

# Visual Assistance for Object-Rearrangement Tasks in Augmented Reality

Anonymous Authors




Fig. 1. (a) We introduce a novel application of augmented reality (AR): providing optimal assistance to users completing quotidian object-rearrangement tasks, with house-cleaning as our exemplar. By considering the user and AR device as a ‘hybrid’ agent and solving capacitated vehicle routing problems (CVRP) we compute time-optimal paths that are displayed to users in AR. (b) We develop a novel AR simulator using Habitat [39]. This serves as the test-bed of our experiments, an approach that allows us to run remote, large-scale, controlled user-studies for evaluating AR assistance. (Image courtesy [41]). (c) We deploy our framework on Amazon Mechanical Turk (AMT) [3] to study the effect of our proposed AR assistance on users’ task performance and sense of agency over a range of task difficulties.

**Abstract**—Augmented-reality (AR) glasses—which will have access to real-time, high-fidelity data regarding a user’s environment via onboard sensors, as well as an ability to seamlessly display real-time information to the user—present a unique opportunity to provide users with assistance in completing quotidian tasks. Many such tasks—such as house cleaning, packing for a trip, or organizing a living space—can be characterized as *object-rearrangement tasks* defined by users navigating through an environment, picking up objects, and placing them in different locations. We introduce a novel framework for computing and displaying AR assistance for object-rearrangement tasks that consists of (1) associating an optimal action sequence with the policy of an *embodied agent* and (2) presenting this optimal action sequence to the user as *suggestion notifications* in the AR system’s heads-up display. We make modeling choices that allow us to formalize the problem of computing AR assistance for any object-rearrangement task as a planning problem that reduces to solving a capacitated vehicle-routing problem (CVRP), a variant of the classical traveling salesman problem (TSP). In addition, we introduce a novel AR simulator that can enable web-based evaluation of AR-like assistance and associated large-scale data collection. Finally, we present the results of a user study that evaluates the effect of the proposed AR assistance on an exemplar object-rearrangement task—house cleaning—on users’ task performance and sense of agency over a range of task difficulties. Our results indicate that the proposed form of AR assistance improves overall user performance. Additionally, even though users report a negative impact to their agency, they may prefer this when presented with a system that aids them in task completion.

**Index Terms**—digital assistance, vehicle routing problem, 3D simulator, crowdsourcing, augmented reality

## 1 INTRODUCTION

Compared with current personal computing devices, always-on AR devices (1) have access to a much larger volume and more diverse set of sensor data and (2) are able to display real-time information to the user in a much lower friction manner [24]. This exposes the exciting potential for such AR devices to provide users with continual, contextually relevant assistance toward achieving their personalized goals. Consequently, assistive AR systems have been used quite extensively in specialized applications such as maintenance and manufacturing [15, 37], education [23], tourism [49], and surgery [46]. However, AR devices hold the promise of providing assistance to users on a much broader and less specialized class of commonly occurring quotidian activities. Advancements in AI and sensing technologies (e.g. Facebook’s Replica [41]) will enable AR devices to reason about

the structure and state of quotidian tasks (e.g. cooking, cleaning, and organizing) and environments. In turn, this advanced AR system can be leveraged to benefit users completing daily tasks by improving performance, reducing physical and cognitive effort while preserving a user’s sense of agency. We envision such AR assistance as providing “*superpowers for everyday tasks*”. Our goal is to develop a framework for such AR assistance applicable to an important class of everyday tasks—those involving object rearrangement—and to evaluate its value to users at scale. AI-driven technology has been shown to suffer from algorithmic bias arising from unrepresentative data [28]. At-scale data collection and evaluation can mitigate such issues for AI-powered AR assistance [28] and can also enable personalization of assistance [36].

Making progress towards the objective of pervasive AR superpowers is challenging for myriad reasons. First, there has been limited work towards formalizing the problem of computing and displaying AR assistance that can lead to improved task performance, reduced effort, and preserved sense of agency for users; while this has been investigated for specialized tasks such as AR-assisted assembly [43, 48], it has not yet been pursued for a broad class of quotidian tasks. Second, no widely available consumer AR device currently exists; as such there is no large AR user base or associated infrastructure that can support the evaluation

• *Anonymous Authors are with Anonymous Organization*

*Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx*

of AR assistance on users at scale. The field of human–robot interaction (HRI) faces a similar challenge; to overcome it, they have leveraged web-based studies that ask users to react to videos of humans and robots interacting [21, 31]. This reliance on a third-person perspective lacks immersion and limits the types of interactions that can be studied; further, no analogous approach to AR assistance would be viable as the AR assistance is not directly observable from third parties. Finally, how real users will respond to AR assistance in everyday tasks—in particular, how it affects their task performance, effort, and sense of agency—remains an open question [2, 8].

To address the first challenge in the context of object-rearrangement tasks, we formalize the problem of computing and displaying AR assistance by (1) associating an optimal action sequence with the policy of an *embodied agent* and (2) presenting this optimal action sequence to the user as *suggestion notifications* in the AR system’s heads-up display. In our formulation, the embodied agent comprises a user–AR-system ‘hybrid’ in that it has the observation space (i.e., sensor measurements) of the AR system and the action space (i.e., task-execution actions) of the user, and its policy is learned by minimizing the time to complete the task. In this initial study, we assume that the AR-system has full observability of the environment, which includes a map and real-time localization of the objects and user within that map. These modeling choices allow us to formalize the problem of computing AR assistance for any object-rearrangement task as a planning task that reduces to solving a capacitated vehicle-routing problem (CVRP) [12, 18] from combinatorial optimization. Because the optimal action sequence comprises a sequence of location visits along shortest paths, we present this action sequence by displaying the next shortest path to the user in the form of world-locked digital breadcrumbs in the heads-up display. If the user ignores the AR-system’s suggestion notifications and deviates from the optimal action sequence by visiting an alternative pickup or delivery location, we replan on the fly.

To address the second challenge, we propose a novel AR simulator that can enable large-scale web-based evaluation of AR assistance and associated data collection. Scaled data collection is difficult in AR tasks, as they require specialized hardware and environmentally situated interactions. We propose interactive web-based studies as a way to collect large scale data that can then be used in conjunction with data-driven techniques to improve AR interactions. The current work focuses on the development and deployment of an AR simulator. The simulator is based on Habitat [39] for embodied AI and satisfies the key criteria we have in an AR simulator: (1) it support the observations and actions of the proposed embodied-agent policy, (2) it emulates a first-person view through an AR device, including an ability to display suggestion-notifications in a heads-up display (HUD) in the form of digital objects and information, (3) it enables a user in the loop to autonomously perform the task-execution actions, and (4) it is deployable on the web at scale via integration with Amazon Mechanical Turk (AMT) and supports data collection related to task performance and sense of agency via psiTurk [20] integration.

To address the third challenge, we define house cleaning as an instance of an object-rearrangement task, implement the task and the proposed CVRP-based assistance using OR-Tools [19] in the proposed AR simulator, and evaluate it at scale using AMT. We collect user data across a range of task difficulties and types of AR assistance in order to evaluate how the proposed form of AR assistance affects users’ task performance, effort, and sense of agency. We find that by following the optimal assistance, users are able to decrease their total distance traveled though this comes at a cost of feeling less in control over their own actions. This cost may be one users are willing to pay, however, as we also find that users report preferring the optimal assistance to a system that does not provide them with the optimal solution. Additionally, we find that users are not consistent in their willingness to follow assistance.

Our contributions are: (1) framing daily tasks as object-rearrangement tasks, formulated and solved as a CVRP using a ‘hybrid’ single agent system (i.e., the user–AR-system), and viewing assistance as the solution to this formulation (2) a web-deployable, simulated AR experience that supports real-time interaction and displays in HUD

assistance (a visualization of the solution to the prior formulation) that can enable large-scale web-based evaluations, and (3) the results of a large-scale web-based study that assesses how users respond to the proposed form of AR assistance in a house-cleaning task over a range of task difficulties. In particular, we envision our web-deployable framework as being useful beyond the domain of AR assistance, as it can provide a framework for at-scale user-in-the-loop evaluation of different kinds of digital assistance. To the best of our knowledge, this is the first at-scale study of AR assistance for quotidian tasks.

## 2 RELATED WORK

We review existing work in AR assistance in Sect. 2.1, simulation frameworks for training and evaluating embodied-agent policies in Sect. 2.2, and assessing users’ response to digital assistance in Sect. 2.3.

### 2.1 AR assistance

To date, most work on employing AR devices to assist users has focused on displaying *predefined information overlays* that can be useful in completing a prescribed task; the information overlays are often spatially registered to objects or locations relevant to the task. For example, researchers have investigated approaches wherein the AR device overlays information or instructions on parts to be assembled or maintained [43, 44], retrieves and displays maintenance documentation using object recognition [13], overlaying medical imaging data on patients in real time [4], or provides location and activity-based, world-locked information to the users during indoor navigation [33].

Critically, none of the above approaches are truly robust for complex multi-step tasks, as they do not leverage the device’s on-board sensors to infer the current task state; as such, they are unable to provide the user with up-to-date assistance toward optimal task completion or properly adapt when users deviate from the system’s suggested steps. To fill this gap, planning-based<sup>1</sup> approaches that track task state in real time and update assistance accordingly have been proposed for a variety of AR applications such as robot tele-operation [17, 45], assistive surgery [9], assembly and manufacturing [1], and even an quotidian task of sandwich assembly [22]. In particular, Abramovici et al. [1] enable AR assistance for collaborative manufacturing using a complex back-end infrastructure where smart devices broadcast their status to a centralized planner with a predefined task dependency graph.

Instead, our proposed approach for AR assistance applicable to complex object-rearrangement tasks described in Sect. 3 leverages only the sensor measurements of the AR device, and executes a policy whose state is informed only by these sensor measurements, thus enabling up-to-date assistance toward optimal task completion. In this initial study, we assume the device has ‘full environment knowledge’ via its sensors and can support mapping and localization of the user and objects [35]. This allows us to compute the policy using a planner based on a capacitated vehicle routing problem. Perhaps the most closely related work to ours is the planning-based cognitive assistant developed by Hu et al. [22], which employs a perception module driven by computer vision to identify task state in a sandwich-assembly task and re-plan using a finite state machine. In contrast to this work, we consider a much broader category of object-rearrangement tasks and consider planners based on a general capacitated vehicle routing problem as described in Sect. 3.2. Further, our embodied AI formulation enables straightforward extensions that go beyond the strict environment-knowledge requirements of planning to learning policies that can employ perception-based partial environment observations.

### 2.2 Embodied-agent simulation frameworks

The robotics and embodied AI communities have leveraged simulation frameworks to train and evaluate embodied-agent policies in lieu of expensive real-world experimentation infrastructure [16, 25, 26, 39]. Most efforts consider fully autonomous agents (e.g. robots) whose actions directly change the physical state of the environment.

Our AR-assistance application is fundamentally different from these, in that actions that can modify the physical state of the environment be-

<sup>1</sup>Planning refers to the process of computing an agent’s policy using a model for the transition dynamics of the environment, i.e., a model that predicts the effect of the agent’s actions on the environment states [42].

long to the human agent. The AR device (our parallel to the autonomous system) can only influence the physical state of the environment indirectly via suggestion-notification actions displayed to the human. Thus, for AR assistance, any realistic *evaluation* of the how the embodied-agent policy affects task performance must involve a human in the loop—or a high-fidelity model of the human’s behavior. Evaluating how assistance affects important qualitative aspects of the user experience (e.g., effort, sense of agency) necessitates collecting data from real users at scale. On the other hand, *training* of the embodied-agent policy can be done in a variety of ways that need not rely on a human in the loop; Sect. 3 describes in detail our proposed formulation for computing the policy that does not require a human in the loop—only a planner coupled to a simulator that can compute the geodesic distances between locations and track task state—although future work will consider human-in-the-loop training loop analogously to Refs. [29, 50]. Thus, for learning assistive AR policies, rather different demands are placed on the environments used for *evaluation* and *training*; the only strong compatibility requirement is that the evaluation environment must support the states and actions associated with the trained policy to deploy it in real time. Here, we focus on the evaluation environment.

Arguably the closest related work is in the field of human–robot interaction (HRI), where researchers are interested in evaluating users’ response to robotic assistance [10]. In this field, the two most common evaluation environments comprise either (1) in-person studies where users are asked to respond to robotic assistance [14], and (2) web-based studies where users are asked to view third-person videos or photos of humans and robots interacting and react accordingly [21, 31]. The former category suffers from lack of scalability, while the latter category lacks immersion due to its reliance on a third-person perspective.

Thus, we believe that our proposed evaluation framework—which is based on an AR simulator and is described in Sect. 4—combines the attractive attributes of each of these paradigms, as it enables scalable, first-person interaction of a user with an embodied assistant and thus we hypothesize can capture more realistic user experiences—and thus yield higher quality user-experience data—at scale.

### 2.3 Evaluating user response to digital assistance

It is generally challenging to characterize the user experience for digital assistance owing to a plethora of factors involved and owing to the difficulty of measuring some of these factors with high fidelity [2]. Past studies on evaluating users’ response to digital assistance have focused on evaluating factors such as usability, utility, effectiveness, and agency [8]. In our study, we focus primarily on evaluating effectiveness as measured by the user’s task performance and sense of agency. We focus on task performance as it is easy to measure and directly addresses the key value proposition of AR assistance for complex tasks. We consider agency due to its central role in human–computer interaction (HCI) research and its relatively uncharacterized role in immersive AR-like assistive scenarios.

The sense of agency refers to the feeling of being in the driver’s seat when it comes to selecting one’s actions [32]. HCI research has long recognized the sense of agency as a key factor in characterizing how people experience interactions with technology [30, 32, 40]. In particular, one of the eight classic rules of interface design places emphasis on designing interfaces that support the user’s sense of agency [40]. Further, the sense of agency may also influence a user’s acceptance of technology [6, 27]. Despite the importance of agency, it has received limited attention in the domain of AR. It is important to address this topic early in developing novel AR technologies, as enabling user’s sense of agency is indeed challenging in assistive and immersive systems; for example, previous research has found a reduction in sense of agency with increase in automation [7].

The study presented in Sect. 5 evaluates task performance and the user’s sense of agency on the proposed cleaning-house task. Akin to other studies on AR assistance that have shown improved task performance for instance in assembly tasks [43, 48], we show shorter completion distances with the use of the proposed AR assistance. Further, we study the effectiveness over a range of task-difficulty and assistance-fidelity levels. Our study also suggests that while reports of user agency may be negatively affected by increased assistance in this

scenario, this may be a cost participants are willing to pay in order to realize improved task performance; these are promising results, and can serve as a baseline for alternative interfaces for displaying the computed AR assistance to the user.

## 3 AR-ASSISTANCE MODEL

The goal of our AR assistance model is to formalize the problem of computing and displaying AR assistance for object-rearrangement tasks. We develop this formalization for this general category of tasks because many daily activities such as following a recipe, packing for a trip, or grocery shopping could be framed in this way. Quotidian object rearrangement tasks can be repetitive, mundane, and time-consuming making them ripe for opportunity to explore how AR assistance could accelerate the task, or remove the burden of completion by decreasing a user’s cognitive burden. For the purposes of this work, we consider house cleaning as our exemplar object-rearrangement task. The specifics of the house cleaning task are described in Sect. 5.1, while in the following section we describe our generalized object-rearrangement problem formalization.

To provide assistance in a generalized object-rearrangement tasks, we adopt the perspective of embodied AI and (1) associate an optimal action sequence with the policy of an embodied agent, and (2) present this optimal action sequence to the user as suggestion notifications in the AR system’s heads-up display. To particularize this setup, we must define the embodied agent and associated partially observable Markov decision process (POMDP) ingredients: the states, observations, actions, and rewards characterizing the environment, agent, and task.

### 3.1 Embodied AI formulation

We begin by defining the objective (and associated reward) of the task as moving each object in question from its initial position to its final desired position in as little time as possible. Minimal time to task completion is not only an intuitive choice for an objective and reward, it also draws inspiration from the long history of AI assistants for task and time management [34].

Regarding the choice of embodied agent, one could adopt a fully multiagent perspective and consider the user and AR system to be independent but cooperating agents with the shared aforementioned objective but with different observation and action spaces. In this case, one could learn the AR system’s policy and present its action sequence to the user as suggestion notifications. Unfortunately, this approach introduces significant challenges, as the user is included in the AR system’s environment and thus any simulation-based learning algorithm for the AR-system’s policy would require modeling the user’s policy acting on an observation space augmented by the AR-system’s displayed information [11]. Instead, we simplify this setup and consider a single-agent formulation with an agent comprising a ‘hybrid’ between the AR system and the user. Namely, it has the observation space (i.e., sensor measurements) of the AR system and the actions (i.e., physical task-execution actions) of the user; in the case of object-rearrangement tasks, the latter corresponds to navigation and object pick/place actions. In this case, learning- or planning-based approaches for computing the embodied agent’s policy require modeling only the physical dynamics of object rearrangement.

As mentioned above, the embodied-agent’s observations correspond to those of the AR system. While many current AR head-mounted devices are equipped with only RGB video, future devices are likely to be equipped with much more high-fidelity observations. Indeed, it is likely that—at least for familiar environments—the AR device will have access to a complete map and will be able to perform localization of the user and objects [35]. As such, for this initial study, we assume that the AR device can has complete observations that include a map of the environment, the current position of all objects in question, and the position of the user. Thus, the observations and state coincide in this case, and the POMDP reduces to an MDP with deterministic transition dynamics, exposing the use for a deterministic planner as described in Sect. 3.2.

Finally, we assume that the user can carry only two objects at once, the AR device has knowledge of the desired final location of all objects

in question (i.e., the ‘goal state’ of the environment), and the AR device can calculate the shortest path between any two points on the map. These assumptions allow us to compute the policy of the embodied agent by a planner that solves a capacitated vehicle routing problem (CVRP), which we describe in Sect. 3.2 below. We remark that future work will relax the above assumptions and consider multiagent formulations, partial observations, and model-free learning of embodied-agent policies.

### 3.2 Object-rearrangement planner: capacitated vehicle routing problem

We now formulate the (single-vehicle) CVRP [12, 18] for object-rearrangement tasks. We assume that we are rearranging  $n$  objects such that each object has an initial location and final (desired) location, thus yielding  $2n + 1$  total locations of interest (including the initial position of the user). Given any arbitrary enumeration of these locations with the zeroth location corresponding to the user’s initial position (i.e., the *depot*), we decompose these locations into *pickup locations*  $\mathcal{P} \subset \{1, \dots, 2n\}$  associating with the objects’ current locations and *dropoff locations*  $\mathcal{D} \subset \{1, \dots, 2n\}$  associated with the final locations such that  $|\mathcal{P}| = |\mathcal{D}| = n$ ,  $\mathcal{P} \cap \mathcal{D} = \emptyset$ , and  $\mathcal{P} \cup \mathcal{D} = \{1, \dots, 2n\}$ . We assume that *transportation costs* can be calculated from an operator  $d: \{0, \dots, 2n\} \times \{0, \dots, 2n\} \rightarrow \mathbb{R}_+$  that calculates the geodesic distance between any two locations and satisfies  $d(i, i) = 0$ ,  $i \in \{0, \dots, 2n\}$ . If any two pickup or dropoff locations  $i$  and  $j$  coincide, we treat them as separate locations with zero separating distance such that  $d(i, j) = 0$ . This assumes the user’s time to complete the task is proportional to their total path length. We assume the user has a *capacity constraint* of  $c \in \mathbb{N}$  (which we set to be  $c = 2$  in the experiments, which assumes a user can carry one object per hand). We introduce a delivery operator  $f: \mathcal{P} \rightarrow \mathcal{D}$  that maps each pickup location to its corresponding dropoff location. We associate any solution to the problem with an invertible operator  $x: \{0, \dots, 2n\} \rightarrow \{0, \dots, 2n\}$  that maps the step number to location index; note that invertibility of this operator assumes that each location is visited exactly once and a location visit is associated with the pickup or dropoff of the appropriate object. Assuming the user has already executed  $m (\leq 2n)$  steps of the task by visiting locations  $\bar{x}(0), \dots, \bar{x}(m)$  with  $\bar{x}(0) = 0$  and the task initialized at  $m = 0$ , the planner defines the optimal sequence of location visits  $(x^*(0), \dots, x^*(2n))$  as the solution to the combinatorial optimization problem

$$\begin{aligned} & \underset{(x(0), \dots, x(2n))}{\text{minimize}} && \sum_{i=1}^{2n} d(x(i), x(i-1)) \\ & \text{subject to} && \sum_{i=1}^j 1_{\mathcal{P}}(x(i)) - 1_{\mathcal{D}}(x(i)) \leq c, \quad \forall j \in \{1, \dots, 2n\} \quad (1) \\ & && x^{-1}(f(i)) > x^{-1}(i), \quad \forall i \in \mathcal{D} \\ & && x(i) = \bar{x}(i), \quad \forall i \in \{0, \dots, m\}, \end{aligned}$$

where  $1_A$  is the indicator function that evaluates to 1 if its argument is in the set  $A$  and evaluates to zero otherwise.

The objective function in problem (1) corresponds to the transportation costs (i.e., total distance traveled); the first set of constraints corresponds to the capacity constraints; the second set of constraints ensures that each object’s pickup location is visited before its dropoff location; the third set of constraints enforces that the solution is computed from the current state of the task at step  $m$ . Given the solution  $(x^*(0), \dots, x^*(2n))$  to problem (1), at step  $p (\geq m)$  of the task, the AR device provides assistance by displaying to the user the shortest path between locations  $x^*(p) (= \bar{x}(p))$  and  $x^*(p+1)$  in the form of world-locked digital breadcrumbs in the heads-up display. If the user ‘violates’ assistance and instead visits a feasible alternative location  $\bar{x}(p+1) (\neq x^*(p+1))$ , then the system replans by solving problem (1) with  $m \leftarrow p+1$  and resumes. We emphasize that replanning is essential to ensure robustness to realistic user behavior in complex multi-step object-rearrangement tasks. In practice, we solve planning problem (1) using OR-Tools [19].

## 4 AR SIMULATOR AND DEPLOYMENT AT-SCALE

To evaluate how the proposed AR-assistance model described in Sect. 3 affects the user experience, an AR simulator must (1) support the observations and actions of the embodied-agent policy, (2) emulate the first-person view through an AR device, including support of the display of suggestion-notifications in the form of digital objects and information in a heads-up display (HUD), (3) enable a user in the loop to autonomously perform the task-execution actions, and (4) be deployable on the web at scale and support data collection related to task performance and sense of agency.

To satisfy the above criteria, we build our AR simulator upon Habitat [39]. We make this choice because Habitat satisfies all four of these criteria. First, it supports a suite of embodied-agent observations that mimic the on-board sensors of an AR device (e.g., RGBD video, compass) including those relevant to our current formulation (i.e., mapping and localization of objects and the user). Second, it can display world-locked digital objects and other information that can mimic suggestion notifications in an AR system’s HUD, thus supporting the actions of the embodied-agent policy. Third, it supports 3D environments based on real-world reconstructions (e.g., the Replica dataset [41]) rather than synthetically generated environments. Fourth, with modifications, it enables users to perform the task-execution actions required for object-rearrangement tasks. In particular, it natively supports navigation, and we implemented a rudimentary pick/place action as described below. Fifth, Habitat supports the WebGL JavaScript API that enables web-based deployment at scale. Finally, we note that Habitat supports reinforcement-learning training algorithms (e.g., PPO) [47] that we can employ to train embodied-agent policies in future studies that assume reduced observation fidelity (e.g., assume only RGBD video); this obviates the need to employ different platforms for policy training and evaluation in future work.

We make several modifications to the design and functionality of Habitat in order to generate our AR simulator. The final interface design is shown in Fig. 2. We iterated over this design with graphical and user experience designers to develop an interface sufficient for measuring our intended effects. Descriptions of the key elements of the graphical and interaction design follow. First, we implement a virtual HUD to mimic the first-person view of an AR device, Fig. 2(a); the HUD supports the display of information relevant to the proposed AR assistance, as well as the ability to display world-locked digital objects in the environment. Surrounding the HUD are text based descriptions that inform the user of the drop-off bins and their locations (left of the HUD), important interactions such as when the user places an item successfully or attempts an infeasible action (above the HUD), and the objects needing to be picked up (right of the HUD). At the suggestion of designers, we were careful to imbue the breadcrumb trail with a sense of direction (by utilizing a variable height for each breadcrumb beginning high at the paths starting location and descending towards the next object). Additionally, we eliminated the use of a 2D overhead room display, which would draw user attention away from the HUD and detract from the situated AR simulation experience.

Second, we introduce a ‘virtual knapsack’ of capacity two that represents the user’s current inventory; we display the current contents of this backpack as the user’s inventory in the virtual HUD. Third, we introduce an object pick and place action that either (1) given capacity, places an object in the virtual knapsack, or (2) removes an object from the knapsack and places it in a dropoff location matching the object type. In both cases, the user must be within a preset radius to either the object or dropoff location for the action to execute. Fourth, we integrate Habitat with the OR-Tools planner to support real-time replanning as described in Sect. 3.2.

By leveraging Habitat’s existing WebGL JavaScript API, we can deploy our AR simulator at scale on AMT and collect data related to task performance and sense of agency by leveraging psiTurk [20]. Appendix B describes the details of this setup.

## 5 STUDY SETUP

The overarching goal of our study is to evaluate user responses to AR assistance generated by the proposed framework on a house cleaning task

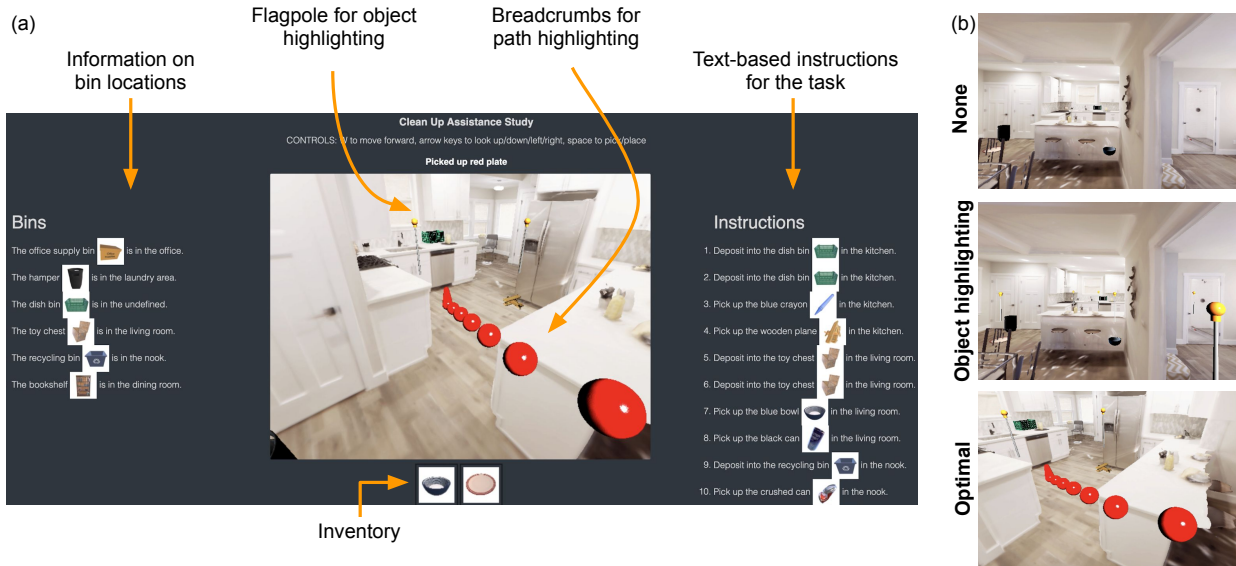


Fig. 2. (a) AR simulator interface with various elements of assistance in the designed virtual HUD. (b) Three assistance conditions from the study are shown. Optimal Assistance condition uses both object highlighting via flagpoles and path highlighting via breadcrumbs.

using our AR simulator. We describe the task and objective (Sect. 5.1), study conditions (Sect. 5.2), and the measures employed for qualitative and quantitative evaluation (Sect. 5.3).

### 5.1 House-cleaning task

The house cleaning task takes place in our AR simulator. The study consists of four phases: (1) task setup, (2) house familiarization and navigation-controls training, (3) cleaning-task execution, and (4) survey. We first situate the user into the study by explaining the purpose of the study and anchoring them onto our desired situation: a task that they are familiar enough with to do confidently, but unfamiliar enough with in order to need the help of the AR assistant. To meet both of these criteria, we decided on a house cleaning task with familiar objects in a relatively unfamiliar environment (e.g. a short term rental). This is explained to them through the following study prompt:

In this HIT, you are a guest at a short-term rental. You have been staying at the house for the past several days and checkout time is fast approaching. You must clean up this house according to the host’s instructions in as little time as possible. In order to avoid a late fee, you must navigate through the house to pick up the misplaced items and place them in the appropriate bins before checkout. For instance, you will be asked to place socks in a laundry hamper, books in a bookshelf, dishes in a dish bin, etc... You will be performing the task in a 3D virtual environment using keyboard controls, described later.

To acquaint the user with the environment as if they had been using it as a short term rental, they are shown a pre-recorded fly-through video. This video displays the room name the video frame currently inhabits and which receptacles occupy the room, if any. Additionally, the video takes care to stop in front of each receptacle as they come into view. We consider six semantic object categories and employ a specific receptacle bin for each category: dishes (dish bin), toys (toy box), books (bookshelf), laundry (laundry hamper), office supplies (office-supply box), recycling (recycling bin). Fig. 1(a) illustrates some of the objects and bins used in the study. We note that this formulation leads to a capacitated vehicle routing problem as described in Sect. 3.2, where  $n$  denotes the number objects to be cleaned up, and delivery locations are repeated when more than one object is associated with the same bin. Appendix C describes how the bins, objects, and starting location are determined for a given experiment.

After being familiarized with the environment, the user was given the opportunity to complete a shortened version of the task, with the goal

of further familiarization with the environment and keyboard controls. The objective was to pick a single item (a book) and place it into the appropriate bin (the bookshelf). This was the only bin present for the training task. Participants use the up arrow key to navigate forward, the left and right arrow keys to look left and right, and the space bar to pick and place the book. No time limits nor incentives were employed in the training task. Once the book was correctly placed, the experiment progressed to the next stage.

The next phase of the task was the experimental phase, where users completed the main house cleaning task. This phase executed similarly to the training phase, but now the participant had to organize a number of objects dictated by the difficulty condition to which they had been randomly assigned. Furthermore, they were now provided with one of the three versions of assistance described in Sect. 5.2.1. As in the training, no time limits nor incentives were employed in the task.

Finally, the participant was given a survey that collects demographic information (e.g., age, gender), responses to the subjective questions described in Sect. 5.3.2, and a free-form response to capture anything else relevant to their experience. Except for the actual task in the third phase of the HIT—which varied across participants depending on the study conditions described in Sect. 5.2—all other phases were consistent across participants.

### 5.2 Study conditions: assistance fidelity and task difficulty

We vary two key variables across experiments: *assistance fidelity*, which represents how much assistance the AR simulator provides to the user; and *task difficulty* measured by the number of objects needing replacement. We hypothesize that these two variables will be the key drivers of the user experience and task performance.

#### 5.2.1 Assistance Fidelity

We consider three different levels of assistance fidelity: no assistance, object-highlighting assistance, and optimal assistance. Each assistance level is characterized by both a world-locked digital-object component and a text-information component; see Fig. 2. In all conditions, text-information assistance includes a list of bins including their picture and semantic location.

*No assistance (None).* Participants receive no assistance from the system in this condition, which serves as our control. The egocentric frame contains only the scene, rendered objects and rendered bins; there are no additional visual cues. The text-information assistance provides a (randomized) list of objects the participant must reorganize in order to complete the task. Each item in this list contains the name of the object, a picture of the object, and the bin in which it should be



placed. Once a participant picks up an item, the text corresponding to the selected item is crossed out and moved to the bottom of the list.

*Object-highlighting assistance.* This form of assistance is designed to provide assistance to the participant under the assumption that the AR device knows the location of the objects and bins salient to the house-cleaning task and can highlight them. Such assistance—which does not rely on knowledge of traversable paths in the environment nor a real-time planner—would be especially helpful in situations where certain items are obstructed from view or are difficult to spot. This form of assistance enables participants to understand the rough locations of all objects at once and form a plan themselves. We implement this form of assistance by placing a digital flagpole over each object as depicted in Fig. 2; the corresponding text-information assistance includes all of the information from the No-assistance condition but with the addition of the name of the room in which the object can be found. Again, this list is randomized per participant and list items are crossed off as they are completed.

*Optimal assistance.* Optimal assistance corresponds to the form of AR assistance proposed in Sect. 3. A solution to this problem for the lowest difficulty setting is shown in Fig. 3. To display this information to the participant, we display the next segment of the optimal path in the egocentric frame as a trail of digital breadcrumbs, which we set to red spheres. After the user executes a feasible pick/place action, we display the next segment of the optimal path (which may involve replanning as described in Sect. 3. To prevent the participant from losing their orientation with respect to these start and end positions of the optimal path segment, the  $z$ -coordinate of the path’s breadcrumbs start at the participant’s chest level and end at the floor level (see Fig. 2). In contrast to other assistance conditions, the text-information assistance is now ordered according to the optimal path: each step lists the action the participant should take, the object they should perform it on, a picture of this object, and the object’s location. Additionally, each step is numbered in order to emphasize the importance of the list’s order. As before, items are crossed off the list as they are completed.

## 5.2.2 Task Difficulty

We consider four levels of task difficulty, where we define difficulty in terms of the number of misplaced objects that must be cleaned up. To control for certain objects being easier to reorganize (due to visual salience or bin location), the number of each object type remained fixed for each difficulty setting. A single difficulty level can thus be defined by the ratio of the number of objects to the number of bins. We use four task-difficulty levels, where this ratio corresponds to 1:1 (6 total objects), 2:1 (12 total objects), 3:1 (18 total objects), and 4:1 (24 total objects), respectively. Appendix C describes our approach for generating the bin, object, and initial user locations for any of these conditions. Fig. 3 shows the bin locations as squares in each difficulty setting, with each color representing a specific semantic object category.

## 5.3 Metrics

To determine the effect of assistance fidelity and task difficulty on participants’ task performance and experience, we collect an array of objective and subjective participant-response data.

### 5.3.1 Objective metrics

We employ four metrics to evaluate task performance: (1) *Normalized Deviations*: the number of deviations from the optimal location ordering (accounting for replanning) normalized by the total number of possible deviations, (2) *Inverse Path Length (IPL)*: the ratio of the minimal possible path length for to the sequence of location visits taken by the participant<sup>2</sup> to the total distance traveled by the participant, (3) *Task Distance*: the total distance traveled by the participant, and (4) *Task Completion Time*: the ratio of the total time taken by a participant to complete the task to the time taken by a participant to complete the fly-through familiarization phase; this ratio accounts for any system-dependent latency that may affect completion time.

<sup>2</sup>That is,  $\sum_{i=1}^{2n} d(\bar{x}(i), \bar{x}(i-1))$  in the notation of Sect. 3.2.

### 5.3.2 Subjective metrics

We employ subjective metrics that are measured using a five-point Likert scale, which focus on the following two categories<sup>3</sup>:

1. *Agency*, which is defined as the feeling of being in control [32].

“I am in charge of deciding what step I complete next during the house cleaning task” (*Control what to do*)

“I am responsible for the speed at which I completed the task” (*Control of speed*)

“I feel that I need to follow the suggestions given to me by the system” (*Need to follow*)

“I prefer that the system show me what to do next rather than figure it out myself during the house cleaning task” (*Prefer to show*)

2. *Utility and Usability*, which is aimed at measuring usefulness, user-friendliness, and acceptability; it is inspired by the System Usability Scale (SUS) [5].

“The assistance provided to me by the system during the house cleaning task helped me complete the task faster than if I had used the help provided to me during the training task” (*Usable*)

“I found the help given to me by the system to be useful” (*Useful*)

## 6 USER STUDY

We conducted a  $3 \times 4$  between-subjects user study ( $n = 447$  participants, male = 265, female = 177, other = 1, no answer = 4; three assistance-fidelity levels; four task-difficulty levels) on AMT using our AR simulator setup described in Sect. 4. Participants were compensated \$7.50 for completing the 25-minute-long study. Appendix A reports details on the number of participants in each condition. We first present and evaluate the hypotheses related to different assistance fidelity in easiest task difficulty condition using objective and subjective metrics defined in Sect. 5. We then show the interaction effects between assistance and task difficulty.

### 6.1 Effects of varying assistance fidelity

We now conduct a study that varies the assistance fidelity and fixes task difficulty to the easiest level (i.e., 6 total objects).

#### 6.1.1 Hypotheses.

Our hypotheses with respect to varying the assistance fidelity are:

**H1** Participants will follow optimal assistance when presented.

**H2** Presenting optimal assistance will lead to higher performance.

**H3** Participant agency will be unaffected by assistance fidelity.

**H4** Participants will perceive the AR system equipped with optimal assistance as more usable and useful than the AR system equipped with no assistance or object-highlighting assistance.

#### 6.1.2 Results summary

**Objective metrics.** We performed four one-way ANOVA tests to measure the effect of assistance on normalized deviations, IPL, task completion distance, and task completion time. The effect of assistance was statistically significant on normalized deviations ( $F(2, 116) = 41.48$ ,  $p = 0.001$ ), IPL ( $F(2, 116) = 15.39$ ,  $p = 0.000$ ) and task-completion distance ( $F(2, 116) = 13.05$ ,  $p = 0.000$ ) (see Figs. 4–6). No statistically significant effect of assistance was found on task-completion time ( $F(2, 116) = 0.652$ ,  $p = 0.523$ ). We performed post hoc Tukey honest significant difference tests to measure differences between each group.

In the case of optimal assistance, we observe that participants are more likely to pick and place items in the optimal order than they are in either the none or object-highlighting assistance conditions as measured

<sup>3</sup>These questions only comprise a subset of the questions that we asked in the original study. We limit our analysis here to this subset due to space constraints; see the Appendix D for the full list.




Fig. 3. The four task-difficulty levels (increasing in difficulty going from left to right), with the optimal path (corresponding to the solution of the associated CVRP) shown in the leftmost panel. Colored polygons represent approximate bounding boxes of the apartment’s floor plan, including the boundaries for each rooms; circular markers denote objects; square markers denote bins; marker color denotes semantic object category (e.g., books); the star represents the initial position.




Fig. 4. On average participants in the optimal condition deviated from the optimal ordering less frequently than participants in object-highlighting ( $p < 0.001$ ) and no assistance conditions ( $p < 0.001$ ).

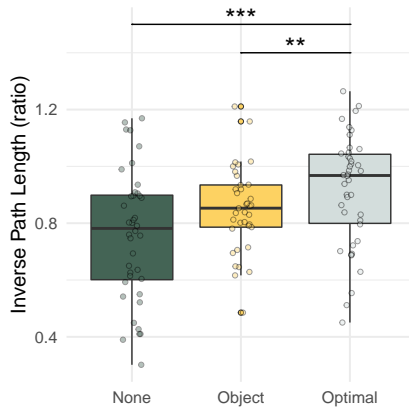


Fig. 5. Participants provided with optimal assistance followed paths that more closely resembled the optimal paths than in object-highlighting ( $p < 0.01$ ) and no assistance conditions ( $p < 0.001$ ).




Fig. 6. Participants provided with optimal assistance ( $p < 0.001$ ) and object-highlighting assistance ( $p < 0.01$ ) were more likely to find shorter paths than participants in the no assistance condition.

by normalized deviations (Fig. 4). *This suggests that people may not be able to compute independently the optimal ordering of location visits for object-rearrangement tasks based on only a first-person perspective, even with object highlighting.* Further, in the case of optimal assistance, participants follow the shortest-path trajectory between points more closely than they do in either the no assistance or object-highlighting assistance conditions (see IPL in Fig. 5). *This suggests that either a) the optimal assistance reduces a user’s burden of planning (the next object is chosen for them by way of the breadcrumb path) or b) when left to navigate freely, users do not naturally tend to the shortest path between objects.* Further studies to disambiguate this result would provide for interesting research questions that could help determine the focus of AR assistance systems in navigation contexts, but is outside the scope of the current work. Taken together, these two findings support our hypothesis H1 that participants will follow optimal assistance when presented with it.

Even though participants tend to follow optimal assistance, this does not necessarily translate to improvements in all performance metrics. We measure task performance in two additional ways: total distance traveled and total task completion time. We observe that participants presented with object-highlighting and optimal assistance generated significantly shorter total paths than those generated in the no-assistance

case, but the average path length traveled with these two forms of assistance was not substantially different (Fig. 6). *This indicates that it is possible to get users to follow shorter paths than those they might find on their own, but that simply highlighting objects may be sufficient for decreasing user path distance. Interestingly, even though participants found shorter paths in the optimal and object-highlighting conditions, there was no significant difference in the total task completion time between the three conditions.* This further indicates that even though we are able to shorten path length, this may come at some cost to the speed at which a user completes the task (potentially in the interpretation of the interface). Ultimately, navigating this trade-off is likely user or task specific, and can likely be made more favorable with alternative and personalized interface design. Taken together, these two findings partially confirm our hypothesis H2 that participants will have higher task performance when presented with optimal assistance.

**Subjective metrics.** We performed six one-way ANOVA tests to test for a main effect of assistance type within each question (Sec. 5.3). We found a statistically significant effect of assistance on Control what to do ( $F(2, 116) = 11.38, p = 0.000$ ), Need to follow ( $F(2, 116) = 7.59, p = 0.000$ ), Prefer to show ( $F(2, 116) = 15.56, p = 0.000$ ), Useful ( $F(2, 116) = 7.83, p = 0.000$ ) and Usable ( $F(2, 116) = 12.72, p = 0.000$ ). Following each ANOVA, we performed a post hoc Tukey test

when a main effect was found; Fig. 7 summarizes these results.

**Agency.** We found that participants generally felt in control of what they should do next and how quickly they completed the task. Even so, participants provided with optimal assistance, while still generally agreeing with feeling in control, reported that they felt less in control than in the other assistance conditions. Participants provided with no assistance and object-highlighting assistance were neutral about their feelings of needing to follow the assistance; participants provided with optimal assistance rated that they did feel the need to follow the assistance. Finally, participants provided with no assistance and object-highlighting assistance disagreed that they would prefer to have the system show them what to do next and where to go next. Since they were not exposed to a condition where they were provided this information, they are likely rating this against their idea of what such a system might look like. *Participants who actually were exposed to optimal assistance agreed that they preferred this to not having this assistive information. So, even though participants in the optimal assistance condition felt less in control overall, they seemed to prefer this than to an alternative.* Overall, this does not support our hypothesis H3 that participants' sense of agency would remain unaffected by assistance fidelity; however, it seems that despite feeling a slight loss in their sense of agency, they may actually prefer this sacrifice in order to obtain useful information for optimal task completion.

**Usability and Utility.** We found that participants generally felt that they were faster with assistance than without assistance. *Participants provided with both optimal assistance and object-highlighting assistance believed that they completed the task faster with the additional information than they would have without it (i.e., with no assistance).* While participants in all assistance-fidelity conditions perceived the assistance to be useful, *those provided with optimal assistance perceived it to be more useful than in either the object-highlighting or no assistance conditions.* These results are especially interesting given that no significant difference was found in total task completion time between assistance-fidelity levels. Overall, these findings partially support our hypothesis H4 that participants will be more likely to perceive the optimal assistance to be usable and useful than other assistance types.

## 6.2 Interactions between assistance and task difficulty

We had an additional hypothesis related to the interaction between assistance and task difficulty:

**H5** As the task difficulty increases, participants will be more willing to accept the assistance.

To test this hypothesis, we conducted two individual two-way factorial ANOVA tests to measure the main effects of assistance type and task difficulty on normalized deviations and IPL, and any combined effects of assistance type and difficulty (see Fig. 8). Assistance type ( $F(2, 435) = 265.75, p = 0.000$ ) and difficulty ( $F(3, 435) = 76.62, p = 0.000$ ) had statistically significant effects on normalized deviations. The interaction between the two independent variables was also statistically significant ( $F(6, 435) = 2.865, p = 0.001$ ). The effect of assistance type on IPL was statistically significant ( $F(2, 435) = 30.71, p = 0.000$ ). There was no significant effect of task difficulty on IPL, and the interaction effect was not statistically significant. We performed post hoc Tukey honest significant difference tests, where effects were found.

Across the board, we observed that participants deviated more from the optimal ordering as task-difficulty increased (Fig. 8). In the cases of no assistance and object-highlighting assistance, this trend conforms to our expectation that people have difficulties in finding the optimal solution in object-rearrangement tasks. We still observe this trend, however, when people are explicitly given the optimal ordering as in the optimal-assistance condition. Interestingly, though, the variance in the optimal-assistance condition is consistently much greater than those in either the no assistance or object-highlighting assistance cases. This suggests that the underlying distribution of assistance acceptance is likely multi-modal, and future work can investigate mechanisms to increase acceptance among various sub-populations of users.

*We also observed that participants were more likely to take shorter paths in the optimal-assistance condition than in the other two con-*

*ditions, and that participants provided with object-highlighting assistance were more likely to take shorter paths than those provided with no assistance (Fig. 8).* However, there was no change in IPL as the task-difficulty level increased. Taken together, this refutes our hypothesis H5 that users would be more likely to accept assistance as the task became harder. In fact, the data suggest that they are either less likely or just as likely to accept the assistance as the task difficulty increases.

## 7 DISCUSSION, LIMITATIONS, AND FUTURE WORK

While completing our work, we made several decisions that lead to specific limitations of our work. Each limitation, though, leads to interesting discussion and directions for future work. We group these limitations into two categories and discuss them as well as the future work they imply in Sect. 7.1 and Sect. 7.2.

### 7.1 Scale versus ecological validity

One important limitation in the current work is the lack of a study verifying the ecological validity of our AR simulation. Unfortunately, we were unable to run this validation study due to the challenges that arose for in-person user studies due to the COVID-19 pandemic, when the development and execution of this study took place. In developing our AR simulator, we made special care to ensure that it is compatible with AR/VR devices, and could be deployed in an ecological validation study in the future. Performing this study and comparing the user behavior in our simulator with behavior in an ecologically valid environment is an important part of our future research direction.

We believe that work into scaling AR interactions is important for the future progress of daily AR interactions. Even with limited ecological validity, there is a benefit to the data that can be collected. Specifically, the data can provide important insights about user preferences related to AR-assistance. Furthermore, such data could support the training of data-driven, personalized AR assistance models. For example, a future direction for our work is to personalize the assistance model from an task-optimal model to one that incorporates user path preferences. By collecting large volumes of data, we would be able to train assistance algorithms that understand and consider the types of paths users feel comfortable navigating (e.g., paths without stairs, paths with fewer turns).

Another important future direction of our work is to relax the constraint that our AR system have access to extremely detailed environmental information. Pursuing this work also benefits from a method of simulated interaction, again, in order to employ data-driven techniques to train algorithms that are robust to uncertain environments and noisy perception.

### 7.2 User interaction with optimal assistance

Our experiments consider a single exemplar task with a singular layout. Our aims were to test the usefulness of our AR simulator setup and the effects of the proposed assistance paradigm in controlled experiments. User acceptance of assistance, though, is likely dependent on a variety of factors including environment familiarity, prior task knowledge or experience, or task enjoyment. Some of these factors may not even be measurable at a population level, and require individual analyses. Given the positive results from our experiments, alongside the ability to rapidly prototype and deploy new object rearrangement tasks at scale, we are excited to explore the effects of optimal assistance in new contexts and different tasks in order to explore these questions in the future.

Another facet of our experiment was the reliance on users' ability and desire to follow both written (the 'Instruction' list) and visual (the breadcrumbs) instructions. While we refined our interface with the help expert graphic and user experience designers to develop a task appropriate interface, this design process was not the main focus of our study. Further iteration on the interface design as well as inclusion of additional instruction delivery modalities (e.g., audio or social) would greatly benefit our understanding of how to best deliver assistance.

Finally, one of our main aims was to determine if users would follow optimal assistance, when given, and if this assistance would improve their performance (H1 and H2). While both of these hypotheses were confirmed, there are still open questions regarding user acceptance of various assistance paradigms. For example, would users blindly






Fig. 7. **Left**, When compared against no assistance and object-highlighting assistance, participants in the optimal assistance condition were less likely to feel in control (both  $p < 0.001$ ), more likely to feel the need to follow the assistance (both  $p < 0.01$ ), and more likely to prefer being told what to do (both  $p < 0.001$ ). **Right**, Participants in the optimal assistance were more likely to rate assistance as usable ( $p < 0.001$ ) and useful than those provided with no-assistance ( $p < 0.001$ ). Those in object-highlighting more strongly agreed that the assistance was useful when compared to no-assistance ( $p < 0.05$ ). Those in optimal assistance more strongly agreed that the assistance was useful as compared to the object-highlighting assistance ( $p < 0.05$ ).




Fig. 8. **Left**, participants provided with optimal assistance were less likely to deviate from the optimal ordering than participants provided with either no assistance or object-highlighting assistance ( $p < 0.001$ , both). As the difficulty increased, participants were more likely to deviate in each group. **Right**, participants provided with optimal assistance were more likely to follow shortest paths (higher IPL) than those provided with both object-highlighting ( $p < 0.01$ ) and no ( $p < 0.001$ ) assistance. Participants provided with object-highlighting assistance took shorter paths than those in none ( $p < 0.001$ ).

follow random, non-optimal, or even worst-case assistance should it be provided? Does the user’s knowledge of the task affect how they view the AR assistance system? Would users more readily accept assistance from a sub-optimal but personalized paradigm? Studies relating to the user’s trust of the assistance were outside the scope of our own study, but they are a fruitful and important direction for future work in AR assistance.

## 8 CONCLUSION

This work has presented (1) a novel framework for computing and displaying AR assistance for object-rearrangement tasks, which characterize a broad category of quotidian tasks, (2) a novel AR simulator that can enable web-based evaluation of AR-like assistance and large-scale data collection, and (3) a study that assesses how users respond to the proposed AR assistance in the AR simulator on a specific object-rearrangement task: house cleaning.

The study illustrated several salient trends. First, by following the optimal assistance, participants were able to reduce the overall distance travelled, suggesting that people could benefit from a system like the one we propose. Second, though participants’ reported feeling less in control over their own actions when following the optimal assistance they were also more likely to agree that they preferred a system that told them what to do than users in either other group. This indicates that users may be willing to sacrifice a small amount of agency in favor of a system that provides useful assistance. Finally, users were less

likely or equally likely to accept the assistance as the difficulty of the task increased, though the population of users in the optimal assistance condition exhibited a much wider variability of assistance acceptance. This indicates that there are potential subgroups in the user population, and that future work should be conducted to discover these groups and develop personalized assistance systems.

Finally, future work will explore extensions of the current framework and study, including performing ecological validation studies, developing and assessing perception-based learned policies that assume lower fidelity of the embodied AR agent’s observations, considering multi-agent formulations that incorporate models of the user’s behavior and preferences, and extending the current assistance framework and AR simulator to other quotidian object-rearrangement tasks.

## ACKNOWLEDGMENTS

The authors wish to thank A, B, and C. This work was supported in part by a grant from XYZ (# 12345-67890).

## REFERENCES

- [1] M. Abramovici, M. Wolf, S. Adwernat, and M. Neges. Context-aware maintenance support for augmented reality assistance and synchronous multi-user collaboration. *Procedia CIRP*, 59(1):18–22, 2017.
- [2] W. Albert and T. Tullis. *Measuring the user experience: collecting, analyzing, and presenting usability metrics*. Newnes, 2013.
- [3] Amazon. *Mechanical Turk*.

- [4] M. Bajura, H. Fuchs, and R. Ohbuchi. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *ACM SIGGRAPH Computer Graphics*, 26(2):203–210, 1992.
- [5] A. Bangor, P. T. Kortum, and J. T. Miller. An empirical evaluation of the system usability scale. *Intl. Journal of Human-Computer Interaction*, 24(6):574–594, 2008.
- [6] A.-M. K. Baronas and M. R. Louis. Restoring a sense of control during implementation: how user involvement leads to system acceptance. *Mis Quarterly*, pp. 111–124, 1988.
- [7] B. Berberian, J.-C. Sarrazin, P. Le Blaye, and P. Haggard. Automation technology and sense of control: a window on human agency. *PLoS One*, 7(3):e34075, 2012.
- [8] P. M. Berry, T. Donneau-Golencer, K. Duong, M. T. Gervasio, B. Peintner, and N. Yorke-Smith. Evaluating user-adaptive systems: Lessons from experiences with a personalized meeting scheduling assistant. In *IAAI*, vol. 9, pp. 40–46. Citeseer, 2009.
- [9] H. Bourquain, A. Schenk, F. Link, B. Preim, G. Prause, and H.-O. Peitgen. Hepavision2—a software assistant for preoperative planning in living-related liver transplantation and oncologic liver surgery. In *CARS 2002 Computer Assisted Radiology and Surgery*, pp. 341–346. Springer, 2002.
- [10] R. Brooks. A brave, creative, and happy hri, 2018.
- [11] L. Buşoniu, R. Babuška, and B. De Schutter. Multi-agent reinforcement learning: An overview. In *Innovations in multi-agent systems and applications-1*, pp. 183–221. Springer, 2010.
- [12] G. B. Dantzig and J. H. Ramser. The truck dispatching problem. *Management science*, 6(1):80–91, 1959.
- [13] J.-Y. Didier, D. Roussel, M. Malle, S. Otmane, S. Naudet, Q.-C. Pham, S. Bourgeois, C. Mégard, C. Leroux, and A. Hocquard. Amra: augmented reality assistance for train maintenance tasks. 2005.
- [14] A. D. Dragan, K. C. Lee, and S. S. Srinivasa. Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 301–308. IEEE, 2013.
- [15] J. Egger and T. Masood. Augmented reality in support of intelligent manufacturing—a systematic literature review. *Computers & Industrial Engineering*, 140:106195, 2020.
- [16] Z. Erickson, V. Gangaram, A. Kapusta, C. K. Liu, and C. C. Kemp. Assistive gym: A physics simulation framework for assistive robotics. *arXiv preprint arXiv:1910.04700*, 2019.
- [17] H. Fang, S. Ong, and A. Nee. A novel augmented reality-based interface for robot path planning. *International Journal on Interactive Design and Manufacturing (IJIDeM)*, 8(1):33–42, 2014.
- [18] B. L. Golden, S. Raghavan, and E. A. Wasil. *The vehicle routing problem: latest advances and new challenges*, vol. 43. Springer Science & Business Media, 2008.
- [19] Google. *OR-Tools*.
- [20] T. M. Gureckis, J. Martin, J. McDonnell, A. S. Rich, D. Markant, A. Coenen, D. Halpern, J. B. Hamrick, and P. Chan. psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods*, 48(3):829–842, 2016.
- [21] G. Hoffman. Evaluating fluency in human–robot collaboration. *IEEE Transactions on Human-Machine Systems*, 49(3):209–218, 2019.
- [22] Z. Hu, M. Shvo, A. Jepson, and I. Mohomed. Interactive planning-based cognitive assistance on the edge. In *3rd {USENIX} Workshop on Hot Topics in Edge Computing (HotEdge 20)*, 2020.
- [23] M.-B. Ibáñez and C. Delgado-Kloos. Augmented reality for stem learning: A systematic review. *Computers & Education*, 123:109–123, 2018.
- [24] T. Jonker, R. Desai, K. Carlberg, J. Hillis, S. Keller, and H. Benko. The role of ai in mixed and augmented reality interactions. In *CHI '20 workshop "Artificial Intelligence for HCI: A Modern Approach"*, 2020.
- [25] N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 3, pp. 2149–2154. IEEE, 2004.
- [26] 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*, 2017.
- [27] K. Le Goff, A. Rey, P. Haggard, O. Oullier, and B. Berberian. Agency modulates interactions with automation technologies. *Ergonomics*, 61(9):1282–1297, 2018.
- [28] N. T. Lee, P. Resnick, and G. Barton. Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms. *Brookings Institute: Washington, DC, USA*, 2019.
- [29] J. Li, A. H. Miller, S. Chopra, M. Ranzato, and J. Weston. Dialogue learning with human-in-the-loop. *arXiv preprint arXiv:1611.09823*, 2016.
- [30] H. Limerick, D. Coyle, and J. W. Moore. The experience of agency in human-computer interactions: a review. *Frontiers in human neuroscience*, 8:643, 2014.
- [31] B. F. Malle, M. Scheutz, J. Forlizzi, and J. Voiklis. Which robot am i thinking about? the impact of action and appearance on people’s evaluations of a moral robot. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 125–132. IEEE, 2016.
- [32] J. W. Moore. What is the sense of agency and why does it matter? *Frontiers in psychology*, 7:1272, 2016.
- [33] A. Mulloni, H. Seichter, and D. Schmalstieg. Handheld augmented reality indoor navigation with activity-based instructions. In *Proceedings of the 13th international conference on human computer interaction with mobile devices and services*, pp. 211–220, 2011.
- [34] K. Myers, P. Berry, J. Blythe, K. Conley, M. Gervasio, D. L. McGuinness, D. Morley, A. Pfeffer, M. Pollack, and M. Tambe. An intelligent personal assistant for task and time management. *AI Magazine*, 28(2):47–47, 2007.
- [35] R. A. Newcombe, J. J. Engel, J. Straub, T. J. Whelan, S. J. Lovegrove, and Y. Ren. Augmented reality mapping systems and related methods, July 9 2020. US Patent App. 16/822,828.
- [36] E. Ohn-Bar, K. Kitani, and C. Asakawa. Personalized dynamics models for adaptive assistive navigation systems. *Conference on Robot Learning*, 2018.
- [37] R. Palmarini, J. A. Erkoyuncu, R. Roy, and H. Torabmostaedi. A systematic review of augmented reality applications in maintenance. *Robotics and Computer-Integrated Manufacturing*, 49:215–228, 2018.
- [38] W. Reese. Nginx: the high-performance web server and reverse proxy. *Linux Journal*, 2008(173):2, 2008.
- [39] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 9339–9347, 2019.
- [40] B. Shneiderman and C. Plaisant. *Designing the user interface: strategies for effective human-computer interaction*. Pearson Education India, 2010.
- [41] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.
- [42] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [43] A. Tang, C. Owen, F. Biocca, and W. Mou. Comparative effectiveness of augmented reality in object assembly. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 73–80, 2003.
- [44] P. Thomas and W. David. Augmented reality: An application of heads-up display technology to manual manufacturing processes. In *Hawaii international conference on system sciences*, pp. 659–669, 1992.
- [45] C. S. Tzafestas. Teleplanning by human demonstration for vr-based teleoperation of a mobile robotic assistant. In *Proceedings 10th IEEE International Workshop on Robot and Human Interactive Communication. ROMAN 2001 (Cat. No. 01TH8591)*, pp. 462–467. IEEE, 2001.
- [46] P. Vávra, J. Roman, P. Zonča, P. Ihnát, M. Němec, J. Kumar, N. Habib, and A. El-Gendi. Recent development of augmented reality in surgery: a review. *Journal of healthcare engineering*, 2017, 2017.
- [47] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra. DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames. *arXiv*, pp. arXiv–1911, 2019.
- [48] Z. Yang, J. Shi, W. Jiang, Y. Sui, Y. Wu, S. Ma, C. Kang, and H. Li. Influences of augmented reality assistance on performance and cognitive loads in different stages of assembly task. *Frontiers in Psychology*, 10:1703, 2019.
- [49] R. Yung and C. Khoo-Lattimore. New realities: a systematic literature review on virtual reality and augmented reality in tourism research. *Current Issues in Tourism*, 22(17):2056–2081, 2019.
- [50] R. Zhang, F. Torabi, L. Guan, D. H. Ballard, and P. Stone. Leveraging human guidance for deep reinforcement learning tasks. *arXiv preprint arXiv:1909.09906*, 2019.

# Appendices

## A PARTICIPANTS STATISTICS IN OUR STUDY

Table 1 shows the number of participant in each of our twelve conditions. Participants were uniformly randomly assigned to one of these conditions when they accepted to participate in the study.

	6 objects	12 objects	18 objects	24 objects
None	40	39	44	40
Obj-highlight	37	38	21	30
Optimal	42	42	32	42

Table 1. The number of participants across conditions in our study.

## B WEB APPLICATION IMPLEMENTATION DETAILS

The setup works as follows: we serve a Human Intelligence Task (HIT) to AMT using a psiTurk server, which also allows us to advertise our HIT on through the AMT web portal. Through this portal, participants can view a study description, compensation details, and the estimated completion time. After accepting a HIT, participants consent to study participation, which initiates serving the approximately 8 GB Habitat WebGL application to the user’s web browser using a combination of a psiTurk server and an NGINX server [38]. The majority of the application is loaded directly onto the participant’s computer, with the exception of the OR-Tools replanning module. When a disagreement between the participant’s actual path and the computed optimal path occurs, the participant’s web browser communicates with the psiTurk server to recalculate the optimal path and send this back to the client. After completing the survey, the data collected during the experiment (e.g., keyboard actions, time spent completing each phase, survey responses) are transmitted back to the server where it is stored in a MySQL database for later use. After the participant completes the experiment, we employ psiTurk to approve and disburse payment through AMT.

## C DISORGANIZED HOUSE GENERATION

We now describe how we generated the disorganized house for the study described in Sect. 5.1.

### C.1 Bin placement

Bins were placed manually within the environment in semantically reasonable locations. For example, the dish bin was placed on top of the counter next to the sink and the office supply box was placed on top of the desk in the office. Receptacle locations were kept fixed in each difficulty setting.

### C.2 Object placement

To sample object locations within the scene we randomly sampled a total of 40 navigable scene points using the Habitat simulator. We used rejection sampling in order to ensure that sampled points were at least one unit of distance away from every other sampled point and receptacle location within the scene. We then used the first N points from this list to define the object locations in the scene. For example, the lowest difficulty setting (6 objects) contained the first 6 points from this list, the highest difficulty setting contained the first 30 points. This way, each scene built on top of the previous scene in order to control for any single scene having an outlying dispersion between points. Each point was assigned a semantic object category, as well. This was kept consistent throughout each difficulty setting. The actual object model used for any individual point was held constant within difficulty settings, but was randomly sampled across difficulty settings. Thus, if a sampled point was assigned to the books and magazines category in the lowest difficulty setting, it would be assigned to the books and magazines category in the highest difficulty setting, as well, but it might not be exactly the same book. Object locations for each difficulty setting are shown as circles in Fig. 3.

### C.3 Starting Location

The starting position of each participant was held constant for any individual scene. This position was calculated by first finding the centroid of all of objects placed in the scene. This point was determined to be

navigable using the Habitat simulator. If the point was not navigable, rejection sampling was used to find a navigable point within a small radius surrounding the centroid. This method allowed us to ensure that the participant’s starting position did not bias their solution by their starting location. Additionally, this method allowed us to generate our optimal assistance, discussed in Sec. 3.2.

## D SUBJECTIVE METRICS

The full list of Likert-scale questions used in our study is below (a subset was described in Sec. 5.3 and used for analysis in Sec. 6.1):

### 1. Agency:

“I am in charge of deciding what step I complete next during the house cleaning task” (*Control what to do*)

“I am in charge of deciding where I go next during the house cleaning task” (*Control where to go*)

“I am responsible for the speed at which I completed the task” (*Control of speed*)

“I feel that I need to follow the suggestions given to me by the system” (*Need to follow*)

“I prefer that the system show me what to do next rather than figure it out myself during the house cleaning task” (*Prefer to show what*)

“I prefer that the system show me where to go next rather than figure it out myself during the house cleaning task” (*Prefer to show where*)

### 2. Utility and Self-efficacy:

“I completed the task as quickly as I could”

“I followed the help given to me by the system”

“I took longer than I needed to to complete the task”

“I found the help given to me by the system to be useful”

“I found a better way to complete the task than offered to me by the system”

### 3. Usability:

“I found the house cleaning task more difficult to complete than the training task”

“I thought the system was easy to use”

“I would imagine that most people would learn to use the help provided by the system very quickly”

“I would like to use the help provided to me by system during the house cleaning task frequently in real life house cleaning scenarios”

“The assistance provided to me by the system during the house cleaning task helped me complete the task faster than if I had used the help provided to me during the training task”

“The help provided to me by the system during the house cleaning task was sufficient in order to accomplish the task”.