2012.

H. Nguyen, M. Ciocarlie, J. Hsiao, and C. C. Kemp, "Ros commander (rosco): Behavior creation for home robots," in *ICRA*, 2013.

S. Niekum, S. Chitta, A. Barto, B. Marthi, and S. Osentoski, "Incremental semantically grounded learning from demonstration," in *RSS*, 2013.

J. Nothman, M. Honnibal, B. Hachey, and J. Curran, "Event linking: Grounding event reference in a news archive," in *ACL*, 2012.

H. Poon, "Grounded unsupervised semantic parsing," in *ACL*, 2013.

V. Ramanathan, P. Liang, and L. Fei-Fei, "Video event understanding using natural language descriptions," in *ICCV*, 2013.

N. Ratliff, M. Zucker, D. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *ICRA*, 2009.

J. Rintanen, "Planning as satisfiability: Heuristics," *Artificial Intelligence*, 2012.

R. Ros, S. Lemaignan, E. A. Sisbot, R. Alami, J. Steinwender, K. Hamann, and F. Warneken, "Which one? grounding the referent based on efficient human-robot interaction," in *RO-MAN*, 2010, pp. 570–575.

A. Saxena, A. Jain, O. Sener, A. Jami, D. K. Misra, and H. S. Koppula, "Robobrain: Large-scale knowledge engine for robots," *arXiv preprint arXiv:1412.0691*, 2014.

S. Srinivasa, D. Ferguson, C. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and M. Weghe, "Herb: a home exploring robotic butler," *Autonomous Robots*, vol. 28, no. 1, pp. 5–20, 2010.

M. Steedman, *Surface structure and interpretation*. MIT press, 1996.

——, *The syntactic process*. MIT Press, 2000, vol. 35.

J. Sung, B. Selman, and A. Saxena, "Synthesizing manipulation sequences for under-specified tasks using unrolled markov random fields," in *IROS*, 2014.

S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. G. Banerjee, S. J. Teller, and N. Roy, "Understanding natural language commands for robotic navigation and mobile manipulation." in *AAAI*, 2011.

S. Tellex, R. A. Knepper, A. Li, T. M. Howard, D. Rus, and N. Roy, "Assembling furniture by asking for help from a human partner," *Collaborative Manipulation workshop at Human-Robot Interaction*, 2013.

M. Tenorth, L. Kunze, D. Jain, and M. Beetz, "Knowrob-map-knowledge-linked semantic object maps," in *Humanoids*, 2010.

M. Walter, S. Hemachandra, B. Homberg, S. Tellex, and S. Teller, "Learning semantic maps from natural language descriptions," in *RSS*, 2013.

T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon control for temporal logic specifications." in *HSCC*, 2010.

C. Wu, I. Lenz, and A. Saxena, "Hierarchical semantic labeling for task-relevant rgb-d perception," in *Robotics: Science and Systems (RSS)*, 2014.

H. Yu and J. Siskind, "Grounded language learning from videos described with sentences," in *ACL*, 2013.

L. Zettlemoyer and M. Collins, "Online learning of relaxed ccg grammars for parsing to logical form," in *EMNLP-CoNLL-2007*. Citeseer, 2007.