

Recognizing Hand Gestures ^{*}

James Davis and Mubarak Shah ^{**}

Computer Vision Laboratory, University of Central Florida, Orlando FL 32816, USA

Abstract. This paper presents a method for recognizing human-hand gestures using a model-based approach. A Finite State Machine is used to model four qualitatively distinct phases of a generic gesture. Fingertips are tracked in multiple frames to compute motion trajectories, which are then used for finding the start and stop position of the gesture. Gestures are represented as a list of vectors and are then matched to stored gesture vector models using table lookup based on vector displacements. Results are presented showing recognition of seven gestures using images sampled at 4 Hz on a SPARC-1 without any special hardware. The seven gestures are representatives for actions of Left, Right, Up, Down, Grab, Rotate, and Stop.

1 Introduction

It is essential for computer systems to possess the ability to recognize meaningful gestures if they are to interact naturally with people. Humans use gestures in daily life as a means of communication, e.g., pointing to an object to bring someone's attention to the object, waving "hello" to a friend, requesting n of something by raising n fingers, etc. The best example of communication through gestures is given by sign language. American Sign Language (ASL) incorporates the entire English alphabet along with many gestures representing words and phrases [3], which permits people to exchange information in a non-verbal manner.

Currently, the human-computer interface is through a keyboard and/or mouse. Physically challenged people may have difficulties with such input devices and may require a new means of entering commands or data into the computer. Gesture, speech, and touch inputs are few possible means of addressing such users' needs to solve this problem. Using Computer Vision, a computer can recognize and perform the user's gesture command, thus alleviating the need for a keyboard. Some applications for such a vision system are the remote control of a robotic arm, guiding a computer presentation system, and executing computer operational commands such as opening a window or program.

This paper presents a gesture recognition method using Computer Vision, which permits human users adorned with a specially marked glove to command

^{*} The research reported here was supported by the National Science Foundation grants CDA-9200369 and IRI-9220768.

^{**} For an extended version of this paper, please send e-mail to shah@sono.cs.ucf.edu.

a computer system to carry out predefined gesture action commands. Our system has been created to recognize a sequence of multiple gestures in which a subset of the gestures is comprised of select ASL letters (See Fig. 1). Each gesture begins with the hand in the “hello” position and ends in the recognizable gesture position. The current library of gestures contains seven gestures: *Left*, *Right*, *Up*, *Down*, *Rotate*, *Grab*, and *Stop*. The user must begin in the designated start position and is able to make gestures until the termination gesture (*Stop*) is recognized.

There are several advantages of this system over other methods. First, it uses inexpensive black-and-white video. Incorporating color markers on a glove as interest points [2] requires costly color imaging, whereas a binary marked glove, as used in this research, can be detected in low-cost black-and-white imaging. Second, a simple vision glove is employed, i.e., no mechanical glove with LEDs or bulky wires. Current gesture input devices require the user to be linked to the computer, reducing autonomy [1]. Vision input overcomes this problem. Third, a duration parameter for gestures is incorporated. For example, if this recognition system were connected to a robotic arm and the user makes a gesture for *Left*, the robotic arm would continue to move left until the user moves the hand from the gesture back to the start position. Therefore the user can control the execution duration of the robotic arm. Finally, due to Finite State Machine (FSM) implementation of a generic gesture, no warping of the image sequences is necessary. That is, the number of frames in each gesture can be variable.

2 Related Work

Baudel and Beaudouin-Lafon [1] implemented a system for the remote control of computer-aided presentations using hand gestures. In this system, the user wears a VPL DataGlove which is linked to the computer. The glove can measure the bending of fingers and the position and orientation of the hand in 3-D space. The user issues commands for the presentation by pointing at a predefined active zone and then performing the gesture for the desired command. Gesture models include information pertaining to the start position, arm motion (dynamic phase), and stop position of the gesture. The command set includes such commands as *next page*, *previous page*, *next chapter*, *previous chapter*, *table of contents*, *mark page*, and *highlight area*. Two main types of errors that can occur with this system are system errors and user errors. System errors relate to the difficulties identifying gestures that differ only in the dynamic phase, while user errors correspond to hesitations while issuing a command. With trained users, the recognition rate was 90– to 98%. This system does not use vision to recognize gestures, but instead uses a linked hardware system to track the hand and arm movements, which makes movement less natural for the user.

Cipolla, Okamoto, and Kuno [2] present a real-time structure-from-motion (SFM) method in which the 3-D visual interpretation of hand gestures is used in a man-machine interface. A glove with colored markers attached is used as input to the vision system. Movement of the hand results in motion between the images

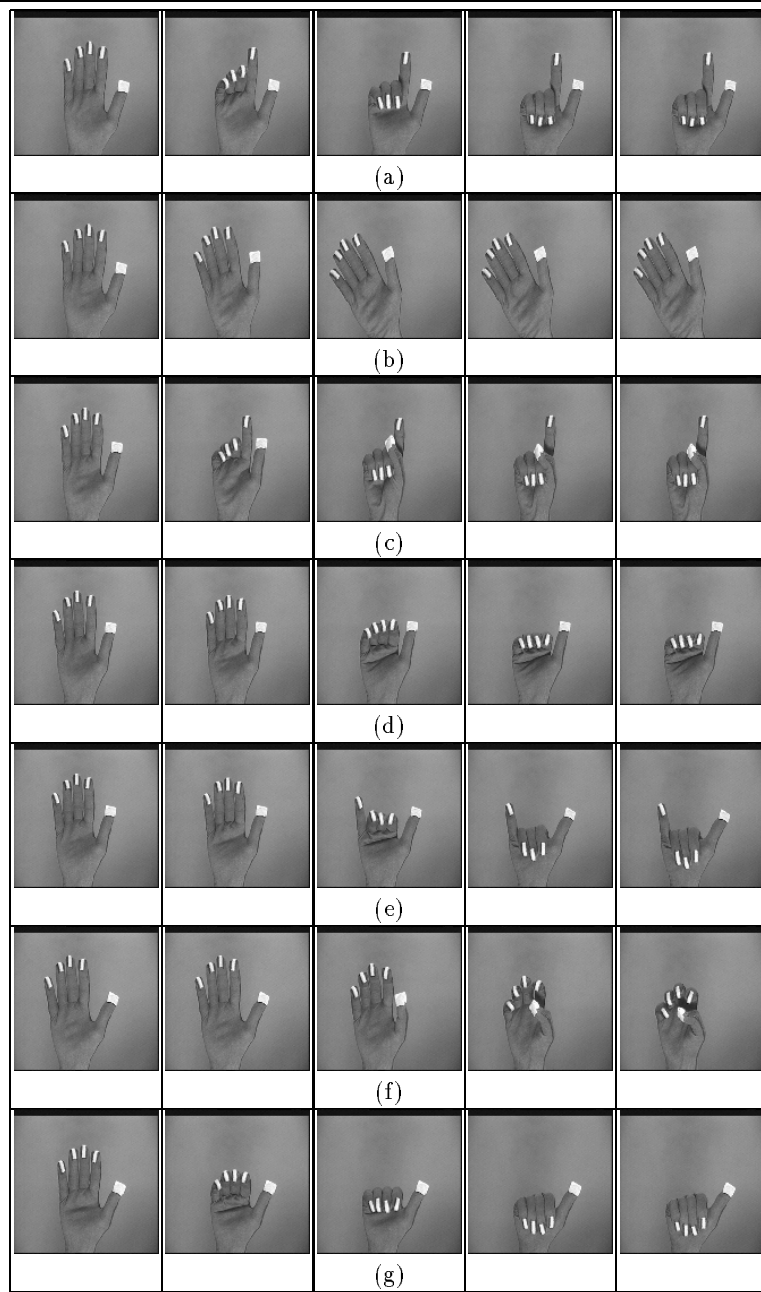


Fig.1. Gestures. (a) Left. (b) Right. (c) Up. (d) Down. (e) Rotate. (f) Grab. (g) Stop.

of the colored markers. The authors use the parallax motion vector, divergence, curl, and deformation components of the affine transformation of an arbitrary triangle, with the colored points at each vertex, to determine the projection of the axis of rotation, change in scale, and cyclotorsion. This information is then used to alter an image of a model. The information extracted from the colored markers does not give the position of the *entire* hand (each finger), it only provides a triangular plane for the SFM algorithm. The SFM method used here assumes rigid objects, which is not true in the case of hand gestures.

Fukumoto, Mase, and Suenaga [5] present a system called *Finger-Pointer* which recognizes pointing actions and simple hand forms in real-time. The system uses stereo image sequences and does not require the operator to wear any special glove. It also requires no special image processing hardware. Using stereo images, their system uses the 3-D location of fingers rather than the 2-D location. The coordinates of the operator's fingertip and the direction it is pointing is determined from the stereo images and then a cursor is displayed in the target location on the opposing screen. The system is robust in that it is able to detect the pointing regardless of the operator's pointing style. Applications of this system can be similar to the gesture controlled computer-aided presentations of Baudel and Beaudouin-Lafon [1] and also can be used in a video browser with a VCR.

Darrell and Pentland [4] have also proposed a glove-free environment approach for gesture recognition. Objects are represented using sets of view models, and then are matched to stored gesture patterns using dynamic time warping. Each gesture is dynamically time-warped to make it of the same length as the longest model. Matching is based upon the normalized correlation between the image and the set of 2-D view models where the view models are comprised of one or more example images of a view of an object. This method requires the use of special-purpose hardware to achieve real-time performance, and uses gray level correlation which can be highly sensitive to noise. Also, their method was only tested in distinguishing between two gestures.

3 Generic Gesture

For a system to recognize a sequence of gestures, it must be able to determine what state the user's hand is in, i.e., whether or not the hand is dormant, moving, or in gesture position. Our approach relies on the qualitatively distinct events (phases) in gestures, rather than on frame by frame correlation. Each gesture the user performs begins with the hand in the start position (all fingers upright, as if one was about to wave "hello" to another person). Next, the user moves the fingers and/or entire hand to the gesture position. Once in position, the system will attempt to recognize and then execute the gesture command until the hand begins moving back to the start position. The system will then wait for the next gesture to occur. Thus, the user is constrained to the following four *phases* for making a gesture.

1. Keep hand still (fixed) in start position until motion to gesture begins.
2. Move fingers smoothly as hand moves to gesture position.
3. Keep hand in gesture position for desired duration of gesture command.
4. Move fingers smoothly as hand moves back to start position.

Since these four phases occur in a fixed order, an FSM can be used to guide the flow and recognition of gestures, based on the motion characteristics of the hand (See Fig. 2). A 1 or 0 in the state diagram represents motion or no motion respectively, between two successive images. A *three frame similarity* constraint, which states that, “at least three consecutive images must have the same motion properties to advance to the next phase,” was found to inhibit premature phase advancement.




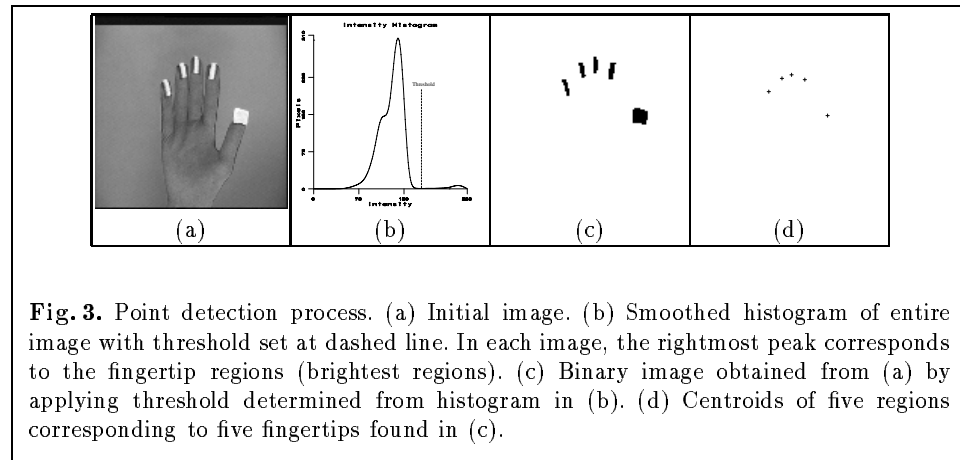
Fig. 2. State diagram. States S0 and S1 depict the initial phase (Phase 1), states S2 and S3 depict the motion to gesture phase (Phase 2), states S4 and S5 depict the gesture recognition phase (Phase 3) where an external device such as a robotic arm could execute the gesture until exiting this phase, and states S6 and S7 depict the motion to initial phase (Phase 4).

Due to the nature of this machine, no *warping* of image sequences is necessary, i.e., it is not required to have a fixed number of images for each gesture sequence. The FSM compensates for varying numbers of images by looping at the current phase as long as the *three frame similarity* constraint is satisfied. The *actual* number of frames which constitute the motion of a gesture yields no information for use with this system. The only useful information is the start and end position of the fingertips. The system does not care *how* the fingers or hand arrive in the gesture position; it wants to know the location of each fingertip before and when the gesture is made. Only the locations and total displacement of the fingertips play a crucial role in gesture recognition, as compared to other motion characteristics such as instantaneous velocity. Therefore, we need only to track each fingertip from the initial position to the final gesture position. The FSM permits the determination of which phase the user is currently executing, and it also tracks the fingertips of a variable-length frame sequence to the gesture position.

In our method, each image in the sequence is analyzed to find the location of the fingertips, and then motion correspondence is used to track these points to the resulting gesture position (Section 4). The trajectories computed by the motion correspondence algorithm are then converted to vector form to be matched with the stored gestures (Sections 5 and 6).

4 Fingertip Detection and Motion Correspondence

The goal of fingertip detection is to identify the 2-D location of the marked fingertips on the vision glove. The location of the fingertips determines the position of the fingers at any time. Since we are using a sequence of images in which the intensity of the fingertips is known a priori to be significantly different from the remaining regions, a multi-modal histogram of the image can be generated and used to segment the image into fingertip regions (See Fig. 3). We represent the five fingertip regions using centroids for ease of calculations, storage, display, etc. and also for motion correspondence (See Fig. 3.d).



Motion correspondence maps points in one image to points in the next image such that no two points are mapped onto the same point. A path, known as a trajectory, is generated for each of the m points, starting with the points in the first image and ending with the points in the n th image.

Rangarajan and Shah's [6] motion correspondence algorithm was chosen for its exploitation of a *proximal uniformity* constraint, which says objects follow smooth paths and cover a small (proximal) distance in a small time. It was stated previously, in the Phase 2 gesture constraint, that the fingers must move smoothly to the gesture position. Additionally, the *three frame similarity* constraint for motion, which requires at least three frames of motion, implies that the fingertips move a small (proximal) distance in each successive frame. Therefore,

the algorithm, using a proximal uniformity constraint, agrees with the previously stated gesture motion constraints.

The authors' algorithm establishes correspondence among points by minimizing a *proximal uniformity function* δ , which prefers the proximal uniform path, such that

$$\delta(X_p^{k-1}, X_q^k, X_r^{k+1}) = \frac{\| \overline{X_p^{k-1} X_q^k} - \overline{X_q^k X_r^{k+1}} \|}{\sum_{x=1}^m \sum_{z=1}^m \| \overline{X_x^{k-1} X_{\Phi^{k-1}(x)}^k} - \overline{X_{\Phi^{k-1}(x)}^k X_z^{k+1}} \|} + \frac{\| \overline{X_q^k X_r^{k+1}} \|}{\sum_{x=1}^m \sum_{z=1}^m \| \overline{X_{\Phi^{k-1}(x)}^k X_z^{k+1}} \|} \quad (1)$$

where Φ^k is one to one onto correspondence between points of image k and image $k+1$, $1 \leq p, q, r \leq m$, $2 \leq k \leq m-1$, $q = \Phi^{k-1}(p)$, $\overline{X_q^k X_r^{k+1}}$ is the vector from point q in image k to point r in image $k+1$, and $\| X \|$ denotes the magnitude of vector X [6]. The first term in the equation represents the smoothness constraint and the second represents the proximity constraint.

5 Gesture Modeling

In general, human finger movements are linear, with extrema moving from an extended position down to palm/wrist area, e.g., from the hand in the "hello" position to the hand making a fist. Even though we have the ability of limited rotational movement in the fingers, we mostly move the fingers up and down to the palm, with the thumb moving left and right over the palm. Since the fingers move relatively linearly (some move curvilinearly at times), we can approximate each fingertip trajectory by a single vector (See Fig. 4). Each vector will originate at the location of the corresponding fingertip before motion to the gesture position, and will terminate at the location of the gesture position. We disregard the actual path each fingertip makes because, as stated previously, we are concerned with only the beginning and ending location of each fingertip. Therefore, if there is some curvature to the fingertip trajectory, it will be disregarded. The motion information leading to the gesture position is not needed. Motion correspondence is used only to map the starting points to the ending points by means of tracking the points in-between in the trajectories. See Fig. 4 for vector representations of the gesture set.

A library model is created from averaging m test models of the same gesture and is represented in a data structure which contains

1. The gesture name.
2. The mean direction and mean magnitude, i.e., mean displacement, for each fingertip vector.
3. The gesture's motion code.

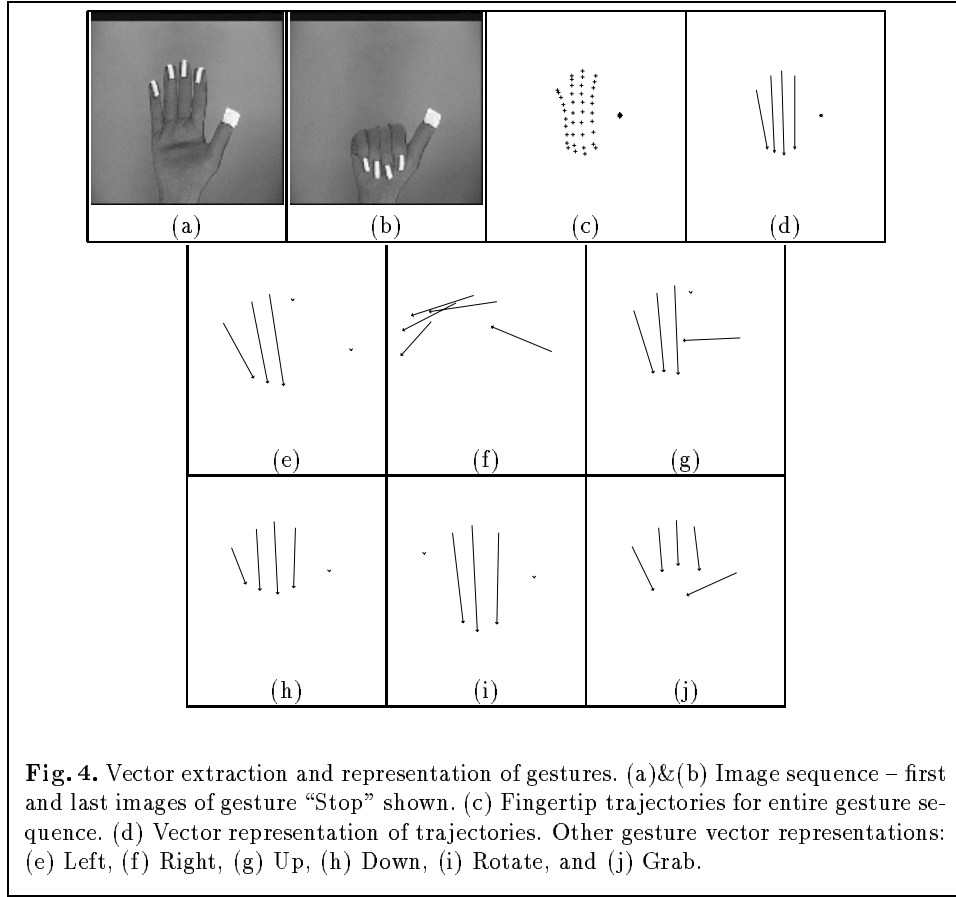


Fig. 4. Vector extraction and representation of gestures. (a)&(b) Image sequence – first and last images of gesture “Stop” shown. (c) Fingertip trajectories for entire gesture sequence. (d) Vector representation of trajectories. Other gesture vector representations: (e) Left, (f) Right, (g) Up, (h) Down, (i) Rotate, and (j) Grab.

The direction, Θ , and magnitude, or displacement, of a fingertip vector is determined from the starting point (x_0, y_0) and stopping point (x_n, y_n) of its trajectory and are easily calculated respectively by

$$\Theta = \arctan \frac{y_n - y_0}{x_n - x_0} , \quad (2)$$

$$Disp = \sqrt{(x_n - x_0)^2 + (y_n - y_0)^2} . \quad (3)$$

We use a five-bit motion code to store the motion activity of the five fingertips for the gesture, which acts as a table key for the model. Each bit of this code corresponds to a specific fingertip vector, with the least significant bit storing finger 1's (thumb's) motion information and progressing to the most significant bit where finger 5's (little finger's) motion information is stored. A bit is set if its respective fingertip vector has motion, i.e., it's fingertip vector magnitude is

above some displacement threshold. Thus, the motion code for a gesture with significant motion in fingertip vectors 3, 4, and 5 only is represented as 11100. This binary number in decimal notation is 28, which is stored as the gesture's motion code.

6 Gesture Matching

Gesture matching consists of comparing the unknown gesture with the models to determine whether the unknown gesture matches with any model gesture in the system vocabulary.

Motion codes enable the matching scheme to consider only those models which have the same motion code as the unknown gesture and also provide information to which motion category the unknown gesture belongs. The library models, when loaded into memory, can be stored in an array of linear linked lists in which the array is indexed by the motion codes (0-31). During the matching stage, the unknown gesture is only compared with the library models that are indexed by the unknown gesture's motion code.

With only a subset of library models to compare to the unknown gesture, we have reduced the search complexity, which is now dependent on the different motion codes of the current library of gestures. A match is then determined by comparison between the stored models and the unknown gesture. A match is made if all vector fields (magnitude and direction for each fingertip) in the unknown gesture are within some threshold of the corresponding model entries.

7 Results

Ten sequences of over 200 frames each were digitized at 4 Hz, stored, and then used for the recognition program. Each run was performed in the same fashion, starting with the gesture for *Left* and progressing to the ending gesture *Stop*, as shown horizontally in Table 1. An image set in which the fixed order shown in Table 1 was altered resulted in perfect recognition, which implies order is not a concern.

The number of images for each sequence depended on the duration of each gesture performed. The overall results on the ten sequences of images yielded almost perfect scores with the exception of run 9, where an error in the sequence caused the remaining gestures to be unrecognizable. A shift of the hand above the threshold limit or occlusion of points due to lighting conditions may cause premature advancement of one phase to another, which in turn may result in the FSM continuing asynchronously with the image sequence.

Recognition of a nine 128×128 frame sequence sampled at 4 Hz took a CPU time of 890 ms on a SPARC-1 (99 ms processing time per frame) with no special hardware. Our experiments show that sampling at a rate of 30 Hz is not necessary for gesture recognition since the processing time needed for our method is small enough for implementation in real-time with images sampled up to 10 Hz.

Table 1. Table 1: Results. \checkmark – Recognized, X – Not Recognized, * – Error in Sequence.

Run	Frames	Left	Right	Up	Down	Rotate	Grab	Stop
1	200	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
2	250	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
3	250	\checkmark	\checkmark	\checkmark	X	\checkmark	\checkmark	\checkmark
4	250	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
5	300	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
6	300	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
7	300	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
8	300	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
9	300	\checkmark	\checkmark	\checkmark	\checkmark	*	*	*
10	300	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark

8 Conclusion

In this paper, we have developed a Computer Vision method for recognizing sequences of human-hand gestures within a gloved environment. A specialized FSM was constructed as an alternative to image sequence warping. We utilize vectors for representing the direction and displacement of the fingertips for the gesture. Modeling gestures as a set of vectors with a motion code allows the reduction of complexity in the model form and matching. We presented the performance of this method on real image sequences. Future work pursued includes extending the gesture vocabulary, removing the glove environment, and relaxing the start/stop requirement.

References

1. Baudel T., Beaudouin-Lafon M.: Charade: Remote control of objects using free-hand gestures. CACM. July (1993) 28–35
2. Cipolla R., Okamoto Y., Kuno Y.: Robust structure from motion using motion parallax. ICCV. (1993) 374–382
3. Costello E.: Signing: How to speak with your hands. Bantam Books, New York (1993)
4. Darrell T., Pentland A.: Space-time gestures. CVPR. (1993) 335–340
5. Fukumoto, M., Mase, K., Suenaga, Y.: Real-time detection of pointing actions for a glove-free interface. IAPR Workshop on Machine Vision Applications. Dec. 7–9 (1992) 473–476
6. Rangarajan, K., Shah, M.: Establishing motion correspondence. CVGIP: Image Understanding. **54** July (1991) 56–73