Available online at www.sciencedirect.com
ScienceDirect
International Journal of Industrial Ergonomics 37 (2007) 589–604
International Journal of INDUSTRIAL ERGONOMICS
www.elsevier.com/locate/ergon

Design guidelines for map-based human–robot interfaces: A colocated workspace perspective

Woohun Lee[a], Hokyoung Ryu[b,*], Gyunghye Yang[c], Hyunjin Kim[c], Youngkuk Park[c], Seokwon Bang[c]

[a]*Department of Industrial Design, Korea Advanced Institute of Science and Technology, South Korea*
[b]*Institute of Information and Mathematical Sciences, Massey University, Auckland 1311, New Zealand*
[c]*Samsung Advanced Institute of Technology, South Korea*

Received 25 May 2006; received in revised form 16 March 2007; accepted 21 March 2007
Available online 4 May 2007

Abstract

To ensure the success of the near future home-service robots, it is essential to develop an affordable and effective instruction mechanism that is well fitted both to the characteristics of the tasks the robots will perform and the work environment where they will operate. As an early exploration of this line of studies, this paper explores a situation where human operators direct a robot to a particular place within a limited workspace using a handheld device. The three experiments revealed that a successful map representation would have significant benefits for the human operator's awareness of both the task and the work environment. As a consequence, several design guidelines for the map representation were empirically attained.

Keywords: Human–robot interaction; Map; Awareness; Interfaces; Hand-held device

1. Introduction

Home-service robot systems has been proposed to assist people in their daily lives with a wide spectrum of practical applications, e.g., cleaning, surveillance and/or search jobs. To ensure the success of these applications, it is essential to have large scale usability testing with both intended user groups and actual home-service robots, so one can develop usable and useful systems. This is a problem, however, in that the design patterns and conventions for home-service robots are still evolving, not yet thoroughly established. Very recently, several researchers (e.g., Kadous et al., 2006; Scholtz, 2005; Yanco et al., 2004) have proposed a systematic design approach for the development of human–robot interaction. This paper therefore takes the same approach in order to explore some design features that one should pay attention to when designing human–robot interfaces. The following sections begin with a brief account as to what features of human–robot interfaces are closely related to, or different from, our understanding from early studies, so that commercial human–robot interface designers may be aware of the issues involved in creating effective interfaces for instructing personal home-service robots.

1.1. Interacting with robots

One of the contrasting characteristics of human–robot interaction (HRI) is that robots are mobile in an open area under a human operator's supervision (Yanco and Drury, 2004). For instance, a synchronous and colocated work context, which is receiving wide attention as an environment for the development of future robotics (Forlizzi, 2005), is said to involve a human operator in order to ensure their own safety (Mynatt et al., 2000). This colocated work context, such as the home environment, would prevent the robot from working solely independently yet.

*Corresponding author. Tel.: +64 9 4140800x9140; fax: +64 9 4418181.
E-mail address: h.ryu@massey.ac.nz (H. Ryu).

0169-8141/$ - see front matter
doi:10.1016/j.ergon.2007.03.004

This issue has led to much research within HRI community as to how an effective collaboration between the robot and the human operator can be made (Scholtz, 2003; Yanco and Drury, 2004; Yanco et al., 2004). The probable primary requirement of this synchronous and colocated HRI is to mix robots and humans in an unstructured and uncontrolled environment in which all manner of obstacles can have unpredictable results.

Fig. 1 illustrates a synchronous and colocated human–robot interaction, delivering the operator's instructions to the robot. Various communication mechanisms have been proposed for this human–robot interaction. Firstly, one can speak to the robot, giving such directions without knowing anything at all about the robot's current location, much like people give directions over the telephone (Perzanowski et al., 2001; Torrey et al., 2006). However, this approach seems to have many obstacles in the way of a commercial application, simply because of the potential lack of accuracy and miscommunication that they inevitably have, particularly when moved out of the laboratory and into the noisy and unpredictable world in which humans typically operate.

Other commercial studies have suggested that map-based interaction would be a cost-effective way to avoid the problems of speech-recognition interfaces, though it could sacrifice ease-of-use in instructing the robot. For instance, Huttenrauch and Norman's (2001) PocketCERO proposed a human–robot interface that presents a drop-down list box to select a task from the task list and specifies the appropriate objects from the object list, as shown in Fig. 2. However, the maps used in PocketCERO did not clearly present all the physical cues available during local coordination of the robot when operating it. Therefore, while it makes sure that human operators can easily perform some collaboration tasks with it, the coordination of the robot would still be heavily reliant on expensive sensing technologies that could make such a system less cost-effective. It is very probable that enriched map use would make an effective commercial case, in conjunction with the approach of PocketCERO. However, one of the challenges is how to provide relevant environmental cues for the human–robot collaboration.

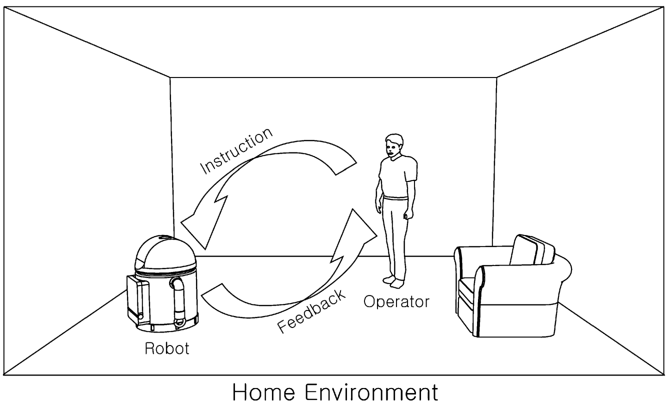

Fig. 1. A typical synchronous and colocated human–robot interaction in the home environment.

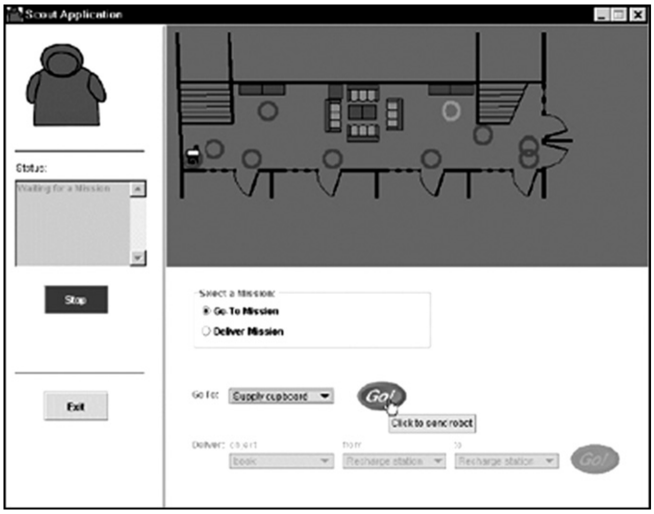

Fig. 2. A human–robot interface, reprinted from Hutternrauch and Norman (2001).

Several researchers (e.g., Fong, 2001; Fong et al., 2001; Fong and Nourbakhsh, 2005; Nourbakhsh et al., 2005; Yanco et al., 2004) have considered a *three dimensional (3D) image-based instruction* mode for providing relevant cues in this workspace model. It captures the surrounding images using a robot-mounted camera, sending the data back to the human operator (Fig. 3(a)). It allows the human operator to accurately recognise the local cues from a robot's forward-field-of-view (FFOV), and to guide the robot, avoiding potential obstacles. However, lack of global awareness of the whole environment is inevitable in the 3D instruction mode, so it is difficult for the human operator to build up and maintain effective situational awareness and/or teamwork plans (Yanco et al., 2004).

By contrast, *two-dimensional (2D) map-based instruction*, as shown in Fig. 3(b), cannot provide such realistic environmental cues, but it can remedy the problems by simply being able to present all the relevant environmental cues that are useful for the human operator to perceive the global context (Yanco et al., 2004). In practice, the 2D map-based instruction mode has been favoured in many human–robot applications, e.g., Jacoff et al. (2000, 2001), Perzanowski et al. (2001), and Skubic (2005).

Yet, a thorough investigation has not yet been made as to whether the map-based instruction mode would be an effective interaction style and if it could adequately support human–robot interaction tasks under the home context. This paper therefore explores this issue, by focusing on what characteristics should be taken into consideration in the map design. The following sections describe some underlying challenges of map design for human–robot interaction.

1.2. Map design for human–robot interfaces: coordination

There are certain map design aspects that provide relevant physical cues of the collaborative work and environment. One can identify many CSCW (Computer-Supported

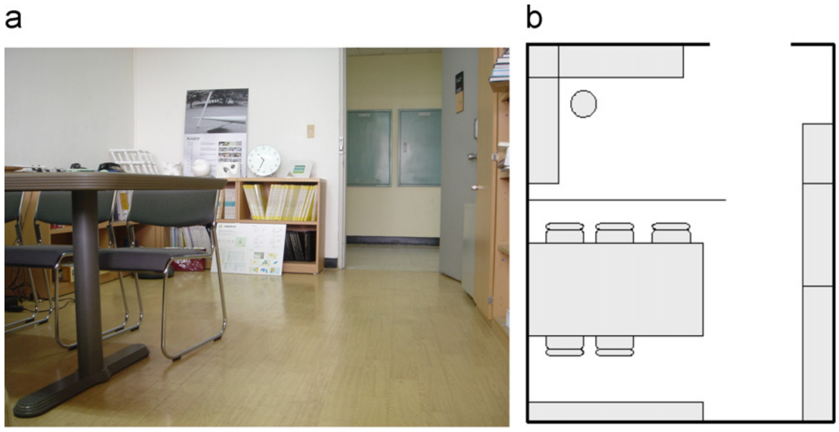

Fig. 3. Two instruction modes. (a) Three-dimensional image-based instruction; (b) two-dimensional map-based instruction.

Collaborative Work) studies that have already set out these aspects of map design for human–robot interfaces. For instance, *peripheral awareness* (Gutwin et al., 1996), which is being aware of all the collaborative participants' existence or their current location, provides information on who else is working together and where they are. *Workspace awareness*, which is about who is working on what, allows collaborative participants to have up-to-the-second knowledge of other participants' interactions with the elements in the environment (Gutwin et al., 1995). There is also the perception of the elements in the shared workspace and the comprehension of their meaning, which is *situation awareness* (Endsley, 1988). Hence, it can be thought that a map for coordinating a robot should be possible to present the workspace, all the elements that both the human operator and the robot are working on and what they must avoid, and how their collaborative tasks can be accomplished (Drury et al., 2003).

With respect to workspace and situation awareness, several HRI applications, especially tele-operating robots (Fong et al., 2001), prefer 3D maps to 2D maps, thanks to the realism and the self-sufficiency that they can offer. Others employ simple 2D map for its accurate perception.

As a compromise, one can also consider a small elevation above the surface of the 2D map (Fig. 4(b)). The elevated two-dimensional ($2\text{-}_{1/2}$ D) map has been successfully exploited in Geographical Information Systems (GIS) design, allowing the user to perceive the depth or the relative volume of each element via the "surface space" (Chin and Dyer, 1986). Of course, any benefits of the elevated 2D map would be heavily reliant on the tasks that the human–robot collaborative activities are intended to perform. For instance, the elevated 2D map would be of little value with a simple locating task when the robot is near to the object, as opposed to when more demanding locating jobs are required of the robot, e.g., moving the robot to the door of the fridge. This issue should be empirically validated, however.

There is also a quite intuitive way, often overlooked in the map design, of extending awareness of the workspace and the objects by adding legends (or labels) to the objects.

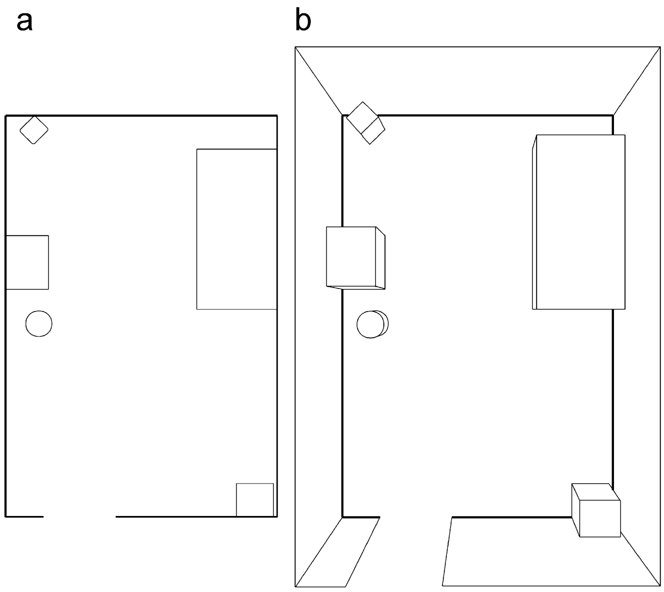

Fig. 4. Two-dimensional map (a) vs. elevated two-dimensional map (b).

Consider Fig. 4 again. Here, the comprehension of the objects would depend only on the skeletal drawings of them, so that the human operator should connect the sketches using his or her local perception. In contrast, the maps as shown in Fig. 5 do not require this extra cognitive process. In fact, many Virtual Reality (VR) studies have long adopted this map design convention, though this should also be empirically examined in HRI situations.

In conjunction with awareness of the workspace and the objects, the human–robot collaborative activities may also ask where the collaborative participants are now. Many CSCW studies, e.g., Baecker et al. (1993) and Gutwin et al. (1995), demonstrated that awareness of the collaborative participants would be highly beneficial to their collaborative task performance. However, current HRI studies (e.g., Borenstein et al., 1997; Perzanowski et al., 2001) have claimed that only the current position of the robot is sufficient. In this respect, it would be worth empirically comparing the two possible versions of the map

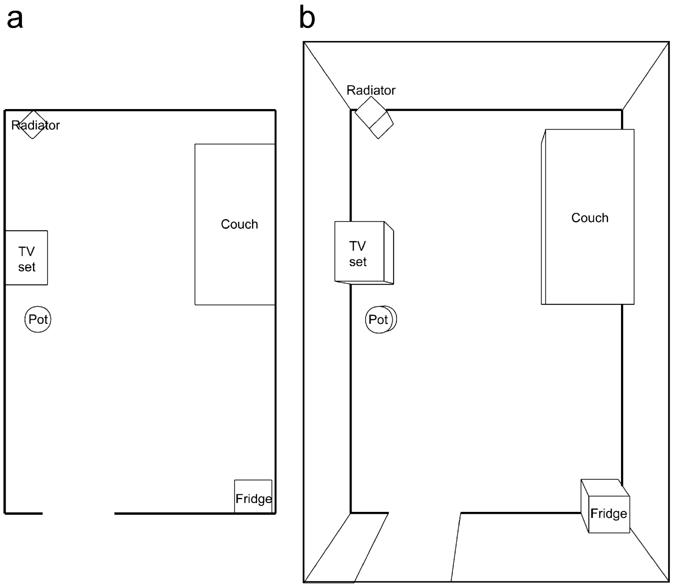

Fig. 5. Labelled maps. Two-dimensional model with labels (a); elevated two-dimensional model with labels (b).

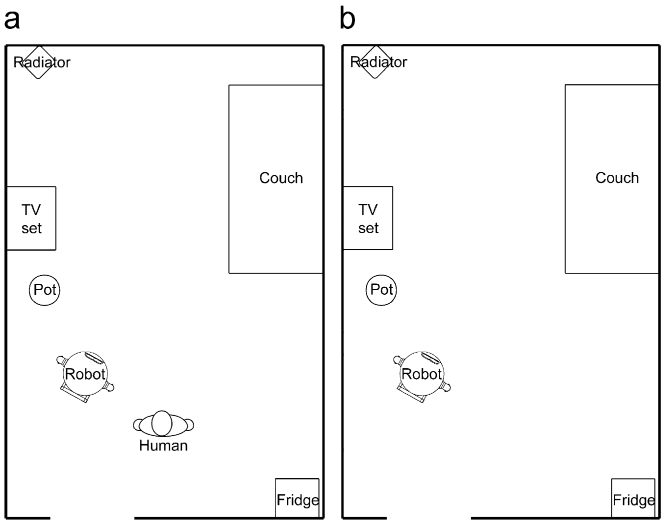

Fig. 6. Representation of collaborative participants. (a) Both the robot and the human operator are presented; (b) only the robot position.

representation, as shown in Fig. 6, for this peripheral awareness support.

In summary, this section discussed what issues would be relevant to the human operator's awareness for a human–robot interface. In particular, we considered best practices in the literature on spatial cognition, identifying that many (though not all) design issues in map-based human–robot interfaces might be answered by them. An empirical understanding of these issues is central to this paper.

1.3. Devices for human–robot interfaces: communication

Apart from the coordination issues previously discussed, there are several communication issues that should be considered in map design. Natural language communication used with a map has attracted many researchers (e.g., Perzanowski et al., 2001), in order to maximise the benefits of both the map and verbal intercommunication. However, current technology for natural language processing cannot avoid certain levels of ambiguity. For instance, if a human operator instructs a robot to search a certain 'area' of the room, natural language communication mechanisms would be an obvious challenge. Instead, direct manipulation with a map has been proposed as a more pragmatic approach. For example, drawing a route (Skubic, 2005) or tapping a position on the map (Fong et al., 2003) assumes that a handheld device with a stylus pen would be a cost-effective communication medium.

The usefulness of such a system would be subject to several drawbacks. Firstly, it does not show a cursor that provides the current position of the possible point selection, so there is no opportunity to figure out misjudgements of the point selection. Second, it cannot avoid optical distortion such as parallax error (Tian et al., 2002), which means that even when a user precisely taps a point that is believed to be correct, the point hit is generally several millimetres away from the one that they want to select. Finally, the physical specification of handheld devices is also limited, e.g., the size of the tip of the stylus pen and touch-sense resolution. The tip size may vary, but is generally approximately 0.5 mm wide. As a consequence, it cannot provide precise pointing performance beyond this. The tip size issue is also closely related to the touch-sense resolution of the screen. For instance, common handheld devices have one sensor for about every 5 mm, so that they can provide around 0.2 mm touch-sense resolution (Ramachandra, 2002). In Huttenrauch and Norman's (2001) interface design, for instance, the handheld device (screen size 57 (width) $\times$ 76 (height) mm with 240×320 pixels screen resolution) was used to display a large office area. It meant that roughly one pixel, i.e., around 0.24 mm of the screen, represents about 10 in in the office environment. Comparing this setting with the common tip size (more or less 0.5 mm) and the touch-sense resolution (0.2 mm) of commercial handheld devices, their interface cannot avoid some distance-related errors if the human operator instructs the robot using screen taps with the stylus pen.

Nonetheless, this paper sees a handheld device with stylus-pen input as the most plausible commercial case for colocated HRI tasks, thanks to its accuracy relative to natural language-based intercommunication, and the portability that the handheld device offers. Further, a relatively small workspace, where this paper is aiming for, would also justify using a handheld device for human–robot interaction. An empirical testing of this type of human–robot interface is central to this paper.

2. Experimental task

The experimental task was undertaken in a real room (3.9 m $\times$ 6.0 m) that was determined by a common living room space in Korea. To coordinate a robot in a particular space, in all the experiments in this paper, participants used a map on the screen of a handheld device, and tapped it with a stylus pen to direct a robot. The map was presented on a Compaq™ T1000, of which the real map size was 39×60 mm with 195×300 pixels resolution. Therefore, one pixel on the handheld device was equivalent to $4\,cm^2$ in the actual size of the room. The tip size of the stylus pen was around 0.3 mm. To partially lessen the optical distortion that comes from the angle of the stylus pen and the human operator's angle of vision, the participants were asked to use a consistent posture when using the stylus pen, and always place the handheld device in the same position.

As the experimenter indicated a location on the floor with a laser pointer, the participants were asked to tap the corresponding point of the destination on the map interface. The robot in all the experiments did not actually move to the places indicated by the participant's instruction, so there was little feedback to the user as to what he or she had done. This is not, of course, the full human–robot interaction envisioned in the near future. Furthermore, it seems to be both unnatural in interpersonal communication and unlikely as a successful interface for human–robot communication. Nevertheless, this experimental setting (i.e., although there is a robot in the room, the robot does not move and could just as well be a piece of furniture) has two advantages. Firstly, arguably, it is more natural in that this setting removes all the distractions that may be triggered by the robot's movement while the human operator perceives the workspace. In fact, the actual robot movement *per se* is barely of value in perceiving a destination. Rather, the current location of the robot as an object in the shared workspace would be more useful in perceiving the destination. Second, the primary concern of these experiments is to see how the human operator would be aware of the shared workspace via the map interface, so the best map would provide better pointing performance for a particular location where they were asked to direct the robot to that location. Experiment 1 for example, assessed whether the involvement of the human operator's position on the map would enhance their awareness of the destination, and we thought that this could be achieved without any movement from the robot. That is, this limitation would not cloud the interpretations of the experimental results.

There is another concern about this experimental task. In all the experiments, the participants were asked to tap on the screen where the robot was to be directed to. This seems to be closely related to each participant's ability to do spatial reasoning. To reduce the effect of any individual differences, all the participants were asked to perform two or three practices before they carried out their main experiment. Also, their pointing performances in these learning trials were instantly reported so that they had the opportunity to modify their task performances. This practice was intended to ensure as much as possible that all the participants had similar spatial reasoning ability in the main experiments.

3. Experiment 1: Peripheral awareness in a limited space

Collaborative activities among people demand appropriate peripheral awareness with regard to who is engaging in the collaboration. In a similar fashion, human–robot interaction in a limited workspace may require information about where the collaborative participants are. However, most of the current human–robot interfaces (e.g., Perzanowski et al., 2001) only present the current position of the robot. Experiment 1 was designed to empirically investigate this design convention.

To emphasise the effects of awareness of the two collaborative participants, all possible distractions, such as surrounding objects in the room, were removed in this experiment. There were therefore no landmarks shown on the map. This may appear to make the task difficult, in that participants basically had to use only the robot, the operator, and/or both locations to mark a point on the map relative to a point shown in the room. However, this experimental setting is very likely to reveal whether peripheral awareness of both the robot and the human operator in the "limited space" makes a significant difference.

The way of representing the two collaborative participants was the main manipulation (independent variable) in this experiment. The measure taken (dependent variable) follows from the practical implications of users having difficulty with the locating task, that is, they will lead to a larger error distance between the to-be-located places and the actual points on the map.

The three experimental groups were formed by the three types of representation of the two entities as shown in Fig. 7. The first map provided both the robot and the human operator. The others represented either the robot or the operator on the map, respectively. If the absence of any entities in the map representation had an effect on task performance, it would lead to an unequal task performance against the map with the locations of both the human operator and the robot.

3.1. Method

3.1.1. Participants

30 participants (12 females, 18 males) were all undergraduate students (aged 18–26 yrs.) at the University where the authors are working. Upon completing Experiment 1, the participants were paid two dollars for their participation.

3.1.2. Design

The experimental design was a two-way (map type-points) mixed design. The different map presentations, as shown in Fig. 7, were the between-subject independent

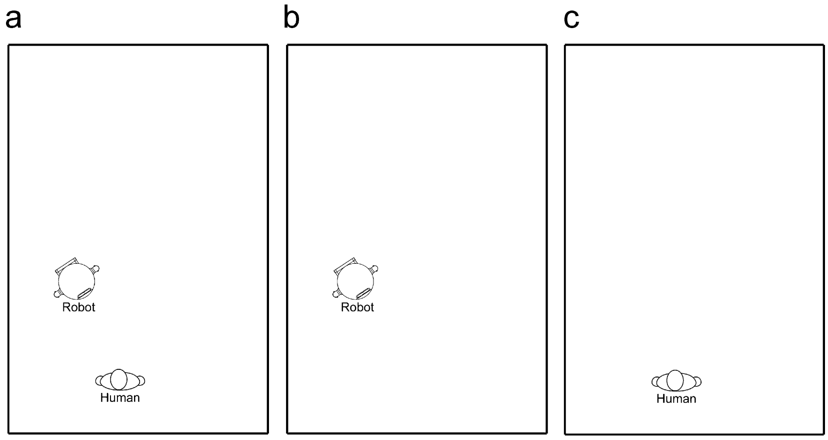

Fig. 7. The three map representations in Experiment 1. (a) Both the human operator and the robot were on the map, (b) only the robot and (c) only the human operator.

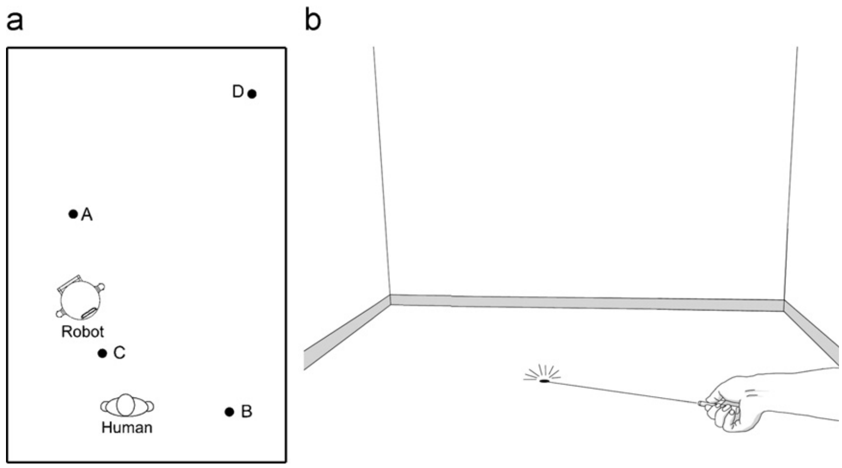

Fig. 8. (a) The four destinations. (b) The experimenter indicated the destinations on the floor of the room.

variable. The four destinations, as shown in Fig. 8(a), served as the other within-subject independent variable. The sequence of the destinations shown to the participants was counterbalanced using a Latin square. The dependent variable, the Euclidian error distance between the to-be-located points on the floor and the points that the participants tapped on the handheld device, was used to assess the effect of the independent variables.

3.1.3. Apparatus

The workspace contained no surrounding objects (i.e., no landmarks), except a desk and a desk chair on which participants sat and performed their pointing tasks. The current position of the human operator on the map was thus specified by the position of the desk.

The four destinations were determined to repeat the effects of awareness of the two collaborative participants (i.e., the human operator and the robot). They were marked with a red sticker on the floor. All the destinations, except Point D, were within a 1 m range of either the current position of the human operator or the current position of the robot. This 1 m range was empirically chosen by the authors, virtually ensuring that destinations within that range were more easily targeted. Therefore, Point A was considered to be relatively close to the robot, but relatively distant from the human operator. By contrast, Point B was located close to the human operator, but far from the robot's current position. Both the human operator and the robot are close to Point C; and Point D was the farthest one from both the robot and the human operator.

3.1.4. Procedure

The participants were first provided with the instructions regarding the experiment. These gave information about the experiment, the purpose of the study, and the data protection policy. They were then randomly assigned into one of the three different map representations as depicted in Fig. 7.

Before the main experiment, the participants were allowed to become familiar with this task and the apparatus, performing two or three trials. As the experimenter indicated a point placed on the floor of the room (see Fig. 8(b)), they were first asked to look at the destination on the floor, and then tap the point on the handheld device with the stylus pen. The destinations in the practice session were not the same ones used in the main experiment. The procedure followed in the main experiment was the same as that in the practice session. The participants were asked to tap the four points on the map, of which the sequence was counterbalanced using a Latin square.

3.2. Results

Table 1 gives the mean pointing accuracy between the to-be-located places (i.e., Points A, B, C, and D) and the points tapped on the map for each point. Looking at Table 1, we firstly noted that the overall task performance (mean error distance = 15.35 cm) was not so poor, when compared with the actual size of the room (3.9×6 m) which might indicate the potential for this type of interface in a commercial context. Second, comparing the mean error distances for each point, it appeared to be dramatically reduced when the points were close to either the robot (Point A = 13.99), the human operator (Point B = 14.49) or both (Point C = 11.60), which should be true for pointing tasks. Finally, the entities presented on the map did not explicitly enhance the mean task performance (14.65 in Map A, 16.21 in Map B, and 15.18 in Map C).

A two-way (maps × points) mixed analysis of variance was conducted on the task performance, revealing that there was no significant main effect of the map representation as to the existence of the collaborative participants ($F_{2,27} = 0.97$, n.s.); but there was a significant main effect of the destinations on the task performance ($F_{3,81} = 23.88$, $p < 0.01$). Tukey tests (at $p \leqslant 0.05$) were performed to further examine the effect of the destinations. The error distance was significantly greater at Point D (mean 21.31) and smaller at Point C (mean 11.60) rather than both Point A (mean 13.99) and B (mean 14.49), which were not significantly different from each other. However, there was no further higher level interaction effect between the map representation and the destination ($F_{6,81} = 0.92$, n.s.).

3.3. Summary and discussion

The main research question concerned in Experiment 1 was whether all the collaborative participants should be explicitly presented on the map. Our original hypothesis was that it would increase the task performance, as it is generally expected in human-to-human collaboration. However, the results of Experiment 1 showed that although Map A, which presents all the collaborative entities, had the smallest error distance, the statistical analysis exhibited that this was not the case. A possible explanation for this result could be that the workspace used in this experiment was not very large, so all the destinations were able to be instantly determined by either the current robot location (i.e., Map B) or the human position on the map (i.e., Map C). That means that the location information of one of the entities (rather than both) is sufficient for coordinating the robot in such a "limited space". This finding seems to be worthy of attention, given the potential of either Map B or C to not compromise locating task performance while reducing the resource requirement of representing both entities as considered in Map A. Indeed, it raises a logical question as to which entity (either the human operator or the robot) should be present on the map. Even though this should be answered with more thorough tests, the location of the robot should probably be present in the sense that the object of the coordination is not the human operator but the robot. Furthermore, it appeared that the location of the human operator in the limited space would be locally perceived based on their relative distance from the robot or self-referenced. Many human–robot interfaces (e.g., Perzanowski et al., 2001) have adopted this approach (i.e., Map B), but Experiment 1 empirically supported this design decision.

Apart from this finding, the results of Experiment 1 provided another practical contribution to the map-based human–robot interface for a limited space, which has not been empirically shown in the previous literature. The error distance with a handheld device does not seem to be large when compared with the actual size of the room. In fact, at most, the largest error distance (mean 21.31 cm for Point D) was only equivalent to 2.13 mm on the map, which suggests this is a realistic way of coordinating a robot in a limited space. Of course, although this benefit seems to be subject to both the handheld device with a 195×300 pixel

Table 1
Task performance

	Mean error distance (s.d.)				
	A	B	C	D	**Total**
Robot, Human (Map A)	12.44 (3.48)	13.36 (3.95)	10.40 (2.04)	22.39 (8.56)	**14.65**
Robot (Map B)	14.29 (6.09)	16.70 (4.24)	12.13 (3.71)	21.71 (5.64)	**16.21**
Human (Map C)	15.23 (3.97)	13.41 (4.26)	12.27 (3.32)	19.81 (4.85)	**15.18**
Total	**13.99**	**14.49**	**11.60**	**21.31**	**15.35**

Unit: Centimetres.

display and the workspace (3.9 m × 6.0 m) used in this experiment, undoubtedly this practical advantage can be equally applied to current handheld devices that generally employ at least 240 × 320 pixel screens.

It should be noted, however, that this experimental setting intentionally overlooked the common home context that includes many household goods. While this artificial setting helped us to explore the issue discussed above, it could not entirely address the actual design challenges of map-based interface. This issue will be further investigated in Experiment 2.

4. Experiment 2: Workspace awareness in a limited space

Experiment 1 demonstrated the commercial applicability of a human–robot interface using screen taps. However, it has some limitations for direct application to the design of a human–robot interface in the general home context. Indeed, the objects or elements in the shared workspace play an important role in determining destinations. Experiment 2, therefore, intends to present these objects on the map, and explores whether they can support the locating task performance as the human operator instructs the robot in a shared workspace.

Here, as a way of presenting the elements or objects in the environment, two modelling features were considered: *dimensionality*, and *legends*. The dimensionality issue of map representations has been widely dealt with in spatial cognition studies, as a means of extending workspace awareness. For instance, Rate and Wickens (1993) showed 2D maps would provide better awareness of lateral and vertical positioning over their 3D counterpart, providing a more accurate response. By contrast, if being asked to report their current position on the map, most users responded faster with the 3D map that provided an additional depth dimension. These early studies strongly implied that a 2D map in human–robot interfaces would be of value for accurate lateral and vertical positioning of the robot, though many robotics applications still prefer 3D modelling to the 2D model.

As a compromise between these two map design conventions, the environment can also be modelled with a small elevation above the surface. Consider Fig. 9(b). The small-scale elevation of each object was determined by the relative size of the robot's stature (70 cm) in the map representation; therefore, it can convey information regarding the volume of each object relative to the robot. It is generally believed that a more accurate locating task seems to justify the use of the $2_{1/2}$D model. However, the elevated 2D map requires more 'graphical space' for drawing the objects. As a consequence, the resolution of the space of the map where the user should point may decrease, thus probably causing a reduction in accuracy. To address this issue, both maps have the same floor size (shown by the bold lines in Fig. 9). Even though the elevated representation of the objects would occupy more space on the map, the floor spaces which the users would tap are equal.

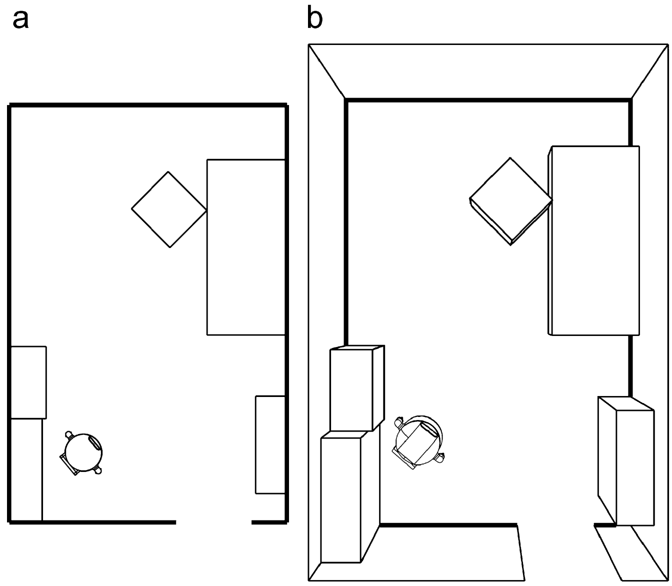

Fig. 9. Dimensionality of the map. (a) Two-dimensional representation; (b) the elevated two-dimensional ($2_{1/2}$) representation.

The comprehension of elements in the environment can also improve workspace awareness (Gutwin et al., 1995). Of course, the objects in the colocated workspace would be locally perceived by the human operator. However, the explicit description of objects would be useful to instantly specify appropriate references (Poole and Wickens, 1998). In practice, many VR studies adopted the explicit description of the objects on their maps, in order to extend awareness of the landmarks.

Two types of legend could be attached to the objects—*text label* and *picture image*, as shown in Fig. 10. Fig. 10(c) shows the 2D representation with the photo image of each object, and Fig. 10(a) with the text labels. Both Fig. 10(b) and (d) were depicted in the elevated 2D format, but different legends were used. In particular, the photo images in Fig. 10(d) were added on the face which the human operator was supposed to see. In addition, the two maps from Fig. 9, which did not have any legends for the objects, were considered as a control condition to highlight the effects of the legends used in the other maps.

A within-subjects experimental design was developed to reduce the number of participants over the six experimental conditions, in which every participant served under all combinations of both variables. The six experimental treatments were formed by the three types of codification (none, text label, and picture) and the two types of dimensionality (2D and elevated 2D) of both the environment and the objects.

4.1. Method

4.1.1. Participants

Sixty participants were recruited. Half of them were females (mean 25.26 yrs) and the others males (mean

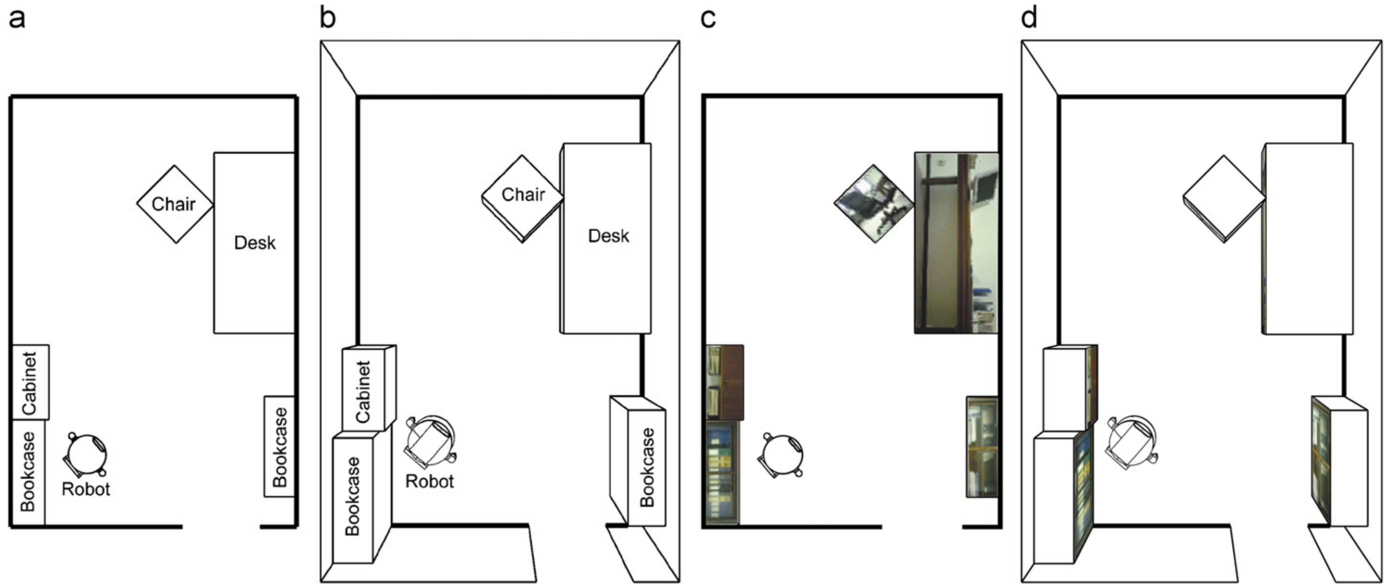

Fig. 10. The map representation considered in Experiment 2. (a) 2D representation with the text labels of the objects; (b) elevated 2D representation with text labels; (c) 2D representation with photo images; (d) elevated 2D representation with photo images.

28.10 yrs). Upon completing Experiment 2, the subjects received a ten-dollar voucher for their participation.

4.1.2. Design

Two independent variables—dimensionality (2D vs. elevated 2D) and legend types (none, text, and picture) form the six map types as shown in Figs. 9 and 10. Also, the 12 destinations as shown in Fig. 11 served as another independent variable. They were simply categorised into three types, i.e., A, B and C. A destinations (A1, A2, A3, and A4) were all within a 1 m range of the closest wall. All B destinations were located within a 1 m range from the closest object. By contrast, C destinations were more than 1 m from both the walls and the objects. Therefore, the experimental design was a 2 (dimensionality)2×3 (legend) $\times$ 3 (destination types) within-subjects design. The performing sequences of both the six map types and the 12 destinations were counterbalanced using a Latin square. The dependent variable was the Euclidian error distance between the to-be-located point that was and the actual point that the participants tapped on the map.

4.1.3. Apparatus and procedure

The same apparatus as Experiment 1 was used, except that the six different maps of the environment were used (Figs. 9 and 10). The elevated 2D maps were designed by RHINOCEROS^{TM} v.2.0 and the small-scale elevation was chosen to best support the realism of the environment, relative to the actual robot stature, i.e., 70 cm. In addition, only the current position of the robot was displayed on the map, given the results from Experiment 1. The procedure in Experiment 2 was almost the same as Experiment 1, except that participants were asked to tap the 12 destinations under each of the six different map representations.

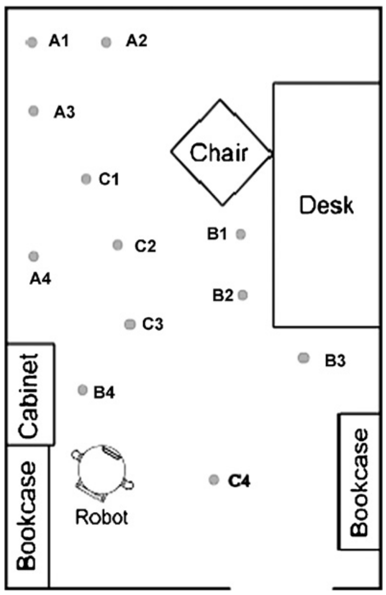

Fig. 11. The 12 destinations used in Experiment 2.

4.2. Results

Table 2 gives the mean error distance over each destination under the six experimental conditions. The overall task performance (mean = 20.64) roughly replicated the result from Experiment 1. A three-way (dimensionality, legends, and locations) within-subjects analysis of variance was carried out on the error distance. It was significantly decreased in the 2D maps over elevated 2D ($F_{1,59} = 21.20$, $p < 0.01$). However, it showed that our participants were not particularly sensitive to the different labels of the objects on the map ($F_{2,118} = 0.18$, n.s.).

Table 2
Task performance

	Mean error distance (s.d.)												
	Text			Picture			None			Sub-total			**Total**
	A	B	C	A	B	C	A	B	C	A	B	C	
2D	15.72 (10.27)	13.26 (4.32)	31.47 (17.99)	17.22 (9.53)	15.81 (6.41)	27.06 (16.52)	17.45 (8.39)	14.19 (7.50)	28.99 (15.16)	16.80 (9.40)	14.42 (6.08)	29.17 (16.56)	**20.13 (10.66)**
Elevated 2D	16.28 (10.12)	15.46 (8.32)	31.80 (22.57)	20.71 (12.23)	10.51 (9.12)	32.77 (13.36)	20.69 (10.61)	18.81 (8.34)	23.23 (15.37)	19.23 (10.99)	14.93 (8.59)	29.27 (17.10)	**21.14 (12.25)**
Sub-total	16.00 (10.20)	14.36 (6.32)	31.64 (20.28)	18.87 (10.88)	13.16 (7.77)	30.19 (14.94)	19.07 (9.50)	16.50 (7.92)	26.11 (15.27)	17.98 (10.19)	14.67 (7.34)	29.22 (16.83)	
Total	**20.66 (12.34)**			**20.68 (11.21)**			**20.56 (10.90)**						**20.64 (11.49)**

Unit: Centimetres.

Furthermore, the locations strongly influenced the performances ($F_{2,118} = 32.19$, $p<0.1$), which followed a Tukey test (at $p \leqslant 0.05$ level), revealing that C locations had higher error distances than both A and B locations, which were not significantly different from each other. There was no further higher level interaction effect among the dimensionality, the types of legends and the locations.

4.3. Summaries and discussion

The main concerns in Experiment 2 were, firstly, to replicate Experiment 1 in the general home environment which includes several household items; and secondly, to explore what map representation would dictate the task performance, which would imply important design guidelines for a successful commercial application of the map-based human–robot interface in the home context.

Firstly, we reconfirmed that the overall task performance (mean error distance = 20.64 cm) with the handheld device was adequate, compared with the actual room size of 3.9×6.0 m. Of course, being acceptable as a commercial case for the human–robot interface demands more rigorous validity tests in the other contexts, such as more than one room space that is a common home environment. However, the results of Experiment 2 simply indicate that this map-based human–robot interface with a handheld device may be a practical way of coordinating a robot in the home environment, the potential of which has not been demonstrated empirically before.

Secondly, regarding the map representation, the appropriate landmark for each destination, e.g., the chair, the desk, the cabinet, and even the surrounding walls, would be very likely to help our participants to precisely locate the destinations (say A destinations and B destinations), which are equivalent to the findings from the earlier literature (e.g., Wickens, 1999; Yates, 1990). In particular, the lower performance in locating C destinations can be explained by the fact that they were placed relatively distant from the objects. However, it is difficult to say clearly what spatial relationships between landmarks and destinations should be considered in map design, a question that will be further investigated in Experiment 3. In addition, the task performance was better with the 2D map than with the elevated 2D map, probably because of the exact lateral and vertical location awareness from the given 2D map, which are in line with the previous navigation studies (e.g., Barfield and Rosenberg, 1995; Rate and Wickens, 1993; Yeh and Silverstein, 1992). Yet, as opposed to our hypothesis, the legends themselves had no effect on the task performance. Indeed, what difference the legends could make in this experiment was not evident for our participants, as they rarely mapped from actual objects in the room to objects on the map. In fact, the weak influence of legends on the task performance has already been identified in some HCI-related studies (*Room effect*; Colle and Reid, 1998, 2003) when the space is relatively small and the user can directly view the objects, irrespective of the representation of the workspace and the objects.

In effect, Experiment 2 legitimises a design convention of the map-based interface for the colocated environment, i.e., the planar 2D map with no labels. It is an empirical contribution of this paper, given that the following experiment further investigates the spatial relationship issue between the destination and the landmark, which was raised in this experiment.

5. Experiment 3: Relation between landmark and destination

Both Experiments 1 and 2 contributed to establishing some design guidelines for map-based interfaces for a colocated workspace, e.g., the 2D map without labels would simply best serve the task performance in this context. Also, we identified that the landmarks, e.g., the objects, the surrounding walls, and the robot itself, would play an important role in enhancing the locating task performance. To explore the effectiveness of the map-based interface proposed by the previous experiments, and to further examine the spatial relationship between landmarks and destinations, a more intensive empirical study should be carried out.

We developed a different experimental setting for this part of the study. Consider Fig. 12(a) first. Based on the

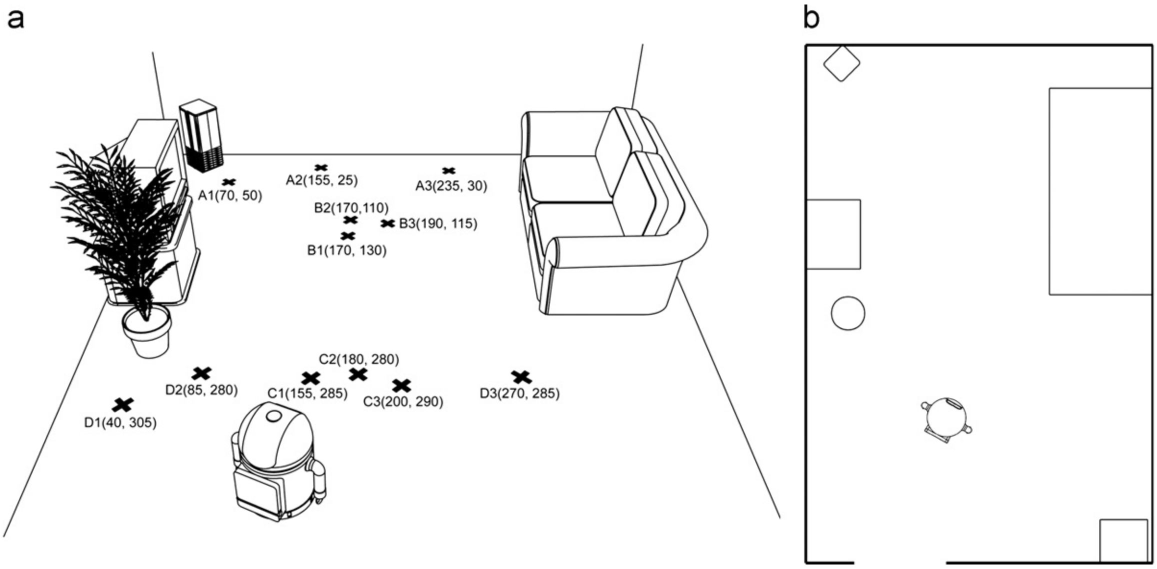

Fig. 12. The environment for Experiment 3 (a); the map used in the experiment (b).

results from both Experiments 1 and 2, it is most likely that Destination A1 is quickly and precisely targeted, because it has a very obvious landmark—the radiator. Similarly, when the human operators intend to move the robot to Destination A3, the couch may be frequently employed as a reference to the destination. The wall near Destination A3 may also be a possible landmark in this case. In contrast, only the front wall can be used as a unique landmark for moving the robot to Destination A2. Based on the results from Experiment 2, the surrounding walls of the limited workspace could be used as relevant landmarks for the destinations, but their usefulness might be less explicit against that of the closest objects. Destination A2 was considered for this issue.

In contrast, destinations B1–B3 would have to select from multiple possible landmarks, i.e., the radiator, the couch, the TV set, and/or even the front wall, in that these points are between these landmarks but their distance from the landmarks is relatively greater than that of the A destinations. Comparing the task performance in the two sets of points (i.e., A's and B's), one can understand what spatial relations, i.e., the relative distance from the landmarks and/or the number of landmarks, would have effects on the locating task.

The current robot's position can itself be the reference point for some destinations, as demonstrated in Experiment 1, such as C destinations (C1–C3) and D destinations (D1–D3). C destinations are located very close to the robot, but relatively distant from the other objects, such as the pot and the couch. Therefore, it is very likely that C destinations would have only one salient landmark, i.e., the robot. By contrast, both the objects and the robot can be used as the possible landmarks to D destinations. By using these two sets of points, one can identify if the robot itself might be a landmark for such destinations in the home context and what reference, i.e., the robot or the objects,

Table 3
Specification of the 12 points in terms of the relative distance from the robot and the closest object

Destination		Closest object	Distance from the closest object	Distance from the robot
A	A1	Radiator	0.31 m (close)	3.12 m (distant)
	A2	Wall	0.25 m (close)	2.92 m (distant)
	A3	Wall	0.30 m (close)	3.25 m (distant)
B	B1	TV set	1.00 m (distant)	1.95 m (distant)
	B2	TV set	1.00 m (distant)	2.25 m (distant)
	B3	Couch	1.00 m (distant)	2.35 m (distant)
C	C1	Pot	1.10 m (distant)	0.65 m (close)
	C2	Couch	1.20 m (distant)	0.72 m (close)
	C3	Couch	1.08 m (distant)	0.80 m (close)
D	D1	Wall	0.40 m (close)	0.74 m (close)
	D2	Pot	0.72 m (close)	0.75 m (close)
	D3	Couch	0.58 m (close)	0.99 m (close)

Note: See texts for details.

would be better used as the appropriate landmark. The spatial relationships in this experimental setting can be categorised as in Table 3. Fig. 12(b) shows the map representation used in this experiment, following the results from Experiment 2.

A note of the 1 m criterion is need here. As the 12 points were initially located in the same room environment used in both Experiments 1 and 2, B destinations were lined up on the y-axis with the same x-coordinate (180), in order to make 1.1 m range to halve the spatial relationship between the destination and the landmark. This let B destinations have the same distance to both the TV set and the couch, given the same contingency to use either the TV set or the couch. However, a pilot test with two participants demonstrated that these three points (B1–B3) were too difficult to be distinguished when the experimenter indicated one of them by the laser pointer. As a slight modification of this initial experimental setting, Destination B3 was moved to the right as shown in Fig. 12(a). As a result, the original criterion—1.1 m—cannot be guaranteed in the new setting, alternatively the 1 m criterion can only encompass all the 12 points in an exclusive way, as both Experiments 1 and 2 did.

In effect, an underlying difference in Experiment 3 against both Experiments 1 and 2 was the twelve different destinations, which were designed to address the spatial relationship in performing the HRI tasks. To simplify the analysis of the spatial relation between the destination and the possible landmark, all the destinations were characterised along with the 1 m criterion. For instance, Destination A1 was considered to be relatively close to the radiator, and relatively distant from the current robot position. Therefore, the operator would be very likely to use the radiator as the landmark for the destination, if a better locating performance is seen. Similarly, Destination A3 can be seen as close to the wall, but distant from the robot. By contrast, the human operator would use the TV set to locate the robot into Destination B1 and B2, but the couch for Destination B3. Furthermore, they can be equally referred by either the TV set or the couch, which inevitably requires the human operator to select the appropriate landmark (Warren, 1994; Wickens, 1999), so it may take more time to point these destinations. On the other hand, C destinations (C1–C3) do not have any objects close at hand, instead, they have the current robot position as a possible landmark, compared with D destinations that have both the closest objects and the current robot position as possible landmarks.

5.1. Method

5.1.1. Participants

Twenty participants who took part in Experiment 2 were reinvited. It was intended to form a more homogeneous participant group and reduce the experimental efforts without further training. In particular, this recruitment ensured that the time stamped log data could be collected. Upon completing Experiment 3, they were also given a 10-dollar voucher for their participation.

5.1.2. Design

The experimental design was a 2×2 within-subject design. Distance from the robot (close and distant), and distance from the closest object (close and distant) served as the independent variables. Twelve destinations were predefined as shown in Fig. 12(a). The sequence of the 12 trials was counterbalanced using a Latin square. The dependent variables, the Euclidian error distance and time taken to tap on the map, were used to assess the effects of the independent variables.

5.1.3. Apparatus and procedure

The same apparatus from the previous experiments was also used here, except for the 12 different destinations and the map. The same procedure as Experiment 2 was used, except that the participants were asked to tap the points on the map as quickly as possible when the experimenter indicated a location on the floor of the room.

Table 4
Task performance

Destination		Closet objects	Mean Euclidian error distance (s.d.) (unit: cm)	Mean completion time (s.d.) (unit: sec)	Distance from the closest object	Distance from the robot
A	A1	Radiator	9.30 (3.64)	9.85 (9.48)	Close	Distant
	A2	Wall	16.14 (6.19)	11.26 (12.41)		
	A3	Wall	13.62 (5.97)	9.46 (9.31)		
	Total		**13.10 (3.84)**	**10.19 (10.40)**		
B	B1	TV set	27.39 (9.65)	10.11 (6.71)	Distant	Distant
	B2	TV set	19.89 (6.88)	11.49 (8.95)		
	B3	Couch	23.69 (8.42)	10.11 (7.02)		
	Total		**23.66 (8.31)**	**10.57 (7.56)**		
C	C1	Pot	30.74 (25.27)	5.32 (4.12)	Distant	Close
	C2	Couch	29.44 (24.36)	9.37 (6.27)		
	C3	Couch	30.80 (25.07)	6.13 (4.31)		
	Total		**30.33 (24.90)**	**6.94 (4.90)**		
D	D1	Wall	18.05 (14.17)	7.17 (4.43)	Close	Close
	D2	Pot	15.01 (14.67)	5.73 (2.58)		
	D3	Couch	13.90 (6.35)	5.91 (3.01)		
	Total		**15.66 (11.73)**	**6.27 (3.34)**		

5.2. Results

The mean error distances and mean task completion time at each point are shown in Table 4. Two additional columns were added to help the reader to understand the spatial relation scheme that was used in this experiment. Comparing the figures of the mean error distance at every point, it can be seen that there was a consistent effect of the distance from the closest objects. In A destinations (mean 13.10) and D destinations (mean 15.66), the mean error distances were almost half of the corresponding counterparts, i.e., B's (mean 23.66) and C's (mean 30.33), respectively. The distance from the robot seemed to be opposite as our participants would slightly outperform in the situations in which the destinations were distantly located from the robot (mean 13.10 in Destination A's, and mean 23.66 in Destination B's), rather than the corresponding counterparts, i.e., Destination D's (mean 15.66), and Destination C's (mean 30.33), respectively.

These observations were firstly analysed by a two-way within-subjects analysis of variance on the error distance. With regard to the spatial relation between the object and the destination, the error distance was significantly decreased when the destination was placed near the closest object ($F_{1,19} = 6.75$, $p < 0.01$). The analysis of each point under the same set of the points was followed, respectively. A Tukey test (at $p \leqslant 0.05$ level) showed that Destination A1 was significantly less error-prone than both Destinations A2 and A3, which were not significantly different from each other. However, the other destinations were not significantly different from one after another within each set of the destinations, i.e., Destination B's, C's, and D's. This observation will be further discussed in Section 5.3.

Interestingly, our participants made significantly less error distance when the destinations were distant from the robot ($F_{1,19} = 95.97$, $p < 0.01$). This result, that is that the error distance reduced when the destination was distant from the robot, seems to be against what Experiment 1 demonstrated, which revealed that the robot itself could be a landmark for the locating task. It can be explained in two ways. Firstly, the human operators tended to select the most obvious landmarks first. In Experiment 1, there were no other objects except the robot, so the human operator used the current robot's position as a possible landmark. However, this experiment had other obvious landmarks available, so the robot's position might not be preferred to the objects. This selection process of landmarks was identified in the early studies (e.g., Warren, 1994; Wickens, 1999). Second, it might result from the fact that our participants were more careful to point to the destinations (i.e., A's and B's) where they were away from the current robot position. In particular, this experimental setting ensured that the A and B destinations were also distant from the human operator. Therefore, the error distances of these destinations, i.e., A's and B's, could be consequently less. The mean completion time of this task supported this interpretation. That is, as the destinations were closer to the robot (or the operator), i.e., Destination C's (mean 6.94 s) and D's (mean 6.27 s), our participants seemed to quickly decide to tap the points on the map rather than the more distant locations, i.e., B's (mean 10.57 s) and A's (mean 10.19 s), respectively. It implied that our participants took more time to carefully tap the points as they were asked to direct the robot to Destinations A's and B's. A two-way within-subjects analysis of variance of task completion time revealed that the completion times were significantly affected by the distance of the destinations from the robot ($F_{1,19} = 12.80$, $p < 0.01$), not from the object ($F_{1,19} = 0.58$, n.s.).

5.3. Summaries and discussion

The conclusion to be drawn from this experiment was that the human operator's spatial interaction would be highly affected by whether the destinations could be referred to by salient landmarks (e.g., Point A's and D's). The importance of landmarks was reviewed in the early literature (e.g., Colle and Reid, 1998; Siegel, 1981; Thorndyke and Hayes-Roth, 1982; Warren, 1994; Wickens, 1992, 1999; Wickens et al., 1996) which emphasised that the locating task performance would be improved where obvious landmarks were provided. However, they only considered a wide open area, so there are limitations on applying their findings to the limited workspace considered in this experiment. This experiment empirically established the case for the colocated workspace.

Another finding from this experiment was the characteristics of the landmarks. In Experiment 2, we simply assumed that both the surrounding walls and the objects in the shared workspace would be appropriate landmark for the destinations. This was also partially supported in this experiment, but in a slightly different way. Consider the destinations which have the surrounding walls as landmarks. The three points (A2, A3, and D1) were placed close to the walls, so we assumed that our participants would have a similar task performance with the other destinations in each set of the destinations, respectively. The task performances for both destination A2 and A3 were not as good as destination A1. This can be interpreted in two ways. Firstly, it was probable that our participants were aiming to tap the points near to, but partly offset from, the surrounding walls (drawn in bold lines) on the map. They already perceived the size of the robot in the shared workspace, so they might have intended to avoid colliding with the walls as they were asked to direct the robot to the points near them. Therefore, the task performance in Destination A2 and A3 might be not good as the destinations that were near to the objects. Second, the surrounding walls themselves would not be such an effective landmark as a household item that occupies a visible space within the environment. Because an object has its own area in the space, it can provide a supplementary reference for the destination via the location of the object itself. Consider Destination A1. The destination would be firstly referenced by the radiator itself (the direct landmark of Destination A1), and then the destination could also be referred by the location of the radiator that would be perceived by the relative distance from the surrounding walls. Therefore, Destination A1 has both a primary landmark (i.e., the radiator) and a supplementary landmark (i.e., the surrounding walls). In contrast, both Destination A2 and A3 have only one landmark, i.e., the front wall.

6. General discussion and future work

Taken together the three experiments presented here demonstrated that the map-based interface on a handheld device could be of practical value in instructing a robot in a limited workspace. Neither of these possibilities has been demonstrated empirically before.

6.1. Using the results: guidelines of the map-based interface

The first conclusion to be drawn was that the explicit representation of the objects in the home context is critical, but the locations of both the human operator and the robot are less crucial as one can easily see where the collaborative human participants are in the colocated environment. It should be noted that there is no direct test of this issue; however, it seems reasonable that the sequential negotiation of the experimental settings from Experiments 1 to 3 would prove the interpretation. A possible guideline therefore is that a cost-effective human–robot interface would be either the human operator or the robot position present on the map, given that providing the locations of both the human operator and the robot is not crucial. Furthermore, it can be said that the location of the robot should be present, since the results from both Experiments 1 and 3 showed that the location of the human operator in a limited space would be easily self-referenced. This design convention has proved successful in many human–robot interfaces (e.g., Perzanowski et al., 2001). As to the representation of the other objects, Experiment 3 demonstrated that the task performance was better in those destinations where close objects acted as appropriate landmarks. This provides another design guideline for an effective map-based human–robot interface in situations where there are no landmarks available for the destination. As Experiment 1 identified, the location of the robot itself can be a possible landmark; but undoubtedly appropriate technical supports are necessary. In fact, magnifying the area of the destination where there is no landmarks available is being investigated by the authors.

The second conclusion directly follows from Experiment 2, that the 2D map without any legends best served the task performance. In effect, the *room effect* (Colle and Reid, 1998, 2003) should be considered in map design for the collocated HRI situation. That is, a compact and concise representation of both the objects and the workspace would prove useful. However, these conclusions are of course not to override the benefits of the other types of human–robot interface, e.g., image-based human–robot interfaces or speech-based interfaces. They only suggested that spatial cognition support and/or awareness support via the map representation should be the essential design challenge in the map-based human–robot interaction. Fig. 13 summarises the above conclusions as heuristics to be applied in a map-based human–robot interface for pointing tasks.

Now consider industrial contributions that can be drawn from these three experiments. The aim of this study was to prove the applicability of the map-based human–robot interface with a handheld device. From these three experiments, one can firstly see that a certain level of distance errors is inevitable as the human operator locates the home-service robots with the small handheld devices. In this respect, the three experiments established a practical "baseline" study of the locating task performance of the home-service robots with a small handheld device (195×300 pixel display), at most 30.80 cm (7.88% error distance) in the $3.9\,\text{m} \times 6.0\,\text{m}$ space. This empirical finding

A. *The location information of the human operator and robot (peripheral awareness)*

The current position of the human operator may not be able to enhance peripheral awareness, in cases where the workspace can be locally perceived by the human operator. However, the current position of the robot should be present at any time, because it is the entitiy to be located by the human operator, and the position information is very likely to be used as supplementary information for pointing tasks (see B-1 below)

B. *Representation of objects in the workspance (workspace awareness)*

The human operator will use the objects on the map as landmarks for the destination. Thus, the objects should be explicitly represented on the map for the operator to recognise them easily.

B-1. In cases where there is no landmark available, either the current position of the robot or the closest object is likely to be used.
B-2. In cases where there is no landmark available, appropriate supports should be considered.

C. *Representation of the workspace (Workspace awareness)*

A compact and concise representation of the workspace would best serve the task performance. The two dimensional map without any legends of the objects may be cost-effective for representing limited workspaces, such as home environments.

Fig. 13. Heuristics for designing map-based human–robot interfaces in a limited space.

can also contribute to the map-based human–robot interface design, assuming that such task performance can be accommodated in the commercial handheld devices, the screen resolution of which is at least 240×320 pixels.

6.2. Limitations of this study

There are many limitations of this study. Firstly, the interfaces tested in this article have not been externally validated with different users in different contexts. The participants used in this paper might not be regarded as the primary user group of this type of home-service robot, and the workspace considered in this experiment might not be the same as that from other cultural contexts. While these concerns should be addressed within this area of research, this paper focused first on map design issues, with the other issues planned as further work.

Secondly, although there was a robot in a room, the robot did not move and could just as well be a piece of furniture in this study. Thus, it may be unclear what the value of the design recommendations discussed above would be. However, this approach proved a very effective way of collecting large amounts of data for identifying the design factors of the map-based interface. It would require extra effort to implement such a more realistic experimental setting where the robot actually followed the instructions of the human operator. This study is thus worthy of attention in that other researchers can easily set up the same type of test for commercial product testing.

Finally, this paper only considered a one-off locating task, which is not sequential process, so it seems both unnatural and unlikely to present the values of the design recommendation discussed above. However, arguably, this one-off locating task is based on the assumption that this is the initial action that should be taken by the human operator, so that the understandings from this simple task would be equally applicable to more complex human–robot communication.

6.3. Further work

The results in this paper raised several questions that could be pursued in a future study. For example, the three experiments in this paper only used a single room environment. However, the home context generally has many rooms, and the human operator may not be in the same room with the robot. This may require different map design guidelines from what this paper established. As well as considering a more realistic home context, the situation in which both the human operator and the robot is on the move is worth mentioning. If the collaborating participants are mobile, the congruence issue between the human operator's local perception and the map representation is inevitable. All the experiments demonstrated here assumed that both the human operator and the robot were not on the move; thereby their FFOV is the same as that of the map representation. Indeed, this issue is being investigated under a new experimental setting, along with the robot on the move.

Acknowledgments

The authors owe much of these works to the anonymous reviewers who thoroughly commented the first draft of the paper. Also, the authors specially thank Dr. David Parsons for his thorough comments and helpful suggestions on the second draft of this chapter. Major funding for this work was provided by Samsung Advanced Institute of Technology, titled as 'Interaction design for Home-service robot using a map-based interface'.

References

Baecker, R.M., Nastos, D., Posner, I.R., Maywby, K.L., 1993. The user-centered iterative design of collaborative writing software. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems, Amsterdam, The Netherlands.

Barfield, W., Rosenberg, C., 1995. Judgments of azimuth and elevation as a function of monoscopic and binocular depth cues using a perspective display. Human Factors 37, 173–181.

Borenstein, J., Everett, H.R., Feng, L., Wehe, D., 1997. Mobile robot positioning—sensors and techniques. Journal of Robotics Systems 14 (4), 231–249.

Chin, R.T., Dyer, C.R., 1986. Model-based recognition in Robot vision. ACM Computing surveys 18 (1), 67–108.

Colle, H.A., Reid, G.B., 1998. The room effect: metric spatial knowledge of local and separated regions. Presence: Teleoperation and Virtual Environments 7, 116–128.

Colle, H.A., Reid, G.B., 2003. Spatial orientation in 3D desktop displays: using rooms for organizing information. Human Factors 45 (3), 424–439.

Drury, J.L., Scholtz, J., Yanco, H.A., 2003. Awareness in human–robot interactions. Paper presented at the IEEE Conference on Systems, Man, and Cybernetics, Washington, DC.

Endsley, M.R., 1988. Design and evaluation for situation awareness enhancement. Paper presented at the Human Factors Society 32nd Annual Meeting, Santa Monica, CA.

Fong, T., 2001. Collaborative control: a robot-centric model for vehicle teleoperation. Unpublished Ph.D., Carnegie Mellon University.

Fong, T., Nourbakhsh, I., 2005. Interaction challenges in human–robot space exploration. ACM Interactions March–April, 42–45.

Fong, T., Cabrol, N., Thrope, C., Baur, C., 2001. A personal user interface for collaborative human–robot exploration. Paper presented at the International Symposium on Artificial Intelligence Robotics, and Automation in space, Montreal, Canada.

Fong, T., Thrope, C., Glass, B., 2003. PdaDriver: a handheld system for remote driving. Paper presented at the IEEE International Conference on Advanced Robotics.

Forlizzi, J., 2005. Robotic products to assist the aging population. ACM Interactions March–April, 16–18.

Gutwin, C., Stark, G., Greenberg, S., 1995. Support for workspace awareness in educational groupware. Paper presented at the Computer Supported Collaborative Learning, Bloomington, IN.

Gutwin, C., Greenberg, S., Roseman, M., 1996. Workspace awareness support with radar views. Paper presented at the SIGCHI Conference on Human Factors in Computing Systems, Vancouver, British Columbia.

Huttenrauch, H., Norman, M., 2001. PocketCERO-mobile interfaces fro service robots. Paper presented at the Mobile HCI, Lille, France.

Jacoff, A., Messina, E., Evans, J., 2000. A standard test course for urban search and rescue robots. Paper presented at the Performance Metrics for Intelligent System Workshop, Gaithersburg, MD.

Jacoff, A., Messina, E., Evans, J., 2001. A reference test course for autonomous mobile robots. Paper presented at the SPIE-AeroSense Conference, Orlando, FL.

Kadous, M., Sheh, R., Sammut, C., 2006. Effective User Interface Design for Rescue Robotics. Paper presented at the Human–Robot Interaction, Salt Lake, UA.

Mynatt, E.D., Essa, I., Rogers, W.A., 2000. Increasing the opportunities for aging in place. Paper presented at the Universal Usability, Arlington, Virginia.

Nourbakhsh, I., Sycara, K., Koes, M., Yong, M., Lewis, M., Burion, S., 2005. Human–robot teaming for search and rescue. Pervasive Computing, 72–78.

Perzanowski, D., Schultz, A.C., Adams, W., Marsh, E., Bugajska, M., 2001. Building a multimodal human–robot interface. IEEE Intelligent Systems Jan./Feb., 16–21.

Poole, P.E., Wickens, C.D., 1998. Frames of reference for electronic map displays: their effect on local guidance and global situation awareness during low altitude rotorcraft operations (ARL-98-7/NASA-98-2): University of Illinois Institute of Aviation.

Ramachandra, P., 2002. Information at your fingertips [URL]. Retrieved 30. August, 2005, from the World Wide Web ⟨http://www.pcquest.com/content/search/showarticle1.asp?arid = 36169⟩.

Rate, C., Wickens, C.D., 1993. Map dimensionality and frame of reference for terminal area navigation display: where do we go from here? (ARL-93-5/NASA-93-1): University Illinois Institutite of Aviation.

Scholtz, J., 2003. Human–robot interactions: creating synergistic cyber forces. Paper presented at the International Conference on System Science, Hawaii.

Scholtz, J., 2005. Have robots, need interaction with humans!. ACM Interactions March–April, 13–14.

Siegel, A.W., 1981. The externalization of cognitive maps by children and adults: in search of ways to ask better questions. In: Liben, L.S., Patterson, A., Newcombe, N. (Eds.), Spatial representation and behavior across the life span: theory and application. Academic Press, New York, pp. 167–194.

Skubic, M., 2005. Qualitative spatial referencing for natural human–robot interfaces. ACM Interactions March–April, 27–30.

Thorndyke, P., Hayes-Roth, B., 1982. Differences in spatial knowledge obtained from maps and navigation. Cognitive Psychology 14, 560–589.

Tian, Z.Z., Kyte, M.D., Messer, C.J., 2002. Parallax error in video-image systems. Journal Of Transportation Engineering 128 (3), 218–223.

Torrey, C., Powers, A., Marge, M., Fussell, S., Kiesler, S., 2006. Effects of adaptive robot dialogue on information exchange and social relations. Paper presented at the HRI, Salt Lake City, Utah.

Warren, D.H., 1994. Self-localization on plan and oblique maps. Environment and Behavior 26, 71–98.

Wickens, C.D., 1992. Engineering Psychology and Human Performance. HarperCollins Publishers Inc., New York.

Wickens, C.D., 1999. Frames of reference for navigation. In: Gopher, D., Koriat, A. (Eds.), Attention and Performance XVII. MIT Press, Cambridge, MA, pp. 112–144.

Wickens, C.D., Liang, C.-C., Prevett, T., Olmos, O., 1996. Electronic maps for terminal area navigation: effects of frame of reference and dimensionality. International Journal of Aviation Psychology 6 (3), 241–271.

Yanco, H.A., Drury, J.L., 2004. Classifying human–robot interaction: an updated taxonomy. Paper presented at the IEEE Conference of Systems, Man and Cybernetics, The Hague, The Netherlands.

Yanco, H.A., Drury, J.L., Scholtz, J., 2004. Beyond usability evaluation: analysis of human–robot interaction at a major robotics competition. Human Computer Interaction 19, 117–149.

Yates, J.F., 1990. Judgment and Decision Making. Prentice Hall, Englewood Cliffs, NJ.

Yeh, Y.-Y., Silverstein, L.D., 1992. Spatial judgments with monoscopic and stereoscopic presentation of perspective displays. Human Factors 34, 583–600.