

Breaking Down High-level Robot Path-Finding Abstractions in Natural Language Programming

Yue Zhan and Michael S. Hsiao
{zyue91, hsiao}@vt.edu

Bradley Department of Electrical and Computer Engineering,
Virginia Tech, Blacksburg, VA 24060, USA

Abstract. Natural language programming (NLPr) allows people to program in natural language (NL) for specific domains. It poses great potential since it gives non-experts the ability to develop projects without exhaustive training. However, complex descriptions can sometimes have multiple interpretations, making program synthesis difficult. Thus, if the high-level abstractions can be broken down into a sequence of precise low-level steps, existing NLP and NLPr techniques could be adaptable to handle the tasks. In this paper, we present an algorithm for converting high-level task descriptions into low-level specifications by parsing the sentences into sentence frames and using generated low-level NL instructions to generate executable programs for pathfinding tasks in a LEGO Mindstorms EV3 robot. Our analysis shows that breaking down the high-level pathfinding abstractions into a sequence of low-level NL instructions is effective for the majority of collected sentences, and the generated NL texts are detailed, readable, and can easily be processed by the existing NLPr system.

Keywords: Natural language processing · Natural language programming · Program synthesis · LEGO Mindstorms EV3.

1 Introduction

The field of robotics has made significant strides because of the growth of market demands in recent years. However, despite the growing interest in educational robots, the time-consuming learning process and the steep learning curve of programming robots still challenge young robotics enthusiasts. Natural language programming (NLPr) offers a potential way to lower the bar of entry by allowing the users to “program” the robot using natural language (NL). The readability and expressive nature of natural language make it an ideal way to simplify the learning process. Though promising for this use case, NLPr has several challenges of its own. First, NL texts used to give instructions are typically low-level (LL) specifications to ensure precision and completeness. For example, the movement specifications used in the NLPr system for LEGO Mindstorms EV3 robot in the work [22] are categorized as a controlled natural language (CNL) [8]. The movement sentences used in the system are object-oriented sentences

Copyright ©2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

like “*The robot goes forward/backward/left/right ...*”. The requirement to use such low-level specifications makes the process of directing the robot more difficult for novice users, as they would rather give a high-level instruction such as “*The robot moves from point A to point B*” than to list out every individual step the robot must take. Unconstrained NL texts are highly flexible and expressive but can sometimes be ambiguous. Designing a language model for NLPr to cover all of the language structures in NL is extremely difficult, if not impossible [3]. As such, it would be a huge benefit for NLPr tasks if the information in a higher-level abstraction can be effectively extracted and used to generate a sequence of precise, unambiguous lower-level sentences that explain the intention and plans the proper actions. Suppose the information related to the robot tasks can be extracted. In that case, the language structures that need to be covered in the domain-specific function library and lexicon can be simplified, and the existing NLPr system can be directly adapted with fewer necessary modifications to handle the high-level (HL) NL abstractions, as shown in Figure 1.

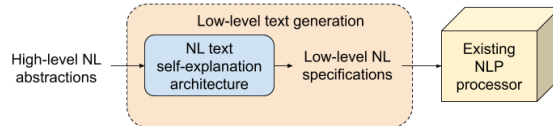


Fig. 1: High-level NL to low-level NL transformation

The key challenge addressed in this paper is effectively extracting semantic information from high-level sentences to synthesize low-level pathfinding NL instructions. In order to demonstrate our proposed low-level text generation process, we use a robot pathfinding task, in which a robot must find an optimal path between two points while avoiding obstacles along the way. To succeed at this task, our system must generate a sequence of low-level instructions that take the robot to its goal while minimizing the time cost and the number of actions taken by the robot. Once our system identifies a path based on the high-level input, it outputs a sequence of low-level NL to an existing NLPr system [22], which then generates the executable program for the LEGO Mindstorms EV3 robot.

2 Previous Work

Due to its promise of better ease of use and improved human-computer interaction, the foundations of NL based programming for robotics have been well established. In [1,13], an NLPr system that navigates a vision-based robot with an instruction-based learning method is presented. In these systems, robot-understandable procedures are generated from command-like NL instructions based on a set of pre-programmed sensory-motor action primitives and routes in a miniature city map. Users can give instructions based on available primitives to the robot when facing an unknown route with a human-robot interaction interface. In the work [21], a Vision-language robot navigation (VLN) system that combines the vision information and descriptive NL commands reasoning using a data-based model is proposed for in-home environments. When the NL instructions are given, a sequence of actions is generated by the reasoning navigator. Gathering sufficient data on various environments for model training purposes could be

costly. In the work [14], a probabilistic combinatory categorial grammar (PCCG) based parser is used to translate procedural NL commands into logic-based Robot Control Language (RCL) for robot navigation. In [22], a grammar-based Object-Oriented Programming Controlled Natural Language (OOP-CNL) model is used to translate NL sentences into executable code for LEGO robots. In this work, the NLPr program synthesis system utilizes contextual and grammatical information to derive desired robot functionalities with a domain-specific function library and lexicon. While the language model used here can process more complex sentence structures, such as conditional statements, the sentences used for navigating the robot are still at a lower level.

There has been a significant amount of work done in the field of NLPr program synthesis, and most of this work has been focused on solving domain-specific problems. The work in [4] emphasizes the importance of NLP techniques in analyzing textual contents in software programs. The authors propose a system called Toradocu, which they developed using Stanford parser and a pattern and lexical similarity matching that converts Javadoc comments into assertions, and a system called Tellina, which is trained with an RNN [10] to generate bash commands and use these systems to illustrate the potential of program synthesis with NL texts. The Metafor platform [11,17] is a descriptive NLPr system that can convert NL components into class descriptions with associated objects and methods. This work takes advantage of the NL parsing toolkit MontyLingua, mixed-initiative dialog, and programming by example techniques. The authors state that modern parsing techniques and the integration of common sense knowledge can help developers link humans' narrative capacities with traditional programming languages. However, the programs generated by Metafor are not directly executable. Another work, DeepCoder [2] extends the programming by example framework *Learning Inductive Program Synthesis (LIPS)*[16] into a big data problem. DeepCoder generates a sequence of SQL-like function calls for given integer input-output examples by training a neural network to predict possible mathematical properties. However, the generated function calls are basic and low-level. In work [6], an NLPr video game design system translates object-oriented English sentences into JavaScript game code. A hybrid context and grammar analysis is used. Conditional statements also can be handled in this system.

Text generation is a topic of interest in NLP research, and it is also receiving attention in the domain of robotics. A number of systems have worked towards explaining robot behavior, including verbalizing the robot's navigation decisions [18,19] and explaining robot policies by generating behavioral explanations in NL [5]. The idea of generating low-level robot NL specifications based on robot paths presented in this paper is similar to these works: breaking down abstracted robot missions into sequential steps describing robot behaviors. However, instead of being used to explain the navigation to humans, the generated low-level NL texts are used for NLPr program synthesis.

3 Problem Formulation & System Design

3.1 High-level to Low-level (HL2LL) System Overview

Parsing and understanding the semantic meanings of high-level abstractions have been a significant challenge in NLP and NLPr research due to their complex linguistic nature. Just like explaining a complex concept to a child, one needs to break the concept down

to a sequence of discrete, straightforward, and actionable steps for machines to understand. In this work, particularly, the HL2LL mechanism is built upon a domain-specific library; in this case, the LEGO robot functions. The OOP-CNL language model L [22] is used to extract the function information from NL inputs and to match a suitable combination of robot functions in this work. In a nutshell, when the function information extracted from the high-level abstraction contains motion language features that cannot be translated into individual functions in the function library \mathcal{F} , the system would further search for identifying high-level abstractions, like color line tracking or moving to specific mission regions. The high-level abstractions can be explained using a set of low-level specifications. For example, “*The robot moves forward 10 inches.*” is an example of a low-level specification, while “*The robot walks to M4 from M1.*” is considered a high-level abstraction since it can be described using a set of low-level instructions. The transformation process, shown in Figure 2, consists of four steps:

1. Parse the high-level abstraction: Identify the task details.
2. High-level abstraction to path: Find a qualified path from the *source* to the *target* based on the given high-level abstraction using the algorithm in Section 3.3.
3. Path to low-level NL specifications: Generate a set of low-level NL specifications that describe the actions needed for the robot to follow the qualified path.
4. Low-level NL specifications to code: Translate low-level NL specifications into executable codes using the NLPr system.




Fig. 2: System overview

3.2 Map Representation

We model our robot’s task after the First LEGO League (FLL)¹ competition, with an $88'' \times 44''$ *Mission Map* based on the FLL 2018/2019 official competition arena serving as our robot’s environment, shown in Figure 3a. The arena contains eight mission regions, denoted using red blocks and several thick black lines on the map, which can be recognized using the robot’s color sensor. The *Base* located at the bottom left is the required starting point for each run. In this paper, we focus on the task of planning a path for the robot between specified mission regions. Some other actions involving motor and sensor usages can be performed in addition to navigation, as described in the LEGO NLPr system [22].

In order to simplify pathfinding, we break the *Mission Map* into grid squares, as shown in Figure 3b. We denote this grid representation the *Virtual Map*. In the *Virtual Map* asterisks denote the edge of the start region and mission regions are represented by mission blocks. Mission blocks are shaded cells of the form M_n where the n represents the mission index. Each grid block corresponds to a block with size of $4'' \times 4''$. The

¹ <https://www.firstlegoleague.org/>

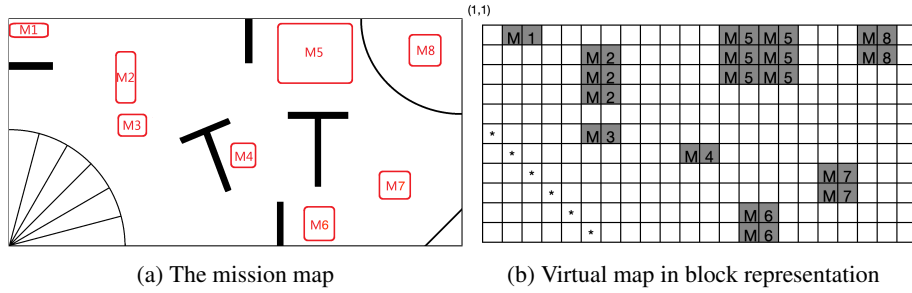


Fig. 3: Virtual game maps

top left corner of the map is initialized with coordinate $(1, 1)$. If the mission regions are treated as block-like obstacles, and the robot is restricted to movement in the cardinal directions, the task of pathfinding in the *Virtual Map* can be treated as a 2D Manhattan pathfinding problem.

3.3 Lee's Algorithm and Its Adaption

Lee's Algorithm [9] is one of the most effective breadth-first search (BFS) based single-layer routing methods for finding the shortest paths in a Manhattan graph. Lee's Algorithm searches for the target from the source using a wave-like propagation. With a source block S and a target set of adjacent blocks T , there are two main phases in Lee's Algorithm:

1. **Search:** Begin by labelling block S as k , where $k = 0$. Fill the valid neighbors of blocks labeled k (not already filled and not outside of the map) with label $k + 1$. Proceed to step $k + 1$, repeating the previous process until either the destination T is reached or there are no more valid neighbors.
2. **Retrace:** Once T has been reached, trace backward to build the path from T to S by following the descend of k from k to 0. It is possible that multiple equal-length paths exist between S and T .

Lee's Algorithm can be modified to break ties between equal-length paths in favor of the path with the fewest turns, as shown in Algorithm 1 [15]. By minimizing the number of turns that the robot makes, we reduce the number of NL sentences our system must generate and the accumulation of navigation errors that occur as the robot turns. In the **Search** process, the direction and coordinates are recorded for the **Retrace** phase's reference. An alternative method approach would be to rank paths first by the number of turns taken and only then consider the overall path length, effectively trading off reduced turning time for potentially longer paths [23]. However, FLL players need to finish as many tasks as possible within a given time limit, and as such, we prefer to rank by path length first. Figure 4 shows an example of a grid's state after Algorithm 1 has been executed. Although there are multiple equal-length paths in the grid, the path highlighted in green is chosen by the adapted Lee's Algorithm because it has the fewest turns among the eligible shortest paths.

3.4 Path Information Extraction for NLP

Information extraction [7] in NLP is the process of converting raw text into a form that can be easily processed by machines. A task-driven domain-specific function library \mathcal{F}




Fig. 4: Finding a path from the Base to M2

is used to narrow down the space of function matching for program synthesis in this study. The function library \mathcal{F} includes actions that a LEGO robot can perform with the supported sensors and motors. The key to parsing a sentence’s semantic meaning is to split the sentence into sentence frame components and identify the dependency relations in and between each frame. A grammar-based OOP-CNL model \mathcal{L} [22] is used to construct an intermediate representation for pathfinding based on part-of-speech (POS) tags [20] and parse information using NLTK toolkits [12], defined as:

$$\mathcal{L} = (O, A, P, R) \quad (1)$$

where O stands for the objects in the arena, A represents the robot actions, P indicates the adjectives or adverbs affiliated with the objects and actions, and R represents the hard-coded requirements for the objects or the actions. For the task-driven robot NLPr system, we need to identify three components: the object, the action, and the conditions.

Sentence tokens are categorized based on grammatical tags and dependency relations, and lexicon keywords such as sensor names, port numbers, and mission region names are identified after preprocessing through lemmatization and tokenization. For example, in the input sentence “A happy robot goes to M2.” O is the robot; A is go to M2; P is a Boolean state happy as shown in Figure 5. When mission regions are detected in the sentence, an error-checking step is invoked to detect any underlying errors in the text, as described in Algorithm 2. The object and action pairs identified in this step continue to a function matching process in the function library \mathcal{F} . The object robot and action go match the pathfinding function `find_path(start, target)` instead of the function `move(dir, num, unit)` because of the presence of the target M2 in this example.

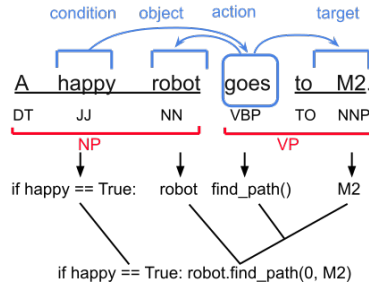


Fig. 5: Parsing a sentence and constructing the intermediate representation.

For our robot application, the number of object and action combinations is finite. For sentences with no ambiguity or errors, each \mathcal{L} should have only one valid match in the finite function library \mathcal{F} . If the system fails to identify such a 1:1 match in the

Algorithm 1 Shortest-and-fewest-turn path: Search and Retrace

```
1: procedure SEARCH(current_point, target_point)
2:   queue.push([source_point, 0])
3:   while queue do
4:     current_point, counter  $\leftarrow$  queue.popleft()
5:     i, j  $\leftarrow$  current_point.i, current_point.j
6:     emap[i][j]  $\leftarrow$  counter ▷ label the current point
7:     queue.push(neighborsOf(i, j), counter + 1) if valid
8:     if any neighbor reaches the target_point then
9:       goal_point  $\leftarrow$  (i, j), break ▷ path found
10:    if goal_point = source_point then
11:      return ▷ no such path exists
12:    else
13:      save current_dir
14:  procedure RETRACE(current_point, source_point)
15:    get_dir  $\leftarrow$  current_dir
16:    while current_point  $\neq$  source_point do
17:      i, j, id  $\leftarrow$  current_point.i, current_point.j, emap[i][j]
18:      L_id, R_id, U_id, D_id  $\leftarrow$  neighbors(emap[i][j]) if exists
19:      if get_dir  $\in$  [L, R, U, D] and (X_id = id - 1) then ▷ X: dir as id ↓ along get_dir
20:        update i, j
21:      else
22:        compare to neighbors in different dirs
23:        update i, j, get_dir
24:      current_point  $\leftarrow$  (i, j)
25:      path.push(current_point)
```

function library, the system will generate an error message with diagnostic information to help users to debug their input. When multiple objects, actions, or interpretations exist, the pre-defined higher priority functions will be chosen to ensure a sample program can be generated. For example, the sentence “*The robot goes straight to M3.*” maps to function `move(forward, 0, 0)` and function `find_path(0, M3)`. As such, the system cannot determine the user’s intention. The system responds to this situation by generating a warning message, notifying users that “straight” is ignored for this conflict. Rather than not produce any low-level output at all, the system produces an output based on the `find_path` function, as it is a higher priority action.

When multiple mission regions are present, the pathfinding process needs to be split into steps. Each input sentence describing robot navigation may contain one midpoint and one avoid-point. For example, in the sentence “*If the robot sees an obstacle in 20 inches, it goes to M7 through M3 but avoids M4*”, the path is parsed into two steps with the midpoint (through (M3)), the target (to (M7)), and the avoid-point (avoids (M4)), under the condition (`if ultrasound_sensor() < 20 inches`).

Multi-conditional statements can be handled in such a language model \mathcal{L} by processing each condition as a Boolean statement and each action separately. For the example shown in Figure 6, the sentence is processed into an *if* statement with 2 conditions:

Algorithm 2 Check for errors in the pathfinding sentence

```
1: procedure CHECK_ERRORS(tokens)
2:   tokens, unknowns = tokens.filter(lexicon)
3:   if unknowns then
4:     Warning: Skipping detected unknown tokens.
5:   obj, act ← tokens.intersection(obj_dict, act_dict)
6:   if obj ≠ robot or act ≠ find_path then                                ▷ mismatch
7:     Error: not valid combination
8:   missions ← tokens.intersection(emap)                                ▷ get all mission regions in the sentence
9:   if len(missions) ≥ 4 then
10:    Error: too many mission regions in one sentence. Consider re-write.
11:   source, target, mid_point ← dependency(tokens)
12:   if !target or any mission ∈ missions unsigned then
13:     Error: no valid target or dangling tokens
14:   else
15:     return [robot.find_path(source, mid_point, target)]
```

condition 1 (NP (color sensor) VP (see black)), condition 2 (NP (robot) VP (is happy)), and action (NP (it) VP (move to M2)). The reference relation between *it* and *robot* is done by contextual analysis on current and previous contents combined with function library restrictions. e.g. *robot* is chosen because of it matches with the action *move* behavior both contextually and functionally.




Fig. 6: Complex sentence example

While an action might have several interpretations, the functions implied in a sentence are limited by the task-driven domain-specific function library. For example, in the sentence “*When the robot sees red at M1, it will speed up and go through M2 to reach M3.*”, the *color* subject indicates a color sensor is needed. Similarly, “see a wall” indicates the ultrasound sensor usage and “touch a wall” indicates the touch sensor usage.

3.5 Path to Low-level Sentence

Once a path between *S* and *T* is identified, a sequence of low-level NL sentences describing the corresponding step-by-step actions needed to navigate the LEGO robot is generated. A grammar-based formalization method is used to construct the object-oriented low-level NL sentences. The generated NL texts will be fed to an NLPr system for further translation, as opposed to being intended for humans to read. Our proposed

method does not require a large dataset for training and can be adapted to other high-level abstractions when a suitable function library is available. If the robot starts off




Fig. 7: The robot moves from A to B.

facing North/up, path 1 in Figure 7 is described in low-level NL specifications as:

Path 1: $[4, 1] \rightarrow \dots \rightarrow [1, 1] \rightarrow \dots \rightarrow [1, 4] \Rightarrow$ The robot goes forward 12 inches. The robot turns right 90 degrees. The robot goes forward 12 inches.

The pseudocode in Algorithm 3 illustrates the above path-to-sentence conversion. First, every two neighboring coordinates in the path array are compared to detect turns and step numbers in each turn. The function `compare((pre_row, pre_col), (row, col))` returns `state` that determines if the robot needs to turn. If not, it means the robot still follows the previous direction `pre_state`. The counter records the number of steps in the current direction. Once a turn occurs, a set of NL sentences is generated based on the number of steps, recorded direction, and previous state, i.e., the function `path2NL(pre_state, dir, count)` generates the NL sentences for each turn. We then update the direction and reset the counter for the next steps.

Algorithm 3 Path to NL Generation

```

1: procedure NL TEXT GENERATION(path, direct)
2:   total_step, (pre_row, pre_col)  $\leftarrow$  len(path), path[0]           ▷ total number of steps
3:   counter, state, pre_state  $\leftarrow$  0, 0, 0
4:   for i in range(1, total_step) do
5:     row, col  $\leftarrow$  path[i]
6:     state  $\leftarrow$  compare((pre_row, pre_col), (row, col))
7:     if state = pre_state then
8:       counter += 1
9:     if state  $\neq$  pre_state or i = total_step - 1 then
10:      NL_text  $\leftarrow$  path2NL(pre_state, dir, counter)
11:      update(dir), counter  $\leftarrow$  1
12:      NL2Code(NL_text)                                           ▷ NL texts to code
13:      pre_state, pre_row, pre_col  $\leftarrow$  state, row, col

```

3.6 Generating Code from NL specifications using the NLPr System

The LEGO NLPr program synthesis system [22] generates executable text-based programs directly from the NL input instead of the graph-based programs typical of LEGO robots. It processes the input English Code (EC) into intermediate representations with lemmatization, tokenization, categorization, and function matching. Such representations indicate the desired functions that need to be translated into formal program

snippets. It is used for program synthesis and producing feedback or error information for users. The NL-to-code program synthesis system, `NL2Code(NL_text)` in Algorithm 3, calls the functions that handle the interpretation from generated low-level NL specification into executable programs. A set of robot motion functions in \mathcal{F} are combined to synthesis the programs based on the intermediate representations. For example, the sentence “*The robot goes forward for 12 inches.*” can be represented by `robot.move(forward, 12, inch)`. This representation is translated into 28 lines of code.

4 Experimental Results

We evaluate our system’s performance on a set of 162 robot pathfinding related descriptions. These descriptions were collected by the authors manually, and they collectively describe movements between all eight mission regions. Each description consists of one or more sentences.

Our system successfully translates all 36 descriptions with 2 or fewer mission regions, resulting in programs that navigate the robot on the shortest path with the fewest turns between two terminals. Of the 56 descriptions that navigate between three mission regions, 91.1% of the generated programs are correct. Our system performs worse on descriptions with more complicated structures, which might involve more than three mission regions, with only 68.6% of the 70 such descriptions are translated into programs meeting the semantic meanings of the descriptions. Overall, our system correctly translated 135 (83.3%) of the 162 collected descriptions. Some example descriptions and their corresponding number of lines of code generated are shown in Table 1.

These results show that our proposed NL text transformation system can successfully translate the large majority of the collected high-level robot navigation task sentences into low-level instructions for producing executable programs. This supports our hypothesis that with an effective POS tagging and domain-specific function library and lexicon, the objects, actions, and targets in \mathcal{L} can be effectively identified and useful intermediate representations for further program synthesis can then be generated.

However, despite our system’s strengths, it still struggles with more complicated sentence structures due to the NL’s ambiguous nature and its expressiveness. One such description that poses a challenge for our system is “*The robot wanders through M1 M2 and M3.*” This description cannot be translated properly because there is no clear indication of the robot’s source and target. As this description would be difficult for a human to adapt to low-level instructions, it is understandable that our system fails to translate it. The LEGO NLPr system can handle complex robot navigation NL descriptions by converting them into a sequence of low-level step-by-step instructions. Collected high-level robot navigation abstractions can generate a number of lines of code based on desired functions. When a description includes information that cannot be handled, a program skeleton with the best guess and debugging feedback is generated.

4.1 Case Study

Example 1 109 lines of code are generated for Example description #2 in Table 1 for navigating from M5 to M3, shown in Figure 8a.

Generated path $[1, 2] \rightarrow \dots \rightarrow [7, 2] \rightarrow \dots \rightarrow [7, 8] \rightarrow [6, 8] \rightarrow \dots \rightarrow [6, 17] \rightarrow [5, 17] \dots \rightarrow [2, 17] \dots \rightarrow [2, 19]$

Eg#	English Code Examples	# of lines
1	The robot goes from M5 to M3	99
2	The robot starts with facing to the right. The robot goes to M8 from M1 but avoids M2.	109
3	The robot goes to M1 without going through M2 via M3.	81
4	If the robot is happy, it goes to M2.	67
5	When robot does not see the red line, it goes straight to M3. Otherwise, it follows the red line.	105
6	If the robot sees an obstacle in 20 inches, it goes to M7 through M3 but avoids M4.	108
7	If the ultrasound sensor sees a ball in 5 inches, the robot is happy. A happy robot goes to M7 though M3 but avoids M4	112

Table 1: English Code examples with corresponding number of lines of code generated

Generated low-level instructions “The robot turns right 90 degrees. The robot goes forward 24 inches. The robot stops. The robot turns left 90 degrees. The robot goes forward 24 inches. The robot stops. The robot turns left 90 degrees. The robot goes forward 4 inches. The robot stops. The robot turns right 90 degrees. The robot goes forward 36 inches. The robot stops. The robot turns left 90 degrees. The robot goes forward 16 inches. The robot stops. The robot turns right 90 degrees. The robot goes forward 8 inches. The robot stops.”

Example 2 Paths for sample #7 in Table 1 are shown in Figures 8b and 8c. Sample #7 in Table 1 is a multiple-phase pathfinding task, and as such our system must compute two paths. Note that when the robot reaches mission region #3, the robot is pointing to the East/right. Therefore, the second path starts with turning to the right only 90 degrees rather than turning 180 degrees. The second sentence’s robot movements would only be performed when the state `happy` is true from the last sentence.

Generated path 1 $[11, 1] \rightarrow \dots \rightarrow [6, 1] \rightarrow \dots \rightarrow [6, 6]$

Generated path 2 $[6, 6] \rightarrow \dots \rightarrow [9, 6] \rightarrow \dots \rightarrow [9, 17]$

Generated low-level instructions “The robot goes forward 20 inches. The robot stops. The robot turns right 90 degrees. The robot goes forward 20 inches. The robot stops. The robot turns right 90 degrees. The robot goes forward 12 inches. The robot stops. The robot turns left 90 degrees. The robot goes forward 44 inches. The robot stops.”

5 Future Work

There are two main directions in which we intend to extend this work. First, we found that some input sentences may be invalid as they contain unclear or unknown information, meaning that they cannot be translated into robot functions even by a human, e.g. “*The robot hates moving forward*”. In order to make the system more robust to invalid inputs, it will be necessary to validate inputs with domain-specific formal reasoning and analysis. This will ensure the correctness of the system’s understanding of the users’ intentions and the correctness of the generated programs. Second, the function space contains several basic robot motions. As such, we intend to expand the function space with

References

1. Mobile robot programming using natural language. *Robotics and Autonomous Systems* **38**(3), 171 – 181 (2002), advances in Robot Skill Learning
2. Balog, M., Gaunt, A.L., Brockschmidt, M., Nowozin, S., Tarlow, D.: Deepcoder: Learning to write programs. *CoRR* **abs/1611.01989** (2016)
3. Dijkstra, E.W.: On the foolishness of "natural language programming". In: *Program Construction, International Summer School*. pp. 51–53. Springer-Verlag, London, UK, UK (1979)
4. Ernst, M.D.: Natural language is a programming language: Applying natural language processing to software development. In: *SNAPL 2017: the 2nd Summit on Advances in Programming Languages*. pp. 4:1–4:14. Asilomar, CA, USA (May 2017)
5. Hayes, B., Shah, J.A.: Improving robot controller transparency through autonomous policy explanation. *Association for Computing Machinery*, New York, NY, USA (2017)
6. Hsiao, M.S.: Automated program synthesis from object-oriented natural language for computer games. In: *Controlled Natural Language - Proceedings of the Sixth International Workshop, CNL 2018, Maynooth, Co. Kildare, Ireland, August 27-28, 2018*. pp. 71–74 (2018)
7. Jurafsky, D., Martin, J.H.: *Speech and Language Processing (2nd Edition)*. Prentice-Hall, Inc., USA (2009)
8. Kuhn, T.: A survey and classification of controlled natural languages. *Comput. Linguist.* **40**(1), 121–170 (Mar 2014)
9. Lee, C.Y.: An algorithm for path connections and its applications. *IRE Transactions on Electronic Computers* **EC-10**(3), 346–365 (1961)
10. Lin, X.V., Wang, C., Pang, D., Vu, K., Zettlemoyer, L., Ernst, M.D.: Program synthesis from natural language using recurrent neural networks. *Tech. Rep. UW-CSE-17-03-01*, University of Washington Department of Computer Science and Engineering, Seattle, WA, USA (Mar 2017)
11. Liu, H.: Metafor: Visualizing stories as code. In: *10th International Conference on Intelligent User Interfaces*. pp. 305–307. ACM Press (2005)
12. Loper, E., Bird, S.: *Nltk: The natural language toolkit*. Association for Computational Linguistics, USA (2002)
13. Lopes, L.S., Teixeira, A.: Human-robot interaction through spoken language dialogue. In: *Proceedings. 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000)* (Cat. No.00CH37113). vol. 1, pp. 528–534 vol.1 (2000)
14. Matuszek, C., Herbst, E., Zettlemoyer, L., Fox, D.: Learning to parse natural language commands to a robot control system. In: *Experimental Robotics*
15. Maxemchuk, N.: Routing in the manhattan street network. *IEEE Transactions on Communications* **35**(5), 503–512 (1987)
16. Menon, A.K., Tamuz, O., Gulwani, S., Lampson, B., Kalai, A.T.: A machine learning framework for programming by example. In: *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*. pp. I–187–I–195. ICML'13, JMLR.org (2013)
17. Mihalcea, R., Liu, H., Lieberman, H.: Nlp (natural language processing) for nlp (natural language programming). In: Gelbukh, A. (ed.) *Computational Linguistics and Intelligent Text Processing*. pp. 319–330. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
18. Perera, V., Selvaraj, S.P., Rosenthal, S., Veloso, M.: Dynamic generation and refinement of robot verbalization. In: *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. pp. 212–218 (2016)
19. Rosenthal, S., Selvaraj, S.P., Veloso, M.: Verbalization: Narration of autonomous robot experience. In: *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. pp. 862–868. IJCAI'16, AAAI Press (2016), <http://dl.acm.org/citation.cfm?id=3060621.3060741>

20. Toutanova, K., Klein, D., Manning, C.D., Singer, Y.: Feature-rich part-of-speech tagging with a cyclic dependency network. In: Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics. pp. 252–259 (2003), <https://www.aclweb.org/anthology/N03-1033>
21. Wang, X., Huang, Q., Çelikyilmaz, A., Gao, J., Shen, D., Wang, Y., Wang, W.Y., Zhang, L.: Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. CoRR **abs/1811.10092** (2018)
22. Zhan, Y., Hsiao, M.S.: A natural language programming application for lego mindstorms ev3. In: 2018 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR). pp. 27–34 (2018)
23. Zhou, Y., Wang, W., He, D., Wang, Z.: A fewest-turn-and-shortest path algorithm based on breadth-first search. *Geo-spatial Information Science* **17**(4), 201–207 (2014)