

Instructing Hierarchical Tasks to Robots by Verbal Commands

Péter Telkes¹, Alexandre Angleraud¹ and Roel Pieters¹

Abstract—Natural language is an effective tool for communication, as information can be expressed in different ways and at different levels of complexity. Verbal commands, utilized for instructing robot tasks, can therefore replace traditional robot programming techniques, and provide a more expressive means to assign actions and enable collaboration. However, the challenge of utilizing speech for robot programming is how actions and targets can be grounded to physical entities in the world. In addition, to be time-efficient, a balance needs to be found between fine- and course-grained commands and natural language phrases. In this work we provide a framework for instructing tasks to robots by verbal commands. The framework includes functionalities for single commands to actions and targets, as well as longer-term sequences of actions, thereby providing a hierarchical structure to the robot tasks. Experimental evaluation demonstrates the functionalities of the framework by human collaboration with a robot in different tasks, with different levels of complexity. The tools are provided open-source at <https://petim44.github.io/voice-jogger/>

I. INTRODUCTION

Speech and natural language is the most common modality for interaction between humans, due to its expressive nature and its ability for dialogue. Communication between humans and robots can benefit as well from such functionalities, and, as a research field, efforts to achieve this have been ongoing since several decades [1]. The main challenge in instructing robots by verbal commands is how natural language can be easily grounded to the actions of a robot. On the one hand this can be done by directly matching input commands to a single robot action or target location in the world (i.e., lexical grounding), thereby effectively performing teleoperation. On the other hand, learning-based approaches can generalize input commands to desired outcomes in the world, from Large Language Models (LLM), possibly assisted by visual perception (vision-language pre-training, or VLP).

While lexical grounding provides a one-to-one connection from commands to actions, learning-based models do not provide this, making them less reliable to achieve desired outcomes. In both cases basic robot functionalities are needed that act in the world, such as movement primitives in world space, gripper actions and higher-level hierarchical tasks that execute multiple actions in sequence. Most importantly, the grounding of (verbal) natural language commands to robot actions requires manual selection of what robot actions are needed for what tasks. While LLMs and VLP can help in extracting relevant information from language and the scene, they cannot alone solve this grounding problem.

¹Cognitive Robotics group, Unit of Automation Technology and Mechanical Engineering, Tampere University, 33720, Tampere, Finland; firstname.surname@tuni.fi




Fig. 1: Instructing robots by speech requires fine-grained verbal commands (left) that directly refer to actions and targets in the workspace. This work presents methods that extend verbal command instructions to hierarchical tasks.

In this work, we present a framework for verbal commands as tool to instruct tasks to robots, including automatic speech recognition and the grounding of commands to robot actions and targets in the work space (see Fig. 1). Single and multiple command phrases from natural language enable different functionalities to be included such as task hierarchies and task concatenation, without explicitly relying on high-level task planners. These capabilities are of particular importance to ensure smooth and efficient collaboration between human and robot, as continuously using fine-grained language commands is time-consuming (see Fig. 1). Course-grained commands enable high-level functionalities and long-term sequences of actions as demonstrated by our work with two hierarchical tasks for gear assembly. Our developments and tools are integrated in a robotic system and numerous experiments with single commands and long-term tasks demonstrate the results.

The contributions of this work are as follow:

- Framework for instructing robots by verbal commands with basic robot functionalities
- Grounding of commands to robot actions and targets in the scene
- High-level functionalities to support task hierarchies and concatenation
- Experimental demonstration of the framework with representative tasks

II. RELATED WORK

A. Human-Robot Collaboration

Collaboration between human and robot can be supported in many ways by different technologies [2]. While physical interaction (e.g., hand-guiding or learning by demonstration) and visual perception (e.g., for object and gesture detection) are popular modalities [3], also audio perception in the form of natural language and speech is common [1]. Success of the collaboration depends on different factors, such as technical quality of the tools, and how communication between human and robot is enabled [4]. In context to this work, the collaboration between human and robot has the purpose to provide robot instructions via speech. This means that communication is only directed from human to robot, and human verbal commands (i.e., natural language) need to be recognized and interpreted to be utilized for robot task execution. Several earlier works have developed capabilities that enable robots to act on voice commands [1], with different levels of complexity and even with additional sensor modalities to support the understanding of the world and tasks within it.

The related work therefore covers the state of the art in natural language processing and further divides robotics approaches in two directions; lexically-grounded methods and learning-based methods (see Table I).

B. Natural language processing

Speech recognition and natural language processing (NLP) are two related fields, which both have been propelled forward recently due to the application of deep learning models [5], [6]. For this work, we focus on NLP for the understanding of text and assume that a correct recognition of speech (by e.g. Vosk [7]) provides natural language (words or sentences) in textual format. The interpretation of language by Large Language Models (LLM), such as BERT [8] and PaLM [9], have demonstrated capabilities including language understanding and generation, reasoning, and code-related tasks. As such, to enable robot instructions, they still need the explicit grounding of verbal commands to actions of the robot or objects in the scene.

C. Lexical-grounding of robot commands

Commanding robots from natural language by lexical grounding has been demonstrated for single verbal commands [10], command pairs [11] and short phrases [12]. In most of these cases, words to be recognized and tasks to be executed are predefined, limiting the applicability of the methods to new situations and new tasks. Verbal commands have also been utilized to aid in hand-guided task programming [13], thereby effectively enabling multi-modal robot programming. A list of supported single word commands can be commanded and, separately, different parameters can be set via an app to adapt different parameters for robot motion and skills. Speech interfaces have also been utilized for modification of robot skills [14], utilizing bi-directional and high-level communication between human and robot. Interaction scenarios are developed that show long human

TABLE I: Comparison between two approaches of utilizing speech in robotics

	Lexical grounding	Learning-based
Approach	Match words directly to robot actions and targets in the workspace	Use high-level natural language instructions to learn robot actions in the workspace
Commands	Fine-grained	Course-grained
Vocabulary	Limited	Expressive
Actions	Short-time	Mid- to long-term
Response	Fast	Slow
References	Single commands [12], [13] Command pairs [11] Command sentences [10] Hierarchical tasks [15], [16] This work	CLIPort [21] PaLM-E [22] Interactive language [23]

commands and instructions to modify robot knowledge. In [15] the learning of complex hierarchical tasks from natural language instruction is demonstrated, utilizing hierarchical task networks (HTNs) and queries from robot to human in case tasks are unknown. Similarly, in [16] behavior trees are utilized to learn and perform novel complex tasks. As such, both these approaches focused on the planning of complex tasks and the learning of task hierarchies and behaviors. Complex natural language instructions have also been utilized for planning in linear temporal logic formulation [17], with incomplete world knowledge [18] and for the ordering [19] and allocation [20] of tasks to be executed.

What all of the mentioned approaches have in common is that verbal commands directly result in robot actions, or indirectly through task plans. As hand-crafted algorithms extract relevant words from input, this means approaches are typically limited to a predefined vocabularies and action plans.

D. Learning-based robot commands

Natural language as input for robot tasks has seen recent interest by generating task policies that are language-conditioned [24]. Even more recent are the inclusion of other modalities, in addition to natural language, to establish the link between words and perception [25], [21], [26], for tasks such as robot object manipulation and motion trajectory generation. While these methods prove well-capable of transforming natural commands to robot actions, the question remains whether such vision-and-language pretraining improves lexical grounding [27]. This embodiment issue is addressed in [22] and [28] which directly incorporate real-world continuous sensor modalities into a LLM (PaLM) to plan and execute long horizon tasks, i.e., robot planning on table-top and object manipulation and navigation in office kitchen environments, respectively. Another example is embodied BERT (emBERT) [29], which grounds language instructions against visual observations and actions to take in an environment, with the aim of home and office navigation. Finally, real-time (i.e., 5Hz) interaction by speech is presented in [23], which demonstrates the instruction of multiple robots on table-top setting at the same time.

For both lexical grounding and learning-based approaches, still a connection needs to be made to robot actions and the speech commands. While this connection can be learned, as demonstrated by learning-based approaches, only a person can verify whether the learned grounding is correct. In case of high-performance or safety critical tasks (e.g., industry or search-and-rescue, as compared to household actions by a service robot) such uncertainty might not be acceptable. Our lexical-grounding approach is human-coordinated and aims to enable basic and hierarchical tasks to be commanded from human to robot, with both fine and course-grained commands.

III. APPROACH

Our approach divides robot instructions into basic and hierarchical commands (see Table II), enabling both low- and high-level commands to be instructed at the same time. New target poses and tasks can be defined at any time with any combination of basic and hierarchical commands.

A. Grounding of commands

Grounding of verbal commands is done by connecting individual words to actions or targets, when they are detected from automatic speech recognition. Robot actions can be named with any spoken word and multiple words can refer to the same action. This similarly applies to robot pose targets (command `save`) and new defined (hierarchical) tasks (command `record`). Verbal commands are also grounded to robot motion in case of single direction motion (command `<direction> <value>`) and other robot actions (`home`, `open`, `close`, `rotate`, etc.).

B. Basic commands

All relevant commands needed to enable a functional robot are referred to as basic commands. This includes commands to start and stop the robot system and individual motion actions of the different components (single direction step motion, continuous point-to-point motion, gripper motion, etc.). Depending on the functionality of the command, different instruction combinations can be given as arguments, as described in Table II. Basic commands can be utilized at any point during robot operation and can be used for hierarchical tasks in short- or long-term form.

C. Hierarchical task commands

Hierarchical tasks offer the ordering of robot actions in different ways (see Table II). For example, task recording (initiated by command `record`, terminated by command `finish`) enables task hierarchies to be assigned by the human operator. In its simplest form these contain a list of actions and commands that are executed in sequence, as described by few included template commands (i.e., `pick`, `place`, `stack`, `push`). Task concatenation is enabled by the `and` command, by appending following commands to a list. Combining multiple robot motion actions is enabled by the `then` command, which computes the spoken motion phrases to one single motion command outcome. For example, multiple directional commands in sequence will simply

TABLE II: List of available basic and hierarchical commands for instructing robot motion and tasks.

Basic commands	Explanation
<code>start robot</code>	Initialize robot
<code>stop robot</code>	Shut-down robot
<code>set mode step</code>	Set robot to move in steps
<code>set mode continuous</code>	Set robot to move continuous
<code>save position <string></code>	Save robot pose as <string>
<code>position <string></code>	Move robot to pose <string>
<code><direction> <value></code>	Move robot <value> {up, down, left, right, back, front}
<code>stop execution</code>	Stop robot motion
<code>step size <value></code>	Set robot motion step size
<code>open</code>	Open gripper
<code>close</code>	Close gripper
<code>rotate <value></code>	Rotate gripper by <value> deg.
<code>home</code>	Move robot to home position
Hierarchical commands	Explanation
<code>record <string></code>	Start recording commands as <string>
<code>finish</code>	Stop recording command sequence
<code>task <string></code>	Execute command sequence <string>
<code>repeat <task> <string></code>	Repetition of <task>, continuously or specified number of times by <string>
<code>repeat <task> <list></code>	Repetition of <task>, until <list> of positions is empty
<code>pick <string></code>	pick object from pose <string>
<code>place <string></code>	place object to pose <string>
<code>stack <string> <value></code>	Pick and place on pose <string> with vertical offset <value>
<code>push <string> <direction></code>	Planar motion at pose <string> starting from <direction>

be added up as a robot pose target (i.e., down then down then left 200 mm will result in a single robot pose translation target with the spoken offsets). Task repetition is enabled by the commands `again` and `repeat`, given suitable input for executing the motion commands (i.e., list of robot poses and number of repetitions).

New (hierarchical) commands can be easily defined, utilizing existing commands and by providing desired arguments. Several typical robot manipulation commands, such as `pick-and-place`, `wiping` and `polishing` have been developed this way and are demonstrated individually and as integrated in two assembly use cases in Section IV.

D. Robot Control Architecture

The architecture for controlling the robot by verbal commands includes modules for command detection, robot motion planning and robot control as follows. Command detection recognizes speech commands, adds these to a queue and publishes them as suitable messages over ROS topics. This means that the system is continuously listening to audio and adds new robot tasks when they are received. A high priority topic is used for the `stop` command, which allows to bypass the queue and executes a stopping motion immediately, with suitable motion profile. In case of teaching positions or task hierarchies, the system stores them locally in custom XML format with commanded name, to be used when requested. The motion planning module relies on MoveIt2 and is continuously waiting for new motion commands from the command topics. Finally, robot control utilizes the Franka Control Interface (FCI) to achieve real-time execution of control actions.

IV. RESULTS AND DISCUSSION

A. Hardware and implementation

The experimental setup includes a Franka Emika robot with parallel gripper. Robot motion control utilizes Moveit2 motion planner for generating work space motion and ROS (ROS.org) handles all communication. Speech of a person is recorded by a standard microphone (16kHz audio signal), transmitted to the main PC and passed through a voice activity detector [30], to identify if the individual audio packet has human voice or not. A time-delayed filter is applied on the running stream of audio packets to consider the natural pause in human speech while speaking. For speech recognition, Vosk [7] is utilized, which is trained on The People’s Speech dataset [31]. The model is configured by filtering out unnecessary words, which are unsuitable for robot instructions, and by including words and sentence instructions that reflect the content of speech to be expected. All packages, tools and videos are open-source available to replicate our work at <https://petim44.github.io/voice-jogger/>.

B. Speech recognition

The accuracy of speech recognition was tested in a clean and noisy environment with a non-native English speaker (see Table III). In the clean case, verbal commands are recognized without any failures. In the noisy environment the audio contains a background noise similar to a busy restaurant with other people speaking, which deteriorates the detection rate slightly. Speech recognition is also evaluated in the use cases for single and hierarchical tasks, which resulted in a high success rate (see Table V).

TABLE III: Speech recognition experiment results

Command	Clean	Noisy
start robot	10/10	10/10
stop robot	10/10	10/10
move up	10/10	10/10
move down	10/10	10/10
move left	10/10	10/10
move right	10/10	08/10
move back	10/10	10/10
move front	10/10	09/10
stop execution	10/10	10/10
step size <value>	10/10	10/10
open tool	10/10	07/10
close tool	10/10	07/10
rotate tool <value>	10/10	10/10
save position <value>	10/10	10/10
load position <value>	10/10	10/10
home	10/10	08/10

C. Single tasks

Single tasks refer to short robot actions, which by definition, can have internal hierarchical structure. This is demonstrated by a pick and place action as it is composed into hierarchical tasks by verbal instructions (see Table IV). From the basic and hierarchical commands available, first robot end-effector poses are defined (save position <value>) in combination with intermediate robot motion and gripper actions (e.g., move, close, open), see left column in Table II. As a result, these robot poses are utilized

TABLE IV: A pick and place task composed into hierarchical tasks by verbal instructions. The original query required 27 commands and, with tasks composed to replace step-wise instructions, the query is narrowed down to 19 commands and, finally, a single task command.

Commands	Task	Commands	Task
left 300 mm	} pick	position pick	} one
down 300 mm		step size low	
down		record one	
front		down (x2)	
down		back	
save position pick		down	
step size low		close	
down (x3)		step size medium	
back		up	
close		Position place	
step size medium	step size low		
up	down (x3)		
left 500 mm	} place	open	
back 200 mm		position place	
save position place		home	
step size low		finish	
down (x4)			
open			
step size high			
up			
home			

for motion commands and included in a recorded hierarchical task, see right column in Table II. This effectively means that from an initial low-level set of commands (i.e., 27 verbal commands), human instructions can be narrowed down to 19 commands and, finally, a single task command, to achieve the same robot action outcome.

Table V describes the results of several single tasks that are commanded by verbal instructions. During the preparation phase specific robot poses and motion sequences are defined as required for the tasks, in few single and multiple commands. During the execution phase robot actions utilize these poses and motions for specific objects and target applications. Repetition of single tasks towards automated applications is possible as well, by specifying the task and how often it should be repeated. In case suitable, offsets in motion commands can be included to achieve more complex operations, such as wiping and polishing a larger surface. Fig. 2 depicts several snapshots of the single tasks.

D. Hierarchical tasks

To demonstrate further the capabilities of verbal command instructions, experiments were also performed with longer-term tasks with complex assembly steps. This included the assembly of a helical gear system with six assembly steps and a planetary gear system with nine assembly steps. In both cases human verbal instructions guided the assembly procedure and, in the case of the planetary gear, also contained three manual and collaborative tasks. Similar to single tasks, a preparation phase records specific robot poses and motion sequences, by hand-guiding and verbal commands. The execution phase entails the actual assembly task, coordinated by the person with verbal commands. Table V describes the results of these hierarchical tasks averaged over ten trials each.

Helical gear assembly: The helical gear consists of four parts (two gears, front and back plate) and requires six steps for assembly, executed by the robot. This includes five pick-and-place actions with different locations and one pushing action to connect the two gears together. Fig. 4 depicts several snapshots of the helical gear assembly task.

Planetary gear assembly: The planetary gear consists of seven major parts (housing, sun gear, three planet gears and top plate, excluding bolts, nuts and small gears) and requires nine major steps for assembly, six executed by the robot and three by a human operator. This includes five pick-and-place actions with different locations, one holding action and three human assembly actions (i.e., fixing the housing with four bolts, fixing planet gears on the sun gear and fixing the top plate with four bolts). Fig. 3 depicts several snapshots of the planetary gear assembly task.

In the preparation phase of both assembly cases four poses are demonstrated to the robot, by utilizing both hand-guiding and verbal commands. In addition, in the preparation and execution phases, a high number of speech commands are used (see Table V), as the assembly steps require fine adjustments before executing the assembly steps.

E. Discussion

This work demonstrated how verbal commands can instruct robot actions to achieve single and hierarchical tasks. The approach allows for the fast generation of robot actions and tasks, directly instructed by a person and without the need for robot programming or training a model (see Table I and [21], [22]). Compared to offline programming, benefits can be especially identified when changes (e.g., pick-up location, motion sequence) need to be made during or after task execution. Speech commands allow for direct adaptation of robot motion parameters, as compared to the more traditional software practices (i.e., code compilation and evaluation). In addition, the functionalities for task concatenation enable long-term sequences of actions to be scheduled, without explicit task planners and without complex structures for task dependencies.

Natural language interaction has the benefit of expressivity [4], as tasks can be instructed by commands selected by the operator in any desired level of granularity. This means that accurate motions involving contact can be done more carefully and under supervision of the operator (see Fig. 4f), leading to real-time interaction between human and robot. To ensure a responsive system, speech recognition is continuously listening for input commands and all tasks can be interrupted and overwritten by new tasks, as decided by the human operator.

Limitations of our work can be identified as well, and relate to the use of only speech for instructing commands, requiring human actions and feedback for successful robot tasks. For example, objects need to be placed on known or instructed locations and fine-grained verbal instructions are needed for delicate and contact manipulation actions. Vision as additional modality would provide a solution to these limitations, either as single perception model for object

detection [2] or as visual-speech model that combines speech commands with low-level action sequences [23].

One additional limitation is that no formal planner is utilized for constructing a task plan. While this keeps the sequencing of hierarchical tasks simple, it also means that pre-conditions or dependencies between tasks are not included. Moreover, analysis on whether a (hierarchical) task sequence is suitable or optimal needs to be done by the human operator.

Failure cases of our work were mostly caused by the limitations in automatic speech recognition (see Table III and V). While current state-of-the-art language models are very robust, still disturbances occurred in case of verbal accents and words that are similar in sound (e.g., 'four' and 'for', 'tool' and 'two'). These issues were avoided by intentional clear speech pronunciations and excluding words from speech recognition that were often mis-recognized. As can be seen in the videos, few times commands had to be repeated to achieve correct speech recognition.

Future work will investigate how Large Language Models (LLMs) can be utilized to assist in interpreting high-level verbal commands, such that more natural sentences can be instructed to the robot, instead of a predefined vocabulary. In addition, LMMs could provide additional functionalities, such as generating hierarchical tasks or complete assembly sequences automatically in the most effective and efficient manner.

V. CONCLUSION

This work presented a framework for instructing robot action commands by human natural language. Instructions can be both fine-grained and provide the functionalities for single and hierarchical tasks, as specified by a human operator in real-time. Instructing verbal commands includes a preparation phase where all required robot poses are demonstrated to the robot, utilizing both hand-guiding and verbal commands. In addition, hierarchical tasks can be composed to combine and concatenate robot actions. The execution phase entails the actual assembly task, coordinated by the person with verbal commands. Experiments with several single and hierarchical assembly tasks demonstrate that verbal commands can replace traditional robot programming techniques, and provide a more expressive means to assign robot actions and enable human-robot collaboration. To provide more natural language for instructions, future work will investigate how Large Language Models can be utilized in the framework.

ACKNOWLEDGEMENTS

Project funding was received from European Union's Horizon 2020 research and innovation programme, grant agreement no. 871252 (METRICS). The authors thank all students in the Robotics Project Work course at Tampere University that contributed to the developments of the speech tool: Wajeh Ahmad, Ara Jo, Omar Hassan, Tuomas Kiviö, Mikko Kulju, Waleed Rafi, Niklas Sorri and Kalle Tanninen.

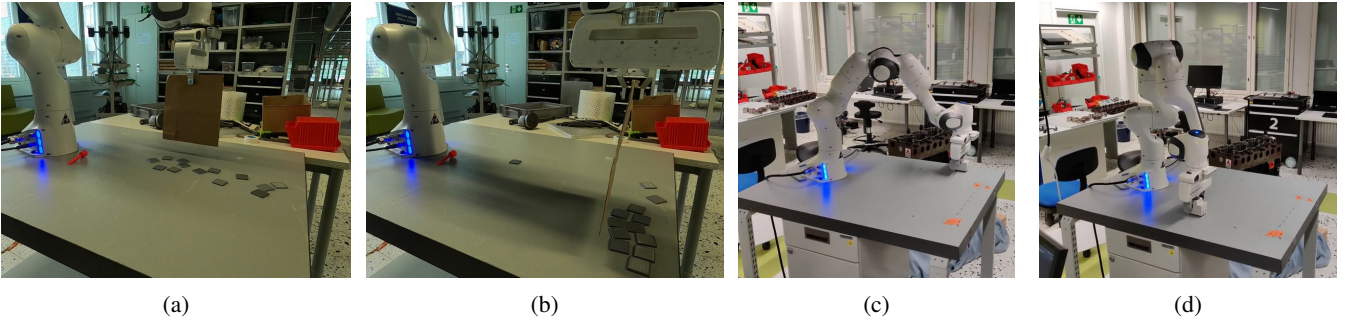


Fig. 2: Snapshots of verbally commanded single tasks: before (a) and after (b) the wiping task, picking (c) and placing (d).




Fig. 3: Snapshots of the verbally commanded helical gear assembly task: (a) pick back plate with pick part, (b) pick gear with pick part, (c) place gear with place gear three, (d) push gears together with forward ninety, (e) pick both gears with close and (f) place top plate with down three.




Fig. 4: Snapshots of the planetary gear assembly task: (a) place motor with hold table distance thirty, (b) pick housing with pick part, (c) place housing on motor with stack table distance fifty, (d) pick sun gear with pick part, (e) operator fixes planet gears on sun gear, (f) place sun gear in housing with hold table distance one hundred, (g) place top plate on housing with place top (h) operator fixes top plate to housing.

TABLE V: Results for speech commanded tasks. The preparation phase aims to define poses and record task sequences, by both verbal commands and hand-guiding. The execution phase entails the actual task, coordinated by the person with verbal commands. Results for both hierarchical tasks were obtained from ten trials.

	Single tasks					Hierarchical tasks	
	Pick	Place	Pick & Place	Wipe	Polish	Helical gear assembly	Planetary gear assembly
Preparation phase							
Prerecorded poses	1	1	3	3	2	4	4
No. of commands (single, multiple)	12 (8/4)	11 (6/5)	27 (16/11)	22 (16/6)	5 (2/3)	44 (15/21)	34 (11/20)
Speech recognition success rate	12/12 100%	11/11 100%	27/27 100%	22/22 100%	5/5 100%	36/44 82%	31/34 91%
Duration	0:45	0:30	1:28	1:26	0:36	6:12	3:14
Execution phase							
No. of automated steps	1/1	1/1	2/2	1/1	1/1	6/6	6/8
No. of commands (single/multiple)	6 (4/2)	6 (4/2)	17 (11/6)	6 (2/3)	3 (1/2)	55 (21/26)	65 (25/28)
Speech recognition success rate	6/6 100%	6/6 100%	17/17 100%	5/6 83%	3/3 100%	46/55 83%	53/65 82%
Duration	0:23	0:22	1:06	2:20	0:30	4:18	8:50

REFERENCES

- [1] S. Tellex, N. Gopalan, H. Kress-Gazit, and C. Matuszek, "Robots that use language," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, no. 1, pp. 25–55, 2020.
- [2] V. Villani, F. Pini, F. Leali, and C. Secchi, "Survey on human-robot collaboration in industrial settings: Safety, intuitive interfaces and applications," *Mechatronics*, vol. 55, pp. 248–266, 2018.
- [3] S. Gross and B. Krenn, "A communicative perspective on human-robot collaboration in industry: Mapping communicative modes on collaborative scenarios," *International Journal of Social Robotics*, pp. 1–18, 2023.
- [4] M. Marge *et al.*, "Spoken language interaction with robots: Recommendations for future research," *Computer Speech & Language*, vol. 71, p. 101255, 2022.
- [5] D. Yu and L. Deng, *Automatic speech recognition*. Springer, 2016, vol. 1.
- [6] D. W. Otter, J. R. Medina, and J. K. Kalita, "A survey of the usages of deep learning for natural language processing," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 2, pp. 604–624, 2021.
- [7] Alpha Cephei, "Vosk Speech Recognition Toolkit," <https://github.com/alphacep/vosk-api>, 2023.
- [8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [9] A. Chowdhery *et al.*, "PaLM: Scaling language modeling with pathways," *arXiv preprint arXiv:2204.02311*, 2022.
- [10] M. Stenmark and J. Malec, "Describing constraint-based assembly tasks in unstructured natural language," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 3056–3061, 2014.
- [11] A. Angleraud, A. Mehman Sefat, M. Netzev, and R. Pieters, "Coordinating shared tasks in human-robot collaboration by commands," *Frontiers in Robotics and AI*, vol. 8, 2021.
- [12] A. C. Tenorio-Gonzalez, E. F. Morales, and L. Villasenor-Pineda, "Teaching a robot to perform tasks with voice commands," in *Mexican International Conference on Artificial Intelligence*. Springer, 2010, pp. 105–116.
- [13] T. B. Ionescu and S. Schlund, "Programming cobots by voice: A human-centered, web-based approach," *Procedia CIRP*, vol. 97, pp. 123–129, 2021.
- [14] T. Frasca, B. Oosterveld, M. Chita-Tegmark, and M. Scheutz, "Enabling fast instruction-based modification of learned robot skills," in *AAAI Conference on Artificial Intelligence*, vol. 35, no. 7, 2021, pp. 6075–6083.
- [15] G. Suddrey, C. Lehnert, M. Eich, F. Maire, and J. Roberts, "Teaching robots generalizable hierarchical tasks through natural language instruction," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 201–208, 2016.
- [16] G. Suddrey, B. Talbot, and F. Maire, "Learning and executing re-usable behaviour trees from natural language instruction," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10643–10650, 2022.
- [17] A. Boteanu, T. Howard, J. Arkin, and H. Kress-Gazit, "A model for verifiable grounding and execution of complex natural language instructions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 2649–2654.
- [18] D. Nyga *et al.*, "Grounding robot plans from natural language instructions with incomplete world knowledge," in *Conference on robot learning*, 2018, pp. 714–723.
- [19] P. Pramanick, H. B. Barua, and C. Sarkar, "Decomplex: Task planning from complex natural instructions by a collocating robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 6894–6901.
- [20] J. K. Behrens, K. Stepanova, R. Lange, and R. Skoviera, "Specifying dual-arm robot planning problems through natural language and demonstration," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2622–2629, 2019.
- [21] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Conference on Robot Learning*, 2022, pp. 894–906.
- [22] D. Driess *et al.*, "PaLM-E: An embodied multimodal language model," *arXiv preprint arXiv:2303.03378*, 2023.
- [23] C. Lynch *et al.*, "Interactive language: Talking to robots in real time," *arXiv preprint arXiv:2210.06407*, 2022.
- [24] J. Luketina *et al.*, "A survey of reinforcement learning informed by natural language," *arXiv preprint arXiv:1906.03926*, 2019.
- [25] L. Shao, T. Migimatsu, Q. Zhang, K. Yang, and J. Bohg, "Concept2robot: Learning manipulation concepts from instructions and human demonstrations," in *Robotics: Science and Systems (RSS)*, 2020.
- [26] P. Sharma *et al.*, "Correcting robot plans with natural language feedback," *arXiv preprint arXiv:2204.05186*, 2022.
- [27] T. Yun, C. Sun, and E. Pavlick, "Does vision-and-language pretraining improve lexical grounding?" 2021.
- [28] M. Ahn *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," 2022.
- [29] A. Suglia, Q. Gao, J. Thomason, G. Thattai, and G. Sukhatme, "Embodied BERT: A transformer model for embodied, language-guided visual task completion," *arXiv preprint arXiv:2108.04927*, 2021.
- [30] Silero Team, "Silero VAD: pre-trained enterprise-grade Voice Activity Detector (VAD), Number Detector and Language Classifier," <https://github.com/snakers4/silero-vad>, 2021.
- [31] D. Galvez *et al.*, "The People's Speech: A Large-Scale Diverse English Speech Recognition Dataset for Commercial Usage," *arXiv preprint arXiv:2111.09344*, 2021.