

Chain-of-Modality: Learning Manipulation Programs from Multimodal Human Videos with Vision-Language-Models

Chen Wang^{1†}, Fei Xia¹, Wenhao Yu¹, Tingnan Zhang¹, Ruohan Zhang²,
C. Karen Liu², Li Fei-Fei², Jie Tan¹, Jacky Liang¹

Abstract—Learning to perform manipulation tasks from human videos is a promising approach for teaching robots. However, many manipulation tasks require changing control parameters during task execution, such as force, which visual data alone cannot capture. In this work, we leverage sensing devices such as armbands that measure human muscle activities and microphones that record sound, to capture the details in the human manipulation process, and enable robots to extract task plans and control parameters to perform the same task. To achieve this, we introduce Chain-of-Modality (CoM), a prompting strategy that enables Vision Language Models to reason about multimodal human demonstration data — videos coupled with muscle or audio signals. By progressively integrating information from each modality, CoM refines a task plan and generates detailed control parameters, enabling robots to perform manipulation tasks based on a single multimodal human video prompt. Our experiments show that CoM delivers a threefold improvement in accuracy for extracting task plans and control parameters compared to baselines, with strong generalization to new task setups and objects in real-world robot experiments. Videos and code are available at [chain-of-modality.github.io](https://github.com/chain-of-modality)

I. INTRODUCTION

Can robots learn to perform physically challenging manipulation tasks (e.g., twisting to open water bottles or drumming) by watching only one human hand video demonstration? One way to enable this capability is through recognizing human task plans from video and then translating them into executable robot skills. While recent advances in video understanding show promising results in action recognition [26], many manipulation skills require precise specification of the control parameters, which cannot easily be inferred from pure visual information, e.g. grasp *lightly* to rotate a key in hand, push *harder* to insert a plug, hit a drum *gently* to produce a soft sound. This limitation restricts robots’ ability to perform diverse manipulation tasks by only watching human videos.

A core challenge here is that extracting task plans from human video data is difficult since vision-only data lacks necessary details in recognizing these plans. A key observation is that many details in human task plans, such as control parameters like force and speed, are better captured through additional signals such as human muscle activities and object interaction sound. For example, when inserting a power plug, humans would first *Grasp* the plug with *low force* to adjust its orientation in hand, then hold it *firmly* while *Insert* it into the socket. In this work, we leverage sensing devices such as modern armbands equipped with muscle sensors and sports cameras with microphones to collect multimodal demonstration videos that include images, muscle activity, and object interaction

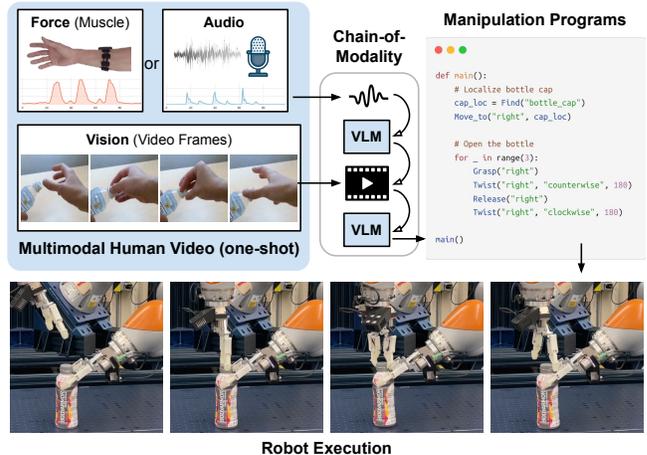


Fig. 1: We introduce Chain-of-Modality (CoM), a prompting strategy that enables VLMs to recognize human task plans from a single multimodal video with force or audio information, and generate corresponding robot control code to reproduce the task.

sounds. They provide additional information about *when* and *how* humans exert physical effort during manipulation. However, effectively utilizing these signals requires new methods for reasoning from multimodal human demonstration videos.

Vision-language models (VLMs) are capable of solving a broad range of practical problems, from visual reasoning to signal processing [63] and even code generation for controlling robots [30]. Recent advancements in long-context input have further enabled VLMs to take videos and long-sequence numerical signals as inputs [10]. This makes us ask: Can VLMs serve as general reasoning models to infer human task plans from multimodal demonstration videos? Most VLM applications still only take one modality as input. To tackle this challenge, we introduce Chain-of-Modality (CoM), a framework that prompts VLMs to analyze each modality one after another, progressively refining the answer by incorporating new information from each modality.

CoM enables VLMs to extract task plans and control parameters by analyzing a single multimodal human video. With CoM, the inclusion of additional modalities helps VLMs better segment subtasks. For instance, when a human is opening a bottle, three peak signals in the force data indicate three *Twist* motions. CoM enables the VLMs to leverage this information to first segment the entire task into a coarse task skeleton and progressively filling in more details by incorporating other modalities. Additionally, force information acquired from additional modalities allows VLMs to generate

¹Google DeepMind ²Stanford University

[†]Work done while interning at Google DeepMind

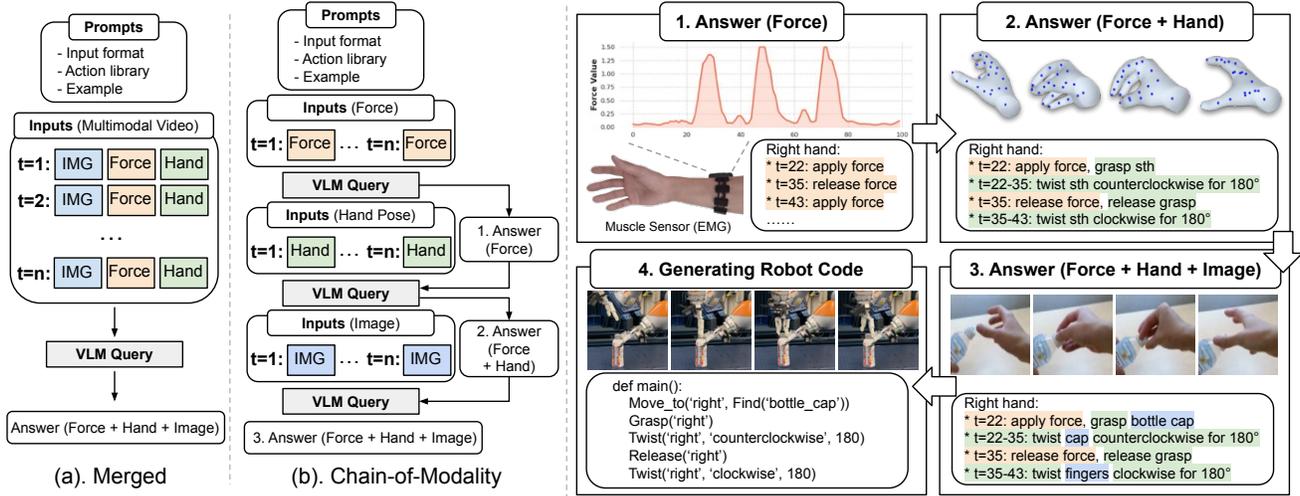


Fig. 2: Overview of Chain-of-Modality (using force as an example). **(a) Baseline method - Merged:** Merges multimodal information (vision, force, and hand pose) into a single input batch and queries the VLM to directly generate the final answer. **(b) Chain-of-Modality (CoM):** Analyzes each modality step-by-step, refining the analysis to produce the final answer. **Example:** First, the VLM uses force data to determine when force is applied. Then, with hand pose information, it infers that the human is grasping and twisting. Next, with image data, the VLM identifies the action as twisting a bottle cap. Finally, VLM transform the CoM analysis into a robot-executable Python program to reproduce the task.

more accurate control parameters for skills like *Grasp* and *Hit* with different level of force. Empirically, we found that CoM has achieved a 60% accuracy in extracting *exact* task plans and control parameters from human video. Methods that rely solely on vision-only data have zero accuracy, and naïve methods that directly query the model with all modalities in a single batch achieve an average accuracy of 17%.

How do we turn these task plans into robot actions? Prior work [30] has demonstrated that foundation models can generate robot-executable API calls based on language instructions. Our setting is different in the sense that we ask foundation models to generate API calls based on a multimodal demonstration video. These API calls provide benefits of cross-embodiment generalization since the code API can abstract away robot embodiment, allowing smooth deployment across different robots. In addition, code API based on advanced vision models further enable the robot to generalize to novel objects and unseen object configurations.

Our main contributions are as follows:

- **Chain-of-Modality (CoM):** A prompting strategy that enables VLMs to reason from multimodal human video demonstration data by progressively incorporating vision and force information.
- **One-shot manipulation program generation:** A pipeline for generating robot control programs from a single multimodal human demonstration video, integrating force information (obtained through muscle or audio signals) to produce fine-grained control parameters of different skills.
- **Generality:** We show the benefit of CoM is consistent across two advanced VLM models, and that our method allows the VLMs to learn from a single human video to write robot code that works on different real-world robot platforms with generalization capabilities.

II. RELATED WORK

Human Activity Understanding from Video. Understanding human activities in video has been a long-standing focus in computer vision [6, 12, 16, 17, 29]. Early works primarily aimed to capture the high-level semantic meaning of videos through classification [56, 61]. In efforts to extract more detailed information, later studies began to focus on deriving task plans from videos [7, 20, 67]. However, these methods tend to be limited by their reliance on specific training datasets, making it challenging to generalize to unseen action categories. In recent years, the development of large vision-language models has led to impressive results in prompting VLMs to understand human activities from videos [25]. Unlike prior works, our work focuses on reasoning multimodal human videos with force or audio information, providing essential information for downstream fine-grained robot manipulation tasks.

Foundation Models for Robotics and Control. In recent years, foundation models have demonstrated significant progress in robotics, spanning high-level reasoning to low-level control [15, 19]. Earlier works primarily focused on language-conditioned robotic reasoning and planning, where tasks were defined using natural language [2, 9, 21, 23, 32, 33, 39, 51, 57, 66]. However, some manipulation tasks—especially those involving spatial ambiguity or requiring fine-grained control—are difficult to specify using language alone. Recent advances in vision-language models (VLMs) have introduced more expressive task specifications, such as visual annotations [18, 22, 36, 54]. Our work, on the other hand, uses single-shot multimodal video as task specifications, enabling robots to extract task plans and control parameters from human demonstrations. To apply foundation models to robot control, several promising approaches have emerged, including subgoal selection for goal-conditioned policies [11, 46, 50], reward or constraint generation for trajectory optimization [22,

31, 64], and code generation based on perceptual and control primitives [30, 52]. Unlike these approaches, which are conditioned on language inputs, we demonstrate how VLMs can directly reason from *one-shot* human video inputs to generate low-level manipulation programs, providing an alternative way to prompt robot to perform new tasks with rich visual hints.

Learning Manipulation from Human Video. A plethora of recent research has explored leveraging human video data to teach robot manipulation skills [3, 5, 13, 14, 24, 28, 37, 42, 43, 47–49, 53, 55, 58, 60, 62, 65]. These works focus on extracting different information from human video, such as object affordance [4, 27], motion trajectories [55, 59, 68], task dynamics [34, 41], and reward representations [8, 35, 44, 45]. Works like [24, 55, 68] train manipulation policies that are conditioned on human or robot videos, instead of language instructions. Despite their effectiveness, because such methods only learn from videos (sequence of images), they cannot infer important details, such as how much force to apply, required for many manipulation tasks. In this work, we focus on developing methods that can leverage multiple sensing modalities, including images, forces, and sounds, to better understand the subtle details in human demonstrations that are not readily visible, and to enable robots to better perform such tasks.

III. LEARNING FROM MULTIMODAL HUMAN VIDEOS

We introduce our system design, which takes a single multimodal human demonstration video as input, and generates a robot executable code to perform the manipulation task demonstrated in the video. The system has three main components: (1) collecting multimodal human videos; (2) Chain-of-Modality for understanding multimodal human videos; and (3) generating code and controlling the robot. For each component, we first discuss the motivation, followed by examples.

A. Multimodal Human Demonstration Video

Videos often struggle to capture important details of fine-grained details of a human performing a manipulation task, especially those involving force application. For example, when inserting a power plug (Fig. 3), we first apply *light* force to adjust its orientation, then *increase* force for insertion. These varying force levels are critical but hard to observe from video alone, highlighting the need for multimodal data that goes beyond visual information.

To address these challenges, we consider multimodal human video consists of, at each timestep, an RGB image, human muscle signal or object interaction sound, and hand pose (Fig. 1). They collectively provide a more comprehensive view of human task plans. Human muscle signals captured by an armband with muscle sensors (EMG) or object interaction sound captured by a microphone can provide necessary force information, which indicates the timing and the amount of force human applied during the entire task. Moreover, to provide more detailed information on human hand motions, we use a vision-based method [38] to estimate hand pose and treat the pixel locations of the fingertips as another input modality.

B. Chain-of-Modality

Next, we use Vision Language Models (VLMs) to analyze the rich information provided in such multimodal human videos to extract task plan descriptions. The VLM needs to process

signals from all these modalities: recognizing the human actions in the correct temporal order and determining the control parameters of each action (e.g., name of the target object, direction of the motion). One way to use a VLM for this purpose is to directly query the model with all modalities interleaved together in a sequence (Fig. 2(a)). However, we found that state-of-the-art VLMs (e.g. Gemini 1.5 Pro [40], GPT-4o [1]) often struggle to correlate information between modalities, leading to issues like neglecting certain inputs or attempting to extract information from the wrong modality. To improve VLMs’ performance in understanding multimodal human videos, we propose Chain-of-Modality (CoM, Fig. 2(b)), a prompting strategy that queries the VLM to analyze each modality sequentially, extracting key information and progressively aggregating results to produce the final answer.

Prompting Chain-of-Modality. The CoM prompt consists of three parts: (1) descriptions of each modality and their input data format, (2) description of the available action set along with an explanation of action parameters, and (3) one example of video-to-analysis pair introducing how each modality should be analyzed to produce a sequence of recognized actions with parameters.

Examples of Chain-of-Modality. Fig. 2 illustrates an example of using CoM to analyze a multimodal human video. In this video, a person holds the bottle with the left hand and twists the cap with the right hand. CoM sequentially analyzes each input modality and refines the answer based on the prior analysis. In Fig. 2, we highlight the new information contributed by each modality with a separate color. In the first stage, the VLM analyzes force or auditory signals and finds out *when* the person applies and releases force. Then it deduces the number of times the person applied force. However, without hand and image information, it is unclear what the person is doing exactly. In the second stage, the VLM incorporates hand pose information. It now recognizes that the person is *grasping* and *twisting* during force application. The finger position also indicates a *counterclockwise twist* of about 180 degrees while *holding* the bottle and a *clockwise finger rotation* during release. Still, without image data, the objects that appeared in the task remain unknown. In the third stage, the VLM integrates image data. It identifies that the left hand is holding *the bottle* and the right hand is twisting *the cap*. With this information, the VLM generates action functions, specifying detailed action parameters at each time step. Note that, none of the tasks or objects appear in the example prompts. The example prompts are only used to demonstrate the output format of the analysis and the library of available skills.

C. Writing Robot Code

Based on the human video analysis obtained above, the final step is to transform the action sequence into robot-executable code with low-level API calls. We use the same VLMs to perform this code generation [30] to create manipulation programs to accomplish the task. Prompts for code generation include the video analysis as well as descriptions of robot APIs and the required output format e.g.,

```
from skills import Grasp, Release, Twist, Find, Move_to
# Based on video analysis and APIs, generate python code:
```

Examples of Generating High-Level Task Plans. Below is an example of the generated program for the aforementioned bottle-opening task:

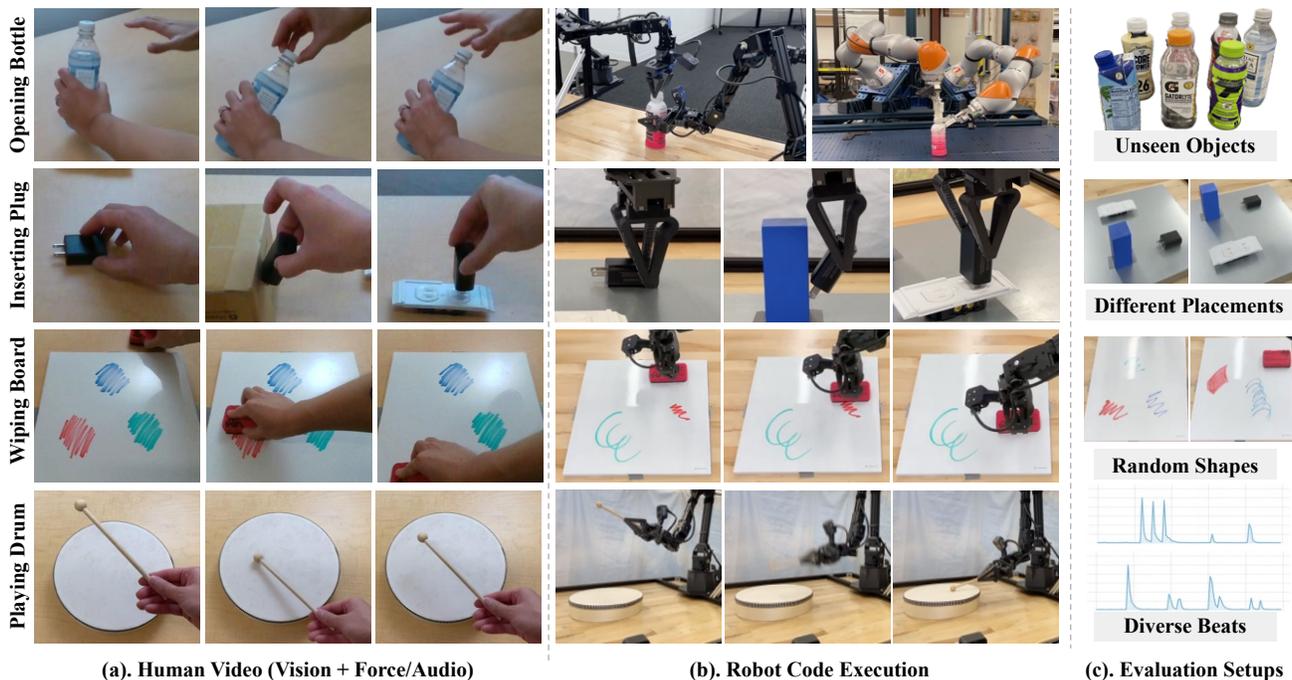


Fig. 3: Overview of Experiment Tasks. **(a) Multimodal Human Video Input:** Our framework processes a single-shot human video with force or audio data, using Chain-of-Modality to extract the task plan and control parameters, then generates a robot control program. **(b) Robot Code Execution:** The robot executes the program to replicate the task observed in the video. **(c) Evaluation Setups:** We evaluate the performance of generated program in various experimental setups.

```

Move_to('left', Find('bottle'))
Grasp('left')
Move_to('right', Find('bottle_cap'))
for _ in range(3):
    Grasp('right')
    Twist('right', 'counterclockwise', 180)
    Release('right')
    Twist('right', 'clockwise', 180)

```

VLM incorporates the video analysis from CoM and generates a detailed task plan for opening a bottle, which includes using the right gripper to `Twist` counterclockwise while holding the cap and `Twist` clockwise without holding the cap. It also generates a for-loop script to specify the periodic twisting motion.

Examples of Generating Control Parameters. Besides generating task plans, in contact-rich tasks such as inserting a plug into a power strip, VLM can also generate control parameters to specify the use of force:

```

from skills import Grasp, Push_towards, Insert
Grasp('right', 'plug', 100) # force range from [0, 100]
Move_to('right', 'box', 20) # rotate plug in-hand
Insert('right', 'power_strip', 100)

```

Leveraging the force information from the multimodal human video, VLM specifies the amount of force to apply during different task stages, allowing the use of `Move_to` to re-orient the plug in hand by pushing towards the wall (force = 20) and to hold the plug *firmly* while `Insert` it into the power strip (force = 100).

D. Implementation Details

Data Collection. The muscle signal (EMG) contains eight channels of data sampled at 200Hz. Since the camera records at 60Hz, we downsample the muscle signals to match the

camera sampling rate and use the maximum value across the eight channels as the force signal at each time step. Similar to audio signals, we compute the loudness of the sound at each time step as the input audio value. For the hand pose estimation, we use HaMeR [38] to localize the pixel locations of the fingertips. More details of the signal processing steps can be found in the Appendix.

Robot Execution. The robot API calls consist of predefined control functions that “ground” the generated programs in a real-world robot system. These APIs benefit significantly from recent advances in perception models. For instance, in our experiments, all object localization is performed by querying Gemini 1.5 Pro with an RGB-D image and the name of the target object (as specified in the generated programs), which directly generates a 2D bounding box around the target object on the RGB image. We then use the depth information and camera parameters to create 3D point clouds of the entities within the detected bounding box and use the average 3D location to represent the object’s position. These open-vocabulary APIs simplify the connection between the generated programs and the robot perception system, directly enhancing the capability of the code-based robot policy. More details of the robot APIs can be found in the Appendix.

IV. EXPERIMENTS

Our experiments aim to answer the following questions:

- Q1: Does Chain-of-Modality improve VLMs’ understanding of multimodal human videos?
- Q2: Does force information help VLMs in reasoning about human task plans?
- Q3: Does hand pose aid fine-grained manipulation?

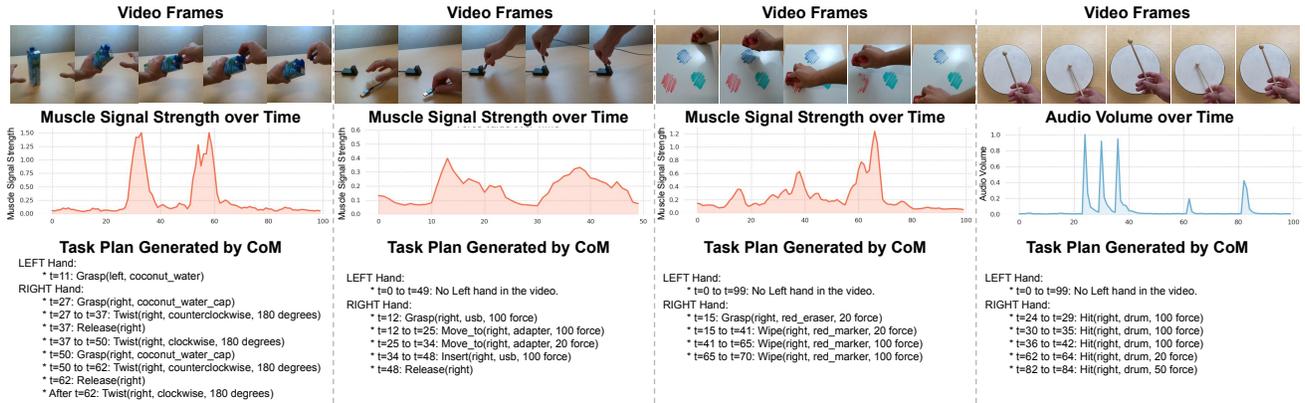


Fig. 4: Qualitative results for Chain-of-Modality. We showcase task plans generated by CoM for four evaluation videos. CoM successfully segments the videos into subtasks, specifying the skills, force, and target objects at each stage.

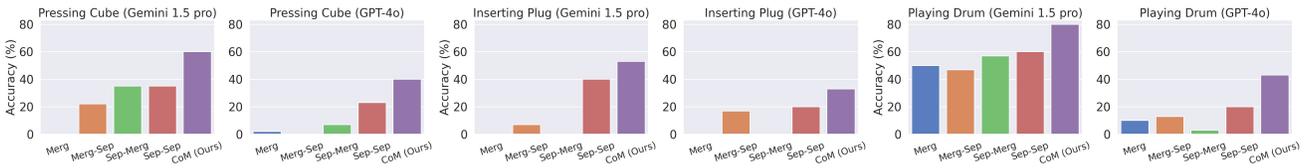


Fig. 5: Quantitative results for Chain-of-Modality. We compare CoM with baselines across three tasks using both Gemini and GPT.

Q4: Can VLMs with CoM extract control parameters from multimodal human videos?

Q5: How does the generated program perform on real robots?

A. Experiment Setups

Baselines. We categorize the baselines into two groups: baselines with different input modalities and baselines with different VLM reasoning procedures. For methods with different input modalities, the baselines include: (1) *Image-only* – uses only video image inputs, (2) *w.o. img* – uses force and hand pose inputs without image inputs, (3) *w.o. force* – excludes force input, (4) *w.o. hand* – excludes hand pose input, and (5) *All* – utilizes all three input modalities: force, hand pose, and image. For methods with different VLM reasoning procedures, the baselines include: (1) *Merg* – merges all modality inputs and directly generates the final answer, (2) *Merg-Sep* – merges all inputs but generates separate answers for each modality, (3) *Sep-Merg* – separate each modality inputs, then directly generates one final answer, (4) *Sep-Sep* – separates each modality input and generates separate answers for each, followed by a final answer, and (5) *Ours* – processes each modality sequentially, analyzing each modality based on the analysis of the previous one until the final answer.

Tasks. We design two categories of tasks: multimodal video analysis and real-world robot evaluation with generated programs. In multimodal video analysis, we test all methods on four tasks: *Pressing Cube*, *Inserting Plug*, *Playing Drum*, and *Opening Bottle*. These tasks feature long task horizons, force sensitivity, and bi-manual manipulation. Each task consists of 10 testing videos with varying objects and camera viewpoints. For each baseline, we query VLM three times per video and report the average success rate in generating the correct human task plans observed in the video, along with the similarity score, which is calculated by finding the

longest common string between the output and the ground truth. For all four tasks, we test with two advanced VLM models: Gemini 1.5 Pro [40] and GPT-4o [1]. In real-world robot evaluation, we test four tasks (Fig. 3): *Opening Bottle*, *Inserting Plug*, *Wiping Board* and *Playing Drum*. Each method is tested with 20 trials, and we report the average success rate.

For testing the generalization capability, as Fig. 3(c) illustrated, we test *Opening Bottle* with 7 types of bottle (6 unseen); we also randomly placed the plug, power strip, and box in *Inserting Plug* task. In *Wiping Board*, we draw marker with different shapes with varying positions on the board. In *Playing Drum* task, we test different drumming beats. For testing the cross-embodiment deployment of our method, we test the *Opening Bottle* task on two robot platforms: the bi-manual ViperX and the bi-manual KUKA.

Prompts. For video analysis tasks, the prompts we use include an explanation of the input data format and one example video. The input format explains the hand pose data, which consists of 2D pixel locations of the thumb and middle fingertips. The muscle or audio signals is normalized to a single float value. The example video involves a human interacting with task-irrelevant objects in random manner. For instance, both the *Pressing Cube* and *Opening Bottle* tasks use the same example video, where the human is pressing and rotating an apple and a can on the table. This video demonstrates only the key features of primitive skills, *which does not include the testing objects or task plans*. We also provide expected output for the example video, helping the VLM generate responses in the same format, which facilitates benchmarking the results. In each evaluation task, all baseline methods use the same example video. After prompting, we provide the testing video without task description and directly let VLM generate its analysis of the input multimodal video. For real-world robot evaluation, the prompts include the video analysis previously generated by

		Pressing Cube	Opening Bottle	Inserting Plug	Playing Drum
Gemini 1.5 Pro	Image-only	0.00/0.68	0.00/0.32	0.00/0.80	0.00/0.31
	w.o. img	0.00/0.00	0.00/0.45	0.00/0.00	0.00/0.00
	w.o. force	0.00/0.68	0.00/0.64	0.00/0.72	0.00/0.03
	w.o. hand	0.70/0.96	0.00/0.49	0.47/0.96	0.57/0.90
	All	0.67/0.92	0.37/0.75	0.53/0.93	0.80/0.96
GPT-4o	Image-only	0.00/0.48	0.00/0.43	0.00/0.77	0.00/0.37
	w.o. img	0.00/0.00	0.00/0.53	0.00/0.00	0.00/0.00
	w.o. force	0.00/0.13	0.00/0.20	0.00/0.55	0.00/0.22
	w.o. hand	0.43/0.76	0.00/0.54	0.30/0.81	0.34/0.84
	All	0.40/0.84	0.00/0.46	0.33/0.79	0.43/0.86

TABLE I: Multimodal eval (Accuracy / Similarity Score).

	Opening Bottle (ViperX)	Opening Bottle (KUKA)	Insert Plug	Wiping Board (red marker)	Wiping Board (blue marker)	Playing Drum
Ours	12/20	15/20	15/20	16/20	14/20	16/20
Oracle	16/20	20/20	18/20	20/20	16/20	20/20

TABLE II: On-robot evaluation results. All results are obtained in the generalization settings introduced in the **Task** section.

the our CoM pipeline, along with the definitions of robot API calls in Python format. We then ask the VLM to generate the main function to reproduce human task plans from the video analysis. After the VLM generates the program, we use *exec()* to execute the generated code on the real-world robot.

B. Results

Chain-of-Modality helps in understanding multimodal human videos. In Fig. 5, we compare baseline methods and Chain-of-Modality across three tasks using both Gemini 1.5 Pro and GPT-4o. We first observe that processing and analyzing each modality separately consistently outperforms other baselines that either merge the modality inputs or generate a single merged answer. Both *Sep-Sep* and *CoM* achieve the best performance in all tasks with different VLMs. *CoM* further outperforms *Sep-Sep* by more than 19% with Gemini 1.5 Pro and 17% with GPT-4o. This result indicates that although advanced VLMs have long-context inputs, directly reasoning over long-context inputs remains challenging. *CoM*’s per-modality analysis and progressively fulfilling the final answer is more suitable for current VLMs to reason from multimodal human videos.

Force information benefits learning from human videos. In Tab. I, we compare methods with different input modalities. We observe that images play a crucial role in recognizing task objects, as *w.o. img* baseline achieve no success. For tasks like *Opening bottle*, the main skill *Twist* doesn’t rely on object names, yielding a decent similarity score, but in general, images are essential. Force information greatly enhances the understanding of human task plans. Methods with force inputs (*w.o. hand* and *All*) significantly outperform those without. This indicates that extracting key information for manipulation purely from images is still challenging. Force inputs help VLMs better segment the video into different stages, leading to an average of 42% improvement in the similarity score between the extracted task plan and the ground truth (*All* v.s. *w.o. force*). Without force information, the baselines achieve no success in predicting the correct task plans and skill parameters from human videos.

Hand pose helps understand fine-grained manipulation. In the *Opening Bottle* task, only the method with all modalities as input (*All*) achieves a non-zero success rate. This task requires extracting fine-grained information about fingertip

rotation direction, along with multiple instances of grasping and releasing the bottle cap (as illustrated in the qualitative results, Fig. 5, first column). Hand pose plays a crucial role in this task, indicating that current VLM models are not yet proficient at extracting human hand motions. The hand pose estimated by specialized vision models [38] provides significant assistance.

CoM is capable of extracting control parameters from multimodal human videos. In the qualitative results shown in Fig. 5, CoM successfully (1) identifies task plans and target objects for single-arm or bi-manual tasks from multimodal human videos and (2) extracts detailed control parameters, including motion direction, force intensity, and timing of each action. This information paves the way for generating robot executable programs to perform the same task.

VLMs can generate manipulation programs based on CoM analysis. Although CoM can generate task plans from a single human video, these plans still lack detailed robot API calls, such as *Find* the target object and *Move_to* the object’s location before *Grasp*. VLMs can fill in these missing sub-steps and generate final robot-executable Python API calls. In Tab. II and our video submission, we demonstrate how the final generated program controls the robot to perform the task with an average success rate of 73% across all tasks, showing strong generalization to novel objects (e.g., opening 8 unseen bottles, wiping a board, and inserting a power plug with randomized object placements). This video demonstrates only the key features of primitive skills, *which does not include the testing objects or task plans*. The *Oracle* results (92%) showcases the upper bound of the robot system with manually-created robot code for each task. Notably, the *Find* API call we use is implemented by directly querying Gemini 1.5 pro, which generates 2D bounding boxes from free-form language inputs and further processes them into 3D locations using RGB-D images. This highlights the potential of leveraging VLMs for straightforward implementation of low-level robot API calls. In the *Opening Bottle* task, we tested our framework on two types of bi-manual robots (ViperX and KUKA), demonstrating the potential for cross-embodiment deployment of our method.

V. CONCLUSION AND LIMITATIONS

We present Chain-of-Modality (CoM), a prompting strategy that enables Vision-Language Models (VLMs) to understand multimodal human video demonstration data by combining video with force or audio inputs. By progressively refining task plans and control parameters through each modality, CoM enhances robots’ ability to perform one-shot imitation from human videos in fine-grained manipulation tasks. Our experiments show that CoM significantly improves the accuracy of task plan recognition and control parameter extraction compared to baselines, with strong generalization to new task setups and objects in real-world robot experiments.

The limitations of this work include: (1) The audio modality we use focuses only on the volume of impact sounds, without fully capturing other aspects such as frequency and pitch. Future work could explore the broader use of audio in foundation models for multimodal human video reasoning. (2) In this work, we focus on extracting task plans and control parameters from human videos and executing them on the robot in an open-loop fashion. In future work, we plan to explore generating closed-loop control programs that can adapt to unexpected situations.

REFERENCES

- [1] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [2] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” *arXiv preprint arXiv:2204.01691*, 2022.
- [3] S. Bahl, A. Gupta, and D. Pathak, “Human-to-robot imitation in the wild,” *arXiv preprint arXiv:2207.09450*, 2022.
- [4] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, “Affordances from human videos as a versatile representation for robotics,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 778–13 790.
- [5] H. Bharadhwaj, A. Gupta, V. Kumar, and S. Tulsiani, “Towards generalizable zero-shot manipulation via translating human interaction plans,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024, pp. 6904–6911.
- [6] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles, “Activitynet: A large-scale video benchmark for human activity understanding,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 961–970.
- [7] C.-Y. Chang, D.-A. Huang, D. Xu, E. Adeli, L. Fei-Fei, and J. C. Niebles, “Procedure planning in instructional videos,” in *European Conference on Computer Vision*, Springer, 2020, pp. 334–350.
- [8] A. S. Chen, S. Nair, and C. Finn, “Learning generalizable robotic reward functions from” in-the-wild” human videos,” *Robotics: Science and Systems (RSS)*, 2021.
- [9] B. Chen, Z. Xu, S. Kirmani, B. Ichter, D. Sadigh, L. Guibas, and F. Xia, “Spatialvlm: Endowing vision-language models with spatial reasoning capabilities,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 14 455–14 465.
- [10] H.-T. L. Chiang, Z. Xu, Z. Fu, M. G. Jacob, T. Zhang, T.-W. E. Lee, W. Yu, C. Schenck, D. Rendleman, D. Shah, *et al.*, “Mobility vla: Multimodal instruction navigation with long-context vlms and topological graphs,” *arXiv preprint arXiv:2407.07775*, 2024.
- [11] Y. Cui, S. Niekum, A. Gupta, V. Kumar, and A. Rajeswaran, “Can foundation models perform zero-shot task specification for robot manipulation?” In *Learning for dynamics and control conference*, PMLR, 2022, pp. 893–905.
- [12] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray, “Scaling egocentric vision: The epic-kitchens dataset,” in *European Conference on Computer Vision (ECCV)*, 2018.
- [13] N. Das, S. Bechtel, T. Davchev, D. Jayaraman, A. Rai, and F. Meier, “Model-based inverse reinforcement learning from visual demonstrations,” in *Conference on Robot Learning*, PMLR, 2021, pp. 1930–1942.
- [14] A. D. Edwards and C. L. Isbell, “Perceptual values from observation,” *arXiv preprint arXiv:1905.07861*, 2019.
- [15] R. Firoozi, J. Tucker, S. Tian, A. Majumdar, J. Sun, W. Liu, Y. Zhu, S. Song, A. Kapoor, K. Hausman, *et al.*, “Foundation models in robotics: Applications, challenges, and the future,” *arXiv preprint arXiv:2312.07843*, 2023.
- [16] R. Goyal, S. Ebrahimi Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag, *et al.*, “The” something something” video database for learning and evaluating visual common sense,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5842–5850.
- [17] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Hamburger, H. Jiang, M. Liu, X. Liu, *et al.*, “Ego4d: Around the world in 3,000 hours of egocentric video,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 995–19 012.
- [18] J. Gu, S. Kirmani, P. Wohlhart, Y. Lu, M. G. Arenas, K. Rao, W. Yu, C. Fu, K. Gopalakrishnan, Z. Xu, P. Sundaresan, P. Xu, H. Su, K. Hausman, C. Finn, Q. Vuong, and T. Xiao, *Rt-trajectory: Robotic task generalization via hindsight trajectory sketches*, 2023. arXiv: 2311.01977 [cs.LG].
- [19] Y. Hu, Q. Xie, V. Jain, J. Francis, J. Patrikar, N. Keetha, S. Kim, Y. Xie, T. Zhang, Z. Zhao, *et al.*, “Toward general-purpose robots via foundation models: A survey and meta-analysis,” *arXiv preprint arXiv:2312.08782*, 2023.
- [20] D.-A. Huang, S. Nair, D. Xu, Y. Zhu, A. Garg, L. Fei-Fei, S. Savarese, and J. C. Niebles, “Neural task graphs: Generalizing to unseen tasks from a single video demonstration,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 8565–8574.
- [21] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, “Language models as zero-shot planners: Extracting actionable knowledge for embodied agents,” in *International conference on machine learning*, PMLR, 2022, pp. 9118–9147.
- [22] W. Huang, C. Wang, Y. Li, R. Zhang, and L. Fei-Fei, “Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation,” *arXiv preprint arXiv:2409.01652*, 2024.
- [23] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar, *et al.*, “Inner monologue: Embodied reasoning through planning with language models,” *arXiv preprint arXiv:2207.05608*, 2022.
- [24] V. Jain, M. Attarian, N. J. Joshi, A. Wahid, D. Driess, Q. Vuong, P. R. Sanketi, P. Sermanet, S. Welker, C. Chan, *et al.*, “Vid2robot: End-to-end video-conditioned policy learning with cross-attention transformers,” *arXiv preprint arXiv:2403.12943*, 2024.
- [25] C. Ju, T. Han, K. Zheng, Y. Zhang, and W. Xie, “Prompting visual-language models for efficient video understanding,” in *European Conference on Computer Vision*, Springer, 2022, pp. 105–124.

- [26] Y. Kong and Y. Fu, "Human action recognition and prediction: A survey," *International Journal of Computer Vision*, vol. 130, no. 5, pp. 1366–1401, 2022.
- [27] H. S. Koppula and A. Saxena, "Anticipating human activities using object affordances for reactive robotic response," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 1, pp. 14–29, 2015.
- [28] S. Kumar, J. Zamora, N. Hansen, R. Jangir, and X. Wang, "Graph inverse reinforcement learning from diverse videos," in *Conference on Robot Learning*, PMLR, 2023, pp. 55–66.
- [29] Y.-L. Li, L. Xu, X. Liu, X. Huang, Y. Xu, M. Chen, Z. Ma, S. Wang, H.-S. Fang, and C. Lu, "Hake: Human activity knowledge engine," *arXiv preprint arXiv:1904.06539*, 2019.
- [30] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, "Code as policies: Language model programs for embodied control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2023, pp. 9493–9500.
- [31] J. Liang, F. Xia, W. Yu, A. Zeng, M. G. Arenas, M. Attarian, M. Bauza, M. Bennice, A. Bewley, A. Dostmohamed, *et al.*, "Learning to learn faster from human feedback with language model predictive control," *arXiv preprint arXiv:2402.11450*, 2024.
- [32] K. Lin, C. Agia, T. Migimatsu, M. Pavone, and J. Bohg, "Text2motion: From natural language instructions to feasible plans," *Autonomous Robots*, vol. 47, no. 8, pp. 1345–1365, 2023.
- [33] B. Liu, Y. Jiang, X. Zhang, Q. Liu, S. Zhang, J. Biswas, and P. Stone, "Llm+ p: Empowering large language models with optimal planning proficiency," *arXiv preprint arXiv:2304.11477*, 2023.
- [34] Y. Liu, A. Gupta, P. Abbeel, and S. Levine, "Imitation from observation: Learning to imitate behaviors from raw video via context translation," in *2018 IEEE international conference on robotics and automation (ICRA)*, IEEE, 2018, pp. 1118–1125.
- [35] Y. J. Ma, S. Sodhani, D. Jayaraman, O. Bastani, V. Kumar, and A. Zhang, "Vip: Towards universal visual reward and representation via value-implicit pre-training," *arXiv preprint arXiv:2210.00030*, 2022.
- [36] S. Nasiriany, F. Xia, W. Yu, T. Xiao, J. Liang, I. Dasgupta, A. Xie, D. Driess, A. Wahid, Z. Xu, *et al.*, "Pivot: Iterative visual prompting elicits actionable knowledge for vlms," *arXiv preprint arXiv:2402.07872*, 2024.
- [37] G. Papagiannis, N. Di Palo, P. Vitiello, and E. Johns, "R+ x: Retrieval and execution from everyday human videos," *arXiv preprint arXiv:2407.12957*, 2024.
- [38] G. Pavlakos, D. Shan, I. Radosavovic, A. Kanazawa, D. Fouhey, and J. Malik, "Reconstructing hands in 3D with transformers," in *CVPR*, 2024.
- [39] S. S. Raman, V. Cohen, E. Rosen, I. Idrees, D. Paulius, and S. Tellex, "Planning with large language models via corrective re-prompting," in *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.
- [40] M. Reid, N. Savinov, D. Teplyashin, D. Lepikhin, T. Lillicrap, J.-b. Alayrac, R. Soricut, A. Lazaridou, O. Firat, J. Schrittwieser, *et al.*, "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," *arXiv preprint arXiv:2403.05530*, 2024.
- [41] N. Rhinehart and K. M. Kitani, "First-person activity forecasting with online inverse reinforcement learning," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3696–3705.
- [42] K. Schmeckpeper, O. Rybkin, K. Daniilidis, S. Levine, and C. Finn, "Reinforcement learning with videos: Combining offline observations with interaction," *arXiv preprint arXiv:2011.06507*, 2020.
- [43] K. Schmeckpeper, A. Xie, O. Rybkin, S. Tian, K. Daniilidis, S. Levine, and C. Finn, "Learning predictive models from observation and interaction," in *European Conference on Computer Vision*, Springer, 2020, pp. 708–725.
- [44] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, S. Levine, and G. Brain, "Time-contrastive networks: Self-supervised learning from video," in *2018 IEEE international conference on robotics and automation (ICRA)*, IEEE, 2018, pp. 1134–1141.
- [45] P. Sermanet, K. Xu, and S. Levine, "Unsupervised perceptual rewards for imitation learning," *arXiv preprint arXiv:1612.06699*, 2016.
- [46] D. Shah, B. Osiński, S. Levine, *et al.*, "Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *Conference on robot learning*, PMLR, 2023, pp. 492–504.
- [47] L. Shao, T. Migimatsu, Q. Zhang, K. Yang, and J. Bohg, "Concept2robot: Learning manipulation concepts from instructions and human demonstrations," *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1419–1434, 2021.
- [48] P. Sharma, D. Pathak, and A. Gupta, "Third-person visual imitation learning via decoupled hierarchical controller," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [49] K. Shaw, S. Bahl, and D. Pathak, "Videx: Learning dexterity from internet videos," *CoRL*, 2022.
- [50] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Conference on robot learning*, PMLR, 2022, pp. 894–906.
- [51] T. Silver, S. Dan, K. Srinivas, J. B. Tenenbaum, L. Kaelbling, and M. Katz, "Generalized planning in pddl domains with pretrained large language models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, 2024, pp. 20 256–20 264.
- [52] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, "Progprompt: Generating situated robot task plans using large language models," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2023, pp. 11 523–11 530.
- [53] L. Smith, N. Dhawan, M. Zhang, P. Abbeel, and S. Levine, "Avid: Learning multi-stage tasks via pixel-level translation of human videos," *arXiv preprint arXiv:1912.04443*, 2019.
- [54] P. Sundaresan, Q. Vuong, J. Gu, P. Xu, T. Xiao, S. Kirmani, T. Yu, M. Stark, A. Jain, K. Hausman, *et al.*,

“Rt-sketch: Goal-conditioned imitation learning from hand-drawn sketches,” 2024.

- [55] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar, “Mimicplay: Long-horizon imitation learning by watching human play,” *arXiv preprint arXiv:2302.12422*, 2023.
- [56] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool, “Temporal segment networks for action recognition in videos,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 11, pp. 2740–2755, 2018.
- [57] Z. Wang, S. Cai, G. Chen, A. Liu, X. Ma, and Y. Liang, “Describe, explain, plan and select: Interactive planning with large language models enables open-world multi-task agents,” *arXiv preprint arXiv:2302.01560*, 2023.
- [58] B. Wen, W. Lian, K. Bekris, and S. Schaal, “You only demonstrate once: Category-level manipulation from single visual demonstration,” *arXiv preprint arXiv:2201.12716*, 2022.
- [59] C. Wen, X. Lin, J. So, K. Chen, Q. Dou, Y. Gao, and P. Abbeel, *Any-point trajectory modeling for policy learning*, 2023. arXiv: 2401.00025 [cs.LG].
- [60] H. Xiong, Q. Li, Y.-C. Chen, H. Bharadhwaj, S. Sinha, and A. Garg, “Learning by watching: Physical imitation of manipulation skills from human videos,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2021, pp. 7827–7834.
- [61] H. Xu, A. Das, and K. Saenko, “R-c3d: Region convolutional 3d network for temporal activity detection,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5783–5792.
- [62] M. Xu, Z. Xu, C. Chi, M. Veloso, and S. Song, “Xskill: Cross embodiment skill discovery,” in *Conference on Robot Learning*, PMLR, 2023, pp. 3536–3555.
- [63] Z. Yang, L. Li, K. Lin, J. Wang, C.-C. Lin, Z. Liu, and L. Wang, “The dawn of lmms: Preliminary explorations with gpt-4v (ision),” *arXiv preprint arXiv:2309.17421*, vol. 9, no. 1, p. 1, 2023.
- [64] W. Yu, N. Gileadi, C. Fu, S. Kirmani, K.-H. Lee, M. G. Arenas, H.-T. L. Chiang, T. Erez, L. Hasenclever, J. Humplik, *et al.*, “Language to rewards for robotic skill synthesis,” *arXiv preprint arXiv:2306.08647*, 2023.
- [65] K. Zakka, A. Zeng, P. Florence, J. Tompson, J. Bohg, and D. Dwibedi, “Xirl: Cross-embodiment inverse reinforcement learning,” in *Conference on Robot Learning*, PMLR, 2022, pp. 537–546.
- [66] A. Zeng, M. Attarian, B. Ichter, K. Choromanski, A. Wong, S. Welker, F. Tombari, A. Purohit, M. Ryoo, V. Sindhwani, *et al.*, “Socratic models: Composing zero-shot multimodal reasoning with language,” *arXiv preprint arXiv:2204.00598*, 2022.
- [67] C.-L. Zhang, J. Wu, and Y. Li, “Actionformer: Localizing moments of actions with transformers,” in *European Conference on Computer Vision*, Springer, 2022, pp. 492–510.
- [68] Y. Zhu, A. Lim, P. Stone, and Y. Zhu, “Vision-based manipulation from single human video with open-world object graphs,” *arXiv preprint arXiv:2405.20321*, 2024.