The Privacy-Utility Tradeoff for Remotely Teleoperated Robots

Daniel J. Butler, Justin Huang, Franziska Roesner, and Maya Cakmak
University of Washington. Computer Science & Engineering
185 Stevens Way, Seattle, Washington, USA
{djbutler,jstn,franzi,mcakmak}@cs.washington.edu

ABSTRACT

Though teleoperated robots have become common for more extreme tasks such as bomb diffusion, search-and-rescue, and space exploration, they are not commonly used in human-populated environments for more ordinary tasks such as house cleaning or cooking. This presents near-term opportunities for teleoperated robots in the home. However, a teleoperator's remote presence in a consumer's home presents serious security and privacy risks, and the concerns of end-users about these risks may hinder the adoption of such in-home robots. In this paper, we define and explore the *privacy-utility* tradeoff for remotely teleoperated robots: as we reduce the quantity or fidelity of visual information received by the teleoperator to preserve the end-user's privacy, we must balance this against the teleoperator's need for sufficient information to successfully carry out tasks. We explore this tradeoff with two surveys that provide a framework for understanding the privacy attitudes of end-users, and with a user study that empirically examines the effect of different filters of visual information on the ability of a teleoperator to carry out a task. Our findings include that respondents do desire privacy protective measures from teleoperators, that respondents prefer certain visual filters from a privacy perspective, and that, for the studied task, we can identify a filter that balances privacy with utility. We make recommendations for in-home teleoperation based on these findings.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems—*human factors, software psychology*

General Terms

Design; Human Factors

Keywords

Privacy; Remote teleoperation

HRI '15 March 02 - 05 2015, Portland, OR, USA

ACM 978-1-4503-2883-8/15/03 ...$15.00.
http://dx.doi.org/10.1145/2696454.2696484

1. INTRODUCTION

While full autonomy in unstructured environments remains highly challenging for robots, many complex tasks can be performed reliably with human supervision or direct human control of robots. Indeed, there are already commercially available systems for remote teleoperation, such as the iRobot Packbot.[1] Though teleoperated robots have become common in extreme environments, they are not commonly used in human-populated environments for more ordinary tasks. Thus ordinary tasks such as house cleaning or cooking present unexploited opportunities for robot teleoperation, which can allow remote operators to work anywhere at any time, shifting night jobs to day time zones and avoiding transportation costs for workers, or improving productivity through partial automation.

Unfortunately, the introduction of teleoperated robots into human-populated environments presents serious privacy, security, and safety risks. These risks present a hurdle to making in-home teleoperated robots attractive to more people. In this paper, we focus primarily on privacy risks: a worker operating a robot remotely in a customer's home can learn significant information about that customer (e.g., their financial information, personal habits, medical conditions, and political or religious views). Such concerns may be greater for a teleoperator than a physical worker in the home due to the anonymity and de-personalization created by the physical distance. Furthermore, digital recordings of people's homes are inherently vulnerable to being intentionally or accidentally revealed to a public audience.

To reduce such privacy concerns, one might suggest that the information (e.g., video) provided to remote teleoperators should be limited. However, providing teleoperators with too little information may interfere with their proper execution of tasks, raising concerns not only about their effectiveness but also about potential physical harm caused by poor task execution (e.g., breaking items in the home). Thus, we are faced with a tradeoff.

In this paper, we define and explore this *privacy-utility tradeoff* for remotely teleoperated robots. Different tasks require different types of information, and likewise, different users have different privacy preferences. As a result, it is not obvious *a priori* how to strike a balance. To begin characterizing this complex tradeoff, we contribute:

- a framework for specifying privacy issues in a teleoperated robot scenario, based on a survey that reveals people's privacy attitudes in this context;
- a sample set of 2D and 3D filtering techniques for vi-

[1] `www.irobot.com/us/learn/defense/packbot.aspx`

sual information provided to teleoperators;
- empirical results from a second survey revealing people's preferences towards these filtering techniques applied in different contexts; and
- a user study that investigates a teleoperator's ability to perform a specific task with different privacy filters.

We report on both qualitative and quantitative results from our studies. From these results, we distill recommendations for understanding and balancing the privacy-utility tradeoff. For example, we observe that end-users may not anticipate all of their privacy concerns without sufficient context in which to consider them; that certain visual filters do indeed meet end-users' privacy preferences; that a small loss in utility can result in a large privacy gain; and furthermore that the performance hit of a high-privacy filter decreases as the teleoperator gains experience. Our empirical results and characterization of the privacy-utility tradeoff lay the groundwork for enabling in-home teleoperated robots to become socially acceptable and useful.

2. RELATED WORK

Remote teleoperation. Remote teleoperation has become a subject of interest both commercially and in the research community [9]. Most existing teleoperated systems target extreme conditions, such as bomb diffusion or search and rescue [2]. More recently, however, researchers have started to look into teleoperation in human-populated environments, such as homes or offices [16, 17]. We target such everyday, human-populated environments in this work as well.

Privacy in robotics. Privacy has increasingly become a topic in robotics. For example, Feil-Seifer *et al.* [10] consider privacy for socially assistive robotics, and Kahn *et al.* [14] consider bystander privacy around humanoid robots. Others have discovered that anthropomorphic robots naturally deliver privacy notice [6] reducing the privacy-enhancing behaviors of older adults compared to a camera [5]. Telepresence systems [23] naturally mitigate some privacy concerns by displaying the person controlling the robot; nevertheless, privacy is a major concerns for older adults considering a telepresence robot in their home [3, 4]. Drones have also recently raised significant privacy concerns [7]. Willow Garage's Heaphy project[2] involving robots teleoperated by Mechanical Turk workers was shut due in part to privacy concerns. In this work, we study how to better balance privacy and utility for teleoperated robots to make them more acceptable to end-users and ultimately more useful.

Other related work in privacy. Beyond robots, many researchers have studied privacy issues with video surveillance and wearable cameras. Solutions generally involve explicit opt-outs of various kinds for bystanders and objects [11, 18–21], and/or more automatic video filtering techniques [12, 13, 22, 24]. These previous approaches assume that sensitive objects can be explicitly detected via computer vision techniques or rely on expensive instrumentation of the world. However, this assumption conflicts with a major motivation behind teleoperation: namely, that human teleoperators can identify and manipulate objects that are not currently recognizable by computer vision. In this work, we thus develop generic filters that are widely applicable to a large class of unknown objects. Nevertheless, more targeted computer vision and/or explicit opt-outs can supplement our blanket approach.

[2]The Heaphy project: `http://youtu.be/0aqghgoeCWk`

3. HOME PRIVACY FRAMEWORK

Our motivating scenario involves remotely teleoperated robots in the home that can carry out ordinary tasks such as cleaning, organizing, and cooking. The workers teleoperating the robot may be located anywhere, but we envision that they are vetted by the service company, that their performance may be rated by end-users, and that their actions through the teleoperation interface may be audited.

Although such robots can provide great benefits to both end-users and to workers, their success hinges on the willingness of end-users to allow such robots into their home. End-users are likely to have privacy concerns about allowing unknown workers to view their home through the robot's sensory feed. We thus begin by considering the privacy concerns of end-users in their homes.

3.1 Privacy concerns

To characterize in-home privacy concerns, we developed a set of dimensions that may affect a person's level of concern, based in part on relevant privacy literature (e.g., [13, 24]). We generated the following (overlapping) dimensions for the evaluation of privacy concerns:

1. *Locations:* Different in-home locations — such as the bedroom, bathroom, living room, or kitchen — may present inherently different levels of privacy concern. For example, the bedroom may be more likely to contain private or sensitive objects than the living room.
2. *Objects:* Different in-home objects may be more sensitive than others, and this sensitivity may vary among users. For example, keys may be sensitive, because photos of keys can be used to replicate them.[3]
3. *Information:* Finally, we can classify privacy concerns according to the higher-level information revealed through objects and/or locations. Potentially sensitive information may include financial information, medical information, information about a person's identity, personal habits, political or religious views, etc.

The relative concern of end-users along each dimension and the variability of concern between different end-users will inform the design of privacy filters or other approaches for limiting the information provided to teleoperators.

3.2 Survey design

To better understand people's concerns in the teleoperated robot scenario, and to empirically validate the above framework for evaluating privacy concerns in particular, we conducted a web-based user survey using Google Forms.[4]

The first page of the survey described the in-home teleoperated robot scenario alongside an image of a UBR-1 robot[5] for context. The second and third pages each asked a general free-response question of the form: "In this scenario, what are some X you would be concerned about?", where X was

[3]`https://keysduplicated.com/`

[4]Google Forms is a free service for creating web-based surveys. `http://www.google.com/forms/about/`

[5]UBR-1 is a state-of-the-art mobile manipulator with a circular omni-directional base and one 7-DoF arm. `http://unboundedrobotics.com/ubr-1/`

replaced with "things" and "privacy-related issues" respectively. This ordering was chosen to find out if privacy would come up naturally as a concern, before the survey revealed that its main focus was privacy.

Each of the next four pages consisted of 5-point Likert-scale questions that asked about objects, rooms, information types, and threat types respectively (see the previous section for our rationale). For example, in the case of objects (e.g., keys, pants, pills), the questions were of the form: "If this object was present in the robot's environment, I would be ..." with 1 indicating "Not at all concerned about privacy" and 5 indicating "Extremely concerned about privacy".

The last page consisted of demographic questions and general privacy-related questions to allow us to categorize respondents by their level of privacy concern according to the Westin Privacy Index [15].

3.3 Findings

Demographics. Our survey respondents were 25 male and 25 female volunteers recruited via email at the University of Washington. Ages ranged from 18 to 71 years old (mean = 28.4, standard deviation = 10.3). An analysis of our Westin Privacy Index questions (coded as described in [15]) revealed 21 of 50 respondents as *Privacy Fundamentalists*, 25 as *Privacy Pragmatists*, and 4 as *Privacy Unconcerned*. Compared to historical Westin Index data [15], our respondents may therefore be slightly, but not dramatically, skewed towards privacy concerned.

Finding 1: Privacy and harm are major concerns. The first question of the survey asked in free-response (qualitative) form about general concerns with the teleoperated robot scenario. Though this question explicitly did not yet mention privacy, many respondents voiced privacy-related concerns. Specifically, two authors independently coded 10 concerns commonly mentioned by respondents, and then attempted to reach consensus wherever there was disagreement (Table 1). Respondents' most common concerns were privacy (22 of 50), harm to people or property (18 of 50), and "other" (things that did not fit into any other category, e.g. size, expense) (13 of 50). As an additional check, we noted that 17 of respondents specifically used the words "privacy" or "private" in their response. In total, 26 of 50 of respondents mentioned concerns about either privacy issues or leakage of sensitive information, suggesting that sensitive visual information collected by a robot is a major issue to address for teleoperated robots in the home. We observe that privacy concerns may be in tension with concerns about physical harm: for a well-intentioned teleoperator, this is precisely the privacy-utility tradeoff.

Finding 2: Privacy concerns vary by context and are greatest for tangible harms. Next, we consider a quantitative measure: Likert scale ratings of respondents' level of privacy concern for different objects (Fig. 1a), types of information (Fig. 1b), and locations (Fig. 1c). Wilcoxon signed rank tests were used for computing significance. We find that respondents are more concerned about some contexts than others. For example, they are significantly more concerned about privacy in the bedroom or the bathroom than in the living room or the kitchen; they are more concerned about bank statements than about jewelry, and more concerned about jewelry than about deodorants; and they are more concerned about financial and personal identification information than about personal habits or gender information. In general, respondents were more concerned about information that they could imagine concretely causing them harm (e.g., financial or home security harm) than less tangible privacy violations (e.g., learning their gender for targeted advertising). These impressions were borne out in respondents' ratings of their concern about specific threats (Fig. 1d): in general, respondents were concerned about data thefts and embarrassing information getting out onto the web, but not about targeted advertising.

Concern	**% answers** (N, κ)
Privacy	**44%** (22, 0.92)
Harm to people or property	**36%** (18, 0.91)
Other (e.g., size, expense)	**28%** (14, 1.00)
Home security (break-in, theft)	**26%** (13, 1.00)
Inability to perform tasks well	**24%** (13, 0.89)
Leakage of sensitive information (e.g., financial, identity)	**14%** (7, 1.00)
Operator actions (nonspecific)	**13%** (6.5, 0.91)
Pets	**12%** (6, 1.00)
Who is liable for damage / harm	**10%** (5, 1.00)
Hackers	**10%** (5, 1.00)

Table 1: Percentage of 50 survey respondents mentioning each concern, given the prompt: "In the described scenario with a teleoperated robot in the home, what are some things you would be concerned about?" Percentages are averaged from two authors' codings of free responses. Average raw count and inter-coder agreement, as measured by Cohen's κ, are shown in parentheses.

Finding 3: Respondents may not always anticipate threats. We hypothesize that respondents did not always imagine the full context and/or anticipate the resulting threats when rating their privacy sensitivities. For example, in free response answers, respondents expressed concern about embarrassing information getting out on the web, but may not have considered this threat when rating their (lower) sensitivity towards information about their personal habits or messiness. As another example, one respondent (female, age 29) said explicitly: *Before these questions came up, I honestly have not considered things like credit cards or mail or other things being a privacy concern. But after going through the survey, I see how the Robot could "steal" your identity (or someone could hack into the Robot and steal your identity or personal information).* We return to this observation when we discuss the results of our second survey (Section 4.3), where we find even stronger evidence that respondents' stated privacy preferences may vary by the degree of context provided.

4. PRIVACY FILTERING FRAMEWORK

Our first survey provides a clearer understanding of the types of information that a privacy filter ought to remove. Next, we set out to determine candidate filters and evaluate them in terms of their effectiveness in enhancing privacy. Though robots may have many sensors that capture different types of information, we focus on sensors that capture visual information, as these are most intuitive for a human viewer.

Mapping the findings of our survey to filter designs is not straightforward. The human eye performs highly complex transformations on visual information in an image, in order to extract meaningful information from it. It is impractical to reverse engineer these transformations so as to create filters such as a "political information filter" or an

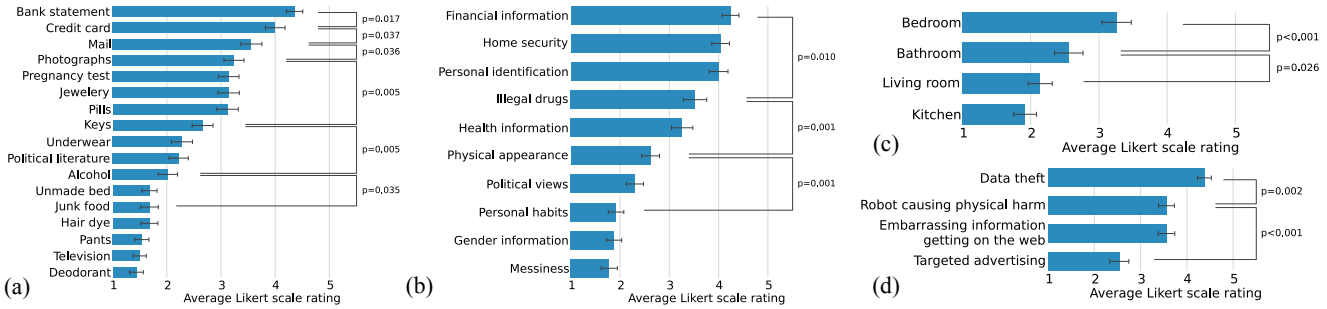

Figure 1: Survey #1 results about in-home privacy sensitivity towards (a) objects, (b) information types and (c) rooms; (d) concern about specific threats, where 1 on the scale corresponds to "Not at all concerned about privacy" and 5 corresponds to "Extremely concerned about privacy." We indicate p values where differences are significant.

"illegal activity filter." Instead, filters operate a much lower level where they manipulate properties such as edges and color. Nonetheless, our findings indicate that filters that remove *text*, would have high impact in improving privacy, as information rated as highly sensitive is primarily revealed through text (e.g., bank statement, credit card, pregnancy test). Focusing on this observation, we explore four filters.

4.1 Image filters

Blur (Fig. 2a). The simplest approach for removing fine details like text from an image is to apply a Gaussian blur filter. Gaussian blur removes image features smaller than a certain scale, controlled by the width parameter σ. Larger values of σ remove larger text but reduce the utility of the image for executing tasks. We found that $\sigma = 5$px was approximately the minimum size that made most of the text in our image set illegible, so we used this value.

Edge (Fig. 2b). The distribution of color and intensity in an image reveals information about the identity of objects, the type of material a surface is made of, the 3D shape of surfaces, and cleanliness. In order to hide color and intensity, we used the Canny edge detection algorithm [8] to remove all information except the edges between visually similar regions. Canny edges are often present along the outlines of objects, which may improve utility for manipulation tasks.

Superpixel (Fig. 2c). Blur not only removes fine details like text but also distorts objects boundaries. To mitigate the latter effect, we used the SLIC superpixel algorithm [1] to cluster pixels that are close in 2D space and color space, and then replaced each cluster with its average value. This process acts like a non-linear filter that removes fine details while preserving the boundaries of objects.

Color-skewed superpixel (Fig. 2d). Superpixels have the disadvantage that they preserve color regions, which may allow familiar objects and brands associated with particular colors to be identified. To conceal color information, we rotated hue by 180°. This helps hide identifying colors while preserving shading useful for perceiving 3D shape.

4.2 Survey design

We designed a second user survey to better understand how these filters interact with end-user privacy preferences, again aiming for both qualitative and quantitative results. We presented respondents with the same in-home teleoperated robot scenario as in the first survey (Section 3.2). We then showed respondents images or short videos of nine in-home objects, displayed in a real context but with minimal additional clutter (e.g., keys on a table, pants on the floor). Based on the results of our first survey, we selected three high-sensitivity objects (credit card, photograph, pregnancy test), three medium-sensitivity objects (pills, keys, underwear), and three low-sensitivity objects (unmade bed, hair dye, pants). For each object, we asked respondents to rate:

1. their level of privacy concern related to the object's unmodified image or video (5-point Likert scale), and
2. their level of privacy concern related to the object's image modified with four different filters (Fig. 2) described in Sec. 4.1 (5-point Likert scale for each filter, with free response explanation).

Additionally, we asked about respondents' level of comfort with human workers versus teleoperators, with known teleoperators, and with being around the robot. We also asked respondents' specifically about whether they would be willing to give more information to the teleoperator in exchange for performance. Finally, we asked demographic and Westin Privacy Index [15] questions.

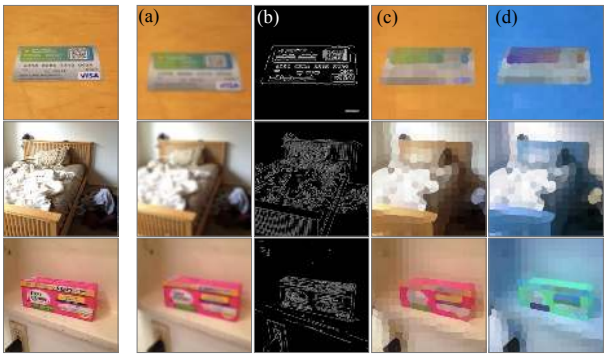

Figure 2: The effect of four filters on sample images from survey #2: (a) *blur*, (b) *edge*, (c) *superpixel*, (d) *color-skewed superpixel*.

4.3 Findings

Demographics. The respondents to the second survey were 25 male, 21 female, and 1 other or unknown gender volunteers recruited via email at our institution. Ages ranged from 18 to 57 years old (mean = 26.9, standard deviation = 8.7). As in the first survey, we categorized respondents by Westin Privacy Index [15], finding 22 of 47 *Privacy Fun-*

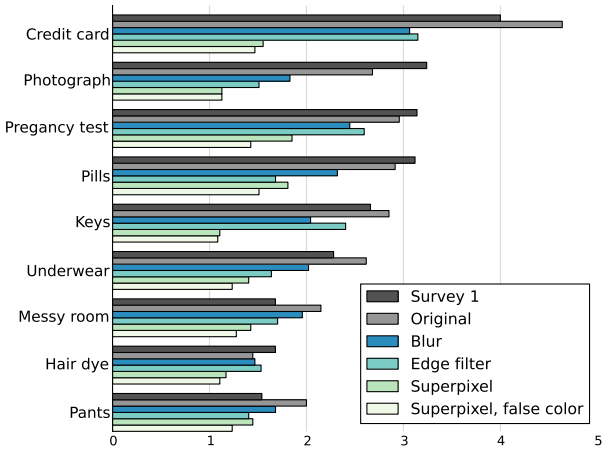

Figure 3: Survey #2 results about privacy sensitivity towards objects modified by various filters, where 1 on the scale corresponds to "Not at all concerned about privacy" and 5 corresponds to "Extremely concerned about privacy."

damentalists, 22 *Privacy Pragmatists*, and 3 *Privacy Unconcerned*. Again, our respondents are thus skewed slightly towards being more privacy-sensitive than historical Westin Index respondents [15].

Finding 1: The superpixel filters were preferred for preserving privacy. Figure 3 shows respondents' reported privacy sensitivity for objects whose images or videos were modified by each filter, as well as the original image or video. This is a quantitative measure using a Likert scale. The Wilcoxon signed rank test was used to test significance. Lower bars on the graph represent greater comfort with a remote teleoperator viewing the image or video. Among the four tested filters (blur, edge, superpixel, and color-skewed superpixel), we find that respondents were more comfortable with images or video modified with a superpixel filter. For example, one respondent observed that the superpixel filters "let the robot know it is a credit card, but nothing else." Specifically, for all objects, the color-skewed superpixel filter preserved privacy better than every filter other than superpixel in a statistically significant way ($p < 0.05$).

Finding 2: Context affects responses about privacy sensitivity. In our first survey (Section 3.3) we found that respondents sometimes gave seemingly inconsistent responses about their privacy sensitivity. Our second survey strengthens this observation: when shown an in-content image or video of some objects in the second survey, respondents were statistically more concerned about privacy than when the same objects were only mentioned verbally in the first survey. Specifically, between objects in the first survey and unfiltered objects in the second survey, increases in privacy concern are significant for the credit card (p=0.006), pants (p = 0.020), and messy room (p=0.019). We used the Kruskal-Wallis test as our test of significance. See Figure 3. Thus, again we find that the context in which respondents are asked about their privacy sensitivity affects their responses, suggesting that our results are a lower bound: we would expect respondents to have been even more sensitive about objects in their own, real homes. We return to this lesson when we make recommendations in Section 6.

Finding 3: Respondents desire a tradeoff based on context. We asked respondents directly about the privacy-utility tradeoff, i.e., whether they would be "willing to show clearer images to teleoperators to improve their performance." We coded their free-text answers (two coders; Cohen's κ for inter-coder agreement was 0.82) and report average percentages across both coders. We find that a plurality of respondents (59% of 47) would be willing to show clearer images in some cases, as long as certain objects or rooms remain obscured, or other conditions (e.g., asking permission) are met. For example, one respondent (male, age 22) wrote: *Some things I wouldn't want them to see at all (pregnancy test), some it doesn't matter the clarity (clothes).* By contrast, only 19% of respondents answered with an unconditional "yes" and 16% answered with an unconditional "no" (with 6% unsure). These responses underscore the need for a solution that trades off privacy with utility, where the tradeoff may vary depending on the specific context.

5. PRIVACY-UTILITY TRADEOFF

So far we have focused on privacy, considering only the end-user's perspective. We now turn to the perspective of the teleoperator, who needs as much information as possible to perform a task. We hypothesize that for a given task, not *all* sensory information is relevant, and that some methods of filtering the video input received from the robot's surroundings will not affect the teleoperator's performance on that task. To balance the privacy-utility tradeoff, we would ideally like to find a filter that is acceptable for an end-user's privacy and that minimally impacts the teleoperator's performance. Following up on our first two surveys, we conducted a user study with the goal of identifying at least one such filter. We describe this user study next.

5.1 Platform

Our user study involved the PR2 research robot. PR2 is a wheeled robot with two arms (7 degrees of freedom) and parallel grippers that can manipulate everyday objects. For the teleoperation of the PR2 we used an open source graphical system developed by Leeper *et al.* called *Interactive Manipulation* (IM) [16]. IM allows users to click on parts of the robot to manipulate them. To manipulate objects the user clicks on the end-effector of the a 3D rendering of the robot on their screen, which reveals a 6-dimensional control with 3 arrows for translation in each direction and 3 wheels for rotation around each direction. The user right clicks on the gripper and selects *open* or *close* to grasp or place items.

IM allows users to customize the sensors that are overlaid in the 3D world of the robot or displayed in a separate panel on the same window. Our study involved a typical configuration with a point cloud obtained from the Kinect sensor rendered in the 3D view side-by-side with a camera image view from the robot's pan-tilt head [16].

5.2 Filters

The second survey revealed that the color-skewed superpixel filter (Sec. 4.1) was superior in preserving privacy. Therefore we are interested in knowing how it ranks in terms of utility. For comparison we wanted to design two additional settings that would be in either extreme of the privacy-utility spectrum. In addition to the 2D image filtering, the IM interface required designing filters for the 3D point cloud as well. The three views we designed are ex-

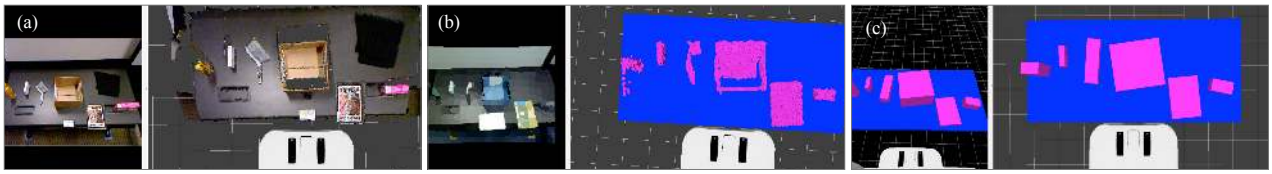

Figure 4: Three teleoperator views designed for the user study: (a) *clean*, (b) *obscured*, and (c) *box* views.

plained in the following (Fig. 4). We assumed a scenario in which the robot manipulates objects on a table.

Clean view. On one extreme, we provide an unfiltered image and point cloud that contains all available information.

Obscured view. The second view contains the 2D image filtered with the color-skewed superpixel method. As a companion 3D filter for this view we removed all points from the point cloud that do not correspond to objects on a table and removed color information from the rest of the points. In addition, using the distance channel RBGD image from the Kinect sensor, we blacked out all pixels in the 2D image that approached the sensor more than a certain distance.

Box view. On the other extreme, we wanted to push privacy as far as possible, leaving minimal task information. To that end, we fit a bounding box to all detected clusters on the table and only displayed these bounding boxes, leaving the point cloud completely out. For the 2D view we provided a rendering of these boxes as seen from the camera.

5.3 User study design

Our user study has two aims. The first is to compare the three filters and assess their ranking on the privacy-utility spectrum. The second is to better understand the risks associated with an *active* adversarial who does not just observe the information captured by the robot, but also can control the robot to gather further information. To investigate these issues, our study included two tasks.

Functional task. The first task aimed to measure the utility of the filtered views for teleoperation. Participants had four minutes to use the PR2 to pick up three objects on a table and place them into a box also on the table. The objects were arranged to the left of the box and were pointed out to participants on the teleoperation views. To reflect common challenges in teleoperation and robot perception, one of the objects was selected to have an irregular shape (e.g., a brush with a handle) and another was selected to be semi-transparent (e.g., a water bottle). After the functional task, participants were asked to respond to a *privacy quiz* asking questions about which objects were present. Specifically, participants were asked to identify whether the robot's environment contained a men's product, political literature, medication, clothing, etc., and to indicate their certainty in their answer. They were also asked to identify individual objects and transcribe any text that they might have seen.

Adversarial task. In the second task, the goal was to use the robot to gather as much information as possible in order to complete the privacy quiz. Participants did not have to place objects into the box in this task.

5.3.1 Protocol

For convenience, the robot was located in the same room as the participant, and remote teleoperation was simulated by hiding the robot behind a barrier. To mask sounds of the robot in operation, participants wore noise-isolating headphones playing background noise. The overall structure of the experiment was as follows:

1. **Introduction.** Participants gave informed consent and were seated at a computer workstation. Experimenter explained the goal of the study as the development of privacy-preserving interfaces.
2. **Tutorial.** Experimenter provided a step-by-step explanation of the teleoperation interface used to control the robots. Participant demonstrated understanding by successfully using each interface element.
3. **Practice task.** Participant teleoperated the robot to pick up an object, moved the arm to its fullest extent in the horizontal and vertical directions, and placed the object back down. Participants familiarized themselves with the privacy quiz.
4. **Experiment.**
   (a) Flight 1 functional task, followed by privacy quiz.
   (b) Flight 1 adversarial task, followed by privacy quiz.
   Similarly for Flights 2 and 3.
5. **Questionnaire.** Participant answered questions about the difficulty of the functional and adversarial tasks under each view, as well as questions about level of privacy concern under each view, demographic questions, the Westin index questions.

The flight numbers correspond to three sets of objects that were presented in the same order to all participants, with each flight displayed under a different view (clean, obscured, or box). The views were presented in a different order for each participant. Over the 18 participants, each of the 6 possible orderings was repeated 3 times.

5.4 Findings

Finding 1: Sacrificing a little utility can significantly improve privacy. In terms of utility, the obscured view was about as good as the clean view, but it provided much better privacy. In particular, teleoperators rated the obscured view just -0.67 Likert points lower than the clean view on ease-of-use (Figure 5b), but rated it $+2.06$ Likert points better in terms of privacy for the adversarial task (-0.63 and $+2.62$ standard deviations respectively). See Figure 5d. One user (male, age 25) summed up the tradeoff this way: *I could tell the general shape of the objects, and could tell what some of them were (box, book etc.), but I couldn't get any details.* We see the same trend for objective measures: the obscured and clean views did not exhibit a significant difference in the average number of objects successfully placed into the box (paired t-test: $p = 0.72$, see Figure 6), while adversarial teleoperators using the obscured view were significantly worse at answering questions about

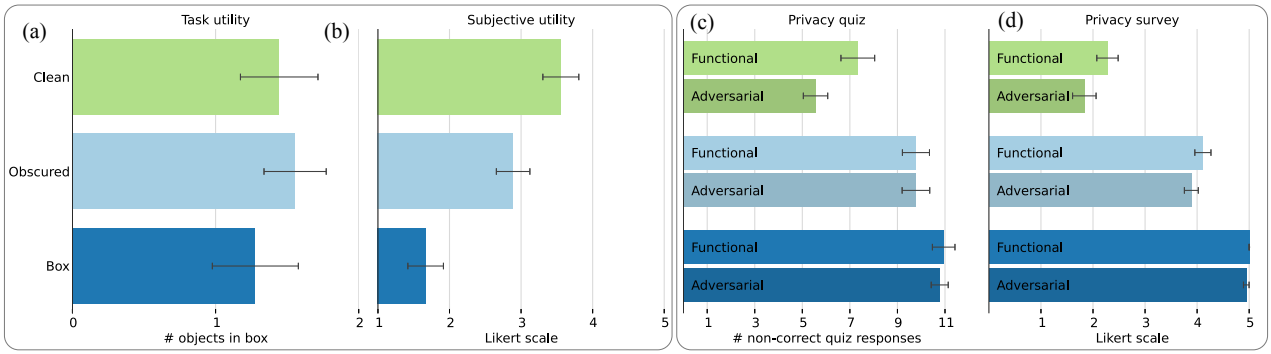

Figure 5: Summary of user study results: objective and subjective measures of task utility (left) and privacy (right).

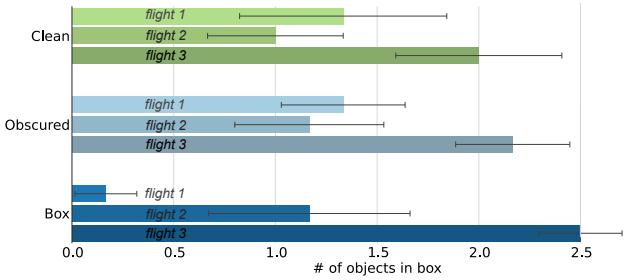

Figure 6: The number of objects successfully placed into the box in 4 minutes, broken down by flight and filter. Sect. 5.4 discusses the stronger learning effect for the box view.

the objects in the scene than those using the clean view (paired t-test: $p < 0.001$).

Finding 2: Practice can mitigate lower-utility views. We observed that practice greatly improved teleoperator performance on the subjectively difficult-to-use box view, and moderately improved performance on the other views (Figure 6). While this learning effect may have been due in part to additional information about the object set gleaned by teleoperators over the course of the study, this finding nevertheless suggests that the performance hit of a filter is not fixed but may in fact diminish or disappear over time.

Finding 3: Active information gathering reduces privacy, but filters can limit this effect. During the adversarial task, teleoperators used a variety of strategies to glean information from the scene. For example, they pushed objects around to get a better view of labels and text, brought objects closer to the camera to try to overcome the effects of image filters, and even tested object rigidity to help determine which objects were made of fabric. In the clean view, adversaries scored twice as high on the privacy quiz (paired t-test: $p < 0.0001$) and rated the difficulty of the quiz significantly lower (Wilcoxon signed-rank test: $p < 0.001$). However, in the obscured and box views, the gains of adversaries were not significant, suggesting that these views were somewhat resilient to active threats.

6. DISCUSSION

Finally, we step back to discuss limitations of our work, make recommendations, and outline avenues for future work.

Limitations. Our work has several limitations. First, our user study evaluated only a single, specific task (placing objects in a box). Though we were able to identify a visual filter that balanced privacy and utility in an acceptable way for this task, this choice of filter does not necessarily generalize to other tasks. Second, several aspects of our user study limited our ability to draw conclusions from the data, including a strong learning effect among participants and the lack of fine granularity data about what participants did at what time. For example, we did not study the degree to which different filters led to inadvertent physical harm, i.e., the disruption of the scene. Third, we have not studied the perspective of human bystander near the robot; knowing that their image is filtered may not be sufficient to make bystanders comfortable around the robot, and further investigation of bystander attitudes are needed. Finally, our survey and study populations were limited in demographic diversity (e.g., it did not include a large older adult or disabled population, who might be early adopters of the studied technology). Nevertheless, our work presents an important first step in understanding and managing the privacy-utility tradeoff for remotely teleoperated robots.

Recommendations and future work. Based on our findings in two surveys and a user study, we make the following recommendations for the design of services for in-home teleoperated robots and beyond:

- Users express different privacy preferences as details and context emerge. Thus, privacy preferences should be elicited from users with as much context as possible. For example, a user could be shown images of their own home (as in [24]), rather than an abstract list of objects, when making preference decisions.
- Users were most comfortable with our two superpixel filters. We recommend future empirical study of these and similar filters for different teleoperation tasks.
- Users recognize that the optimal point on the privacy-utility spectrum may vary by task, by object, and by user. Future work should explore how to balance this tradeoff dynamically as these contexts change. For example, the filter parameters or even the choice of filter could change in real time. To aid this process, users could explicitly mark sensitive objects [19, 20].
- We were surprised how much a filter's performance hit diminishes with practice. Thus, low-utility filters may ultimately prove more valuable than expected.
- While this paper has studied visual filters not specific to particular context, there are other possible techniques for balancing privacy and utility that must be

studied. For example, in-home robots may be restricted from certain rooms rather than certain objects. As computer vision and robotic autonomy improve, this tradeoff can perhaps also be balanced by reducing the involvement of the teleoperator.

- Privacy filters have robotic applications beyond just teleoperation. Autonomous robots could similarly store information in a filtered form in order to be less vulnerable to unintended security breaches.

7. CONCLUSION

This paper has defined and explored the privacy-utility tradeoff for remotely teleoperated robots in the home. Although such robots present tremendous near-term opportunities, their success depends on the willingness of end-users to allow them into their home. We conducted two surveys to characterize qualitatively and quantitatively the privacy concerns and preferences of end-users, finding that respondents are concerned both about privacy and physical harm from teleoperated robots. We observed that respondents were not always able to anticipate all threats, and thus recommend that end-users be asked about their privacy preferences with as much context as possible. We also found that privacy concerns vary by specific context, but that most respondents were comfortable with the level of privacy provided by one of our visual filters (color-skewed superpixel). Finally, in a user study in which participants manipulated a robot, we found that an intermediate filter provided a good privacy-utility balance for the studied task: participants were able to carry out the task with reasonable accuracy and only moderate difficulty, but they were not able to answer privacy-invasive questions. We also found that the performance hit of a privacy-preserving filter diminishes with practice. Though the optimal point in the privacy-utility spectrum varies by task, by context, and by end-user, our findings suggest how these properties can be traded off in acceptable ways. Our characterization of in-home privacy concerns and our empirical exploration of the privacy-utility tradeoff thus lays a foundation for future designs of remotely teleoperated robots in the home.

References

[1] Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Susstrunk, S. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence 34*, 11 (2012), 2274–2282.

[2] Baker, M., Casey, R., Keyes, B., and Yanco, H. A. Improved interfaces for human-robot interaction in urban search and rescue. In *SMC (3)* (2004), pp. 2960–2965.

[3] Beer, J. M., and Takayama, L. Mobile remote presence systems for older adults: acceptance, benefits, and concerns. In *6th International Conference on Human-Robot Interaction* (2011), ACM, pp. 19–26.

[4] Boissy, P., Corriveau, H., Michaud, F., Labonté, D., and Royer, M. A qualitative study of in-home robotic telepresence for home care of community-living elderly subjects. *Journal of Telemedicine & Telecare 13*, 2 (2007), 79–84.

[5] Caine, K., Sabanovic, S., and Carter, M. The effect of monitoring by cameras and robots on the privacy enhancing behaviors of older adults. In *ACM/IEEE International Conf. on Human-Robot Interaction* (2012).

[6] Calo, R. The drone as privacy catalyst. *Stanford Law Review Online 64* (2011).

[7] Calo, R. Against notice skepticism in privacy (and elsewhere). *Notre Dame Law Review 87* (2012).

[8] Canny, J. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6 (1986), 679–698.

[9] Chen, J. Y., Haas, E. C., and Barnes, M. J. Human performance issues and user interface design for teleoperated robots. *IEEE Transactions on Systems, Man, and Cybernetics 37*, 6 (2007), 1231–1245.

[10] Feil-Seifer, D., Skinner, K., and Matarić, M. J. Benchmarks for evaluating socially assistive robotics. *Interaction Studies 8*, 3 (2007), 423–439.

[11] Halderman, J. A., Waters, B., and Felten, E. W. Privacy Management for Portable Recording Devices. In *Workshop on Privacy in Electronic Society* (2004).

[12] Jana, S., Molnar, D., Moshchuk, A., Dunn, A., Livshits, B., Wang, H. J., and Ofek, E. Enabling Fine-Grained Permissions for Augmented Reality Applications with Recognizers. In *USENIX Security Symposium* (2013).

[13] Jana, S., Narayanan, A., and Shmatikov, V. A Scanner Darkly: Protecting User Privacy from Perceptual Applications. In *IEEE Symposium on Security and Privacy* (2013).

[14] Kahn Jr, P. H., Ishiguro, H., Friedman, B., Kanda, T., Freier, N. G., Severson, R. L., and Miller, J. What is a human?: Toward psychological benchmarks in the field of human–robot interaction. *Interaction Studies 8*, 3 (2007), 363–390.

[15] Kumaraguru, P., and Cranor, L. F. Privacy Indexes: A Survey of Westin's Studies. Tech. Rep. CMU-ISRI-5-138, Carnegie Mellon University, 2005.

[16] Leeper, A. E., Hsiao, K., Ciocarlie, M., Takayama, L., and Gossow, D. Strategies for human-in-the-loop robotic grasping. In *ACM/IEEE International Conference on Human-Robot Interaction* (2012), ACM, pp. 1–8.

[17] Mast, M., Španěl, M., Arbeiter, G., Štancl, V., Materna, Z., Weisshardt, F., Burmester, M., Smrž, P., and Graf, B. Teleoperation of Domestic Service Robots: Effects of Global 3D Environment Maps in the User Interface on Operators' Cognitive and Performance Metrics. In *Social Robotics*. Springer, 2013, pp. 392–401.

[18] Patel, S. N., Summet, J. W., and Truong, K. N. BlindSpot: Creating Capture-Resistant Spaces. In *Protecting Privacy in Video Surveillance*, A. Senior, Ed. Springer-Verlag, 2009, pp. 185–201.

[19] Raval, N., Srivastava, A., Lebeck, K., Cox, L. P., and Machanavajjhala, A. MarkIt: Privacy Markers for Protecting Visual Secrets. In *UPSIDE* (2014).

[20] Roesner, F., Molnar, D., Moshchuk, A., Kohno, T., and Wang, H. J. World-Driven Access Control for Continuous Sensing Applications. In *ACM Conference on Computer and Communications Security* (2014).

[21] Schiff, J., Meingast, M., Mulligan, D. K., Sastry, S., and Goldberg, K. Y. Respectful Cameras: Detecting Visual Markers in Real-Time to Address Privacy Concerns. In *Proceedings of the International Conference on Intelligent Robots and Systems* (2007).

[22] Senior, A., Pankanti, S., Hampapur, A., Brown, L., li Tian, Y., and Ekin, A. Blinkering surveillance: Enabling video privacy through computer vision. *IBM Research Report 22886* (2003).

[23] Sheridan, T. B. Teleoperation, telerobotics and telepresence: A progress report. *Control Engineering Practice 3*, 2 (1995), 205–214.

[24] Templeman, R., Korayem, M., Crandall, D., and Kapadia, A. PlaceAvoider: Steering first-person cameras away from sensitive spaces. In *Network and Distributed System Security Symposium* (2014).