

Gesture control interface for immersive panoramic displays

Marcel Alcoverro · Xavier Suau · Josep R. Morros ·
Adolfo López-Méndez · Albert Gil ·
Javier Ruiz-Hidalgo · Josep R. Casas

© Springer Science+Business Media New York 2013

Abstract In this paper, we propose a gesture-based interface designed to interact with panoramic scenes. The system combines novel static gestures with a fast hand tracking method. Our proposal is to use static gestures as shortcuts to activate functionalities of the system (i.e. volume up/down, mute, pause, etc.), and hand tracking to freely explore the panoramic video. The overall system is multi-user, and incorporates a user identification module based on face recognition, which is able both to recognize returning users and to add new users online. The system exploits depth data, making it robust to challenging illumination conditions. We show through experimental results the performance of every component of the system compared

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 248138. This work has been partially supported by the Spanish Ministerio de Ciencia e Innovación, under project TEC2010-18094.

M. Alcoverro · X. Suau (✉) · J. R. Morros · A. López-Méndez
A. Gil · J. Ruiz-Hidalgo · J. R. Casas
Department of Signal Theory and Communications,
Universitat Politècnica de Catalunya, Barcelona, Spain
e-mail: xavier.suau@upc.edu

M. Alcoverro
e-mail: marcel.alcoverro@upc.edu

J. R. Morros
e-mail: ramon.morros@upc.edu

A. Gil
e-mail: albert.gil@upc.edu

J. Ruiz-Hidalgo
e-mail: j.ruiz@upc.edu

J. R. Casas
e-mail: josep.ramon.casas@upc.edu

to the state of the art. We also show the results of a usability study performed with several untrained users.

Keywords Interactivity · Panoramic display · Human-machine interfaces

1 Introduction

The need of gesture control interfaces emerges due to the increasing complexity of the functionalities of information systems and the current demand of more natural user interfaces. Commercially available display sets, for instance, offer larger degrees of freedom to manipulate the appearance of the displayed content. As the rendering capabilities of the terminal evolve, users tend to feel constrained by traditional interfaces, such as remote controls. This compromises the principle of a natural user interface [19, 20], which should become effectively invisible to its users as a communication modality.

A human gesture is a non-verbal, natural communication made with a part of the body [2]. Gestures may have several phases, and different ways for kinematic and functional coding [15], that can be tracked [33] and detected [17]. Gesture interfaces may not require physical buttons, multi-touch devices or second screens that could tangle up the users in the interaction process. In the particular case of user interaction with TV sets, the need for enabling interaction with new types of data, such as high-quality panoramas or 3D/FVV video, has fostered the study of new control modalities. Remote controls have two major drawbacks: they can be misplaced or out of reach of the user, and they require a lot of attention by the user as its control functionalities increase. Commercial innovations in the field of TV sets is focusing on natural user interfaces based on gesture and speech recognition [24] to overcome such drawbacks. Such commercial innovation follows studies from the generic field of Multimodal Human Computer Interaction (MMHCI).

In the last years, video transmission has experienced significant transformations. Internet technologies and new workflows are currently challenging traditional business models for the broadcast industry. New paradigms in the audiovisual sectors aim at offering new possibilities to the viewers from professional content providers.

Currently, viewers expect to enjoy video content in a variety of displays, such as high-resolution cinemas, TVs, tablet, laptops, smart-phones, etc. Furthermore, they are increasingly expecting to be able to control their audio-visual experience by themselves. This opens the path to going away from a fixed production and selecting one of several suggested regions of interest (ROIs), or even by freely exploring the audio-visual scene.

Combining advanced panoramic rendering and gesture recognition tools allows the viewer a larger degree of freedom to select the way the contents are displayed and to freely exploit the scene in a natural and non-intrusive way. The European project FascinatE [7] investigates a Format-agnostic Approach for Production, Delivery and Rendering of Immersive Media. Such approach requires a level of interactivity at the user end that may hardly be satisfied with remotes, speech commands or second screens, and will be used as experimental setup for the gesture control interface introduced in this paper.

With this purpose, we propose a novel interaction paradigm exploiting the combination of shortcuts to activate control functions, and a grasp-release navigation to explore a panoramic scene. In our approach, shortcuts are implemented with robust static gestures, whilst navigation is based on a fast hand tracking system. We show that the use of intuitive and easy-to-learn static gestures allows to include more functionalities to the system, while keeping the interface simple. In addition, navigation and static gestures are easily combinable, adding flexibility to the system design.

The paper is organized as follows. Next section gives an overview of the related work in the field of user interaction, commercial sensors and gesture interfaces. A proposed system design for a gesture control interface, with its component architecture and state diagram, is presented in Section 3. The technological system components are further detailed in Section 4, including Head Tracking, Hand Tracking, Gesture Localization and Face Identification. Section 5 introduces a setup for the user evaluation of the proposed system. Experimental results are provided for every main component of the system, evaluating accuracy and performance. Furthermore, overall results are provided by means of a usability study with untrained users. Finally, Section 6 outlines the conclusions and possible extensions of the proposed system.

2 Related work

MMHCI is a widely studied field and some surveys have been previously published [12, 30]. Moreover, MMHCI is at the crossroad of several related areas extensively studied such as face detection and identification [38], facial expression analysis [23], eye tracking [13], gesture recognition [34], human motion analysis [25] or audio-visual automatic speech recognition [26].

The recent commercialization of new game console controllers as Kinect or Wiimote, has been rapidly followed by the release of proprietary or third part drivers and SDKs suitable for implementing new forms of 3D user interfaces based on gestures [8]. On the one side, several authors propose gesture control interfaces based on accelerometers as the Wii controller [16, 29]. On the other side, the Kinect depth sensor allows for device-less interfaces. Some solutions as ZigFu [39] or GesturePak [11] use the skeleton tracking SDKs [14, 22] as input for gesture recognition based on skeletal poses. Thus, such approaches require a human pose estimation step, which is a complex task with high computational cost, often prone to errors in presence of clutter. Alternatively, several approaches make use of raw Kinect depth data as input for hand pose and gesture recognition methods [27, 28, 35]. In contrast to 2D color video, the use of depth data makes such methods robust to illumination changes and suitable for dark environments. Moreover, depth data provides 3D information valuable to account for scale invariance of human body parts. In general, these features make depth based methods perform better than its color-based counter part.

As introduced by Wachs et al. [37], gesture-based interfaces for entertainment applications, such as TV control, must address to major issues: *intuitiveness* and *gesture spotting*.

By *intuitiveness*, one means that the types of gesture selected should have a clear cognitive association with the functions they perform. However, a gesture natural to one user may be unnatural to others, due to the strong associations with cultural background and experience. Stern et al. [32] proposed a method to design and evaluate gesture vocabularies taking into account psycho-physiological measures (e.g. intuitiveness, comfort) together with machine factors (recognition accuracy). Alternatively, Nielsen et al. [18] propose a procedure to design the gesture vocabulary based on the Wizard-of-Oz paradigm. In a Wizard-of-Oz experiment, user interact with the system, but the response of the system is simulated by having a person respond to the user commands. In this way, the users are the ones who decide which gesture best represent its intentions.

Gesture spotting [37] consists of distinguishing useful gestures from unintentional movement. This problem may be afforded by the recognition technique, by performing a temporal segmentation to determine where the gesture start and ends. However, this is a difficult problem and often the recognition methods assume temporally segmented actions [5]. Also the recognition may require spatial segmentation of the body parts (e.g hands) which may be also a task prone to errors. To overcome this problem, López-Méndez et al. [17] focus on the problem of localization of static gestures on depth data. Their method learns the local appearance of gestures, and neither temporal nor spatial segmentation are required.

3 System design

The proposed system works in a device-less and marker-less manner allowing users to control the content on their TV sets only using their hands (Fig. 1). It is responsible of locating the people that want to interact with the system and of understanding and recognizing their gestures.

The current implementation is focused on controlling the content on a high definition TV screen situated in a home scenario. The content is composed of a panorama image (a high resolution view of the scene) together with several audio tracks and Regions of Interest (ROIs) associated with the panorama.

Fig. 1 Several users interacting with the proposed gestural interface



The current setup of the system allows the user to be standing or seating in a chair or coach. Although a single user is able to interact with the system at any given time, the interface is multi-user in the sense that several users can ask for control of the system and interact with it while the others might still be present in the scene (Fig. 1).

In order to allow the user to interact with the TV content, the system supports the following functionalities:

Menu selection	A menu is overlaid on the current screen and the user can select any button of the menus by pointing at it (Fig. 2).
Navigation	The user is able to navigate through the panorama scene by panning, tilting and zooming in the content.
ROI selection	The system informs the user of the available ROIs in the current view and the user is able to change between them.
Pause/resume	The user is able to pause / play the content at any time.
Audio volume	The user is able to increase or decrease the volume and to mute it completely.
Take control	The control of the system can be passed between users.

3.1 User experience and design aspects

The system design is a compromise between the best user experience and a feasible solution from a technical perspective.

To ensure that the system is always providing a responsive, convenient and intuitive experience for the user, the following design decisions have been adopted:

- Control functionalities of the system are activated with static gestures. These gestures should be easy to learn and perform, providing a fast way to do simple tasks (*shortcuts*).
- Navigation is performed in a grab-release manner. More precisely, the panoramic scene acts as a sheet which can be grabbed (closing one hand), moved (moving a closed hand) and released (opening the hand). This way, the interface simulates a virtual tablet in front of the user.

Fig. 2 Detail of the displayed content and the user feedback. The icon on the *bottom centre* (above the bar) is used to give feedback about static gesture commands (the icon shows that the user is lowering the volume). The bar on the *bottom* shows the 5 available static gestures



- The proposed system is touch-less and the user stands far from the display. For that reason, the interface must provide visual feedback to the user, providing information about the system state (e.g. the identified user, the detected gestures, or where the detected hands are located and pointing).
- To allow the social aspects of interaction with TV sets, the system provides an easy mechanism to release, take/switch the control between multiple users. The active user may freely move within the capture range.
- Panoramic navigation and gestures recognition modes are separated to improve user experience and at the same time reduce false/wrong detections.

3.2 Architecture

To fulfill all the functionalities, the current system implements the architecture depicted in Fig. 3. It is divided in three main layers. In the first layer, the **Capture** component is responsible of communicating with a single Kinect camera and feeding the images into the system. The Kinect sensors provide color and depth video with VGA resolution that are feed into all other components of the system. The second layer is the core of the interface. Here the processing, detection and recognition algorithms take place. It is composed of the following components:

The **Head tracking** (Section 4.1), where the depth image obtained in the capture module is analyzed to detect heads (oval areas of the same depth) [33]. The position of the heads are used in all subsequent components to locate possible users of the system. Other persons in the field of view of the Kinect sensor are not tracked and their gestures do not interfere with the system.

The **Face identification** component (Section 4.4) recognizes users. Faces are detected using a modified Viola-Jones detector [36] on the color images and recognized using a temporal fusion of single image identifications. The recognized users will determine the number of people able to control the system.

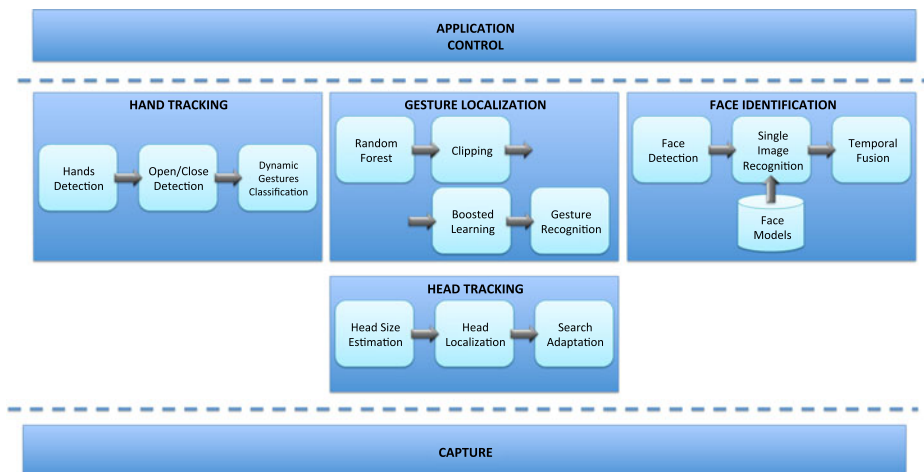


Fig. 3 Architecture of the gesture control interface

Once users are detected, hands are tracked by the **Hand Tracking** component (Section 4.2) using a 3D virtual box in front of the head of the user with control of the system. 3D blobs in the virtual box are segmented and treated as hands [33].

The **Gesture localization** component (Section 4.3) is responsible of detecting and classifying static gestures (gestures performed with still hands). In this case, the classification of the gesture is purely done using the shape and position of the hand, extracted using only the depth data provided by the Capture component. The classification is based on random forests [4], aiming at accurately localizing gesture and object classes in highly unbalanced problems.

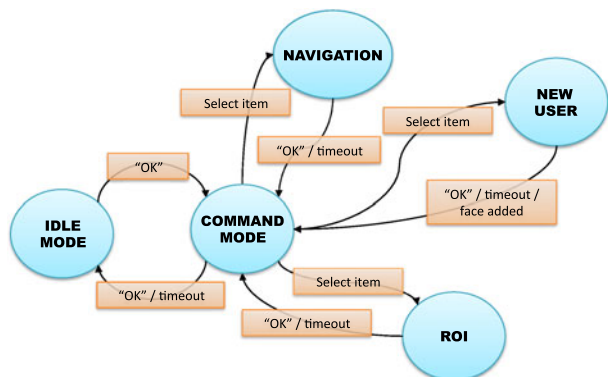
We remark that the functional distance range between the kinect and users is between 1m and 3.5m. If placed closer, it is difficult to have the user's head and hands framed. If placed farther, the Kinect resolution is not enough to ensure a good performance of the proposed system. Nevertheless, the proposed algorithms adapt to scale, providing a consistent performance within the suggested distance range.

Finally the third layer of the architecture contains the **Application Control** component. This module is responsible of acquiring all the detections and tracking information obtained by the components in the middle block (head, hands, user recognized and classified gestures) and mapping them into the functionality listed in the previous section. It communicates with the TV to perform all the needed interactions (select ROIs, change volume, pan-tilt-zoom the panorama, etc.). It controls the user interface and provides all the needed feedback through the GUI overlaid in the TV (Fig. 2).

3.3 State diagram

This section describes the state diagram of the current system and further explains the relation between the different components in the system architecture. Each state in the diagram allows users to perform specific gestures to control the system. The **Application Control** module is responsible of controlling the general status of the system and performing the transition between states. The end user does not need to be aware of the general state diagram or the current state as the **Application Control** component informs the user about the available commands at each stage through feedback in the GUI overlay. Figure 4 shows the different states possible. The arrows

Fig. 4 State diagram for the gestural interface



and yellow boxes indicate the gesture or timeout that will transition the interface from one state to the other one.

- The system always starts in *idle mode*. In this state no menu is shown in the screen and the only gesture available is the *OK* or *take control* gesture (Fig. 9). In this state, the Gesture Localization component only recognizes the *OK* gesture and all other components are suspended. If the *OK* gesture is detected by the Gesture Localization component, the user doing the gesture takes control and the interface transitions to the next state. In all following states only the user in control is recognized and other users might be present in the scene without affecting the system.
- The *command mode* is the principal state where users interact with the system. In this state all the components (Head and Hand Tracking, Gesture Localization and Face Identification) are started and report to the Application Control. The Head Tracking ensures that the user is followed and that all sub-sequent components only focus on the user controlling the system. The Hand Tracking tracks the position of a single hand which is used to access the menus of the GUI overlay. The selection is done by pointing a single hand to the menu item, wait for a few seconds until an arrow is shown, and then moving the hand in the direction of the arrow (like grabbing the menu item to the centre of the screen). The Face Identification recognize users and allows them to access specific options or interactions with the system. Finally, the Gesture Localization component recognizes 5 static gestures (Fig. 9) and the Application Controls maps them to the following functionalities:
 - The *OK* gesture (the same one the user used to take control of the system) is used to go back to *idle mode*.
 - The volume can be increased by locating the finger on top of the mouth and lowered by putting the hand on the ear.
 - The volume can be completely muted by doing a cross sign in front of the mouth.
 - The video can be paused or resumed by closing the hands together (similar to clapping).

By selecting any of the menus shown in the GUI overlay, the interface can transition to the states of navigation, ROI selection or new user. From all these states the user can go back to command mode by doing the *OK* gesture or by waiting several seconds without doing anything to trigger a timeout.

- The *navigation mode* allows the user to navigate freely through the panorama. In this mode no menus are shown in the screen and only the hands of the user are overlaid in the screen. The user can then pan and tilt the panorama by using one hand and grabbing the scene and zoom in or out by using both hands (such as done in maps applications on tablets). The Hand Tracking component is responsible of tracking this movements and, in this state, no Gesture Localization is performed as experimental results showed an increase of false/positive detections.
- The *ROI mode* allows the user to navigate through the available ROIs by moving your hand left to right or right to left. Also, the Hand Tracking algorithms is

- responsible to track the hands of the user to detect grabbings and movements in this state with no possibility of doing static gestures.
- Finally, the *new user mode* is used to add new users to the system. This mode creates new user models to be used in the Face Identification component.

4 System components

This section further describe each of the components (Head and Hand tracking, Gesture Localization and Face Identification) presented in the architecture of the system.

4.1 Head tracking

The head tracking algorithm is composed of three steps: head size estimation, head localization and a search area resizing (Fig. 5).

Head size estimation For every foreground pixel, an elliptical template (\mathcal{E}) of the size of a regular adult head (about 17×23 cm) is placed at the pixel's depth level d and projected onto the camera image plane, obtaining an ellipse of the apparent head size (H_x, H_y) in pixel dimensions (see Fig. 6). This ellipse is called *template* or \mathcal{E} is then used to find a head location estimate.

We assume that people interacting with the TV will either stand-up or be seated, keeping their head more or less vertical.

Head localization Aiming at finding the image region which better matches (\mathcal{E}), a matching score between (\mathcal{E}) and the global foreground mask (\mathcal{F}) is calculated at every pixel position (m, n) of the image. Matching is calculated within a rectangular search area of size $R_x \times R_y$, by sliding \mathcal{E} across the image. The matching score is calculated according to conditions C^k presented in (1), where $b = \text{background}$ and $w = \text{foreground}$. Conditions are checked at every pixel position $(u, v) \in \mathcal{E}$ of the template, which is itself centered at (m, n) . When a condition C^k is satisfied, $C^k = 1$, otherwise $C^k = 0$. The final matching score for the pixel (m, n) is




Fig. 5 Head matching score values using a Kinect sensor. Different situations are presented (from left to right): back view, side view (with slight head tilt), far person and long-haired person. In all these cases, the matching score presents a maximum in the head zone




Fig. 6 Head tracking snapshots. Head templates (*ellipses*), search areas (*rectangles*) and obtained head locations (*crosses*)

calculated as the sum of all the scores obtained on the template pixels, as shown in (1) (Fig. 5).

$$\begin{aligned}
 C_{u,v}^1 &: (\mathcal{E}_{u,v} = b) \wedge (\mathcal{F}_{u,v} = b) \\
 C_{u,v}^2 &: (\mathcal{E}_{u,v} = w) \wedge (\mathcal{F}_{u,v} = w) \wedge (|d_{u,v} - d_{H_i}| < d_{\max}) \\
 C_{u,v}^3 &: (\mathcal{E}_{u,v} = b) \wedge (\mathcal{F}_{u,v} = w) \wedge (|d_{u,v} - d_{H_i}| > d_{\max}) \\
 \mathcal{M}_{m,n} &= \sum_{\forall (u,v) \in \mathcal{E}} (C_{u,v}^1 + C_{u,v}^2 + C_{u,v}^3)
 \end{aligned} \quad (1)$$

The pixel (m, n) with a higher score \mathcal{M}^H is selected as the best head position estimation \hat{p}_H in a given search area. Conditions C^2 and C^3 provide robustness against clutter and partial occlusions, incorporating depth information to the matching score.

Search adaptation The position and size of the search area (R_x, R_y) is adapted to the head position variance (σ_x, σ_y) , and also to the confidence on the estimation $\bar{\mathcal{M}} = \frac{\mathcal{M}^H}{\mathcal{M}^{\max}} \in [0, 1]$, as described in (2).

$$\begin{aligned}
 R_x &= \sigma_x + (1 + \mu) \cdot H_x \\
 R_y &= \sigma_y + (1 + \mu) \cdot H_y
 \end{aligned} \quad \text{with} \quad \mu = e^{\frac{-\bar{\mathcal{M}}}{1-\bar{\mathcal{M}}}} \quad (2)$$

Such resizing is effective against fast head movements. For example, horizontal movements will enlarge the search area along the horizontal axis. Furthermore, including the matching score in (2) makes the system robust against noisy estimations.

4.2 Hand tracking

Hand detection The hand tracking system relies on the robust head estimation presented in Section 4.1. Hands are supposed to be *active* in a space placed in front of the body. We define a HandBox Ω as a virtual 3D box, which is attached to the head position \hat{p}_H so that it follows the user's head at every time instant, as shown in Fig. 7.

Dense clusters are searched in Ω by means of a *kd-tree* structure, which allows fast neighbor queries in 3D point clouds [9]. A list of candidate clusters is obtained and filtered according to the following criteria:

- Merging** Two clusters are merged as a single cluster if the Hausdorff distance between them $< \delta_{\min}$
- Size filtering** The resulting merged clusters are filtered by size, keeping the largest ones ($size > s_{\min}$) as hand candidates.

Depth filtering Clusters that fulfill the previous criteria are sorted by depth, keeping those closer to the camera (up to two).

Thresholds δ_{\min} and s_{\min} should be tuned depending on the type of camera and scene. For example, two hands are being detected in Fig. 7.

Open/closed hand detection The interactivity of the viewer with the rendering node may be increased by determining whether the user opens or closes hands.

We propose to compute the area of the detected hands (Section 4.2) and threshold it to quickly decide if it is closed or not. Such strategy assumes that the observed area is reasonably perpendicular to the camera. Therefore, the depth of each pixel may be replaced by the mean depth of the observed region, simplifying the physical area calculation.


The physical area of an open hand perpendicular to the camera is about 70–90 cm², whilst that of a closed hand is about 20–30 cm². It seems reasonable to set a threshold of 50 cm² to determine the hand status.

Dynamic gestures The pan, tilt and zoom angles of the panoramic viewport are controlled through gesturing. Navigation across the panoramic video is performed by a series of gestures consisting of grabbing (close hand), moving (while hand closed) and releasing, as if the screen was a piece of tissue. The pan and tilt angles are calculated depending on the current position of the hand in the HandBox Ω .

For the zooming case, once grabbed (both hands closed), the distance between hands is calculated. The zoom angle is proportional to that distance.

In addition, a family of dynamic gestures based on trajectory is included. In this setup, for example, horizontal movements with a considerable speed (user-defined) are utilized to shift between ROIs in the ROI mode. Temporal segmentation of these gestures is performed using the entry and exit points in the HandBox, and also assuming a low initial speed if the hand was already in the HandBox.

Fig. 7 Caption of the proposed approach, where both hands are detected inside the (green) HandBox



4.3 Gesture localization

The gesture localization builds upon the method proposed by [17], which is based on class-specific (one-vs-all) random forests using depth data. Random forest localization in depth data renders an efficient localization algorithm that can operate in real-time under realistic conditions. This implies detecting gestures with high accuracy in the presence of clutter and unintentional motion.

The localization problem is challenging due to two main reasons. Firstly, there is a high unbalance between gesture and non-gesture classes, since non-gesture classes contains the set of other gestures, as well as clutter and unintentional poses. Secondly, without accurate segmentation, it is difficult to accurately model the appearance of clutter using a limited number of training samples. The gesture localization approach used in our system addresses these issues with two main components:

- **Clipping:** In order to better capture the local appearance of static gestures, we automatically clip the depth data in a local vicinity of the training samples. In this manner, training is less prone to over-fitting problems caused by learning the specific backgrounds of training data.
- **Boosted learning:** We use a specific learning approach to deal with the high unbalance of training data [17]. We take advantage of the learning strategy on random forests to efficiently mine the most meaningful samples of the negative class.

For the sake of completeness, in the following of this section we summarize some details of the approach proposed in [17] (extended version currently submitted to IEEE Trans. MM).

Clipping binary tests Random forests [4] are an ensemble of m randomized trees, which are binary trees. Each tree is trained separately with a small subset of the training data obtained by sampling with replacement. Learning is based on the recursive splitting of training data into 2 subsets, according some binary test f and a threshold θ . The binary test is a function of the feature vector \mathbf{v} obtained from each training example. At each node of the tree, a test f and a threshold θ are randomly generated, and the one that maximizes an information gain criteria is selected.

We define our binary tests based on [31], but we introduce an auxiliary parameter that clips the depth of the available training examples. In such manner, tests are more robust to changes in background, while avoiding a segmentation step. Specifically, the clipping parameter is a value that represents the maximum and minimum relative depth with respect to the depth value of the center pixel. Formally, let κ denote the clipping parameter. Then, for a given pixel \mathbf{x} the test f has the following expression:

$$f_{\theta}(\mathbf{I}, \mathbf{x}) = \max \left(\min \left(d_{\mathbf{I}} \left(\mathbf{x} + \frac{\mathbf{u}}{d_{\mathbf{I}}(\mathbf{x})} \right), d_{\mathbf{I}}(\mathbf{x}) + \kappa \right), d_{\mathbf{I}}(\mathbf{x}) - \kappa \right) \\ - \max \left(\min \left(d_{\mathbf{I}} \left(\mathbf{x} + \frac{\mathbf{v}}{d_{\mathbf{I}}(\mathbf{x})} \right), d_{\mathbf{I}}(\mathbf{x}) + \kappa \right), d_{\mathbf{I}}(\mathbf{x}) - \kappa \right) \quad (3)$$

where $d_{\mathbf{I}}$ is the depth map associated to image \mathbf{I} and \mathbf{u} and \mathbf{v} are two randomly generated pixel displacements that fall within a patch size. Pixel displacements are normalized with the depth evaluated at pixel \mathbf{x} in order to make the test features invariant to depth changes.

Boosted learning of random forests In the localization problem posed in this paper, gesture and non-gesture classes are naturally unbalanced. On the one hand, in real applications users are not constantly performing gestures. On the other hand, the actual appearance of a gesture may be represented by a relatively low number of pixels. Summing up, the distribution of gesture classes (*positive*) with respect to non-gesture (*negative*) is biased towards the latter. This unbalance makes that low false positive rates constitute actually a large number of false positive votes. Taking into account this phenomenon is important during the training phase of a random forest since, under such unbalance, it will be difficult to optimize the information gain.

In order to overcome this problem, we adopt a boosted learning scheme, designed in a tree-wise manner as follows:

We train the first tree with a balanced set of samples from each class. Once the tree is trained, we evaluate it against the out-of-bag set [4]. The wrongly classified samples are added to the training set of the second tree (up to a maximum number of training samples). This new training set is completed by sampling with replacement from the full training set until balance is achieved. We train the second tree with this training subset and we repeat the process until the forest is fully trained.

Gesture localization For gesture detection and localization, a set of patches are provided to the detection forest, which casts a vote whenever a *positive* class has more probability than the *negative* class and other *positive* classes. Figure 8 illustrates the casted votes for a positive class in a class-specific learning example. To detect a gesture, we first estimate a probability density using the votes within a frame and we take into account temporal consistency by recursively updating this distribution with votes aggregated from past time instants. In order to construct the probability density, we use a Parzen estimator with Gaussian kernel. In order to account for the time component of the approximated density, we sequentially update such density $p(c|\mathbf{I}_t)$ as follows:

$$p'(c|\mathbf{I}_t) = \alpha p(c|\mathbf{I}_t) + (1 - \alpha) p'(c|\mathbf{I}_{t-1}) \quad (4)$$




Fig. 8 Gesture localization with random forests (best viewed in color). **a** A number of votes (*green dots*) are casted for the target gesture **b** votes are aggregated to estimate a probability density (overlaid in *red* on the input depth map) and a localization is estimated (*green square*)

This is a simple yet effective method to keep temporal consistency of the casted votes, as it requires storing a single probability map. An adaptation rate $\alpha = 0.8$ works well in practice, as it prevents several false positives while avoiding a delayed response.

Finally, we compute the pixel location \mathbf{g}_c of a gesture class $c > 0$ as the pixel location with maximum probability. We ensure that such a maximum represents a target gesture by thresholding the probability volume V computed by locally integrating the estimated pseudo-probability measure:

$$V = \sum_{\mathbf{x} \in \mathcal{S}} p'(c|\mathbf{I}_t(\mathbf{x})) \quad (5)$$

where \mathcal{S} is a circular surface element of radius inversely proportional to the depth, and centered at the global maximum. In this way, the localization is depth-invariant.

For efficiency reasons, we approximately segment the scene into relevant and non-relevant pixels by thresholding the depth using two values $z_{\text{near}} = 0.8$ m and $z_{\text{far}} = 3.5$ m. Once a user get the control by performing gesture *OK*, we employ the head tracking algorithm result (Section 4.1) to compute a region of interest, where the other 4 gesture class-specific forests are evaluated. In this way, we can run detection forests in real-time. Besides, the region of interest helps in discarding gestures of interest that other users may perform, thus increasing the detection accuracy (Fig. 9).

4.4 Face identification

The purpose of face identification is to allow users to select preferences based on the identity and to enable the system to establish a hierarchy of users that can be used in applications such as parental control.

The continuous monitoring environment offer the possibility of video-based face recognition. This permits improving the performance of face recognition from single image captures.

The workflow of face ID system is the following: faces are detected in first place and forwarded to the face recognition module. The face ID compares each test image with a set of models consisting of images of each person, created off-line with images from previous recordings. For each test face, a set of scores defining the probability it represents the persons in each model is produced. Finally, the fusion module combine the individual results for a temporal group of images of the same person into a final decision.




Fig. 9 Examples of successful localization results using our approach ($\kappa = 15$ cm) *volume down*, *volume up*, *Take Control*, *mute* and *pause*

In the following, we will analyze each module (detection, still image identification and temporal fusion).

Face detection The face detection module is cascaded with the head tracking module, described in Section 4.1. Face detection is only performed inside a reduced region defined by the head tracker. This allows reduced computational complexity and a very low probability of false positives.

We use the OpenCV [3] face detection module that relies on the adaboosted cascade of Haar features, i.e. the Viola-Jones algorithm [36]. In our application, the users generally interact with the system facing the camera. The system is thus restricted to detect frontal faces, allowing for some pose variations.

Single image face recognition For face recognition, a set of meaningful face characteristics are first extracted to form a feature vector. This vector should convey all the relevant information in a face. A feature extraction technique based on Local Binary Pattern (LBP) [21] is used. LBP operator is a non-parametric kernel which represents the local spatial structure of an image. It provides high discriminative power for texture classification and a good amount of illumination invariance, because it is unaffected by any monotonic gray-scale transformation which preserves the pixel intensity order in a local neighborhood.

At each pixel position, $LBP_{P,R}$ is defined as an ordered set of binary comparisons between the intensity of the center pixel and the intensities between the center pixel and its P surrounding pixels, taken over a circumference of radius R. The decimal form of the resulting LBP code can be expressed as follows:

$$LBP(x_c, y_c) = \sum_{n=1}^P s(i_n - i_c) 2^n \quad (6)$$

Where i_c corresponds to the grey value of the center pixel (x_c, y_c) , and i_n to the grey values at the P surrounding locations. Function $s(x)$ is defined as:

$$s(x) = \begin{cases} 0 & x < 0 \\ 1 & x \geq 0 \end{cases} \quad (7)$$

Our method is based in [1] For each pixel of the luminance component of the face, its $LBP_{8,2}$ representation is computed, that is, using 8 samples taken over a circumference of radius 2 from the pixel. The resulting transformed image is partitioned into non-overlapping square blocks of 7×7 pixels. The feature vector is formed by concatenation of the histograms of the LBP values of each block. This results in 49 local 64 bins histograms (face image was resized to 49×49). The final feature vector dimensionality is thus 3,136.

A k nearest neighbor (kNN) classifier [6] is used to obtain the final decision. In our system kNN is modified so that a vote is penalized by the inverse of the distance between the test vector \mathbf{t} and the selected neighbor, in such a way that the bigger the distance, the more penalization. The k_i inverse distances are normalized by the overall sum of the k nearest neighbor inverse distances of all classes, leading to the a posteriori probability for each class [6].

These probabilities for each class allow the classifier to give an estimation or score on how confident the classification of the test vector is.

The Chi-Square distance (8) is used to compare feature vectors because it is shown that it performs better than the Euclidean one when using histogram based features.

$$\chi^2(\mathbf{v}_1, \mathbf{v}_2) = \sum_i \frac{(v_{1,i} - v_{2,i})^2}{v_{1,i} + v_{2,i}} \quad (8)$$

where \mathbf{v}_1 and \mathbf{v}_2 are feature vectors.

Models for all individuals in the database should be created off-line, by using sequences of images collected in a different session than the testing recordings. 150 training images per subject are used, extracted automatically from small ad-hoc recordings.

Temporal fusion As face detection is cascaded with head detection and this module performs a temporal tracking of the detected heads, we can ensure that the identity of the person is maintained along the track, so that video based recognition can be used. Temporal fusion combines the information of several images to perform the recognition.

After face detection, each track is composed of T consecutive face images of the same individual. The face identification algorithm is used to compute a vector of scores \mathbf{s}_i for each single frame. A simple score fusion scheme is used, based on averaging the scores along the track. The instantaneous final decision is computed by taking an average of the last N_s score vectors and selecting the class with highest score.

5 Experimental results and user evaluation

The evaluation of the proposed system is divided into two different parts. In the first one, an objective evaluation of the different technical components is performed. In this case, all head and hand tracking, gesture localization and face identification are evaluated using objective metrics against other similar methods from the state of the art. In the second part, the entire system is evaluated subjectively by end users with low or no prior knowledge of the system.

For the experimental results, the proposed system has been implemented and a complete gesture control interface for panoramic TV content has been created as a standalone demonstrator. As described, the demonstrator is composed of multiple algorithms that must be connected in different ways; some can run independently in parallel, others need some signaling and synchronization and others must run serially after another one. Therefore, careful signaling must be shared between them. The demonstrator is implemented in the following hardware:

- A Microsoft Kinect sensor to capture depth and color video
- A laptop with 8 CPUs (at least 6 are needed) and 8GB of RAM

It uses the Debian operating system and the following open-source libraries:

- OpenNI as capture drivers for the Kinect sensor
- SmartFlow as a transport and communication middleware
- OpenCV as a support for the face detection algorithm

- PointCloudLibrary (PCL)
- Boost C++ Libraries

All algorithms described in Section 4 (Head and hand tracking, gesture localization, face recognition and application control) are implemented in C++ in an internal library.

5.1 Objective evaluation of technical components

This section evaluates the performance of each individual component presented in the system.

Head tracking evaluation We analyze the convenience of conditions C^2 and C^3 in (1), and how robustness is increased by including depth data queues. We compare the proposed scheme with a limited version which only verifies condition C^1 . This way, the contribution of C^2 and C^3 is shown. The head location error between a ground-truth (manually marked) head position g^H and the estimated head location is calculated as $\varepsilon = |g^H - \hat{p}^H|$, and presented in Fig. 10 for the two versions of the algorithm. Note that, even if C^1 plays the role of 2D method, depth data is used for the initial head size estimation.

Figure 10 shows the frame by frame error of the $C^1 + C^2 + C^3$ compared to C^1 . The C^1 version loses target twice, a reset of the algorithm being needed (frames 74 and 398). The labeled arrows in Fig. 10 correspond to adverse clutter and occlusion situations. Our proposed algorithm presents an error that does not go above 10 pixels, which is about the head radius.

Hand tracking evaluation Table 1 summarizes the average error for different one-hand and two-hand trajectories, computed as the average 3D error (with respect to

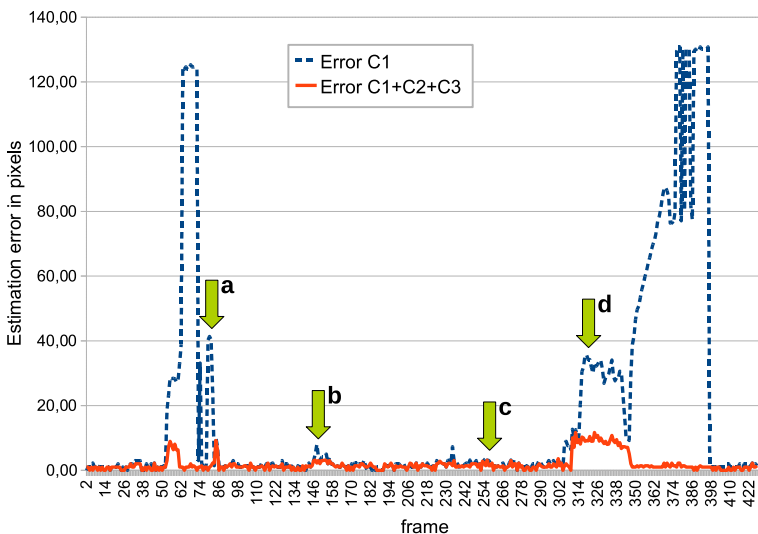


Fig. 10 Error between the obtained head location estimates and the ground-truth

Table 1 Hand detection 3D accuracy on different gestures

Trajectory	# frames	Error R hand (cm)	Error L hand (cm)
Push	30	2.62	–
Circle	30	6.61	–
Replay	35	2.86	–
Hand up-down	115	5.87	–
Separate hands	75	2.36	3.80

the ground-truth hand positions) along the duration of the movements. Ground-truth trajectories have been extracted by hand, selecting a reasonable hand center which may vary in some cm along frames. Despite these adverse noisy conditions, the hand estimation error rarely goes above 10 cm.

The average error is higher for fast movements, resulting in about 6 cm of error. For the other gestures, the error is of about 3 cm, which is fairly adequate given the size of a human hand. For the sake of illustration, a sample of the zooming gesture is included in Fig. 11, showing the detection error between the ground-truth trajectory and the detected one.

Gesture localization evaluation We conduct experiments of the gesture localization method to provide a quantitative performance of the approach and determine optimal clipping parameters. We focus on the gestures defined in Fig. 9. We record 5 training sequences where 5 different actors perform a set of gestures and actions, including the 5 target gesture classes. Additionally, we record 6 test sequences with 4 additional actors that were not included in the training set.

The employed detection forests have 15 trees with maximum depth 20, and each tree is trained with approximately 20,000 examples per class.




Fig. 11 Ground-truth and estimated trajectories of the (R)ight and (L)eft hands. The estimated hand positions are nicely close to the reference ground-truth positions. Only the XY projection of these 3D trajectories is shown

Table 2 Average Area Under the Curve (AUC) for different learning approaches

Method	Volume down	Volume up	OK	Mute	Pause	Average
Boosted [10]	0.17	0.10	0.31	0.14	0.02	0.15
Ours	0.17	0.10	0.31	0.20	0.25	0.21
Ours + κ = 15 cm	0.58	0.53	0.81	0.47	0.11	0.50
Ours + κ = 30 cm	0.69	0.50	0.54	0.27	0.14	0.42
Ours + κ = 50 cm	0.59	0.46	0.61	0.24	0.18	0.41

Bold entries correspond to the best method for each column key

In all cases, we employ squared patches of size 85×85 pixels. We train class-specific forests with the boosted learning method. Additionally, in order to compare the learning method, we implement a boosted approach based on [10].

To measure the accuracy of the proposed methods, we consider a correct localization if the estimated gesture and the actual gesture belong to the same class and if the estimated location is within a radius of 10 pixels. We compute the curves representing 1-precision vs recall to then compute the Area Under the Curve (AUC). Average AUC's per gesture class are shown in Table 2.

The comparison of different training approaches with and without clipping (κ parameter) shows that the proposed boosted learning yields an overall better performance. The proposed learning approach outperforms the boosting scheme proposed in [10].

Face identification evaluation Evaluation of face recognition has been done at two levels. In the first place, the single-frame face ID algorithm has been tested. This measures the ability to recognize persons given a unique test frame. Then, we evaluate the performance using the fusion algorithm, where the use of face tracking allows combining consecutive individual results into a more robust decision.

The system is designed to work in small groups of people that share the access to a TV set (for instance, a family). Having this in mind, 12 individuals have been recorded while using the demo in two sessions: the first one is used to create the models while the second one is used for testing. As some gestures involve putting the hands between the face and the camera, there are a good number of partial face occlusions. Pose is mostly frontal but with variations as the users look at the different corners of the screen. There is also a medium degree of expression variability since the users are asked to behave normally while using the system.

Table 3 shows the results either for single frame and for temporal fusion face ID. Due to occlusions and the restrictive settings, there are good number of frames where no faces are detected. This is not a problem in this setup as a decision is only provided at regular intervals (currently, each one or two seconds). It can be seen that even using single frame, the identification results are very good, with only 37 errors in 6,605 detected faces. When using the temporal fusion, the identity of the person in each segment is always given correctly.

Table 3 Face ID results

	# frames	Detected	Error (%)
Single frame	8,716 frames	6,605 frames	0.56
Temporal fusion	–	266 segments	0.0

Table 4 Usability questionnaire

Questions	Mean	Standard deviation
1. I liked using the interface	4.13	0.96
2. The interface was pleasant to use	3.81	0.98
3. It was simple to use this system	4.00	0.73
4. It was easy to learn to use the system	4.31	0.79
5. The organization of information presented by the system was clear	4.25	0.68
6. The information provided by the system was easy to understand	4.50	0.52
7. Navigation in panoramic video was pleasant to use	3.47	0.99
8. Navigation through interface menu items was comfortable	3.94	1.00
9. The system responded effectively to my gestures	3.19	0.98
10. Task 1 (zoom into a specific part of the scene) was easy to perform	3.88	1.31
11. Task 2 (change audio volume) was easy to perform	4.13	1.36
12. Task 3 (region of interest selection) was easy to perform	4.00	0.89
13. I felt comfortable using the system	3.63	0.89
14. Overall, I am satisfied with this system	3.69	0.87

Sentences and Mean and Standard deviation of the responses of 15 participants (*stronglydisagree* = 1, . . . , *stronglyagree* = 5)

Tests performed with the system with more users show that good results (above 96 % recognition rate) can be obtained for up to 50 people.

5.2 Usability evaluation

In order to validate the design of the interface we conducted a usability study. We selected a set of participants, who were asked to interact with the system and then they answered a questionnaire. This usability study is derived from similar studies that have been recently conducted to evaluate a gesture controlled interface for elderly in [2].

The procedure was the following:

1. First, users watched a short video about how to use the interface.¹
2. Then, they interacted freely with the system to familiarize with its functionalities.
3. They were asked to perform 3 specific conceptual tasks: To change the volume (up, down and mute); To zoom into a defined part of the scene; To select certain ROI.
4. Finally, they answered the questionnaire.

In total, the number of participants were 15. The questions had 5 answer choices according to the degree of agreement with the sentence (*Strongly Disagree* = 1, . . . , *Strongly Agree* = 5). The sentences and the mean and standard deviation results are shown in Table 4.

In general, users were satisfied with the system and they could perform effectively the assigned tasks in a short amount of time. The most successful aspects of the user experience are related to the use of static gestures to activate several control

¹The video is available at <http://vimeo.com/55432492>.

functions. As reported by participants, the system is simple to use, easy to learn and the GUI is clear. From their response and according to our observation of the participants, we can conclude that the selected gestures are intuitive and easy to perform. Moreover, the use of static gestures as *shortcuts* to perform actions reduce the amount of information required in the GUI, improving its clarity. Navigation through panoramic video was not so successful. The reason seems to be due to the fact that the interaction paradigm chosen for zooming and panning (i.e. hand gestures with hands open or closed to simulate touching a virtual tablet in front of the user) results not being comfortable and intuitive by participants. Further research is required to improve the usability of the system and such study offers insights into a better user experience.

6 Conclusion and future work

This paper has presented a gesture recognition interface designed to control a TV set providing immersive viewing functionalities, such as navigating a panorama. It works in a device-less and marker-less manner allowing users to control the content on the display by only using their hands. It automatically locates the people that want to interact with the system and understands and recognizes a small set of natural and easy-to-learn gestures. An objective evaluation of the technologies integrated in the system show results competitive with similar techniques in the state of the art. The complete integrated system has been evaluated subjectively by a group of 15 users which where overall satisfied by the simplicity and intuitiveness of the interface.

Our future work include the extension of the gesture localization part to successfully detect and recognize fingers, in order to explore ways to include finger gestures within the pre-defined gesture database. Performing further user evaluation studies, we plan to improve the usability of the navigation through the panoramic video in order to improve the responsiveness of the system and to limit the fatigue caused to the users.

References

1. Ahonen T, Hadid A, Pietikainen M (2006) Face description with local binary patterns: application to face recognition. *IEEE Trans Pattern Anal Mach Intell* 28(12):2037–2041. doi:[10.1109/TPAMI.2006.244](https://doi.org/10.1109/TPAMI.2006.244)
2. Bhuiyan M, Picking R (2011) A gesture controlled user interface for inclusive design and evaluative study of its usability. *J Softw Eng Appl* 4(9):513–521
3. Bradski G (2000) The OpenCV library. Dr. Dobb's journal of software tools
4. Breiman L (2001) Random forests. *Mach Learn* 45(1):5–32
5. Demirdjian D, Varri C (2009) Recognizing events with temporal random forests. In: *Proceedings of the 2009 international conference on multimodal interfaces, ICMI-MLMI '09*. ACM, New York, pp 293–296. doi:[10.1145/1647314.1647377](https://doi.org/10.1145/1647314.1647377)
6. Duda R, Hart P, Stork D (2001) *Pattern classification*, 2nd edn. Wiley, New York
7. Fascinate: Format-agnostic script-based interactive experience. <http://www.fascinate-project.eu/>. Accessed 28 Feb 2013
8. Francese R, Passero I, Tortora G (2012) Wiimote and kinect: gestural user interfaces add a natural third dimension to hci. In: *Proceedings of the international working conference on advanced visual interfaces, AVI '12*. ACM, New York, pp 116–123. doi:[10.1145/2254556.2254580](https://doi.org/10.1145/2254556.2254580)
9. Friedman JH, Bentley JL, Finkel RA (1977) An algorithm for finding best matches in logarithmic expected time. *Trans Math Softw* 3(3):209–226. doi:[10.1145/355744.355745](https://doi.org/10.1145/355744.355745)

10. Gall J, Yao A, Razavi N, Van Gool L, Lempitsky V (2011) Hough forests for object detection, tracking, and action recognition. *IEEE Trans Pattern Anal Mach Intell* 33(11):2188–2202. doi:[10.1109/TPAMI.2011.70](https://doi.org/10.1109/TPAMI.2011.70)
11. Gesturepak: Gesture recording and recognition toolkit. <http://www.franklins.net/gesturepak.aspx>. Accessed 20 Feb 2013
12. Jaimes A, Sebe N (2007) Multimodal human–computer interaction: a survey. *Comput Vis Image Underst* 108(1–2):116–134. doi:[10.1016/j.cviu.2006.10.019](https://doi.org/10.1016/j.cviu.2006.10.019) (special issue on vision for human-computer interaction)
13. Ji Q, Wechsler H, Duchowski A, Flickner M (2005) Editorial: special issue: eye detection and tracking. *Comput. Vis. Image Underst.* 98(1):1–3. doi:[10.1016/j.cviu.2004.07.006](https://doi.org/10.1016/j.cviu.2004.07.006)
14. Kinect for windows sdk. <http://www.microsoft.com/en-us/kinectforwindows/develop/>. Accessed 20 Feb 2013
15. Lausberg H, Sloetjes H (2009) Coding gestural behavior with the neuroges-elan system. *Behav Res Methods* 41(3):841–849
16. Liu J, Zhong L, Wickramasuriya J, Vasudevan V (2009) uwave: accelerometer-based personalized gesture recognition and its applications. *Pervasive Mobile Comput* 5(6):657–675. doi:[10.1016/j.pmcj.2009.07.007](https://doi.org/10.1016/j.pmcj.2009.07.007) (PerCom 2009)
17. López-Méndez A, Casas JR (2012) Can our tv robustly understand human gestures?: real-time gesture localization in range data. In: *Proceedings of the 9th European conference on visual media production, CVMP '12*. ACM, New York, pp 18–25. doi:[10.1145/2414688.2414691](https://doi.org/10.1145/2414688.2414691)
18. Nielsen M, Störing M, Moeslund T, Granum E (2004) A procedure for developing intuitive and ergonomic gesture interfaces for hci. In: Camurri A, Volpe G (eds) *Gesture-based communication in human-computer interaction*. Lecture notes in computer science, vol 2915, pp 409–420. Springer Berlin Heidelberg
19. Norman DA (2010) Natural user interfaces are not natural. *Interactions* 17(3):6–10
20. Nui Group Community. <http://nuigroup.com>. Accessed 23 Feb 2013
21. Ojala T, Pietikäinen M, Harwood D (1996) A comparative study of texture measures with classification based on featured distributions. *Patt Recogn* 29(1):51–59
22. OpenNI sdk. <http://www.openni.org/openni-sdk/>. Accessed 20 Feb 2013
23. Pantic M, Rothkrantz LJM (2000) Automatic analysis of facial expressions: the state of the art. *IEEE Trans Pattern Anal Mach Intell* 22(12):1424–1445. doi:[10.1109/34.895976](https://doi.org/10.1109/34.895976)
24. Talking and waving to samsung's remote-free tv. <http://spectrum.ieee.org/tech-talk/consumer-electronics/audiovideo/ces-2012-talking-and-waving-to-samsungs-remotefree-tv>. Accessed 24 Feb 2013
25. Poppe R: Vision-based human motion analysis: an overview. *Comp Vision Image Underst* 108(1–2):4–18 (2007). doi:[10.1016/j.cviu.2006.10.016](https://doi.org/10.1016/j.cviu.2006.10.016) (special issue on vision for human-computer interaction)
26. Potamianos G, Neti C, Luetttin J, Matthews I (2004) Audio-visual automatic speech recognition: An overview. *Issues in Visual and Audio-Visual Speech Processing*, pp 356–396
27. Pugeault N, Bowden R (2011) Spelling it out: real-time ASL fingerspelling recognition. In: *ICCV-CDC4CV*
28. Ren Z, Yuan J, Zhang Z (2011) Robust hand gesture recognition based on finger-earth mover's distance with a commodity depth camera. In: *ACM MM, MM '11*. ACM, New York, pp 1093–1096. doi:[10.1145/2072298.2071946](https://doi.org/10.1145/2072298.2071946)
29. Schlömer T, Poppinga B, Henze N, Boll S (2008) Gesture recognition with a wii controller. In: *Proceedings of the 2nd international conference on tangible and embedded interaction, TEI '08*. ACM, New York, pp 11–14. doi:[10.1145/1347390.1347395](https://doi.org/10.1145/1347390.1347395)
30. Sebe N (2009) Multimodal interfaces: challenges and perspectives. *J Ambient Intell Smart Environ* 1(1):23–30
31. Shotton J, Fitzgibbon A, Cook M, Sharp T, Finocchio M, Moore R, Kipman A, Blake A (2011) Real-time human pose recognition in parts from single depth images. In: *CVPR*, pp 1297–1304. doi:[10.1109/CVPR.2011.5995316](https://doi.org/10.1109/CVPR.2011.5995316)
32. Stern HI, Wachs JP, Edan Y (2008) Designing hand gesture vocabularies for natural interaction by combining psycho-physiological and recognition factors. *Int J Semantic Comput* 02(01):137–160. doi:[10.1142/S1793351X08000385](https://doi.org/10.1142/S1793351X08000385)
33. Suau X, Ruiz-Hidalgo J, Casas JR (2012) Real-time head and hand tracking based on 2.5D data. *Trans Multimed* 1(99):1
34. Turk M (2001) Gesture recognition. *Handbook of Virtual Environment Technology*
35. Uebbersax D, Gall J, Van den Bergh M, Van Gool L (2011) Real-time sign language letter and word recognition from depth data. In: *ICCV-HCI*, pp 1–8

36. Viola P, Jones MJ (2004) Robust real-time face detection. *Int J Comput Vision* 57(2):137–154
37. Wachs JP, Kölsch M, Stern H, Edan Y (2011) Vision-based hand-gesture applications. *Commun ACM* 54(2):60–71
38. Zhao W, Chellappa R, Phillips PJ, Rosenfeld A (2003) Face recognition: a literature survey. *ACM Comput Surv* 35(4):399–458. doi:[10.1145/954339.954342](https://doi.org/10.1145/954339.954342)
39. Zigfu. Motion controlled web. <http://zigfu.com>. Accessed 20 Feb 2013



Marcel Alcoverro received his B.S.c. in Telecommunications in 2007 and Master degree in Computing in 2009 from the Technical University of Catalonia (UPC). During 2007 he worked as research assistant at the Equipes Traitement des Images et du Signal (ETIS) and the Institut d'électronique et d'informatique Gaspard-Monge (IGM) in Paris, France, taking part in the EROS-3D project on artwork 3D databases. He is currently working with the Image Processing Group at UPC towards his Ph.D degree where he has been involved in projects of the Spanish Science and Technology System (VISION) and European FP projects (ACTIBIO, FASCINATE). His research interests include markerless motion capture, gesture recognition, 3D video processing and 3D graphics.



Xavier Suau received a degree in Telecommunications Engineering at the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain in 2007. He also received a degree in Aeronautics Engineering at the Ecole Nationale de l'Aéronautique et de l'Espace (SUPAERO), Toulouse, France in 2007.

In September 2007 he received a Master of Research in Automatics, Informatics and Decisional Systems by SUPAERO. His MSc research was carried out in the General Navigation Systems research department of THALES Avionics, Valence, France. During 2008 he joined THALES Avionics as full-time researcher in the Airbus A350-XWB Navigation System project. Since October 2009 he is a PhD student in the Image and Video Processing Group at UPC. He has taken part in the HESPERIA project developing algorithms to compensate illumination in foreground extraction applications. His current work is focused on exploiting range information for feature extraction and body pose estimation, in the framework of his PhD and also in the FP7 FascinatE project.



Josep R. Morros Josep Ramon Morros Rubió received a degree in Physics from the Universitat de Barcelona (UB), Barcelona, Spain, in 1989. He received the Ph.D. from the Universitat Politècnica de Catalunya (UPC) in 2004.

In 1989 he made a short stay at the Research Institute for Technical Physics of the Hungarian Academy of Sciences, in the Division of Semiconductor Microwave Devices, in Budapest, Hungary.

From 1989 to 1990 he worked at Mediterrània d'Enginyeria as a programmer and computer network administrator. From 1991 to 1993 he worked at Instrumentación Electrónica PROMAX, as responsible of the technical documentation department. From October 1994 to February 1997 he was teaching telecommunications at the UPC. In February 1997, he joined the Escola Universitària Politècnica de Mataró EUPMT, Mataró, Spain, where he taught electronics and computer science until June 2001. He also gave courses on PC Hardware at SC2 Formació during several years.

In 2003 he joined the Universitat Politècnica de Catalunya where he is currently Associate Professor. He is lecturing in the areas of communications, signal processing, digital image processing and acoustics. His current research interests include image and sequence coding, segmentation problems and video sequence analysis.

Since 1994 he has been involved in various European Projects (MAVT, MoMuSys, SIMILAR, CHIL, FascinatE ...) as a researcher from the Image and Video Processing Group at the UPC. In 2005 he was in charge of the Halftone Compression Investigation Project in collaboration with Hewlett Packard Española. From 2005 to 2007 he's acting as a leader for the Face Tracking/Detection task in the CHIL project.



Adolfo López-Méndez received his B.S.c. (2007) and Master degree in Signal Theory and Communications (2009) from the Technical University of Catalonia (UPC), where he is currently working towards his Ph.D degree. His research interests include action and gesture recognition, markerless motion capture, 3D video processing and machine learning. Adolfo has been recently involved in the FP7 projects CHIL, ACTIBIO and FASCINATE.



Albert Gil received a degree on Electrical Engineering from the Universitat Politècnica de Catalunya (BarcelonaTECH) while I was working in the Software Design department of SONY in Viladecavalls.

In 2007 I joined the Image Processing Group where I have been designing our software development platform, collaborating in software libraries and applications for different projects and also working as system manager. Because I'm mainly interested in ubiquitous computing, computer vision and software design I also joined the Master in Computing program of the UPC and I have been mainly collaborating in projects related to the SmartRoom.



Javier Ruiz-Hidalgo received a degree in Telecommunications Engineering at the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain in 1997. From 1998 to 1999, he developed an MSc by Research on the field of Computer Vision by the University of East Anglia (UEA) in Norwich, UK. During 1999 he joined the Image Processing Group at UPC working on image and video indexing in the context of the MPEG-7 standard where he obtained his PhD thesis in 2006. Since 1999 he has been involved in various European Projects as a researcher from the Image Processing Group at UPC. During 1999 and 2000 he worked in the ACTS (AC308) DICEMAN project developing new descriptors and representations for image and video sequences. From 2001 to 2003 he is also involved in the IST/FET (2000-26467) project MASCOT developing an efficient compression scheme exploiting metadata information. Since 2006 he is the principal investigator in the HESPERIA (CENIT-2006) project involved in developing new image algorithms in security applications. Since 2001 he is an Associated Professor at the Universitat Politècnica de Catalunya. He is currently lecturing on the area of digital signal and systems and image processing. His current research interests include image segmentation, still image and sequence coding, compression and indexing.



Josep R. Casas is Associate Professor at the Department of Signal Theory and Communication, Technical University of Catalonia (UPC) in Barcelona. He graduated in Telecommunications Engineering in 1990 and received the PhD in 1996, both from UPC, where he is currently teaching Signals and Systems, Image Processing and Television Systems at the School of Telecommunications Engineering (Telecom BCN). He was visiting researcher at CSIRO Mathematics & Information

Sciences in Canberra, Australia from 2000 to 2001. Josep R. Casas is Principal investigator of the project PROVEC (“Video Processing for Controlled Environments”) of the Spanish R&D&I Plan started in 2007, and has led or contributed to a number of industry-sponsored projects, projects of the Spanish Science and Technology System (VISION, HESPERIA) and European FP projects (ACTIBIO, SCHEMA, ADVISOR). In particular, he coordinated UPC contribution to CHIL (“Computers in the Human Interaction Loop”), an IP of the IST/EU 6th Framework Program in the strategic objective of Multimodal Interfaces, involving video, audio and natural language technologies. Josep R. Casas has authored or co-authored over 10 papers in international journals, 12 papers in LNCS, 50 contributions to conferences and 9 book chapters and a teaching book in the areas of video coding, analysis, indexing and image processing.