

TidyBot: Personalized Robot Assistance with Large Language Models

Jimmy Wu,¹ Rika Antonova,² Adam Kan,³ Marion Lepert,² Andy Zeng,⁴ Shuran Song,⁵
Jeannette Bohg,² Szymon Rusinkiewicz,¹ Thomas Funkhouser^{1,4}

Abstract—For a robot to personalize physical assistance effectively, it must learn user preferences that can be generally reapplied to future scenarios. In this work, we investigate personalization of household cleanup with robots that can tidy up rooms by picking up objects and putting them away. A key challenge is determining the proper place to put each object, as people’s preferences can vary greatly depending on personal taste or cultural background. For instance, one person may prefer storing shirts in the drawer, while another may prefer them on the shelf. We aim to build systems that can learn such preferences from just a handful of examples via prior interactions with a particular person. We show that robots can combine language-based planning and perception with the few-shot summarization capabilities of large language models (LLMs) to infer generalized user preferences that are broadly applicable to future interactions. This approach enables fast adaptation and achieves 91.2% accuracy on unseen objects in our benchmark dataset. We also demonstrate our approach on a real-world mobile manipulator called TidyBot, which successfully puts away 85.0% of objects in real-world test scenarios.

I. INTRODUCTION

Building a robot that provides personalized assistance for physical household tasks is a long-standing goal of robotics research. In this paper, we investigate the task of tidying up a room: moving every object on the floor to its “proper place.” One of the challenges in performing this task is determining the correct receptacle (“proper place”) for every object. This is difficult because where objects should go is highly personal, and depends on cultural norms and individual preferences. One person may want to put shirts in a dresser drawer, another may want them on shelves, and a third may want them hanging in a closet. There is no “one size fits all” solution.

Classical approaches to the household cleanup task ask a person to specify a target location for every object [1], [2], which is tedious and impractical in an autonomous setting. Other works learn generic (non-personalized) rules about where objects typically go inside a house by averaging over many users [3], [4], [5]. Works that focus on personalization aim to extrapolate from a few user examples given similar choices made by other users, using methods such as collaborative filtering [6], spatial relationships [7], or learned latent preference vectors [8]. However, all of these approaches require collecting large datasets with user preferences or generating datasets from manually constructed, simulated




Fig. 1. We study the task of household cleanup, where each object on the floor must be picked up and put away while following user preferences.

scenarios. Such datasets can be expensive to acquire and may not generalize well if they are too small.

Our approach is to utilize the summarization capabilities of large language models (LLMs) to provide generalization from a small number of example preferences. We ask a person to provide a few example object placements using textual input (*e.g.*, yellow shirts go in the drawer, dark purple shirts go in the closet, white socks go in the drawer), and then we ask the LLM to summarize these examples (*e.g.*, light-colored clothes go in the drawer and dark-colored clothes go in the closet) to arrive at generalized preferences for this particular person.

The underlying insight is that the summarization capabilities of LLMs are a good match for the generalization requirements of personalized robotics. LLMs demonstrate astonishing abilities to perform generalization through summarization, drawing upon complex object properties and relationships learned from massive text datasets. By using the summarization provided by LLMs for generalization in robotics, we hope to produce generalized rules from a small number of examples, in a form that is human interpretable (text) and is expressed in nouns that can be grounded in images using open-vocabulary image classifiers. Using an off-the-shelf LLM also avoids expensive collection of user preference data and model training.

We investigate the proposed approach in a real-world robotic mobile manipulation system for household cleanup, which we call TidyBot (Fig. 1). Before the robot begins cleanup, we ask the user to provide a handful of example placements for specific objects, which are passed to an LLM

¹Princeton University, ²Stanford University, ³The Nueva School, ⁴Google, ⁵Columbia University

This work was supported in part by the Princeton School of Engineering, Toyota Research Institute, and the National Science Foundation under CCF-2030859, DGE-1656466, and IIS-2132519.

to be summarized into a generalized set of rules (personalized to that user) mapping object categories to receptacles. The nouns of these generalized rules are provided to an open-vocabulary image classifier in order to identify objects on the floor and determine target receptacles for them using the rules. The robot will then carry out the cleanup task by repeatedly picking up objects, identifying them, and moving them to their target receptacles.

We evaluate our approach quantitatively on both a text-based benchmark dataset and our real-world robotic system. On the benchmark, we find that our approach generalizes well, achieving an accuracy of 91.2% on unseen objects across all scenarios in the benchmark. In our real-world test scenarios, we find that TidyBot correctly puts away 85.0% of objects. We also show that our approach can be easily extended to infer generalized rules for manipulation primitive selection (*e.g.*, pick and place vs. pick and toss) in addition to inferring object placements.

Our contributions are: (i) the idea that text summarization with LLMs provides a means for generalization in robotics, (ii) a publicly released benchmark dataset for evaluating generalization of receptacle selection preferences, and (iii) implementation and evaluation of our approach on a real-world mobile manipulation system.

Please see our project page at <https://tidybot.cs.princeton.edu> for additional supplementary material, benchmark dataset and code, and qualitative videos of our real-world system TidyBot in action.

II. RELATED WORK

Household cleanup. Many recent works in Embodied AI have proposed benchmarks or methods for completing household tasks in simulated indoor environments [9], [10], [11], [12], [13], [14]. For household cleanup in particular, the object rearrangement task [10], [15], [12], [16], [17] requires an embodied agent to pick up and move objects so as to bring the environment into a specified state. Household cleanup has also been studied in robotics works, in which instructions for object rearrangement are specified via pointing gestures [1] or target layouts [2]. The drawback of these setups is that a target location must be manually specified for every object to be manipulated, which can require significant human effort. Prior works have addressed this challenge by automatically inferring object placements based on human preferences for where objects typically go inside a house [3], [4], [5], eliminating the need to specify where every individual object goes. However, these works predict human preferences that are generic rather than personalized. To handle the variability in preferences across different users, other works have used collaborative filtering [6], spatial relationships [7], or learned latent preference vectors [8] to predict object placements that are based on personalized user preferences. These methods require the collection of large crowd-sourced datasets for human preferences, which can be expensive. By contrast, our approach uses off-the-shelf LLMs with no additional training or data collection. We are able to directly leverage

the commonsense knowledge and summarization abilities of LLMs to build generalizable personalized preferences for each user.

Object sorting. Object sorting has been studied in robotics using approaches such as clustering [18], active learning [19], [20], metric learning [21], or heuristic search [22], [23]. These setups carry out pre-specified sorting rules using physical properties such as color [24], [18], [19], [20], [22], [25], [23], shape [20], size [18], [20], [25], or material [26]. Notably, they are not able to sort based on semantics or commonsense knowledge, nor are they able to automatically infer sorting rules. More recently, Høeg and Tingelstad [27] studied whether classification of objects into general high-level categories can be improved by using an LLM to take in an object detector’s prediction and output a general category for the object. In our work, we similarly tap into the commonsense knowledge of LLMs to reason about object sorting. However, whereas their setup uses pre-specified sorting rules based on a fixed set of categories, ours is able to infer generalizable sorting rules automatically.

LLMs for robotics. Large language models (LLMs) have been shown to exhibit remarkable commonsense reasoning abilities [28], [29], [30]. As a result, there has been increasing interest in harnessing the capabilities of LLMs to build more commonsense knowledge into robotic systems. Many recent works study how LLM-generated high-level robotic plans (typically produced using the few-shot learning paradigm [28]) can be grounded in the state of the environment. This can be done with value functions [31], [32], semantic translation into admissible actions [33], scene description as context [34], [35], [36], [37], feedback [38], [39], or re-prompting [40]. However, these works assume a setup in which the LLM is expected to output a single generic plan. This is not a good fit for personalized household cleanup, because a “one size fits all” plan would not address the wide variability in user preferences. Instead, our system generates personalized plans that are tailored to the preferences of a particular user. Other works in robotics have used LLMs for PDDL planning [41], code generation for robotic control policies [42], and parsing navigation instructions into textual landmarks [43]. These works all use LLMs as a means of integrating commonsense knowledge into robotic systems, which is also true in our case. However, unlike these works, we additionally show that the summarization ability of LLMs enables generalization in robotics.

III. METHOD

We use the summarization capabilities of an off-the-shelf LLM to generalize user preferences from a small number of examples. Below, we describe how we use the LLM to infer personalized rules for both receptacle selection and manipulation primitive selection, and also how we deploy the approach on a real-world mobile manipulation system for household cleanup.

A. Personalized receptacle selection

Our system first receives a few examples of object placements reflecting the personal preferences of a user. For instance, the user may specify that yellow shirts and white socks go in the drawer, while dark purple shirts and black shirts go in the closet. We provide these examples to an LLM, which then infers personalized rules on where objects belong. Specifically, the LLM (i) summarizes the examples into general rules, and then (ii) uses the summary to determine where to place new objects.

Following recent work [34], [37], we convert the user examples into LLM prompts that are structured as Pythonic code. This prompt form is advantageous because LLMs are trained on large amounts of code, and it also provides a structured output that is easy to parse. To represent the user examples, the prompt first contains a list of objects present in the scene and a list of potential receptacles. This is followed by a series of pick and place commands reflecting where the objects should be placed according to the user. Then, we ask the LLM to complete the last line, which is a code comment summarizing what the preceding code block does. Here is an example LLM completion where the output from the LLM is highlighted:

```
objects = ["yellow shirt", "dark purple shirt", "white socks", "black shirt"]
receptacles = ["drawer", "closet"]
pick_and_place("yellow shirt", "drawer")
pick_and_place("dark purple shirt", "closet")
pick_and_place("white socks", "drawer")
pick_and_place("black shirt", "closet")
# Summary: Put light-colored clothes in the drawer and dark-colored clothes in the closet.
```

In this example, the LLM summarized the provided object placements and inferred that light-colored clothes go in the drawer while dark-colored clothes go in the closet. These examples lead to a generalized rule for where objects belong, personalized to this particular user.

Next, the summary is used by the LLM to generate placements for novel, unseen objects. The prompt consists of the summary from the LLM summarization step (in the form of a code comment), a list of the *unseen objects*, a list of receptacles, and a partial pick and place command for the first object. We then ask the LLM to provide a placement for each object by completing the prompt:

```
# Summary: Put light-colored clothes in the drawer and dark-colored clothes in the closet.
objects = ["black socks", "white shirt", "navy socks", "beige shirt"]
receptacles = ["drawer", "closet"]
pick_and_place("black socks", "closet")
pick_and_place("white shirt", "drawer")
pick_and_place("navy socks", "closet")
pick_and_place("beige shirt", "drawer")
```

The output pick and place commands can then be parsed to determine where each unseen object should be placed.

B. Personalized primitive selection

Similar to the way we infer generalized rules for receptacle selection, we can also infer generalized rules for how to manipulate objects, again leveraging the summarization capabilities of LLMs. First, we provide a few examples of objects along with their user-preferred manipulation primitive to the LLM, and ask it to summarize. Here is an example completion where the output from the LLM is highlighted:

```
objects = ["yellow shirt", "dark purple shirt", "white socks", "black shirt"]
pick_and_place("yellow shirt")
pick_and_place("dark purple shirt")
pick_and_toss("white socks")
pick_and_place("black shirt")
# Summary: Pick and place shirts, pick and toss socks.
```

The summary can then be used as a generalized rule to predict the appropriate primitive to use for *unseen objects*:

```
# Summary: Pick and place shirts, pick and toss socks.
objects = ["black socks", "white shirt", "navy socks", "beige shirt"]
pick_and_toss("black socks")
pick_and_place("white shirt")
pick_and_toss("navy socks")
pick_and_place("beige shirt")
```

C. Real-world robotic system

Given generalized rules from LLM summarization, we can now implement these rules on a robot tasked with tidying up a household environment. To do so, we use a perception system to localize and recognize objects in the environment, and a predetermined set of manipulation primitives to move objects into receptacles. For our setup, we use `pick_and_place` and `pick_and_toss` as our primitives, as they are well-suited for household cleanup. However, other sets of primitives could also be used.

For each new user, the system will receive a set of example preferences and run the previously described LLM pipeline to get personalized rules for the user. The rules contain a set of generalized object categories produced by summarization (e.g., light-colored clothes, dark-colored clothes), each of which is matched to a preferred receptacle and manipulation primitive for that category. The robot will tidy up the environment by iteratively performing the following steps until no more objects remain on the floor: (1) localize the nearest object, (2) classify the object into a generalized category, (3) determine the appropriate receptacle and manipulation primitive for the object using generalized rules produced by the LLM, and (4) use the manipulation primitive to put the object into the receptacle. Fig. 2 provides a conceptual illustration of this procedure, and Algorithm 1 outlines these steps in pseudocode.

One important aspect of our approach is that the LLM summarization automatically provides candidate categories to the perception system. Nouns (or noun phrases) are extracted from the summarization text as categories, and used as the target label set for CLIP [44], the open-vocabulary

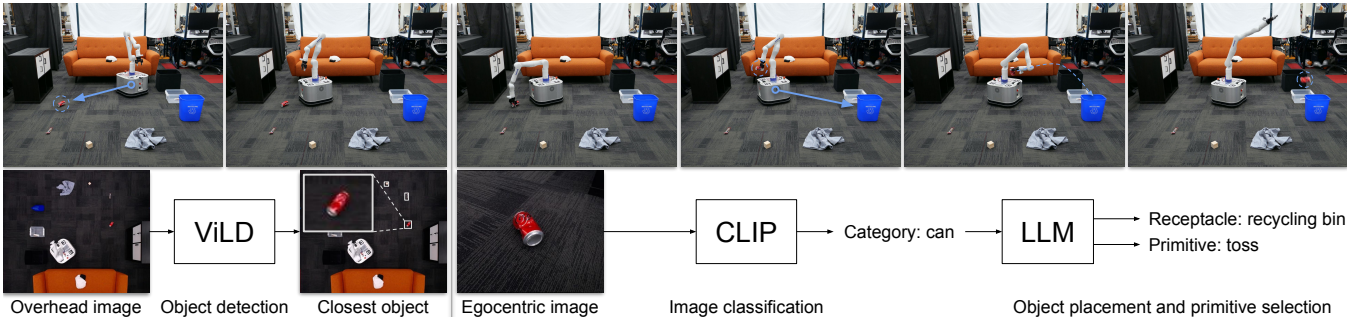


Fig. 2. **System overview.** Once the user’s preferences have been summarized with an LLM, TidyBot will localize the closest object on the floor, move to get a close-up view with its egocentric camera, predict the object’s category using CLIP, use the LLM-summarized rules to select a receptacle and manipulation primitive, and then execute the primitive to put the object into the selected receptacle, repeating this entire process until no more objects can be found on the floor.

Algorithm 1 System pipeline

```

Input:  $E_{\text{receptacle}} = \{(o_1, r_1), (o_2, r_2), \dots\}$ 
Input:  $E_{\text{primitive}} = \{(o_1, p_1), (o_2, p_2), \dots\}$ 
 $S_{\text{receptacle}} = \text{LLM.Summarize}(E_{\text{receptacle}})$ 
 $S_{\text{primitive}} = \text{LLM.Summarize}(E_{\text{primitive}})$ 
 $C = \text{LLM.GetCategories}(S_{\text{receptacle}})$ 
 $\text{robot.Initialize}()$ 
while True do
   $I_{\text{top}} = \text{GetOverheadImage}()$ 
   $o = \text{ViLD.GetClosestObject}(I_{\text{top}})$ 
   $\text{robot.MoveTo}(o)$ 
   $I_{\text{ego}} = \text{robot.GetEgocentricImage}()$ 
   $c = \text{CLIP.GetCategory}(I_{\text{ego}}, C)$ 
   $r = \text{LLM.GetReceptacle}(S_{\text{receptacle}}, c)$ 
   $p = \text{LLM.GetPrimitive}(S_{\text{primitive}}, c)$ 
   $\text{robot.PickUp}(o)$ 
   $\text{robot.MoveTo}(r)$ 
   $\text{robot.ExecutePrimitive}(p)$ 
end while

```

image classification model we use. For example, the following LLM prompt will extract the two general categories in the summary text (light-colored clothing and dark-colored clothing):

```

# Summary: Put light-colored clothes in the drawer and
dark-colored clothes in the closet.
objects = ["light-colored clothing",
"dark-colored clothing"]

```

This combination of summarization and open-vocabulary classification is critical to the autonomy of the system, as it enables the object classifier to work with a small set of generalized object categories. The approach is (i) robust as there are only a small number of categories to differentiate between, and (ii) flexible because it supports arbitrary sets of object categories for different users. In contrast, without LLM summarization, the object classifier would have to be able to recognize all possible fine-grained object classes, which is much more difficult. Alternatively, the user would have to manually specify the list of objects present in each target scene, which would be impractical for an autonomous system.

IV. EXPERIMENTS

We investigate the performance of our proposed approach with two types of evaluation. For the first type of evaluation, we design a benchmark for generalization of receptacle selection using text-based examples, which enables direct comparison to alternative approaches and ablation studies, with quantitative metrics. For the second type of evaluation, we deploy our approach in a real-world mobile manipulation system for tidying up a room based on user preferences. Unless otherwise specified, the LLM we use is text-davinci-003, a variant of GPT-3 [28]. All LLM experiments were run with temperature 0.

A. Benchmark dataset

In order to evaluate the proposed approach and to quantitatively compare it to alternatives, we created a benchmark dataset of object placements. The benchmark is comprised of 96 scenarios, each of which has a set of objects, a set of receptacles, a set of example “seen” object placements (preferences), and a set of “unseen” evaluation placements, all specified as text. The task is to predict the placements in the “unseen” set given the examples in the “seen” set.

The benchmark scenarios are defined in 4 room types (living room, bedroom, kitchen, pantry room), with 24 scenarios per room type. Each scenario contains 2–5 receptacles (potential places to put objects, such as shelves, cabinets, etc.), 4–10 “seen” example object placements provided as input to the task, and an equal number of “unseen” object placements (distinct from the seen examples) provided for evaluation. There are 2 seen and 2 unseen object placements per receptacle. In total, there are 672 seen and 672 unseen object placements, which cumulatively reference 87 unique receptacles and 1,076 unique objects.

Success on this benchmark is measured by the object placement accuracy: the number of objects placed in the correct receptacle divided by the total number of objects. We evaluate accuracy separately for seen and unseen objects, to tease apart memorization versus generalization. For each, we compute the accuracy per scenario, and then average the results across all scenarios to produce the numbers shown in the tables.

TABLE I
COMPARISONS TO BASELINES

Method	Accuracy (unseen)
Examples only	78.5%
WordNet taxonomy	67.5%
RoBERTa embeddings	77.8%
CLIP embeddings	83.7%
Summarization (ours)	91.2%

B. Baseline comparisons

In our first set of experiments, we use the benchmark to provide quantitative evaluation of our approach compared to several alternatives. The results are in Tab. I. Since the main challenge is to generalize from objects in the examples (seen) to those in the evaluation set (unseen), we consider a variety of baseline generalization approaches and report placement accuracy metrics only for unseen objects.

The following paragraphs describe each baseline and provide a discussion of how the performance compares to that of our proposed approach.

Examples only. The first baseline provides a direct comparison to a system like ours if it did not use summarization. The LLM is given a list of objects, receptacles, and example placement preferences, along with a list of unseen objects for a new scene. Then, the LLM is asked to directly infer the proper placements (highlighted text) for unseen objects in the new scene, without summarization as an intermediate step:

```
objects = ["yellow shirt", "dark purple shirt", "white socks", "black shirt"]
receptacles = ["drawer", "closet"]
pick_and_place("yellow shirt", "drawer")
pick_and_place("dark purple shirt", "closet")
pick_and_place("white socks", "drawer")
pick_and_place("black shirt", "closet")

objects = ["black socks", "white shirt", "navy socks", "beige shirt"]
receptacles = ["drawer", "closet"]
pick_and_place("black socks", "drawer")
pick_and_place("white shirt", "closet")
pick_and_place("navy socks", "drawer")
pick_and_place("beige shirt", "closet")
```

The prediction accuracy of this method for unseen objects (78.5%) is significantly worse than that of our method (91.2%). Since the main difference between this method versus ours is that our method leverages summarization, this result presents strong evidence for our main hypothesis — *i.e.*, summarization is useful for generalization. This finding is also consistent with recent work showing that LLMs perform better when they are asked to output intermediate steps of reasoning before the final answer [29], [30]. When looking at the predictions, we find that this baseline approach generally predicts object placements that are sensible but may not be consistent with the user’s preferences.

WordNet taxonomy. This baseline uses a hand-crafted lexical ontology called WordNet [45] to generalize placements

TABLE II
ABLATION STUDIES

Method	Seen	Unseen
Commonsense	45.0%	45.6%
Summarization	91.8%	91.2%
Human summary	97.1%	97.5%

from seen to unseen objects. For each unseen object, we place it in the same receptacle as the most similar seen object, where similarity is measured using the shortest path between two objects in the taxonomy. Since WordNet is a hand-crafted taxonomy, it does not contain all possible object names. For the 694 objects in our benchmark that are missing from WordNet, we manually mapped each of them to the closest WordNet object name. Even with the manual mapping, the performance of this WordNet baseline for unseen objects (67.5%) is far worse than that of our method (91.2%). This shows that LLM summarization provides better generalization than using the hierarchy provided by a hand-crafted ontology.

Text embedding. This baseline uses pretrained text embeddings to assist with generalization. For each unseen object, we place it in the receptacle provided for the most similar seen object, where similarity is defined by cosine similarity between encoded object names in the RoBERTa [46] or CLIP [44] embedding space. For RoBERTa, we use the pretrained Sentence-BERT [47] model from the SentenceTransformers library. Specifically, we use the all-distilroberta-v1 variant which is a distilled version of the RoBERTa [46] model that is fine-tuned on a dataset of 1 billion sentence pairs. For CLIP, we use the pretrained model provided by OpenAI. In either case, the generalization performance for predicting placements of unseen objects does not reach the performance of our proposed summarization approach (77.8% for RoBERTa and 83.7% for CLIP, versus 91.2% for ours). Although text embeddings trained on large datasets encode many types of object similarities, particularly for related object categories, they may not encode the object attributes relevant to the preferences of a particular user (*e.g.*, light objects go here, heavy object go there). In contrast, our summarization approach is able to correctly encode a larger variety of user preferences.

C. Ablation studies

In the second set of experiments, we use the benchmark to evaluate the performance of several variants to our method. The goal of these experiments is to compare its performance to alternatives with far less information (using only common sense, without preferences) or far more information (using human-generated summarizations). We also study the impact of using different LLMs. The benchmark metrics for both seen and unseen objects are provided in Tabs. II and III.

Commonsense. Our first ablation study measures how well an LLM can perform the benchmark tasks using only commonsense reasoning — *i.e.*, without using the preferences at all. For each benchmark scene, we give the LLM the

TABLE III
COMPARISON OF DIFFERENT LLMs

Model	Commonsense		Summarization	
	seen	unseen	seen	unseen
text-davinci-003	45.0%	45.6%	91.8%	91.2%
text-davinci-002	41.8%	37.5%	84.1%	75.7%
code-davinci-002	41.4%	39.4%	88.6%	83.2%
PaLM 540B	45.5%	49.6%	84.6%	75.7%

list of objects and list of receptacles, and then ask it to generate object placements (highlighted text) without using the provided user preferences:

```
# Put objects into their appropriate receptacles.
objects = ["black socks", "white shirt", "navy socks",
"beige shirt"]
receptacles = ["drawer", "closet"]
pick_and_place("black socks", "drawer")
pick_and_place("white shirt", "closet")
pick_and_place("navy socks", "drawer")
pick_and_place("beige shirt", "closet")
```

This baseline performs poorly, even for seen objects (45.0%), due to the high variability of object placement preferences in the benchmark. The predicted object placements are sensible but are not reflective of the particular user’s preferences. In contrast, our method can learn preferences from examples via summarization and performs much better for both seen and unseen objects (91.8% and 91.2%).

Human summary. This ablation studies how the summaries provided by the LLM compare to summaries crafted manually by a human. For each benchmark scenario, a human-written summary was used by the LLM (in place of the LLM-produced summary) to predict object placements for the test objects. The results achieved with this “oracle” summarization are better than the LLM summarization by 6% for both seen and unseen objects. This result suggests that the LLM summarizations are already quite good, and that improvements to LLM summarization could enable further gains for our method in the future.

Different LLMs. Table III reports our performance on the benchmark using different LLMs. We find that text-davinci-002 and code-davinci-002, which are older variants of GPT-3, are not as good as the newest one (text-davinci-003). In particular, there is a much larger gap between seen and unseen objects. This is because the older models are more likely to generate summaries that list out individual objects in the seen set, which does not generalize well to the unseen objects. For PaLM 540B [48], we find that while it shows slightly higher performance on common-sense reasoning, it does not do as well as text-davinci-003 on summarization, particularly in scenarios where there is a larger number of receptacles to choose from.

D. Real-world experiments

In our final set of experiments, we test the proposed approach on a robot performing a cleanup task in the real world (Fig. 1). The robot base is a holonomic vehicle capable of any 3-degree-of-freedom motion on the ground plane. This

maneuverability comes from the vehicle’s Powered-Caster Drive System [49], which consists of four caster wheels that are powered to roll and steer as needed to achieve the desired vehicle motion. The robot manipulator is a Kinova Gen3 7-DoF arm mounted on top of the mobile base with a Robotiq 2F-85 parallel jaw gripper as its end effector.

The robot is placed inside a room with various objects and receptacles on the floor and is then tasked with picking up all the objects and putting them into the correct receptacles according to user preferences. The preferences are provided as a set of textual examples for a particular user (as in the benchmark). As described in Sec. III-C and illustrated in Fig. 2, the robot iteratively locates the closest object on the floor, navigates to it, recognizes its category, picks it up, determines the appropriate receptacle for the object, navigates to the receptacle, and then puts the object inside.

Implementation. The robot uses two overhead cameras for 2D robot pose estimation (x, y, θ) and 2D object localization (x, y). The pose of the robot base is estimated using ArUco fiducial markers [50] mounted on its top plate (see Fig. 1). The object locations are detected in the overhead camera using ViLD [51], while the receptacle locations are hard-coded for each scenario. We found that these design choices work well for our mobile robot system. However, other pose trackers and object detectors could also be used instead.

To navigate in the scene, the robot calculates the shortest collision-free path to the target position using an occupancy map that includes obstacles in the scene such as receptacles. It then uses the pure pursuit algorithm [52] to follow the computed path.

After the robot arrives at the closest object, it uses a camera mounted on its base (and pointed forward at the ground) to take a close-up, centered image of the object, then determines the object category using cosine similarity between text and image features in the CLIP embedding space [44]. The set of object categories in the LLM summary is automatically extracted and used as the target label set for CLIP. Note that without these categories from LLM summarization, a human would have to manually specify a list of fine-grained object classes potentially present in the target scene in order to use CLIP for object classification.

After the object category is identified, the system uses the LLM summarization to predict the appropriate receptacle and manipulation primitive for the object. The robot then moves the object into the receptacle with a sequence of two high-level manipulation primitives: (i) pick and (ii) place or toss. The “pick” primitive uses the gripper to grasp at the center of the detected object. The “place” primitive moves the gripper to a location just above the selected receptacle and drops the grasped object in. The “toss” primitive swings the robot arm and releases the gripper with timing that results in tossing [53] of the grasped object into the selected receptacle.

Real-world evaluation. Using this mobile robot system, we ran tests on 8 real-world scenarios as shown in Fig. 3, each with its own set of 10 objects, 2–5 receptacles, 4–10 “seen” examples indicating preferences for which objects should go



Fig. 3. **Real-world scenarios.** We evaluate our mobile manipulation system in 8 real-world scenarios, encompassing a wide variety of objects and receptacles.

into which receptacles and which primitive should be used to put them there, as well as 10 “unseen” test objects. Across all 8 scenarios, 70 unique “unseen” test objects and 11 unique receptacles are represented.

For each scenario, we asked the robot to perform 3 runs of the cleanup task and measured its success throughout operation. Overall, the system was able to put 85.0% of the objects into the correct receptacle during these tests. For qualitative examples, please refer to the supplementary material and additional videos at <https://tidybot.cs.princeton.edu>.

Looking at the results in more detail, there were 240 objects to be cleaned up in total (8 scenarios, 10 objects per scenario, 3 runs per scenario). We observed that the overhead camera was able to localize 92.5% of the objects, and the object classifier correctly identified the object category for 95.5% of the localized objects. Given the predicted object category, the LLM selected the appropriate receptacle and manipulation primitive for 100% of localized objects. Additionally, the robot succeeded in executing the chosen primitive for 96.2% of the localized objects. In terms of speed, the robot took on average 15–20 seconds to pick up and put away each object.

V. CONCLUSION

In this work, we showed that the summarization capabilities of large language models (LLMs) can be used to generalize user preferences for personalized robotics. Given a handful of example preferences for a particular person, we use LLM summarization to infer a generalized set of rules to manipulate objects according to the user’s preferences. We show that our summarization approach outperforms several strong baselines on our benchmark, and we also evaluate our approach on a real-world mobile manipulator called TidyBot, which can successfully clean up test scenarios with a success rate of 85.0%. Our approach provides a promising direction for developing personalized robotic systems that can learn generalized user preferences quickly and effectively from only a small set of examples. Unlike classical approaches that require costly data collection and model training, we show that LLMs can be directly used off-the-shelf to achieve generalization in robotics, leveraging the powerful summa-

tion capabilities they have learned from vast amounts of text data.

ACKNOWLEDGMENTS

The authors would like to thank William Chong, Kevin Lin, and Jingyun Yang for fruitful technical discussions, and Bob Holmberg for mentorship and support in building up the mobile platforms.

REFERENCES

- [1] R. Rasch, D. Sprute, A. Pörtner, S. Battermann, and M. König, “Tidy up my room: Multi-agent cooperation for service tasks in smart environments,” *Journal of Ambient Intelligence and Smart Environments*, 2019.
- [2] Z. Yan, N. Crombez, J. Buisson, Y. Ruichck, T. Krajnik, and L. Sun, “A quantifiable stratification strategy for tidy-up in service robotics,” in *2021 IEEE International Conference on Advanced Robotics and Its Social Impacts (ARSO)*, 2021.
- [3] A. Taniguchi, S. Isobe, L. El Hafi, Y. Hagiwara, and T. Taniguchi, “Autonomous planning based on spatial concepts to tidy up home environments with service robots,” *Advanced Robotics*, 2021.
- [4] Y. Kant, A. Ramachandran, S. Yenamandra, I. Gilitschenski, D. Batra, A. Szot, and H. Agrawal, “Housekeep: Tidying virtual households using commonsense reasoning,” *arXiv preprint arXiv:2205.10712*, 2022.
- [5] G. Sarch, Z. Fang, A. W. Harley, P. Schydlo, M. J. Tarr, S. Gupta, and K. Fragkiadaki, “Tidee: Tidying up novel rooms using visuo-semantic commonsense priors,” in *European Conference on Computer Vision*, 2022.
- [6] N. Abdo, C. Stachniss, L. Spinello, and W. Burgard, “Robot, organize my shelves! tidying up objects by predicting user preferences,” in *2015 IEEE international conference on robotics and automation (ICRA)*, 2015.
- [7] M. Kang, Y. Kwon, and S.-E. Yoon, “Automated task planning using object arrangement optimization,” in *2018 15th International Conference on Ubiquitous Robots (UR)*. IEEE, 2018.
- [8] I. Kapelyukh and E. Johns, “My house, my rules: Learning tidying preferences with graph neural networks,” in *Conference on Robot Learning*, 2022.
- [9] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi, “Ai2-thor: An interactive 3d environment for visual ai,” *arXiv preprint arXiv:1712.05474*, 2017.
- [10] X. Puig, K. Ra, M. Boben, J. Li, T. Wang, S. Fidler, and A. Torralba, “Virtualhome: Simulating household activities via programs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [11] M. Shridhar, J. Thomason, D. Gordon, Y. Bisk, W. Han, R. Mottaghi, L. Zettlemoyer, and D. Fox, “Alfred: A benchmark for interpreting grounded instructions for everyday tasks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020.

- [12] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. S. Chaplot, O. Maksymets *et al.*, “Habitat 2.0: Training home assistants to rearrange their habitat,” *Advances in Neural Information Processing Systems*, 2021.
- [13] C. Li, F. Xia, R. Martín-Martín, M. Lingelbach, S. Srivastava, B. Shen, K. E. Vainio, C. Gokmen, G. Dharan, T. Jain *et al.*, “igibson 2.0: Object-centric simulation for robot learning of everyday household tasks,” in *Conference on Robot Learning*, 2022.
- [14] S. Srivastava, C. Li, M. Lingelbach, R. Martín-Martín, F. Xia, K. E. Vainio, Z. Lian, C. Gokmen, S. Buch, K. Liu *et al.*, “Behavior: Benchmark for everyday household activities in virtual, interactive, and ecological environments,” in *Conference on Robot Learning*, 2022.
- [15] D. Batra, A. X. Chang, S. Chernova, A. J. Davison, J. Deng, V. Koltun, S. Levine, J. Malik, I. Mordatch, R. Mottaghi *et al.*, “Rearrangement: A challenge for embodied ai,” *arXiv preprint arXiv:2011.01975*, 2020.
- [16] K. Ehsani, W. Han, A. Herrasti, E. VanderBilt, L. Weihs, E. Kolve, A. Kembhavi, and R. Mottaghi, “Manipulator: A framework for visual object manipulation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021.
- [17] L. Weihs, M. Deitke, A. Kembhavi, and R. Mottaghi, “Visual room rearrangement,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021.
- [18] M. Gupta and G. S. Sukhatme, “Using manipulation primitives for brick sorting in clutter,” in *2012 IEEE International Conference on Robotics and Automation*, 2012.
- [19] J. V. Kujala, T. J. Lukka, and H. Holopainen, “Classifying and sorting cluttered piles of unknown objects with robots: a learning approach,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.
- [20] M. Herde, D. Kottke, A. Calma, M. Bieshaar, S. Deist, and B. Sick, “Active sorting—an efficient training of a sorting robot with active learning techniques,” in *2018 International Joint Conference on Neural Networks (IJCNN)*, 2018.
- [21] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo *et al.*, “Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching,” *The International Journal of Robotics Research*, 2022.
- [22] E. Huang, Z. Jia, and M. T. Mason, “Large-scale multi-object rearrangement,” in *2019 International Conference on Robotics and Automation (ICRA)*, 2019.
- [23] H. Song, J. A. Hausstein, W. Yuan, K. Hang, M. Y. Wang, D. Kragic, and J. A. Stork, “Multi-object rearrangement with monte carlo tree search: A case study on planar nonprehensile sorting,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [24] R. Szabo and I. Lie, “Automated colored object sorting application for robotic arms,” in *2012 10th International Symposium on Electronics and Telecommunications*, 2012.
- [25] T. Dewi, P. Risma, and Y. Oktarina, “Fruit sorting robot based on color and size for an agricultural product packaging system,” *Bulletin of Electrical Engineering and Informatics*, 2020.
- [26] T. J. Lukka, T. Tossavainen, J. V. Kujala, and T. Raiko, “Zenrobotics recycler—robotic sorting using machine learning,” in *Proceedings of the International Conference on Sensor-Based Sorting (SBS)*, 2014.
- [27] S. H. Høeg and L. Tingelstad, “More than eleven thousand words: Towards using language models for robotic sorting of unseen objects into arbitrary categories,” in *Workshop on Language and Robotics at CoRL 2022*, 2022.
- [28] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, 2020.
- [29] M. Nye, A. J. Andreassen, G. Gur-Ari, H. Michalewski, J. Austin, D. Bieber, D. Dohan, A. Lewkowycz, M. Bosma, D. Luan *et al.*, “Show your work: Scratchpads for intermediate computation with language models,” *arXiv preprint arXiv:2112.00114*, 2021.
- [30] J. Wei, X. Wang, D. Schuurmans, M. Bosma, E. Chi, Q. Le, and D. Zhou, “Chain of thought prompting elicits reasoning in large language models,” *arXiv preprint arXiv:2201.11903*, 2022.
- [31] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *6th Annual Conference on Robot Learning*, 2022.
- [32] K. Lin, C. Agia, T. Migimatsu, M. Pavone, and J. Bohg, “Text2motion: From natural language instructions to feasible plans,” *arXiv preprint arXiv:2303.12153*, 2023.
- [33] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, “Language models as zero-shot planners: Extracting actionable knowledge for embodied agents,” *arXiv preprint arXiv:2201.07207*, 2022.
- [34] A. Zeng, A. Wong, S. Welker, K. Choromanski, F. Tombari, A. Purohit, M. Ryoo, V. Sindhwani, J. Lee, V. Vanhoucke *et al.*, “Socratic models: Composing zero-shot multimodal reasoning with language,” *arXiv preprint arXiv:2204.00598*, 2022.
- [35] O. Mees, J. Borja-Diaz, and W. Burgard, “Grounding language with visual affordances over unstructured data,” *arXiv preprint arXiv:2210.01911*, 2022.
- [36] B. Chen, F. Xia, B. Ichter, K. Rao, K. Gopalakrishnan, M. S. Ryoo, A. Stone, and D. Kappler, “Open-vocabulary queryable scene representations for real world planning,” *arXiv preprint arXiv:2209.09874*, 2022.
- [37] 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,” *arXiv preprint arXiv:2209.11302*, 2022.
- [38] 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.
- [39] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao, “React: Synergizing reasoning and acting in language models,” *arXiv preprint arXiv:2210.03629*, 2022.
- [40] S. S. Raman, V. Cohen, E. Rosen, I. Idrees, D. Paulius, and S. Tellex, “Planning with large language models via corrective re-prompting,” *arXiv preprint arXiv:2211.09935*, 2022.
- [41] T. Silver, V. Hariprasad, R. S. Shuttlesworth, N. Kumar, T. Lozano-Pérez, and L. P. Kaelbling, “Pddl planning with pretrained large language models,” in *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.
- [42] 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,” *arXiv preprint arXiv:2209.07753*, 2022.
- [43] D. Shah, B. Osinski, B. Ichter, and S. Levine, “LM-Nav: Robotic navigation with large pre-trained models of language, vision, and action,” *arXiv preprint arXiv:2207.04429*, 2022.
- [44] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *International Conference on Machine Learning*, 2021.
- [45] G. A. Miller, “Wordnet: a lexical database for english,” *Communications of the ACM*, 1995.
- [46] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [47] N. Reimers and I. Gurevych, “Sentence-bert: Sentence embeddings using siamese bert-networks,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [48] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, “Palm: Scaling language modeling with pathways,” *arXiv preprint arXiv:2204.02311*, 2022.
- [49] R. Holmberg and O. Khatib, “Development and control of a holonomic mobile robot for mobile manipulation tasks,” *The International Journal of Robotics Research*, 2000.
- [50] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and M. J. Marín-Jiménez, “Automatic generation and detection of highly reliable fiducial markers under occlusion,” *Pattern Recognition*, 2014.
- [51] X. Gu, T.-Y. Lin, W. Kuo, and Y. Cui, “Open-vocabulary object detection via vision and language knowledge distillation,” in *International Conference on Learning Representations*, 2021.
- [52] R. C. Coulter, “Implementation of the pure pursuit path tracking algorithm,” Carnegie-Mellon UNIV Pittsburgh PA Robotics INST, Tech. Rep., 1992.
- [53] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, “Tossing-bot: Learning to throw arbitrary objects with residual physics,” *IEEE Transactions on Robotics*, 2020.