# Fast 3D Cluster Tracking for a Mobile Robot Using 2D Techniques on Depth Images

ARNAUD RAMEY[*], MARÍA MALFAZ, and MIGUEL A. SALICHS

*Robotics Lab, University of Carlos III, Leganés, Spain*

**Abstract**: User simultaneous detection and tracking is an issue at the core of human–robot interaction (HRI). Several methods exist and give good results; many use image processing techniques on images provided by the camera. The increasing presence in mobile robots of range-imaging cameras (such as structured light devices as Microsoft Kinects) allows us to develop image processing on depth maps. In this article, a fast and lightweight algorithm is presented for the detection and tracking of 3D clusters thanks to classic 2D techniques such as edge detection and connected components applied to the depth maps. The recognition of clusters is made using their 2D shape. An algorithm for the compression of depth maps has been specifically developed, allowing the distribution of the whole processing among several computers. The algorithm is then applied to a mobile robot for chasing an object selected by the user. The algorithm is coupled with laser-based tracking to make up for the narrow field of view of the range-imaging camera. The workload created by the method is light enough to enable its use even with processors with limited capabilities. Extensive experimental results are given for verifying the usefulness of the proposed method.

**Keywords**: depth map, human robot interaction, image segmentation, multiuser detection

## INTRODUCTION

Human–robot interaction (HRI) requires the robot to be aware of its environment. It needs to understand what lies in front of it, if users are there, and, if so, whether they want to interact with it. Furthermore, recognizing the objects that were previously seen is an important feature for obtaining a robot that learns from its experiences. We indeed expect it to recognize places, objects, and people previously seen. Such a feature is of paramount importance in social robotics, where interaction with human users and daily objects is the keystone of the robot's activity.

[*] Address correspondence to Arnaud Ramey, Robotics Lab, University of Carlos III, Legane´s, Spain. E-mail: arnaud.a.ramey@gmail.com

In this article, we aim at giving a social robot the possibility of detecting and following objects or people. The user indicates to the robot what a object to track and then the robot plans its motion to stay at a given distance of this object of interest. The method is mainly based on the use of depth maps, which can be used on a wide range of robots. For robust tracking while moving, a fallback method using a laser scanning range finder is used. Such a skill is a very useful feature for the social robot: making it approach a given object or follow a given user is frequently necessary.

This article is structured as follows: In the following section, a comprehensive review of the existing solutions is presented. Next, we will define in further detail the goals and specifications of our research. Then, we will present the algorithmic details of our solution. In the next section, we present the results of a comprehensive series of tests that show the validity of the proposed algorithm. Finally, its usefulness and the future directions of this work are discussed.

## RELATED WORK

A depth map is an image that contains information about the distance of the surfaces that can be viewed from a viewpoint. As such, it directly provides spatial comprehension of the observed scene. Devices that can provide depth maps have existed for many years; for instance, via stereo-vision. However, their presence on mobile robots was limited due to a relatively high price, complex calibration, and an additional computational workload for the robot. This might explain why information from a depth map is not often reported.

Many techniques focus on mixing data from several devices, such as cameras, laser scanners, etc. Among them we can underline the works of Muñoz-Salinas (2007) and Muñoz-Salinas and Aguirre (2009). However, the use of multimodal fusion is often computationally costly and requires several expensive and bulky devices. Therefore, it is not perfectly suitable for small mobile robot bases.

Some articles, however, also provide solutions only using a stereo camera or a range imaging device. The release of the affordable and reliable Kinect device at the end of 2011 led to its soaring use in robotics. In Tomori (2012), a simple motorized support for a Kinect device, called KATE (Kinect Active Tracking Equipment), was presented. It has two degrees of freedom, which enables the camera to look at the object of interest. The latter can be provided by face detection or the center of the image.

Segmentation of the depth images around this center is made by a modified version of the GrabCut method (Rother et al. 2004). However, the KATE platform is not mobile, making segmentation much easier. The patented PrimeSense NITE middleware Berliner and Hendel (2007) allows detecting and tracking human shapes from depth maps. It has been shown to work at a very high frame-rate even with average CPUs. Although the technique employed and the source code are not available, it is likely that motion analysis and clustering techni-ques are at the core. Indeed, human detection is activated by their motion. As such, detecting a motionless audience is challenging. This is especially an issue in HRI, where the user can stand still in front of the robot while speak-ing to it. The software is aimed at playing video games on Microsoft XBox 360 Latta and Tsunoda (2009). As human players are bound to move their bodies to play, this context solves the detection problem.

Stereo-vision devices are also used on mobile platforms (Howard and Matthies 2007; Jia et al. 2011). A complete pedestrian detection system only based on stereo-vision was presented in Howard and Matthies (2007). A polar-perspective map was built and then segmented into regions of inter-est. The system is compatible with classical image processing techniques, such as appearance-based algorithms. Stereo-vision is, according to authors, a reliable solution for navigation and pedestrian detection. However, their system is designed for cameras mounted on outdoor vehicles, with ranges between 5 and 50 m. Thus makes it difficult to use for indoor robotics plat-forms. In Jia et al. (2011), a Point Grey stereo camera mounted on top of a mobile platform provided depth maps. User detection was performed through edge detection with a Canny filter (Canny 1986). Tracking was performed using an extended Kalman filter. However, the process requires a complex calibration process and restricted experimental results do not clearly state the robustness of the method. In a related field, glasses for people with impaired vision presented in Lee (2012) performed obstacle detection based on segmentation using only depth images. The developed framework works with nonconstrained camera position and orientation, and because it discards color data, it also works in the dark. However, it can only detect the objects coming closer and notify the user. There is no recognition or tracking of the objects being detected.

The method proposed in this article aims at correcting some of the limitations of the research previously presented. It focuses on hardware requirements that are easy to meet and a light workload. The code is released under the LGPLv3 license and integrated into a common software platform, allowing other teams to improve it easily.

## PROBLEM STATEMENT

The aim of our work is to provide a mobile robot with a lightweight algorithm for cluster detection and tracking. These clusters can be human users or objects. The only required input data are the stream of depth maps. The RGB data—that is, color images–also supplied by the Kinect, are discarded in this article. This ensures the compatibility of the algorithm with devices such as PMDVision CamCubes or Asus Xtion PRO devices, which do not supply RGB images. It should be able to be processed either onboard or on a remote computer, but the goal is to have it done on-line, in real time. As such, visual tracking of the object of interest is made simultaneously with navigation of the robot.

More accurately, tracking is made at two levels: first, at the sensor processing level, the robot is expected to keep track of the followed cluster. Second, we want the former to physically move itself and keep close to the latter. The whole process results in the robot moving to the tracked cluster and following its trail if it is moving.

## Hardware Specifications

The target robot for this application is the social robot *MOPI* (nondefinitive name). It is a home-made robot of the RoboticsLab of the University Carlos III of Madrid and is shaped like a mobile, car-like platform. It is most notably equipped with a Microsoft Kinect device tilled at 45° and a nontiltable Hokuyo laser scanning range finder. Communication is made through a Wi-Fi connection.

Furthermore, we also make use of a touchscreen computer. It is a TravelMate C110 computer equipped with a touch-sensitive screen and a CPU of modest performance. It displays a graphical user interface that enables the user to choose an object of interest.

## Software Specifications

The robot MOPI works according to the AD paradigm, as presented in Barber and Salichs (2002). This paradigm handles skills relying on primitives. Primitives are in direct communication with the physical devices of the robot and send elementary orders to them. This includes the base motors, the laser sensor, the camera, etc. A skill is the ability of the robot to do a specific action. It relies on the data supplied by the primitives. The actions generated by a skill

can be numerous: move the car to a given point, play games with the user, interact with electric appliances, etc.

Since it is an experimental platform, the robot MOPI has been used as a bridge between the traditional implementation of AD, as seen in Rivas (2007), and a new one relying on the communication mechanisms of *ROS*, the robot operating system (Quigley et al. 2009). The version used for this application is ROS *Electric* running on top of Ubuntu 10.10. The use of ROS most notably provides the possibility of redistributing the computational workload. We can send raw data from the sensors to remote computers for processing via wired or wireless connections. The latter can be more powerful than robot-embodied PCs, also lightening the computation workload of the main computer. Then, the processed data are sent back to the robot. It also embeds the Stage simulator (Gerkey et al. 2003), which gives us the possibility of making first outlines of our algorithms—for instance, the tracking one—before trying them on the real robot.

## APPROACH: DETECTION AND TRACKING PROCESS

The whole processing pipeline is illustrated in Figure 1.

The process is described in the following order: the acquisition of depth images from the depth sensing device is detailed in the next subsection. As
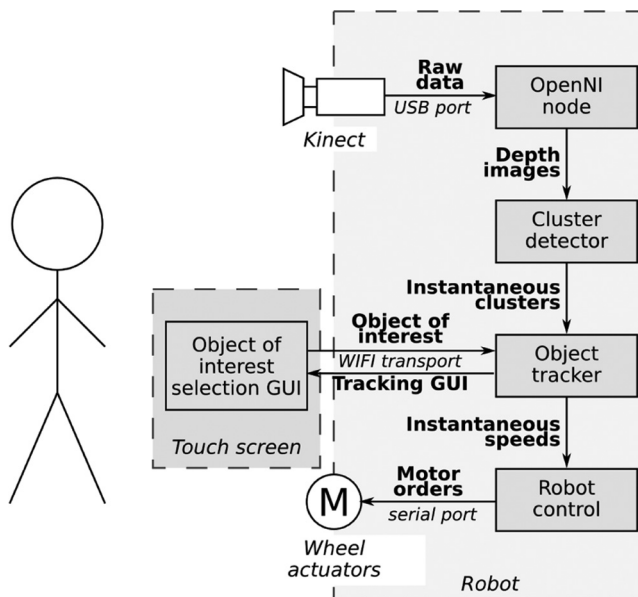


**FIGURE 1** The flowchart of the whole system. It presents all of the components needed for the vision-based algorithm and does not include the laser data fusion presented later in the subsection on robot motion and user tracking.

previously said, the RGB (color) images that can be provided by the Kinect sensor are not used. Thus, the algorithm can also be used with single depth imaging devices, such as CamCubes.

The algorithm developed for converting float depth images into byte images is explained next. The compression of depth images for remote processing follows, and, our tracking algorithm is articulated in three phases:

1. **Cluster detection:** the clusters are found in the current depth map due to 2D image processing techniques.
2. **Cluster matching and tracking:** the clusters of the current depth map are matched to the objects found in the previous depth maps. If the user has selected an object to track, the 3D position of this object is estimated.
3. **Robot control:** depending on the tracking results, the robot motion is controlled. For instance, if the object selected by the user has been found, the robot will move and come closer. Finally, some high-level multimodal fusion is made between the output of our algorithm and the one of another technique based on the data from the range finder.

## Depth Image Acquisition

Acquisition of the depth map is made using the software provided by the ROS architecture. A so-called node is in charge of the communication through the USB port. The stream acquired via that port is transformed into proper depth maps. For each pixel $(x, y)$ in the depth map, the pixel value at $(x, y)$ corresponds to the distance of the closest object intersecting the 3D ray passing by this pixel, in meters.

As such, the values of the acquired depth map are float values, with a range depending on the device. For Microsoft Kinect, they are typically within one to 10 m. Furthermore, the constructed light patterns projected by the device can be reflected for shiny surfaces. The depth map also contains some undefined values, represented by *NaN*.

## Remapping of Float Images to [0,255] Values

Most image processing techniques require the input image to be in the typical RGB space. In this space, grayscale values are represented as byte values (and colors are tuples of bytes):

$$v_{byte} = g \in [0,255] \quad \text{for gray scale images}$$

However, it was seen previously that the values depth maps supplied by the ROS node are decimal numbers representing the physical distance of the object. We then want to convert these images to the [0,255] range while keeping track of the undefined *NaN* points.

It is important not to lose the information about these undefined points. They indeed correspond to physical zones where we have no information. For instance, in the picture visible in Figure 2(a), the black polygons in the ceiling could be, for instance, some reflecting surfaces or some nests in the roof, etc. In this example, they correspond to the neon lights. They reflect the light patterns emitted by the Kinect projector.

For each depth map, this remapping is made in several steps.

1. Detection of the values range of the depth map: we find the minimum $m$ and the maximum $M$ values of the defined pixels in the current depth map.



**FIGURE 2** Detection flowchart. (a) The depth map, as supplied by the range-imaging device (Kinect) and remapped to visible colors; (b) The "clean" depth map, after handling the *NaN* values. The technique used here is *directional propagation* (to the left); (c) Canny edge detection on the clean depth map; (d) Morphological erosion of the Canny edges; (e) Recombination of the opened Canny edges and the *NaN* values; (f) The connected components found and their matches to objects. Each one is drawn with a color corresponding to the object it was matched to, and the index of this object is overlaid.

2. Computation of the affine transform: we determine $\alpha, \beta \in \mathbb{R}$ such that the transform $v_{byte} = \alpha \times v_{float} + \beta$ maps the float values to the byte values:

$$v_{float} \in [m, M] \rightarrow v_{byte} \in [1, 255]$$

Note that the 0 value is discarded. It is used to represent the undefined *NaN* values.

## Data Compression for Remote Image Processing

Float images can come from various origins; for instance, depth maps of a RGBD camera like Kinect. For CPUs with limited capabilities, it is necessary to send these images for processing on a remote computer.

However, ROS does not support compression for float images. We hence developed a package for image transport providing this feature. Float images are remapped to byte images using an affine transform, as seen in the previous subsection. They are then compressed using traditional image compression libraries.

## Clusters Detection

In this section, the method used to find the 3D clusters in the current depth map is explained.

To be able to make use of conventional vision techniques, the float depth map is first converted into usual images with values in the [0,255] span by using the technique seen in the prevoius subsection.

### Undefined values handling

For easier handling of the remapped depth map, *NaN* values are filtered. They correspond to a failure in the depth estimation at the given pixel. For instance, the Canny algorithm explained afterwards would fail with an image containing *NaN* values, as they are seen as a regular 0 value.

Three different and concurrent ways of solving the problem have been tried:

- **Average border**: On each side of the depth map, the average thickness, in pixels, of the *NaN* border is computed. Then the rectangular area corresponding to this average border of the depth map is cropped out. For instance, if on average the first five pixels at each line are undefined, the first five pixels of all lines are removed.
- **Inpainting:** The technique of inpainting, presented in Bertalmio Sapiro (2000), can be used to replace undefined *NaN* values in the remapped depth. Inpainting was primarily designed to remove watermarks or damage to an image. It propagates the color values at the border of the damaged

areas and mixes them. Its use makes sense here as the *NaN* values correspond to unsuccessful depth measurements from the range-imaging device.

- **Directional value propagation:** Another way to fill the missing information is to propagate the values in a direction. It basically consists of setting the value of each undefined pixel to that of its closest defined neighbor to the left. For undefined pixels at the left border, we set them to the closest defined pixel to the right. A sample is shown in Figure 2(b).

EDGE DETECTION

Now the depth images have been transformed into standard byte images and the undefined values have been handled. We can now apply classical image processing algorithms. Because the goal is to find clusters in the depth map, we use a Canny filter on the remapped image (Canny 1986). The Canny filter helps us detect edges in a grayscale image.

It requires two thresholds, a low and a high one. Two edges maps are obtained by passing first a Sobel operator on the grayscale image and then thresholding with the two thresholds. A Sobel operator is a linear filter based on a simple $3 \times 3$ kernel, and it approximates the gradient operator on the grayscale image.

The high threshold edge map contains broken, discontinuous edges, but they are likely to belong to the real contours of the objects. On the other hand, the low threshold edge map contains continuous edges but with many edges that are not useful. The two maps are combined to create an optimal edge map: If a chain of the low threshold map enables to connect two pixels of the high threshold map otherwise disconnected, this chain is added to the final edge map. All of the isolated chains of the low threshold map are then removed. A sample is shown in Figure 2(c).

The two parameters of the Canny filter are defined for values of the RGB space, and their meaning depends on the values of $\alpha$, $\beta$. Hence, they are not consistent between frames. This is why instead of setting the thresholds, we set as constants the values of $\alpha \times low\_threshold$, $\alpha \times high\_threshold$ between frames. Parameter $\beta$ is not taken into account because it is a constant offset factor, and is neutralized by the gradient simulated by the Sobel operator.

MORPHOLOGICAL EROSION

We previously explained how the edges in the depth map are detected. However, some weak local contrasts, such as the feet of the user when close to a wall, can be missed. To compensate for this effect, we apply a morphological transformation that thickens the edges. This can close some edges that had been left open by the Canny filter.

This result is obtained by passing a morphological erosion filter on the image. The erosion filter replaces each value in the image with the maximum

of the values of the surrounding pixels. We used a $3 \times 3$ pixel kernel. It could close gaps in the border that are up to 4 pixels wide. A sample is shown in Figure 2(d).

Then, because we do not want to include the undefined *NaN* values in the objects, we need to restore the undefined values erased in the step of handling of undefined values. They are restored by first computing the minimum of the original remapped and the eroded edge map and then thresholding this minimum with a binary threshold at value 0. A sample is shown in Figure 2(e).

First, we define *connected components* for depth maps: two pixels of a depth map belong to the same connected component if there is a chain from one to the other, such as there is no depth gap between two consecutive elements of the chain.

At this step of the processing pipeline, an edge map has been computed that also includes the undefined *NaN* values returned by the range-imaging device. A 3D object visible in the current depth map corresponds to a cluster without discontinuity in the inside and with a discontinuity at the border with the neighboring pixels. It should then correspond to a connected component in our edge map.

In a previous work (Ramey 2011), the authors presented a lightweight and fast algorithm for connected components fetching in a monochrome image, such as our edge map. It is based on an efficient representation in memory of the connectivity between components due to a *disjoint-sets* forest. The disjoint-sets data structure was first presented by Galler and Fisher in 1964. In Ramey (2012), this algorithm was benchmarked against two popular ones for components labelling, flood-fill and Chang et al. (2010) method. It turned out to be 30% faster on the image collection used in the article. A sample is shown in Figure 2(f).

The method also returns the bounding boxes of the components. We recall here that the bounding box of a set of 2D points is the smallest rectangle that fits all of the points within its surface.

CLUSTER FILTERING

Some filters can be applied to remove some of the nonrelevant clusters found in the current depth map. The actual version of the algorithm discards the too small clusters. This is obtained by checking whether the size of the corresponding connected component is under a given threshold.

## Cluster Matching and Tracking

In the previous section, the detection of the different clusters in the current depth map was explained. However, our aim is to give temporal coherence

to this cluster information: we want to match these clusters to objects detected in previous images. For instance, if the robot is following a person, we want to spot the cluster corresponding to this person in the current image. Having this temporal coherence for objects gives the possibility of a meaningful chasing motion for the robot.

Let us define the concept of *object* in the scope of our detector: an object is a set of connected components of the successive depth maps, each of which corresponds to the same physical entity. There is, at most, one connected component in each depth map corresponding to a given object. However, it might be that some frames do not contain any connected component corresponding to a given object; that is, the object might be occluded or not successfully recognized.

For clarity of the concept, an example issued from real data is presented in Figure 3. It shows the recognition state of an object labeled 3. This object was recognized five times in the previous depth maps. For each frame, we have stored the connected component and the bounding box representing it. On the other hand, the object was not recognized in depth maps 9, 10, 11, 14, and 15.

In order for the computation time to be as short as possible, the recognition is made in two steps.

1. A first rough matching using only the bounding boxes of the components gives us a first estimation of which clusters correspond to which objects. This is explained in the next paragraph.
2. Ambiguities are solved using an analytic distance with strong discrimination properties, the Hausdorff distance.

A definitive match is obtained at the end of this second phase.

ROUGH ESTIMATOR: BOUNDING BOXES CORRESPONDENCES

Let us represent a given connected component $C$ of the current depth map. We want to obtain a quick estimation of what object is the most likely to be matched with $C$.



FIGURE 3 Example to clarify the concept of an object representation. The data are real and come from the tracking sequence of a given user. The object connected component is the filled area and the bounding box is marked as a rectangle.

Let us now consider a given appearance at frame $i$ in the past of an object $O$ in the recognition history, which we call $O^i$. We compare the bounding box $bb_O^i$ of this appearance with the bounding box $bb_C$ of $C$.

Definition of the bounding box distance: Let us define a bounding box distance $d_{bbox}$ such as $\forall X, Y$ bounding box, $d_{bbox}(X, X) = 0$, and the more similar $X$ and $Y$, the smaller $d_{bbox}(X, Y)$.

Let us write $bb_X = \{TL_X, BR_X\}$, $bb_Y = \{TL_Y, BR_Y\}$ where $TL$ refers to the top-left corner of the bounding box, and $BR$ refers to the bottom-right corner. Then, we define $d_{bbox}$ as defined in Eq. (1). This corresponds to the sum of the distance between corresponding corners of both bounding boxes.

$$d_{bbox}(X, Y) = d(TL_X, TL_Y) + d(BR_X, BR_Y) \tag{1}$$

A distance function between 2D points, however, needs to be chosen. Three usual choices are possible, as written in Eq. (2).

$$\begin{cases} d_{L_1}(a, b) = |a \cdot x - b \cdot x| + |a \cdot y - b \cdot y| \\ d_{L_2}(a, b) = \sqrt{(a \cdot x - b \cdot x)^2 + (a \cdot y - b \cdot y)^2} \\ d_{L_\infty}(a, b) = max(|a \cdot x - b \cdot x|, |a \cdot y - b \cdot y|) \end{cases} \tag{2}$$

Keeping in mind the speed of execution, the $L_1$ norm is here chosen. The (Euclidean) $L_2$ norm indeed needs expensive square root computations.

Application to the tracking: The result of $d_{bbox}(bb_O^i, bb_C)$ gives us an idea of the similarity of $O^i$—that is, the appearance of the object $O$ at frame $i$ in the past—and $C$—that is, the current connected component of the depth map. The smaller the value, the more likely it is that $C$ is an appearance of the object $O$ in the current frame. In addition, to take into account the age of this appearance of $O$, we weight the obtained $d_{bbox}(bb_O^i, bb_C)$ by the age (in seconds) of the depth map where $O^i$ belongs.

However, this estimation only takes into consideration the position of the objects in the frame and not their shape. For our given connected component $C$, we compare it to all of the appearances of all objects and store all of the obtained marks in a list. We then sort this list by similarity; that is, with the lowest $d_{bbox}$ distances first. Then, we use the precise estimator presented in the next section iteratively on each element of this sorted list. This gives us the final distance $d_{final}$ of $C$ with each object that appeared in the previous frames.[1]

---

[1]We can thus avoid the computation of $d_{final}$ for the majority of the object appearances. If, for a given object appearance, its $d_{bbox}$ is superior to the smallest $d_{final}$ computed at that time, we can stop comparing all of the following appearances in the list. They will indeed inevitably obtain a higher $d_{final}$ distance than this match.

Using the rough matching with bounding boxes, matching ambiguities can occur. For instance, two connected components, similar in size and close to the last apparition of a given object, could both correspond to the same object and then generate an ambiguous matching. To solve such ambiguities, we need a mathematical tool to properly compare the shapes of components.

Modified Hausdorff distance: In Dubuisson and Jain (1994), a comparison was made between the different ways of computing a distance between two sets of points. According to these findings, we use the distance $d_{22}$, defined as in Eq. (3).

$$\forall A, B \in \mathbb{R}^{2\mathbb{N}}, d_{22}(A, B) = max(d_6(A, B), d_6(B, A))$$
$$\text{with} \begin{cases} d_6(A, B) = \frac{1}{|A|} \sum_{a \in A} d(a, B) \\ d(a, B) = \min_{b \in B} \|a - b\| \end{cases} \quad (3)$$

For a given connected component $C$, $d_{22}(C, C) = 0$, and given two connected components $A$ and $B$, the lower the result returned by $d_{22}(A, B)$, the more similar they are.

$d_{22}$ requires the choice of a norm for estimating the distance between two points.

Accurate component comparator: The Hausdorff distance helps to solve ambiguities. For each frame, all of the components we first scaled to a given size, say $32 \times 32$ pixels.

Then, the different candidates for the same object were compared using $d_{22}$ distance. Similar to how we weighted $d_{bbox}$, $d_{22}$ is weighted by the age of the object appearance. The final distance is then given by:

$$d_{final}(O^i, C) = d_{bbox}(bb_O^i, bb_C) + d_{22}(O^i, C)$$

The best object candidate for the recognition of $C$ is the one that achieves the lowest $d_{final}$. It is considered a positive match if it gets a mark inferior to a given threshold. This mark is empirically determined and can be set through the graphical user interface. Its default value is 0.8.

OBJECT TRACKING: SELECTION OF THE OBJECT OF INTEREST

At this step, the connected components of the current depth map are matched to the objects already found. However, the tracking needs to know what object we aim to track. In other words, the user needs a way to select the object of interest. We will present a graphical user interface (GUI) that we have developed in the next section.

The 3D pose of the object of interest is obtained due to the 3D reprojection of the barycenter of all points of the object of interest. The cluster matcher periodically republishes this 3D pose of the tracked object of interest.

## Robot Motion and User Tracking

### MULTI SENSOR FUSION

As all range-imaging devices, the Kinect is limited by a field of view. Horizontally, it can see objects that belong to an angular domain of 57° and of 43° vertically. It also cannot detect objects at a distance less than to 1.2 m according to the Kinect datasheet. Furthermore, for greater distances, the further the object, the lower the precision. On the other hand, the robot is also equipped with a Hokuyo laser scanning range finder, which can only detect objects on a horizontal plane but with a field of view of 240°. However, its range is limited to approximately 4 m. Both fields of view are compared in Figure 4.

The poor visibility of the Kinect device results in objects that move laterally easily getting out of its view spectrum and hence getting lost by the tracking algorithm, while they remain visible to the Hokuyo laser. On the other hand, distant objects are out of range for the latter.



**FIGURE 4** Field of view of both sensors mounted on the robot: the Kinect depth camera and the Hokuyo range finder.

Thus, to enable robust tracking even with challenging trajectories, including curves with hairpin turns, we chose to combine two different tracking algorithms: the vision-based method presented before and another one based on the laser data only. The latter consists of a simple 2D cluster-tracking algorithm. A cluster is defined by a continuous set of points where each point is separated from its neighbors by a distance inferior to a given threshold, say 35 cm. The tracking is initialized due to a 3D seed point supplied to the tracker.

These two tracking algorithms run simultaneously on the robot. They periodically publish their status and the 3D pose of the object of interest obtained by the tracking (as defined previously). In addition, a *dialog* node decides what algorithm has priority and reinitializes the other due to the 3D pose it returns. When it performs successfully, the vision-based algorithm presented in this article always has priority over the laser-based one. As such, if the tracked object is located in the view frustum of the Kinect device, the former is used before the latter. The latter enables the tracking to continue successfully outside of the view frustum and recover from eventual failures of the vision-based tracking algorithm.

GOAL NAVIGATION

The dialog node republishes the resulting 3D pose of the tracking algorithm that has priority. This is used as a *goal* for the motion planning system: that is, the point the robot should reach. This goal is moved after each iteration of the tracking algorithms. This corresponds to each acquisition of a depth map by the range-imaging device.

The motion control is made through a dynamic window approach, as presented in Fox et al. (1997).

The admissible velocities—that is, the ones where the robot is then able to stop without collision—are determined with the local cost map. This local cost map is obtained via fusion of the laser scanning range finder and the reprojected Kinect point cloud.

When the robot is close enough to the goal, the robot keeps steady until the goal moves again. An example is shown in Figure 5.

EXPERIMENTAL RESULTS

The method explained before was implemented in the robot MOPI. The experiments were performed in several phases. First, the performance of the depth map compression is evaluated. Then, the GUI developed to select the object of interest is presented. The time needed for the algorithm to run on several hardware platforms is presented and discussed. Then, how the workload can be distributed between several computers is explained. Next,

(a)



(b)

**FIGURE 5** Motion planning toward the goal at a given time. (a) The detection status at a given time. The tracked object is the shape with a 0 index on the left of the bottom-right image; (b) Visualization of the plan of the robot. The white spots correspond to the laser scan. The light gray squares represent the local cost map, and the dark ones represent its inflated version (inflation by a radius corresponding to that of the robot). The robot is indicated by its white footprint and its frame axes. The goal is the sphere with the tilted axes on top of the image. The direction indicated by the arrow corresponds to the orientation of the Kinect within the robot. The dotted curved line corresponds to the planned trajectory until the goal.

a discussion on the advantages and limitations of both inpainting methods is provided, the accuracy of the whole tracking algorithm is measured. The success rate of the process is measured as the user goes across a marked path with obstacles and occlusions.

## Depth Image Compression Performance

The compression of depth maps for processing on a remote computer was explained previously. Some experimental results for the compression of depth maps are shown in Figure 6. They were obtained during one of the experimental runs when chasing the user.
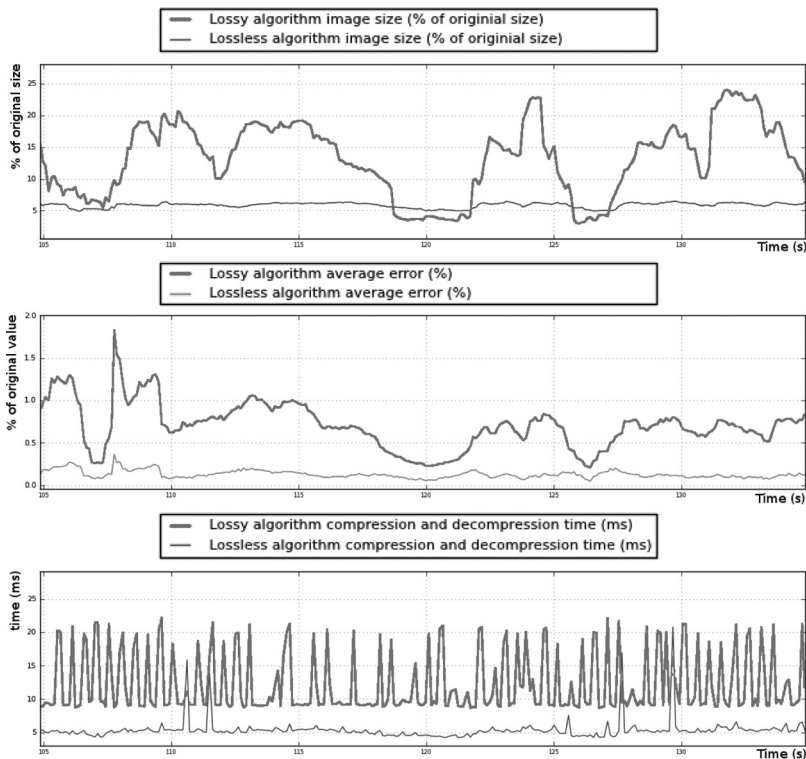
**FIGURE 6** Compression ratio, average error per pixel, and compression time for both lossless (`libpng`, compression factor of 6) and lossy (`libjpeg`, JPEG quality of 85) compressions. All of the curves are plotted against time during a typical tracking phase in an indoor environment.

The bandwidth required by the transfer of the Kinect depth maps to a remote computer typically drops from 9–10 MB/s to less than 300 kB/s. As seen in Figure 6, the ratio of the lossless compressed image to the original one is less than 10%; that is, the image size is reduced more than 10 times. The lossy algorithm requires the storage of the pixel indices with *NaN* value. As such, even if the image itself is smaller than that with the lossless algorithm, the amount of data to transmit is usualy great.

With a lossless algorithm, the compression generates an average relative error less than 1% for floating point values in the range of [1..10], which corresponds to typical values of the Kinect depth map. The results in Figure 6 show an average error on that run of less than 0.5% for the lossless algorithm. The lossy algorithm generates two successive approximations for the depth map: first the remapping to [0,255] values and then the information loss generated by the `libjpeg` compression algorithm. As such, its error rate is higher.

The compression times presented in Figure 6 were obtained with an AMD Athlon 64 X2 Dual Core Processor 5200+. With the embodied

computer of the robot, the compression of a single-channel $320 \times 240$ image requires around 15 ms. This enables a 30 Hz broadcasting (then requiring 30% CPU more or less).

Some input, output, and error values are shown in Figure 7.

## Selection of the Object of Interest Due to a GUI

A GUI interface was developed for the touch screen. It enables the user to select the object of interest by clicking on it. A sample is shown in Figure 8.



(a)

(b)                                    (c)

**FIGURE 7** Error visualization for data compression. No visible difference is appreciable to the human eye between the original image and the ones after compression and decompression. Each pixel corresponds to the error per pixel, as a precentage of the original value at that pixel. A black pixel corresponds to zero error and a white one to the scale indicated in caption. The average error (over the whole frame) is 0.11% for PNG compression vs. 0.65% for JPG. The maximum errors, reached at one pixel, were 0.41 and 18.81% respectively. (a) The original depth map without compression, as supplied by the Kinect device; (b) Visualization of the compression loss per pixel for the lossy compression. A white pixel is a relative error of 18% to the original value; (c) Visualization of the compression loss per pixel for the lossless compression. A white pixel is a relative error of 1.8% to the original value. Note the scale makes errors 30 times more visible than the other image using lossy compression.

**FIGURE 8** A sample picture of a user selecting the object of interest. It is indicated by the shape with a white border. The other objects are drawn with a different color.

## Time Costs for Clusters Detection and Tracking

The tracking algorithm has been tried in several hardware architectures. It was first used on the on-board computer of the MOPI robot, which is an embedded computer with a CPU with limited capabilities (Intel Atom CPU Z530 @ 1.60 GHz). In a second configuration, everything was processed on the touchscreen that was years old (Intel Pentium M @ 1000 MHz). Finally, a desktop PC with a full-power processor was used (AMD Athlon 64 X2 Dual Core Processor 5200 +).

The time needed to run the algorithm is shown in Figure 9. The on-board computer needs around 50 ms to perform a cycle of the algorithm. As such, it can run up to 20 Hz. However, it represents a workload for the CPU with limited capabilities. The use of a remote computer for processing solves this issue.

It would also be possible to run the algorithm using the touchscreen, reducing the number of devices to two. However, it is preferable to have a smooth GUI running at high frequency rather than a choppy display that might cause trouble to the user when selecting the object of interest. This would also reduce the battery autonomy of the touchscreen.

## Workload Distribution Among Several Computers

It is possible to make use of the communication strengths of ROS: several computers can share the same data and the different subtasks can be distributed between them, without affecting the result.

**FIGURE 9** Time needed to run the algorithm on different hardware platforms. The time allocated for each step is also indicated.
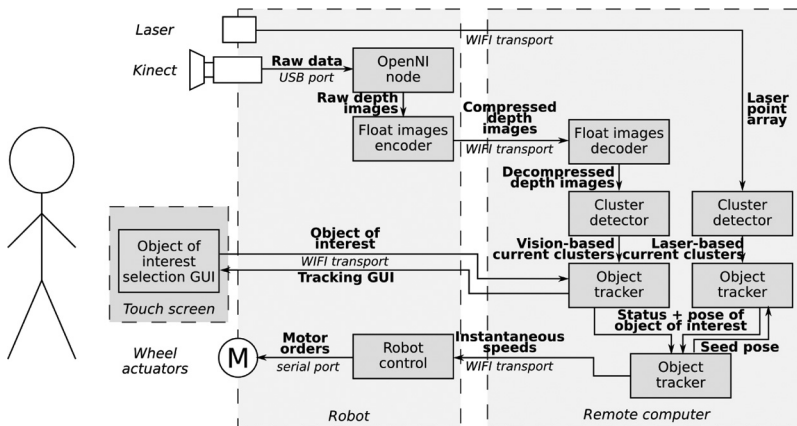


**FIGURE 10** Flowchart of the whole system when distributed between several computers. It also includes the laser data fusion presented in the subsection on robot motion and user tracking.

**FIGURE 11** Different methods for handling undefined *NaN* values. The difference between methods is especially visible on the neon lamp, to the right of the user's head. (a) The original image; (b) *NaN* reduction due to average border removal. The image is identical except that it is slightly more narrow. This is due to the removal of the right undefined border. Note the *NaN* values inside this border were not removed; (c) *NaN* removal due to inpainting; (d) *NaN* removal due to left-value propagation.

As such, it is possible to send the raw depth maps to a remote computer via a wireless connection. It will run the algorithm presented previously and return the 3D pose of the object if found.

The whole architecture for processing is then structured as presented in Figure 10. The task demanding much CPU time is the detection of the clusters. It was moved to a remote computer with a faster processor. The results of the processing are sent back to the robot embedded computer.

## Comparison of the Different Methods for the Undefined Values Handling

Experimental results for methods presented previously are shown in Figure 11.

- The average border removal presents the advantage of being extremely fast while removing most of the undefined values located near the borders

**FIGURE 12** The path followed by the user to test the tracking accuracy. The gray boxes correspond to cardboard boxes aimed at making the path planning harder and preventing the robot from cutting the curves. The robot intends to maintain the chasing distance as close as possible to the goal distance. When the robot enters the circle without losing the track of the user, the test is marked as successful.

of the images. However, it does not solve in any way the undefined zones in the middle of the depth maps, which can lead to objects being cut in two.

- In the inpainting algorithm, the output colors in the undefined zones are obtained by propagation of the values from all along the perimeter of the damaged area. As such, the of the damaged area. turns out to be a soft blending from these values. It especially smoothes the strong contrasts that might occur from one side of the damaged area to the other.
- On the other hand, in the directional value propagation, the same value is propagated. Hence, the strong edges are maintained. Furthermore, it is computationally less expensive than normal inpainting.

These reasons justify the choice of the directional value propagation in the final algorithm.

**TABLE 1** Result of the Tracking Runs along the Complicated Path

| | |
|---|---|
| Number of runs | 20 |
| Success rate | 90% |
| Average time for successful runs | 114 s |
| Standard deviation of average time | 12.8 s |

## Tracking Accuracy on a Complicated Path

The multisensor fusion was presented in the previous section. It gives a high priority to the presented vision-based algorithm and uses a laser-based fall-back tracking to enable successful tracking along a complicated path.

We measured the performance of the tracking system on a complicated path, here an unstructured lab environment. The path followed by the user is shown in Figure 12. The length of the whole path that the robot has to follow, from its starting position to the finish circle, is around 33 m. It includes straight lines, hairpin turns, and narrow passages.



**FIGURE 13** Map generated by a SLAM algorithm (GMapping; Grisettiyz 2005) overlaid with the paths generated during a run. The light gray area corresponds to the free space of the map and the black edges its occupied space. The irregular line is the path of the user, as detected by the algorithm. The smooth line ending in the rectangular shape corresponds to the path of the robot, as determined by its odometry. The remaining symbols are identical to the ones used in Figure 5(b).
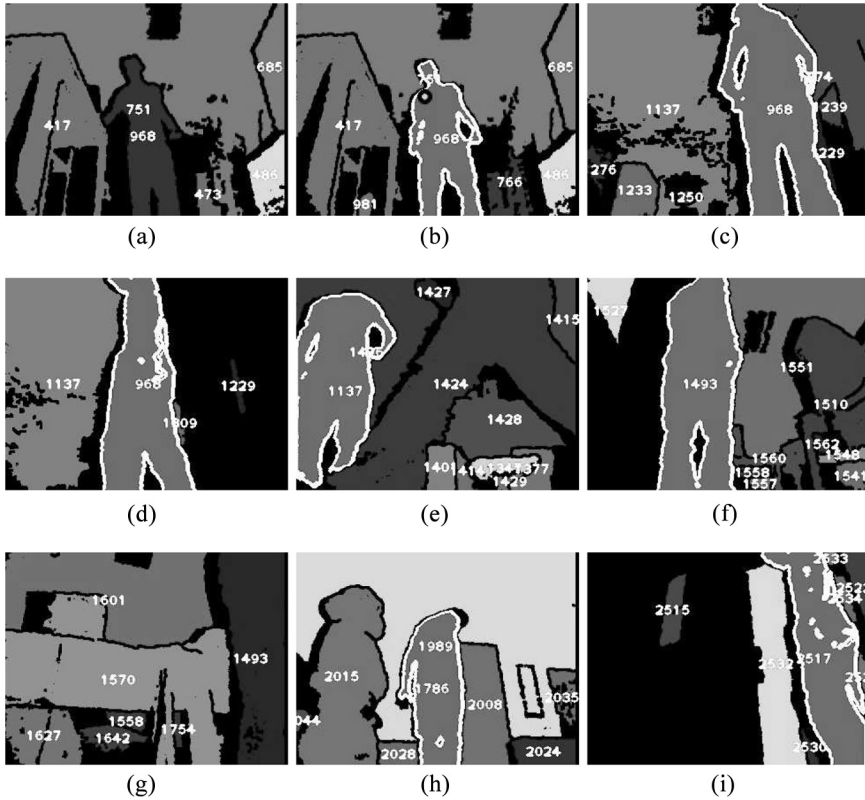
**FIGURE 14** Different frames of the GUI during a tracking sequence on the test path. The numbers stand for the objects, names. The object of interest has a white stroke; (a) GUI just before the user initialization and; (b) just after. The selected user then is marked with a white stroke (cluster 968); (c) Shortly after initialization, the robot has come next to the user (cluster 968), who already started walking along the path; (d) The user starts walking in a transverse direction to the robot motion, which triggers a fast on-place rotation; (e) Another user (cluster 1424) crosses the path of the robot. The tracking is not affected; (f) In the narrow passage between the tables (cluster 1510) and the wall; (g) The vision-based tracking loses track of the user (cluster 1493) after a sharp turn. The tracking goes on thanks to the laser-based tracking; (h) After the sharp S-turns at the end, the final door is visible (cluster 2008); (i) Approaching the final door (cluster 2532).

A user initializes the tracking on him due to the GUI presented earlier and, then follows the path without giving more orders to the robot until it reaches the finish line. A run is declared *successful* if the robot follows the user along the whole path, does not collide with any obstacle, and reaches the finish line. The experience was repeated for 20 runs.

The results are presented in Table 1. Some screenshots of the GUI obtained during a tracking sequence are shown in Figures 13 and 14.[2] Ninety

---

[2]The full output of the algorithm is visible in an on-line video at http://youtu.be/9xclN8ncYSY. It lasts about 2 h.

percent of the runs were finished successfully, on average in under 2 min. The cases of failure were most often due to a wrong match of the user's shape from one frame to the another. For instance, during one of the failed runs, the robot incorrectly recognized a passive observer as the tracked user and started tracking him. A standard deviation of more than 10 s can be observed. The time needed to go along the whole path depends on the speed of the user, which may vary from one run to another.

This high success rate experimentally validates the robustness and usability of the developed tracking system.

## CONCLUSIONS AND FUTURE WORKS

In this article, we presented a lightweight algorithm aimed at detecting and tracking 2D clusters using depth maps. This was applied to a mobile robot to chase a selected object of interest. A robust tracking system was built using a laser fallback when the selected object of interest remained out of the narrow field-of-view of the range-imaging camera.

In a near future, more extended measurement and user experiments were carried out. We especially want to measure the object of interest pose measurement error. This could be obtained by precisely measuring the position of the user due to a camera viewing from above and matching it against the pose computed by the algorithm.

The RGB information given by the Kinect device was discarded. We aim at experimenting some of the methods presented in the articles of the Introduction and measure the improvement obtained by coupling both methods.

## REFERENCES

Barber, R. and M. A. Salichs. "A New Human based Architecture for Intelligent Autonomous Robots." In *Proceedings of the 4th IFAC Symposium on Intelligent Autonomous Vehicles*, Asama, H., Inoue, H., Ed. Elsevier, 85–90, 2002.

Berliner, T., and Z. Hendel. *Modeling of Humanoid Forms From Depth Maps*. U.S. Patent Application 20100034457, 2007.

Bertalmio, M., and G. Sapiro. "Image Inpainting." In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, New York: ACM Press, 417–424, 2000.

Canny, J. "A Computational Approach to Edge Detection." *Pattern Analysis and Machine Intelligence IEEE Trans*. PAM 1-8, no. 6 (1986): 679–698.

Chang, S., S. Ham, and D. Suh. "ROHINI: A Robotic Flower System for Intuitive Smart Home interface." *2010 International Conference on Control Automation and Systems (ICCAS)*, Gyeonggi-do, Korea, October 27–30, 2010.

Dubuisson, M.-P. and A. K. Jain. "A Modified Hausdorff Distance for Object Matching." *Proceedings of the 12th IAPR International Conference on Computer Vision Image Processing*, Vol. 1, 566–8, 1994.

Fox, D., W. Burgard, and S. Thrun. "The Dynamic Window Approach to Collision Avoidance." *Robotics & Automation Magazine.* 4, no. 1 (1997): 23–33.

Galler, B. A. and M. J. Fisher. "An Improved Equivalence Algorithm." *Communication of the ACM.* 7, no. 5 (1964): 301–3.

Gerkey, B., R. T. Vaughan, and A. Howard. "The Player/Stage Project: Tools for Multi-robot and Distributed Sensor Systems." *Paper Presented at the 11th International Conference on Advanced Robotics*, Coimbra, Portugal, 317–323, 2003.

Grisettiyz, G. "Improving Grid-Based Slam with Rao-Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling." *IEEE Transactions on Robotics,* 23, no. 1 (2005): 34–46.

Howard, A. and L. H. Matthies. "Detecting Pedestrians with Stereo Vision: Safe Operation of Autonomous Ground Vehicles in Dynamic Environments." In *Proceedings of the 13th Int. Symp. of Robotics Research. Hiroshima, Japan*, edited by M. Kaneko and Y. Nakamura. Berlin: Springer, 2007.

Jia, S., L. Zhao, X. Li, W. Cui, and J. Sheng. "Autonomous Robot Human Detecting and Tracking based on Stereo Vision." Paper Presented at the 2011 IEEE International Conference on Mechatronics and Automation, 640–645, Beijing, China, August 7–10, 2011.

Latta, S. G. and K. Tsunoda. Gesture Keyboarding. *U.S. Patent Application* 12/391, no. 145 (2009).

Lee, C. H. "An Intelligent Depth-based Obstacle Detection System for Visually-impaired aid Applications." Paper Presented at the 13th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS), Dublin, Ireland, May 23–25, 2012.

Muñoz Salinas, R. "People Detection and Tracking Using Stereo Vision and Color." *Image and Vision Computing* 25, no. 6 (2007).

Muñoz Salinas, R. and E. Aguirre. "Multi-agent System for People Detection and Tracking using Stereo Vision in Mobile Robots." *Robotica* 25, no. 6 (2009): 99–1007.

Quigley, M., B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng. "ROS: An Open-source Robot Operating System." Paper presented at the ICRA Workshop on Open Source Software, Kobe, Japan, November 28–29, 2009.

Ramey, A. "Playzones : A Robust Detector of Game Boards for Playing Visual Games with Robots." Robot 2011—III Workshop de Robótica: Robótica Experimental Paper presented at Sevilla, Spain, November 28–29, 2011.

Ramey, A. "Playzones: A Robust Detector of Game Boards and its Application for Games with Robots." Master thesis, University Carlos III of Madrid, 2012.

Rivas, R. "Robot Skill Abstraction for ad Architecture." *6th IFAC Symposium on Intelligent Autonomous Vehicles* 47, no. 4 (2007): 12–13.

Rother, C., V. Kolmogorov, and A. Blake. "Grabcut: Interactive Foreground Extraction Using Iterated Graph Cuts." *ACM Transactions on Graphics* 23, no. 3 (2004): 309–314.

Tomori, Z. "Active Segmentation in 3D Using Kinect Sensor." Paper Presented at the 20th WSCG International Conference on Computer Graphics, Visualization and Computer Vision, Plzen, Czech Republic, June 25–28, 2012.