ORIGINAL ARTICLE

A location-sensitive visual interface on the palm: interacting with common objects in an augmented space

Seokhwan Kim · Shin Takahashi · Jiro Tanaka

Received: 21 April 2013/Accepted: 13 January 2014/Published online: 9 March 2014 © Springer-Verlag London 2014

Abstract We have created a visual interface using the human palm that is location sensitive and always available. To accomplish this, we constructed an augmented space in an actual workspace by installing several depth cameras. To manage and connect the multiple depth cameras, we constructed a distributed system based on scalable clientserver architecture. By merging depth images from different cameras, the distributed system can track the locations of users within their area of coverage. The system also has a convenient feature that allows users to collect the locations of objects while visualizing the objects via images from the depth cameras. Consequently, the locations of both users and objects are available to the system, thus providing a location-based context for determining which user is close to which object. As a result, the visual interface on the palm becomes location sensitive, which could lead to various applications in daily life. In this paper, we describe the implementation of the aforementioned system and demonstrate its potential applicability.

Keywords Location awareness · Interface on body · Augmented reality · Ubiquitous · System · Interaction

S. Kim (⊠) · S. Takahashi · J. Tanaka Department of Computer Science, University of Tsukuba, Tsukuba, Ibaraki 305-0006, Japan e-mail: skim@iplab.cs.tsukuba.ac.jp

S. Takahashi e-mail: shin@cs.tsukuba.ac.jp

J. Tanaka e-mail: jiro@cs.tsukuba.ac.jp

1 Introduction

The advancement of sensing and display technologies has made it possible to construct augmented spaces. Several systems have already been developed for augmenting all objects in a space and enabling users to interact with those objects. Such systems include Microsofts EasyLiving [1], the Everywhere Display Projector [18], and LightSpace [22]. This study shares the concept of augmented space in such studies: it is assumed that all objects in an augmented space are interactive, and users can potentially interact with all of them.

LightSpace, in particular, enables interactions between augmented objects and a visual interface on the body [22]. Users are able to store digital content on their bodies and drop the content onto other physical objects (cf. the pickand-drop metaphor of Rekimoto's early work [19]).

The present study employs a similar interactive concept (i.e., combining an interface on the body with augmented objects), but applies it to an actual workspace. Our space was constructed by installing depth cameras in a laboratory that is used by students. We installed several depth cameras to track users in the space. Accordingly, we implemented a distributed system that can seamlessly connect the depth cameras and merge their images, and thereby track the locations of users in the space. The installation of the depth cameras and the implementation of the distributed system are described in subsequent sections.

A key feature of our system is that it facilitates the process of registering or deleting the locations of objects. Because the system can track user locations, it provides a location-based context for determining which user is close to which object. A necessary preliminary step to enable this location awareness is the collection (i.e., registration or deletion) of object locations, which generally requires considerable time and effort. To facilitate this process, we developed a support tool that enables such registration or deletion while visualizing objects via the images from the depth cameras. Thus, the system is able to determine the distance between users and registered objects, which establishes the location-based context.

Another key feature is that it provides a visual interface on the human palm. We installed two projected displays in the space. These devices display visual content on the user's palm, and the user can interact with this interface via hand gestures or physical metaphors. Because the projected displays are installed in the space, the user is provided with a simple graphical user interface (GUI) on his/her palm, without temporal or spatial limitations (i.e., it is always available). Hence, we have created a visual interface on the palm that is location sensitive and always available, and could lead to various applications in daily life.

Figure 1 shows an example. Because the system records the locations of objects and users (i.e., user skeletons), it can continuously track the distance between a user body part (e.g., a hand or a head) and an object. Employing this location-based context (i.e., the distance), the system can recognize which user is close to which device. Therefore, the system can intelligently display an appropriate message on a user's palm.

For example, if a user needs to print a document, the system displays a notification of the printing task on the user's palm when the user is close to a printer (i.e., the physical distance is sufficiently small). The user can then complete the printing task by dropping the notification onto the printer. This application allows users to efficiently carry out a task at a physical location. The notification is displayed when the necessary device (the printer in this example) is close enough, and the user can complete the task without carrying any additional devices (i.e., the notification can be displayed on a bare hand).

Our system includes the following features, which are described in detail below: (1) A distributed system to manage the depth cameras and track the locations of users, (2) a support tool to help collect object locations, and (3) projected displays to create a visual interface on the palm.

2 Related work

The contributions of this study can be divided into three related domains: (1) constructing an augmented space, (2) registration of object locations, and (3) creating interfaces on the body. Below, we summarize each of these and clarify the motivations of this study.

2.1 Constructing an augmented space

The augmented space in this study allows interactions among all objects in the space. Here, we summarize some typical techniques for constructing an augmented space.

Computer vision is a popular approach for constructing an augmented space. Computer vision can track or identify both humans and objects. Microsoft's EasyLiving project [1] and the work of Lee et al. [12] are based on this approach. Studies from that era (i.e., about a decade ago) mainly relied on RGB cameras, which were adequate for envisioning futuristic environments in the early stages. However, because RGB image processing is sensitive to different lighting conditions, these techniques may present some practical drawbacks.

Another approach is to use highly capable sensors. The Gator Tech Smart House enables augmentation with various sensors, such as ultrasonic trackers [8]. Such sensors can provide high fidelity, but cost-effectiveness is an issue, because they are still expensive and not readily available.

An approach that has recently attracted attention involves the use of depth cameras. Depth cameras can provide features similar to those of RGB image processing (i.e., they can gather a lot of content) and are much more robust under different lighting conditions. Also, they are now affordably priced. LightSpace was the first system in which an augmented space was constructed using depth

Fig. 1 An application of smart notification: when the user is close to the printer, a message is displayed on his palm regarding a scheduled printing task (a). The user can complete the printing task by dropping the message onto the printer (b)



cameras [22]. The authors augmented all common surfaces (originally non-interactive) in their experimental space, and they successfully demonstrated the plausibility of using depth cameras via promising applications and interactions.

Our environment is implemented using an approach similar to that of LightSpace (i.e., depth cameras) [22]. However, we installed cameras in an actual workspace, whereas the LightSpace was based on an experimental space. Our real-world approach requires additional investigation. For example, it creates more concerns regarding the installation and synchronization of the cameras (The details of the implementation are presented in Sect. 4). Augmented spaces based on actual workspaces may also be necessary for specific research topics. For example, our system can be used to examine the effect of a users familiarity with a space [10], which is difficult to accomplish in an artificially constructed experimental space.

2.2 Registration of object locations

Location-based context means that the system is able to track the locations of users or objects to measure the distances between them. Consequently, the system can determine which users are most likely to manipulate which objects. This is the exact meaning of location-based context in this study. Location-based context can provide many usable features [16], and the registration of object locations is a necessary preliminary step to accomplishing this.

Generally speaking, there are two approaches to registration: automated and manual. The Gator Tech Smart House employs various sensors to track object locations automatically [8, 11]. Devices store their profiles and register themselves automatically, and then, the sensors are able to track the locations of the devices. We agree that this is an elegant approach, but the technology is not yet fully mature, and the cost remains relatively high.

Mobile augmented reality is a manual approach. In CA-MAR [20] and iCam [17], the space is able to track the pose (i.e., location and heading) of a mobile device. With this pose data, it is possible to determine the location of an object using the mobile screen (users can collect the locations from the mobile screen). However, one drawback of this approach is that it can be time-consuming, because the users must be close to the objects, which necessitates physical movement.

The Gator Tech Smart House demonstrated another promising technique [8]. They constructed a sophisticated virtual model that allows users to collect locations while navigating the virtual space via a simple GUI application. This technique is actually cheaper and more efficient than the aforementioned approaches (automatic registration and mobile AR), but considerable time and labor (e.g., 3D modeling) are involved in constructing the virtual model. As noted above, we constructed an augmented environment using multiple depth cameras. By combining images from the depth cameras, it is possible to create a point cloud visualization, which can function as a virtual model without additional cost. Then, users can collect the locations of objects from this point cloud visualization. The details are presented in Sect. 5.1.

2.3 Creating interfaces on the body

There are two different methods for creating interfaces on the body: wearable devices and the infrastructure-dependent approach. In this section, we describe both methods and explain their benefits and shortcomings.

2.3.1 Wearable devices

Skinput, by Harrison et al. [7] attaches analog sound sensors to a user's elbows. When a point on the skin is touched, some vibration occurs, and this vibration is propagated along the skin. The sound sensors are able to detect such vibrations, and recent machine learning techniques allow the point of contact to be recognized. Omnitouch provides similar features, but relies on different sensors (i.e., depth cameras) [6]. The developers of this system tested a shoulder-worn device that included a micro projector and a depth camera. The micro projector displayed some images on a user's palm, and the depth camera was able to track the movement of fingertips on the palm.

2.3.2 Infrastructure-dependent approach

Another technique for creating a visual interface on the body is to exploit the infrastructure of the space. The Palm Display of Kim et al. [9] exploits a projector and camera installed in the space. The projector displays some images on a user's palm, and the user is able to interact with these images via fingertip gestures. LightSpace has a more sophisticated setup [22]. Multiple depth cameras and projectors are used so that every surface in the space becomes a touch-sensitive display, and a user is able to interact with them using the interface on his/her body.

The major advantage of this infrastructure-dependent approach is obvious. It does not require the user to wear any devices. However, it is applicable only to indoor scenarios. On the other hand, wearable devices are applicable to both indoor and outdoor scenarios, but users must wear more than one device.

This study is concerned with interactions in an augmented space, which is an indoor scenario. Accordingly, we developed a visual interface using the infrastructuredependent approach.

3 Applications

As previously noted, we have constructed a space that can record the locations of users and objects. Using this location-based context, the system can determine which user is manipulating which object. The space also creates a visual interface on the palm. In this section, we present some applications. The details of the system are discussed in Sect. 4.

3.1 Data storage on the body

The system can track users' bodies and the locations of devices. Thus, it can sense a touch event between a body part and a device by measuring the distance between them. A possible application arising from this capability is data storage on the body. Figure 2 illustrates this. A user touches a digital picture on a laptop display, and the picture is moved to his hand (Fig. 2a, d). Then, he touches his right pocket (Fig. 2b), and the system stores the picture on his right pocket, and the picture reappears on his hand (i.e., the user extracts the picture from the pocket). Then, the user can drop the picture onto another computer (Fig. 2c, e).

When transferring data from one machine to another, temporary storage (e.g., flash drive) is sometimes

necessary, and this can be rather cumbersome, because the device must be connected to and disconnected from each machine. The present application could be useful in such situations, allowing users to temporarily store data on their bodies and easily transfer the data to another machine.

3.2 Virtual midair input

In the above example, the system recognized a touch event between a body part and a device by measuring the distance between them. Utilizing the same principle, the system can also recognize a touch event between a midair location and a body part. For example, if a location is associated with a specific interface, and a user's hand is close to that location, the system can automatically display the interface on the hand.

Figure 3 illustrates this. The location indicated by the blue rectangle P in Fig. 3a is reserved as the location of an imaginary light switch. The user memorizes this location, and when he wants to control the light, he simply places his hand near the reserved location, and the switch automatically appears, as shown in Fig. 3b, c. The application displays a bidirectional arrow, the two directions corresponding to the on and off commands. The user can control the light by making the appropriate directional hand gestures.



Fig. 2 An example of data stored on the body. The user stores a picture on his hand by touching the screen of a laptop (a, d), stores the picture on his pocket (b), and drops the picture onto another laptop (c, e) Fig. 3 Imaginary input in the air. Users can secure a location in the air for an interface. For example, the user swings his arm around at (P) in **a**. Then, it shows an imaginary interface (**c**) for a light emulator in a large display, i.e., (1) in **b**



Fig. 4 Visual feedback viewer for a hand gesture. When the user makes a gesture (**a**), the shape of the gesture is traced on the other hand (**b**)



This application implies that we can exploit any location in a space as a meaningful location associated with a specific interface. Generally speaking, controllers for electrical devices (e.g., lights or air conditioners) are placed at certain locations, and users must move to those locations. When this type of virtual input is available, users can reserve a convenient location and can access the controllers with less physical movement.

3.3 Complementary interface

Because most interaction schemes are less than perfect on their own, additional modalities generally help to improve their capabilities. In this sense, a visual interface on the palm can be used to provide a visual modality to other interfaces. To illustrate this point, we present a visual feedback viewer for hand gestures.

An advantage of hand gestures is that they can be made without temporal or spatial constraints. Sometimes there may be no means of providing feedback (i.e., devices without visual or auditory components). In this case, an interface on the palm can serve as a visual feedback viewer, as shown in Fig. 4. When a user makes a triangular gesture with one hand (Fig. 4a), the system traces the recognized shape on the other hand (Fig. 4b).

This type of viewer could be particularly useful in certain scenarios where users are manipulating devices without visual components (e.g., printers or microwaves), and for some interaction schemes that have no visual modality (e.g., speech or haptic interfaces).

4 Construction of the augmented space with depth cameras

Thus far, we have presented some applications utilizing augmented space. In this section, we describe the implementation of augmented space in detail. The augmented space in this study was constructed using depth cameras. As mentioned in Sect. 2, depth cameras make it possible to collect a great deal of information, are robust under different lighting conditions, and are cost effective. To employ depth cameras, the first task is obviously their installation.

There are two primary considerations regarding installation: avoiding blind spots and ensuring that each camera covers as much area as possible. With these two concerns in mind, we installed four Microsoft Kinect cameras at points (a)–(d) in Fig. 5. We took into account the cameras' horizontal and vertical fields of view (FOVs), which are 43° and 57° , respectively. The white dotted lines in Fig. 5 indicate the FOVs. (a)–(d) denote the locations of the depth cameras.

Another consideration is efficient tracking range. According to the official specifications, the efficient tracking range for the Kinect sensor is about 0.7–6 m. However, in an informal preliminary test, the device was only able to effectively track users from 1 to 4 m, which is similar to the range supported by Microsoft's official software development kit (SDK). Accordingly, we installed the cameras to cover only this range. As Fig. 5 shows, there were four cameras in a distance of 7.5 m. Hence, these four cameras should be able to cover the entire area if each camera covers a range of about 1–4 m.



Fig. 5 Layout for the depth camera installation. Side view (**a**) and top view (**b**). (a)-(d) indicate the locations of depth cameras, and P1 and P2 indicate the locations of the projected displays. The white dotted lines indicate the fields of view for the depth cameras

4.1.1 Interference between infrared lights

A concern that arises when using multiple Kinect cameras is interference between the infrared (IR) lights. According to an informal experiment reported on a website, this interference can be especially severe when two depth cameras are facing each other directly [4]. However, this did not occur in our installation, because all cameras were pointed toward the floor (see Fig. 5). Moreover, we observed no interference effects in preliminary demonstrations.

4.2 Merging multiple depth cameras

To combine 3D data from multiple depth cameras, synchronization (i.e., aligning the coordinates of each camera) is a necessary step. Here, we describe this process.

4.2.1 Synchronization via image processing

Synchronization (i.e., merging) of multiple depth cameras was successfully demonstrated previously [22]. Because they set up their system in an experimental space, they were able to install retro-reflective dots, which emitted relatively bright light that was easily extracted via simple image processing. However, such retro-reflective dots are generally not available in actual workspaces.

Therefore, we had to develop a new synchronization procedure. Du et al. [2] demonstrated a promising technique. When two images are slightly different, the SIFT algorithm makes it possible to find many shared feature points between them [21]. These shared feature points play the same role as the retro-reflective dots employed in LightSpace [22]. After eliminating outliers using the RANSAC [3] algorithm, a synchronization matrix M can be found by solving Eq. 1:

$$M(X_1, X_2, \dots, X_n) = (Y_1, Y_2, \dots, Y_n)$$
(1)

where $X_1, X_2, ..., X_n$ are inliers in Frame 1 and $Y_1, Y_2, ..., Y_n$ are inliers in Frame 2.

This approach is based on the assumption that there are two slightly different images. Thus, we can only use this method when two cameras are very close to each other, so that their images will differ only slightly. However, as Fig. 5 shows, the cameras are installed at separate locations.

To address this issue, we slowly move a camera from one location to another, capturing many frames in between, and find synchronization matrices between every consecutive frame. Figure 6 illustrates this process. Assume that two cameras are installed at points (a) and (b). The camera at point (a) is manually moved along the indicated path



Fig. 6 Example of synchronization. If two cameras are placed at (a) and (b), they can be synchronized by moving one camera along the indicated path

toward the camera at point (b), and synchronization matrices are found between all consecutive frames along the path. In this way, we can find the desired synchronization matrix by accumulating all intermediate matrices.

Using this technique, we were able to find matrices that synchronized the four cameras in Fig. 5. In preliminary tests, it took about 30 min to complete this procedure.

4.2.2 Manual adjustment

Although we were able develop an automated process, it is very difficult to eliminate all imperfections, because errors can occur when extracting matrices between frames. To deal with such errors, we developed a piece of software that allows the cameras pose (i.e., yaw, pitch, roll, and position) to be adjusted manually.

Figure 7 shows a screenshot of the software. A list (a) of connected depth cameras is displayed. After selecting one

of these cameras, the image panel (b) shows an RGB image from the selected camera. The camera's pose data (c) are also displayed. Users can edit these values, and the software shown in Fig. 8 reflects the changes in real time. This software enabled us to complete the process of finding synchronization matrices for all cameras. Figure 8 shows a visualization of a point cloud from the four depth cameras shown in Fig. 5.

4.3 Distributed system

After synchronizing the four cameras, we developed a distributed system to manage them. Here, we describe the composition of this system.

Figure 9 shows the system composition. Each of the four computers (a) hosts one of the four depth cameras in Fig. 5. The computer (b) functions as a server. Each of the four host computers sends depth data, RGB images, and user skeletons to the server. Then, the server can visualize the environment, manage object locations, and track user locations, because it has access to all data from the four depth cameras. The applications shown in Figs. 7 and 8 run on the server.

We expect that this scalable client—server architecture will be useful when additional cameras are necessary.

4.4 Tracking user skeletons

As Fig. 9 shows, four computers host the depth cameras and send user skeletons to the server. A problem may occur



Fig. 7 Software used for manual adjustment and managing objects in the space. Users can add or delete object locations, and adjust a camera's pose (e.g., for fine synchronization)



Fig. 8 Visualizations of a merged point cloud from four depth cameras, showing the space from different perspectives



Fig. 9 Distributed system. Each of the four computers (a) hosts a depth camera, and all data are sent to the server (b)

if a user is captured by more than one camera, because more than one set of skeleton data will be sent to the server. The server must accept only one set of skeleton data in such a case.

For this purpose, the host computers send not only skeleton data, but also the absolute distance between the user and the camera. Because the closest camera is clearly more reliable in most instances, the system accepts the skeleton data with the shortest distance and ignores data from more distant cameras. Figure 10 shows visualizations of a user (colored green) from different perspectives.



Fig. 10 A user in the space is tracked from different perspectives (a, b)

4.5 Projected displays for creating an interface on the body

In this section, we describe the projected displays, which are located at P1 and P2 in Fig. 5. The purpose of the projected displays is to display appropriate images on the palms of users.

Figure 11 shows the projected displays. As the figure indicates, each projected display is composed of a projector and a depth camera. The depth camera tracks the location of a users hand and provides the location to the projector. Then, the projector displays an appropriate image on the hand (see Figs. 2, 3, 4).

Obviously, it is necessary to calibrate the projector and the depth camera. For this purpose, the projector first displays a chessboard pattern, and two sample images are captured at different heights, as shown in Fig. 11a, b. Then, we can find the directional vectors that connect the corners of the two images, represented by dotted lines (1)-(4) in Fig. 11c. Using these vectors, we can estimate the pose of the projector. This calibration process is required only once [6].

Fig. 11 Calibration of a projector and a depth camera. The camera captures images at two different heights (a, b). Then, the projector's pose is determined by connecting the corners of the two images (c)



Some related studies used the POSIT algorithm [15] for the calibration [6, 22]. However, this algorithm requires several intrinsic parameters of the camera and projector, and these parameters were not available to us. Therefore, we developed this method.



5 The support tool

To enable the applications mentioned in Sect. 3, registration of objects is also necessary. The software shown in Fig. 7 can be used for this purpose. We present the details below.

5.1 Pick what you see: collecting the locations of objects

Although the system provides visualizations of point clouds, the quality is usually not very high. Our support tool is capable of displaying RGB images from each camera. When the depth cameras are connected to the server via the host computers, they appear in list box (see Fig. 7a). Users can select any camera from the list, and the corresponding RGB image will be displayed (Fig. 7b).

This RGB image can be used to collect object locations. When a user clicks on a 2D point in the RGB image, it is converted into a 3D point by adding the relevant depth data, and the results (Fig. 7f) are displayed in the bottom left corner of the window. The user can input the name of an object in a text box (Fig. 7g), and the object will be added to the list (Fig. 7d). Below this list box, there are remove and save buttons. When the user selects an item in the list box and clicks the remove button, the corresponding object is removed. When the user clicks the save button, the current locations of registered items are stored in an extensible markup language (XML) file for later use.

Fig. 12 An example of the XML database. This example shows the positions and names of a printer and a microwave



Fig. 13 The visualization of registered objects in the space. Text labels (a, b) are placed near the objects

Figure 12 shows an example of the database file containing the names and locations of objects.

All of these changes (i.e., registration and deletion) are reflected in real time, as shown in Fig. 13. In the example illustrated, text labels for the two items [printer (a) and microwave (b)] are added.

Because the desktop environment is 2D, we expect that the 2D graphical panel (the RGB image) is easier to use than the 3D visualization (the point cloud). Indeed, the difficulty created by mismatched input and output dimensions has been an issue of note for a long time.

5.2 Visualizing only necessary data to support realtime updating

It is difficult to support real-time updating [i.e., to guarantee more than 15 frames per second (FPS)]. When point clouds from four cameras are merged, the server must process about 17 MB of data per frame, assuming that the cameras have VGA resolution (640×480) , including three floating point numbers per 3D point (12 bytes) and RGB data (3 bytes). Even though the amount of data varies (e.g., missing depth data), it is almost impossible to process such a tremendous amount of data in real time. In preliminary tests, we were only able to obtain 2–3 FPS, and our server has an Intel Core i7 2.8 GHz processor, 8 GB RAM, and an nVidia GTS250 chipset.

To address this, the system has a feature whereby only necessary data are visualized. In this mode, the host computers send only point cloud data corresponding to users, and more than 20 FPS was obtained.

5.3 Summary

Figure 14 shows the overall composition of the system. In this figure, machines (a) and (b) correspond to machines (a) and (b) in Fig. 9. The depth cameras (a) send user skeletons, RGB images, and point cloud data to the server (b).

Then, the server can visualize a point cloud and display an RGB image (see Figs. 7, 8, 10, 12). As mentioned above, users can register objects, and the server can calculate the distances between the objects and user skeletons. These distance data are transferred to the software (c), and the software, in turn, can display an image on a user's palm via the projected display (d).



Fig. 14 Overall composition of the system. Depth cameras (a) send their data to the server (b). The server stores the necessary context (i.e., locations) and provides that context to the software (c) when necessary. The software can turn the projected display (i.e., the visual interface on the palm) on and off

6 Implementation of the applications

In this section, we explain the implementation of the applications described in Sects. 1 and 3.

6.1 Smart notification of a printing task

In Sect. 1, we introduced an application that notifies a user of a printing task when the user is close to a printer. To implement this application, we must first input the location of a printer, and this is done using the software shown in Fig. 7. Then, the server (Fig. 14b) examines the distance between the printer and the user skeleton. The distance data are shared with the software (Fig. 14c). If the distance is relatively small, the software turns on the projected display and causes the notification to be displayed on the user's palm (see Fig. 1).

After an image is displayed on the user's palm, the distance between the printer and the user's palm is continuously tracked. If this distance is relatively small, the software prints a sample page. Communications among the server, software, and projected display are manually established via a network. The system is connected to a printer, and a routine for printing a sample page is included in the software.

6.2 Data storage on the body

In Sect. 3.1, we introduced an application that enables the storage of digital content on some part of the body and the transfer of the content to machines. To implement this application, we must first input the locations of two computers, which can be registered using the software shown in Fig. 7. Then, the distance between the user's right hand and the computers is continuously examined.

When a touch event occurs (i.e., between the user's right hand and a machine or between the hand and the pocket), the system takes a different action depending on the state of



Fig. 15 State diagram for the application that presents data storage on the body

the application. Figure 15 shows the state transition diagram. There are four states, depending on the location (the two machines, the hand, and the pocket) of the data (the image on the computer shown in Fig. 2).

The application starts with the state "picture in machine 1" (state S1 in Fig. 15), and the picture is first shown at that location. When the user touches machine 1, the state changes to picture in hand (state S2, transition (a) in Fig. 15). When the hand is close to the pocket, the state changes to "picture in pocket" (state S3, transition (c) in Fig. 15). When the user touches the pocket again, the state reverts to "picture in hand" (state S2, transition (d) in Fig. 15). Employing the same principle, the user can move the picture to machine 1 (state S1, transition (b) in Fig. 15), to machine 2 (state S4, transition (f) in Fig. 15), or back to his pocket (state S3, transition (c) in Fig. 15).

6.3 Virtual midair input

The implementation of this application is almost the same as that of the smart notification example (see Sect. 6.1). To reserve a midair location, the user places his hand at the location, and the location is secured via the software shown in Fig. 7.

When the user's left hand is near the location, the software turns on the projected display and displays an image. A gesture is recognized by tracking the movement of the right hand. In tests, we constructed a simple light emulator (shown in the display (1) of Fig. 3b) that was manually connected to the system via the network. Recognition of hand gestures is explained in the next section.

6.4 Complementary interface

For the application in this example, we arbitrarily specified the right hand for making gestures and the left hand for the feedback viewer. Using the skeleton data from the depth cameras, the left and right hands are easily distinguished. For gesture recognition, we employed the one dollar gesture recognizer [23]. This created a problem, because the depth camera provides 3D hand movement data, whereas the one dollar recognizer accepts only 2D data. To address this issue, we projected the 3D data onto an appropriate 2D plane, as in a previous study [5]. The system continuously provided the right hand movements to the recognizer, and a visualization was displayed when a gesture was recognized.

7 Limitations and future work

Here, we describe the limitations of the current implementation and suggest some promising future directions.

7.1 Temporary disappearance of a user

One concern with the current system is that a user's location can temporarily be lost. In particular, this occurs when users cross the borders between coverage areas of different cameras (i.e., intersections between the dotted lines in Fig. 5). Currently, the OpenNI implementation takes 1–2 s to recognize a skeleton, depending on the image processing capabilities or sensors. When more sophisticated algorithms and devices are developed, this issue will be addressed.

7.2 Lack of accuracy

The accuracy of the depth camera installation is not high. We are currently using four Kinect cameras to cover a 4–7.5 m space. This setup could be used to cover a broader area, but high accuracy is not easily obtained. Because we employ a distributed system architecture, we could add more depth cameras when higher accuracy is required. However, this issue can be better addressed when more advanced (i.e., higher resolution) depth cameras become available.

7.3 Supporting complex shapes

The current system treats object locations as points. This is because the system is not accurate enough to support complex shapes (e.g., spheres or cubes). Although this was sufficient to implement the applications demonstrated here, a greater variety of applications would be possible if the system supported complex shapes. It would be worthwhile to develop a more advanced version of the software shown in Fig. 7, and resolve the issues related to it.

7.4 Continuous tracking of objects

In this work, we considered only interactions among static objects. Of course, some objects in a space (e.g., books or laptops) can be moved. To support interactions with such movable objects, continuous tracking is necessary and should be investigated in future work.

There are various technologies that can track moving digital devices with good accuracy [14]; however, there are fewer solutions for non-digital devices (e.g., notes or books). We expect that our system may be extended to create an infrastructure for tracking non-digital devices. In our setup, an object is imaged using multiple cameras (see Fig. 5); therefore, when an object is registered using the software (see Fig. 7), the system can obtain multiple images from different angles. Such images are valuable in improving the accuracy of object detection. This is a possible area for future work, when more sophisticated image processing algorithms and higher resolution cameras become available.

7.5 Identifying users

We have demonstrated a system that can track users; however, we did not consider how the system might identify users. Although this is beyond the scope of this work, there is a promising scenario that may be enabled with the current implementation.

In many laboratories and offices, most users undergo some authentication procedure when they enter the workplace. If our system can communicate with the authentication system, it could continuously track a user with a registered profile. Realizing this scenario and addressing new issues related to it is a direction for future work.

8 Conclusions

We have demonstrated applications that can be achieved via a visual interface on the human palm, which is always available and is location sensitive. We constructed an augmented space in a workspace by adding depth cameras. We employed a distributed system to manage the depth cameras and developed a support tool to facilitate the collection of object locations. We have also demonstrated several potential applications.

In future work, we plan to focus on continuous tracking of objects and user identification (see Sects. 7.4, 7.5). Potentially, every surface on objects in a room can be augmented to be interactive. Moreover, it can be fully personalized. Lee et al. [13] demonstrated augmenting various surfaces and, with our system, it is possible to make them fully personalized. For example, when a user reads a newspaper, our system can project a related article onto it, selected based on a user's personal interests; it can also be automatically translated into the user's native language. Another example is that users can write a message using gestures on a book and attach it on a book (or anywhere they wish). The attached message can then only be read by specific people.

We anticipate that this work will contribute to the development of an infrastructure for such services, and may help to create new interactions in such environments.

References

- Brumitt B, Meyers B, Krumm J, Kern A, Shafer SA (2000) Easyliving: technologies for intelligent environments. In: Proceedings of the 2nd international symposium on handheld and ubiquitous computing, HUC'00. Springer, London, UK, pp 12–29
- Du H, Henry P, Ren X, Cheng M, Goldman DB, Seitz SM, Fox D (2011) Interactive 3d modeling of indoor environments with a consumer depth camera. In: Proceedings of the 13th international

conference on ubiquitous computing, UbiComp'11. ACM, New York, NY, USA, pp 75-84

- Erdem CE, Bozkurt E, Erzin E, Erdem AT (2010) Ransac-based training data selection for emotion recognition from spontaneous speech. In: Proceedings of the 3rd international workshop on affective interaction in natural environments, AFFINE '10. ACM, New York, NY, USA, pp 9–14
- Fukushima N (2010) Interference of IR light from multiple kinects@ONLINE. http://youtu.be/p2dokAm1uQg. Accessed 30 Jan 2014
- Gustafson S, Bierwirth D, Baudisch P (2010) Imaginary interfaces: spatial interaction with empty hands and without visual feedback. In: Proceedings of the 23Nd annual ACM symposium on user interface software and technology, UIST '10. ACM, New York, NY, USA, pp 3–12
- Harrison C, Benko H, Wilson AD (2011) Omnitouch: wearable multitouch interaction everywhere. In: Proceedings of the 24th annual ACM symposium on user interface software and technology, UIST '11. ACM, New York, NY, USA, pp 441–450
- Harrison C, Tan D, Morris D (2010) Skinput: appropriating the body as an input surface. In: Proceedings of the SIGCHI conference on human factors in computing systems, CHI '10. ACM, New York, NY, USA, pp 453–462
- Helal S, Mann W, El-Zabadani H, King J, Kaddoura Y, Jansen E (2005) The gator tech smart house: a programmable pervasive space. Computer 38(3):50–60
- Kim S, Takahashi S, Tanaka J (2010) New interface using palm and fingertip without marker for ubiquitous environment. In: Proceedings of the 2010 IEEE/ACIS 9th international conference on computer and information science, ICIS '10. IEEE Computer Society, Washington, DC, USA, pp 819–824
- Kim S, Takahashi S, Tanaka J (2012) Point-tap, tap-tap, and the effect of familiarity: to enhance the usability of see-and-select in smart space. Trans Hum Interface Soc 14(4):445–455
- Lee C, Nordstedt D, Helal S (2003) Enabling smart spaces with OSGI. IEEE Pervasive Comput 2(3):89–94
- Lee J, Choi J, Shin D, Shin D (2006) Multi-user human tracking agent for the smart home. In: Proceedings of the 9th pacific rim international conference on agent computing and multi-agent systems, PRIMA'06. Springer, Berlin, Heidelberg, pp 502–507
- Le JC, Hudson SE, Tse E (2008) Foldable interactive displays. In: Proceedings of the 21st annual ACM symposium on user interface software and technology, UIST '08. ACM, pp 287–290
- Liu H, Darabi H, Banerjee P, Liu J (2007) Survey of wireless indoor positioning techniques and systems. Trans Syst Man Cyber Part C 37(6):1067–1080
- Lowe DG (2004) Distinctive image features from scale-invariant keypoints. Int J Comput Vis 60(2):91–110
- Marquardt N, Diaz-Marino R, Boring S, Greenberg S (2011) The proximity toolkit: prototyping proxemic interactions in ubiquitous computing ecologies. In: Proceedings of the 24th annual ACM symposium on user interface software and technology, UIST '11. ACM, New York, NY, USA, pp 315–326
- Patel SN, Rekimoto J, Abowd GD (2006) icam: precise at-adistance interaction in the physical environment. In: Proceedings of the 4th international conference on pervasive computing, PERVASIVE'06. Springer, Berlin, Heidelberg, pp 272–287
- Pinhanez C (2001) The everywhere displays projector: a device to create ubiquitous graphical interfaces. In: Abowd G, Brumitt B, Shafer S (eds) Ubicomp 2001: ubiquitous computing. Lecture notes in computer science, vol 2201. Springer, Berlin, pp 315–331
- Rekimoto J (1997) Pick-and-drop: a direct manipulation technique for multiple computer environments. In: Proceedings of the 10th annual ACM symposium on user interface software and technology, UIST '97. ACM, New York, NY, USA, pp 31–39

- Shin C, Lee W, Suh Y, Yoon, H, Lee Y, Woo W (2009) Camar
 future direction of context-aware mobile augmented reality. In: Ubiquitous virtual reality, 2009. ISUVR '09. International symposium on, pp 21–24
- Sinha SN, Frahm JM, Pollefeys M, Genc Y (2011) Feature tracking and matching in video using programmable graphics hardware. Mach Vis Appl 22(1):207–217
- 22. Wilson AD, Benko H (2010) Combining multiple depth cameras and projectors for interactions on, above and between surfaces. In: Proceedings of the 23nd annual ACM symposium on user

interface software and technology, UIST '10. ACM, New York, NY, USA, pp 273–282 $\,$

23. Wobbrock JO, Wilson AD, Li Y (2007) Gestures without libraries, toolkits or training: a \$1 recognizer for user interface prototypes. In: Proceedings of the 20th annual ACM symposium on user interface software and technology, UIST '07. ACM, New York, NY, USA, pp 159–168