

Research Article

Adding Image Constraints to Inverse Kinematics for Human Motion Capture

Antoni Jaume-i-Capó, Javier Varona, Manuel González-Hidalgo, and Francisco J. Perales

*Departament de Ciències Matemàtiques i Informàtica, Universitat de les Illes Balears, Ed. A. Turmeda,
Ctra de Valldemossa Km 7.5, 07122 Palma, Spain*

Correspondence should be addressed to Antoni Jaume-i-Capó, antoni.jaume@uib.es

Received 15 May 2009; Accepted 8 July 2009

Academic Editor: João Manuel R. S. Tavares

Copyright © 2010 Antoni Jaume-i-Capó et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

In order to study human motion in biomechanical applications, a critical component is to accurately obtain the 3D joint positions of the user's body. Computer vision and inverse kinematics are used to achieve this objective without markers or special devices attached to the body. The problem of these systems is that the inverse kinematics is "blinded" with respect to the projection of body segments into the images used by the computer vision algorithms. In this paper, we present how to add image constraints to inverse kinematics in order to estimate human motion. Specifically, we explain how to define a criterion to use images in order to guide the posture reconstruction of the articulated chain. Tests with synthetic images show how the scheme performs well in an ideal situation. In order to test its potential in real situations, more experiments with task specific image sequences are also presented. By means of a quantitative study of different sequences, the results obtained show how this approach improves the performance of inverse kinematics in this application.

1. Introduction

In biomechanical applications that aim to study human motion, a critical component is to accurately obtain the 3D joints' positions of the user's body. Usually, the most common methods to obtain the joints' positions require a laboratory environment and the attachment of markers to the body. Modern biomechanical and clinical applications require the accurate capture of normal and pathological human movement without the artifacts associated with standard marker-based motion capture techniques such as soft tissue artifacts and the risk of artificial stimulus of taped on or strapped on markers [1]. Emerging techniques and research in computer vision are leading to the rapid development of the markerless approach to motion capture [2].

In computer vision, algorithms are designed to allow the system to analyze one or multiple image streams in order to recover human motion. However, the images are 2D and the human body representation is in 3D. This fact leads to the presence of ambiguities; there are a number of

possible 3D configurations of the human body that could explain a single image. In addition, these images can be noisy or incomplete (some joints or limbs are not visible). Therefore, we can only estimate the users posture. Inverse kinematics approaches can solve the body posture from their 3D position if we can clearly locate visible body parts such as face and hands. For example, in the work of Zou et al. [3], the angles of joints are estimated by inverse kinematics based on human skeleton constraints, and the coordinates of pixels in the body segments in the scene are determined by forward kinematics. Finally the human motion pose can be reconstructed by histogram matching. Their main drawback is that the algorithm does not handle human motion in the direction perpendicular to the image plane displacement. In the case of multiple cameras, ambiguities appear to be less significant. For example, by using two cameras to recover the user's posture in order to recognize the user's gestures for Human-Computer Interaction applications [4]. This work is also based on computer vision and inverse kinematics in order to recover human body posture.

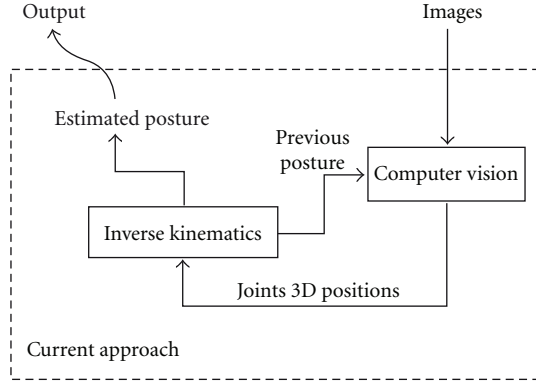


FIGURE 1: Current approach to combine computer vision and inverse kinematics.

However, previous works based on the combination of computer vision and inverse kinematics tend to simplify the approach by combining the results of both techniques that are applied by separate. As it is shown in Figure 1, current approaches are based on the detection of certain joints into the images by using Computer Vision algorithms and, therefore, estimate a plausible body posture by using Inverse Kinematics algorithms. In this way, the Inverse Kinematics algorithms are “blinded” with respect to the projection of body segments into the images. They only use biomechanical constraints to estimate the best position of the non-detected joints in the images.

In this paper we present a new approach where the objective is to include the image information directly onto the inverse kinematics scheme, see Figure 2. The idea is to use image constraints to solve the redundancy of kinematics solvers. In addition, we also explain that it is possible to use a preprocessed image, in other words, the computer vision algorithms could process the input images in order to make the problem more tractable or to enhance a desired image feature for specific applications. Finally, this scheme of posture reconstruction can be used with one or more views, that is, it can work using only one view but the results improve if more views of the performer are applied.

In order to show the viability of this scheme of 3D human posture recognition, different experiments have been conducted. First, using synthetic images in order to show how the scheme works in an ideal situation. This simple case shows how theoretically the scheme performs correctly. Next, in order to test its potential in real situations, experiments with real sequences are also presented. In these experiments an annotated sequence and a known database of human motions that contains motion capture data are used to make a quantitative study of different sequences in order to evaluate the performance of the presented approach.

This paper is organized as follows. In next section, current inverse kinematics approach is reviewed in order to introduce in Section 3 the image constraints. Section 4

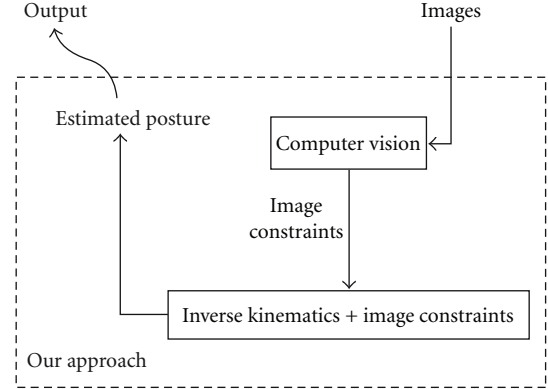


FIGURE 2: Our approach: Adding image constraints to inverse kinematics.

analyzes the obtained results in order to demonstrate the viability of this approach. Finally, conclusions are presented in the last section.

2. Inverse Kinematics

In order to capture human motion, the human body is usually modeled as an articulated chain, which consists of a set of rigid objects, called links, joined together by joints. To control the movement of an articulated chain it is common to use inverse kinematics (IK). IK is exploited to reconstruct an anatomically correct posture of the user (i.e., its joint state) considering the 3D locations of selected end-effectors which are used to constrain the posture.

For the moment, let us consider only a single frame of motion. Write the vector of joint angles as $\boldsymbol{\theta} = \{\theta_1, \dots, \theta_n\}$. Assume that we would like to meet a set of constraints on joints positions as $\mathbf{x} = \{x_1, \dots, x_m\}$ as functions of the joints degrees of freedom $\boldsymbol{\theta}$. The problem of inverse kinematics is to obtain a $\boldsymbol{\theta}$ such that $\mathbf{x} = \mathbf{g}(\boldsymbol{\theta})$. Closed forms solutions are available for at least some parameters, when the limbs of the articulated chain are considered independently [5]. More often, one must see this as a numerical root finding problem based on the linearization of the set of constraints on joint positions, \mathbf{g} , considering small displacements about the current configuration, $\boldsymbol{\theta}$,

$$\Delta \mathbf{x} = \mathbf{J} \Delta \boldsymbol{\theta}, \quad (1)$$

where \mathbf{J} is the $m \times n$ Jacobian matrix

$$\mathbf{J} = \left(\frac{\partial g_i}{\partial \theta_j} \right), \quad i = 1, \dots, m; \quad j = 1, \dots, n. \quad (2)$$

The resulting Jacobian matrix \mathbf{J} is inverted to map the desired constraint variation $\Delta \mathbf{x}$ to a corresponding posture variation $\Delta \boldsymbol{\theta}$. Using the pseudoinverse, noted by \mathbf{J}^+ , the norm of the solution mapped by \mathbf{J}^+ is minimal, that is, it is the




FIGURE 3: (a) Initial configuration of the articulated chain, (b) IK estimation, (c) result by adding the image constraint.



FIGURE 4: Image supports: (a) silhouette; (b) euclidean distance transform.

smallest posture variation realizing the desired constraint variation:

$$\Delta\boldsymbol{\theta} = \mathbf{J}^+ \Delta\mathbf{x}. \quad (3)$$

Since $\text{rank}(\mathbf{J}) = m < n$, there is an infinite number of solutions. For the positioning and animation of articulated figures in computer graphics, the weighting strategy [6] is frequently employed. In the field of robotics however, the strategy is to solve this inverse kinematics redundancy adding a secondary term (usually defined as secondary task) to (3) in order to minimize a criterion $h(\boldsymbol{\theta})$. In this formulation, redundancy solution is accomplished by moving the joints such that the end-effectors are moved in the desired way and the criterion h is always kept at a minimum. This was first exploited by Liégeois [7] who added a secondary task by projecting the negative gradient of $h(\boldsymbol{\theta})$ into the null space of \mathbf{J} , see (4),

$$\Delta\boldsymbol{\theta} = \mathbf{J}^+ \Delta\mathbf{x} - \alpha(\mathbf{I}_n - \mathbf{J}^+ \mathbf{J}) \nabla h(\boldsymbol{\theta}), \quad (4)$$

where \mathbf{I}_n is the $n \times n$ identity matrix, and α is a positive gain factor which is configuration dependent. Definition of the secondary task by means of the criterion $\nabla h(\boldsymbol{\theta})$ depends on the application. In following section, a criterion based on images will be defined in order to capture human motion.

3. The Image-Based constraint

As explained in previous section, it is possible to constrain the solutions of inverse kinematics by adding a scalar

criterion $h(\boldsymbol{\theta})$. Next, we explain how to define this criterion by using the images in order to guide the posture reconstruction of the articulated chain for human motion capture applications.

The definition of the image-based constraint is inspired in the works of Visual Servo Control [8]. Specifically, Marchand and Courty define different secondary tasks for controlling a camera in virtual environments [9]. For motion capture purposes, it should be taken into account that the human structure is highly redundant and, therefore, a large solution space exists. A solution is to generalize (4) to include more tasks by using the priority strategy [10]. In this case, the solution guarantees that a task associated with a high priority will be achieved as much as possible, while a low-priority constraint will be optimized only on the reduced solution space that does not disturb all higher priority tasks. However, for the sake of clarity, we only consider two tasks. It is straightforward to extend to more tasks when the image constraint has low priority. In addition, it is possible to use the Extended Jacobian method [11] in order to give a high priority to the image constraint.

For motion capture applications, we define $h(\boldsymbol{\theta})$ in order to maximize the overlap between the projection of the articulated chain into the images and the human body. Consider the case of the Figure 3(a), where is shown the initial configuration of the articulated chain and the objective is to estimate the elbow's position with the 3D position of the hand as end-effector. By applying IK we obtain the result of Figure 3(b), where the elbow's estimation

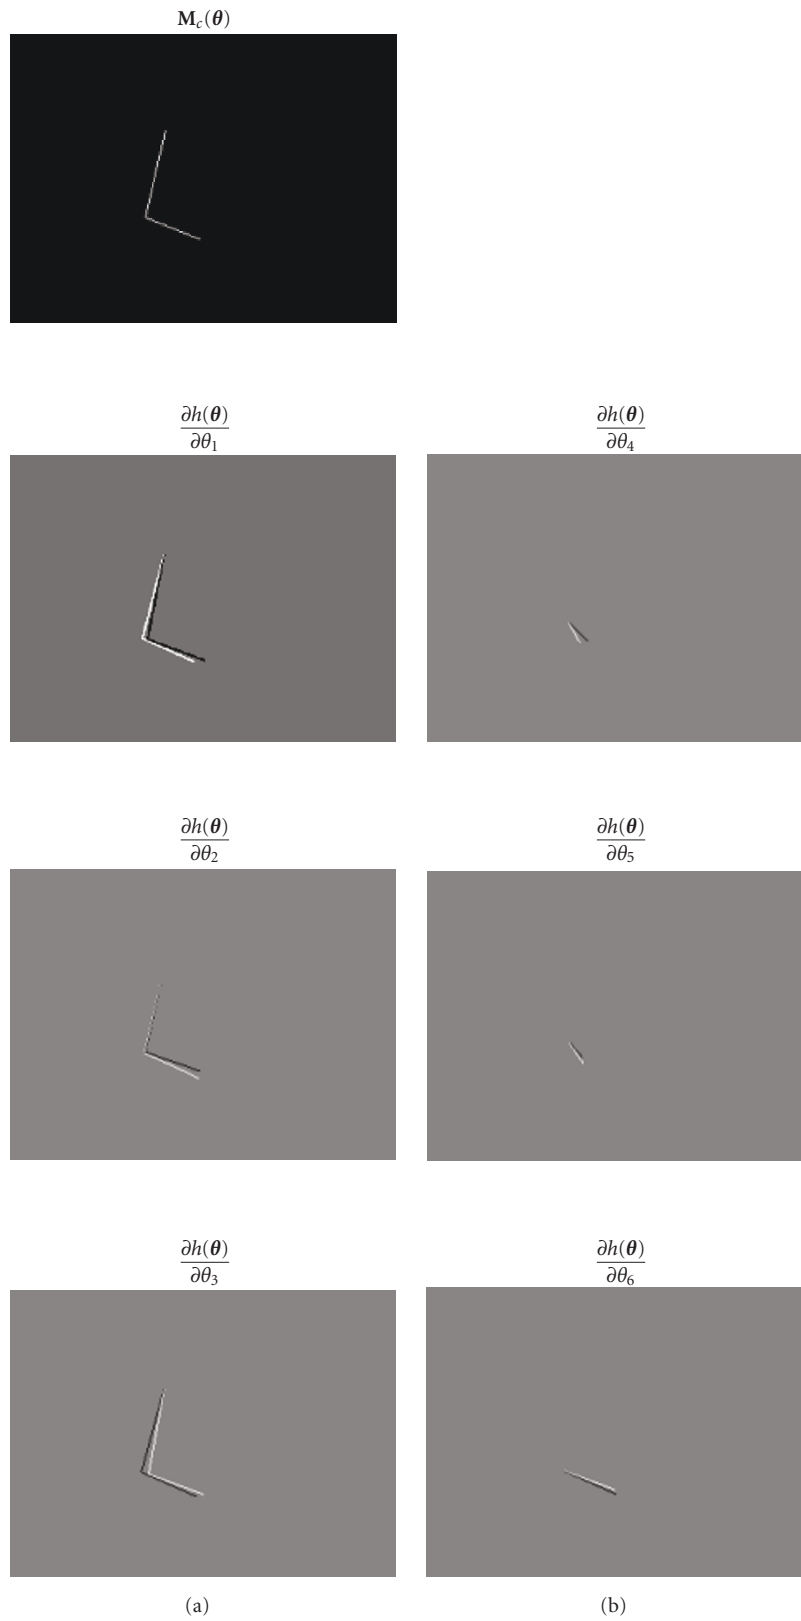


FIGURE 5: The $M_c(\theta)$ function and its partial derivatives for the case of Figure 3.




FIGURE 6: Experiment 1. From the initial configuration of the articulated chain (gray), both approaches try to reach the objective configuration (black). Left column, sequence applying inverse kinematics without image constraint. Right column, sequence applying inverse kinematics with image constraint.

lies outside of the body due its “blind” nature by using only the desired position of the end-effector. In order to solve this problem, we propose a criterion that tries to guide the articulated chain to the body projection into the image, Figure 3(c). Formally, let us to define $h(\boldsymbol{\theta})$ as follows

$$h(\boldsymbol{\theta}) = \frac{1}{n} \sum_c \sum_x \sum_y (\mathbf{I}_c(x, y) \cdot \mathbf{M}_c(x, y, \boldsymbol{\theta})), \quad (5)$$

where $\mathbf{I}_c(x, y)$ represents the intensity of the 2D point (x, y) of the image c ($c \geq 1$), which corresponds to a different view of the user, and n is the number of points (x, y) that belong to the desired image support \mathbf{I}_c . Applying a background subtraction algorithm [12], it is possible to directly use the silhouette as the image support, \mathbf{I} , of the articulated chain, see Figure 4(a). However, in order to get

a smooth surface we apply the euclidean distance transform [13] to the silhouette image, see Figure 4(b). Both operations are fast and do not introduce a significative delay in the algorithm.

In order to complete the definition of (5), let us to define the function $\mathbf{M}_c(\boldsymbol{\theta})$, which is the projection of the articulated chain into the image c . If $\mathbf{X} = (X, Y, Z)$ are the coordinates of the i th joint in the 3D-space, and assuming knowing the calibration data in order to project the 3D coordinates into the 2D images, we define $\mathbf{p}_{c,i} = (x, y)$ such as the 2D image coordinates of the projected i th joint into the image c . Assuming that the joints are ordered in a consecutive way, the $\mathbf{M}_c(\boldsymbol{\theta})$ function is defined as follows:

$$\mathbf{M}_c(x, y, \boldsymbol{\theta}) = \begin{cases} 1, & \text{if } (x, y) \in \overline{\mathbf{p}_{c,i}\mathbf{p}_{c,i+1}} \text{ for some } i \text{ joint,} \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where $\overline{\mathbf{p}_{c,i}\mathbf{p}_{c,i+1}}$ is the segment between the 3D joints’ projection into the image. Figure 5 shows the $\mathbf{M}_c(\boldsymbol{\theta})$ function for the example of Figure 4 and its partial derivatives. Concluding, the image constraint is then given by the gradient of the criterion $h(\boldsymbol{\theta})$ of (7)

$$\nabla h(\boldsymbol{\theta}) = \left(\frac{\partial h(\boldsymbol{\theta})}{\partial \theta_j} \right), \quad j = 1, \dots, n, \quad (7)$$

where the partial derivative of the joint j is defined in (8)

$$\frac{\partial h(\boldsymbol{\theta})}{\partial \theta_j} = \frac{1}{n} \sum_c \sum_x \sum_y \left(\mathbf{I}_c(x, y) \cdot \frac{\partial \mathbf{M}_c(x, y, \boldsymbol{\theta})}{\partial \theta_j} \right). \quad (8)$$

4. Performance Evaluation

The proposed approach is evaluated using three different tests. The first test uses a virtual environment to show how the presented approach runs well in an ideal situation. The second test applies the proposed approach on a sequence of user’s motions to show how the presented approach performs well using real images. Finally, the third test compares the evaluation of the inverse kinematics approach, with and without image constraints, by using HumanEva dataset [14]. This dataset comprises four subjects performing six different types of actions recorded in seven calibrated video sequences from different viewpoints. Additionally, the video sequences are synchronized with their corresponding motion captured 3D pose parameters.

In addition, the complete algorithm (with and without image constraint) has been implemented in Visual C++ using the OpenCV libraries [15] and it has been tested in a realtime interaction context on an Intel Core2 QUAD Q6600 under Windows Vista. First, without the image constraint, we have obtained a performance of 21 frames per second




FIGURE 7: Experiment 2. Initial (gray) and objective (black) configuration with the same end effector position. Left column, sequence applying inverse kinematics without image constraint. Right column, sequence applying inverse kinematics with image constraint.

(with 15 steps of convergence). Second, with the image constraint, we have obtained a performance of 19 frames per second (with 15 steps of convergence). Therefore, its use in human-computer interaction applications is also possible. For other uses in non real-time applications the accuracy could be improved adding more steps of convergence.

4.1. Virtual Environment. First, we test the system in a virtual environment to show how the presented approach works in an ideal situation. We define an articulated chain in 2D space, composed by 4 segments, with a total of





FIGURE 8: Experiment 3. From initial configuration of the chain (gray) the experiment try to avoid an object (black). Left column, sequence applying inverse kinematics without image constraint. Right column, sequence applying inverse kinematics with image constraint.

4 rotational joints of 1 DOF each one (i.e., a rotational joint for each segment). To test the system, we generate an initial configuration and an objective configuration of the articulated chain. Next, we apply the inverse kinematics approach, from initial configuration with and without image constraint, to estimate the objective configuration of the articulated chain.

In the first experiment, displayed in Figure 6, we generate an initial configuration and an objective configuration of the articulated chain, we apply the inverse kinematics approach, with and without image constraint. Without the image constraint, when the articulated chain reaches the end-effector the estimation of the objective configurations stops. On the other hand, by using image constraint it continues to try inserting the articulated chain inside the projection of the objective configuration, even if the articulated chain reaches the end-effector. In Figure 7 we display a second experiment where we generate an initial configuration and an objective configuration of the articulated chain, with the same end-effector. We apply the inverse kinematics approach, from the initial configuration, with and without image constraint. The results show that the inverse kinematics without the image constraint estimation does not change the initial configuration because the end-effector is reached. On the other hand, the inverse kinematics with image constraint enforces the chain to reach the image projection of the objective configuration. Figure 8 shows the last experiment in virtual environment, where the articulated chain tries to avoid an object



(a)



(b)



(c)




(d)

FIGURE 9: Working with real images. (a) Initial configuration of the articulated chain. (b) Arm estimation applying inverse kinematics approach. (c) And (d) arm estimation applying inverse kinematics with image constraints using one and two views respectively.

projected in an image. We define the initial configuration of the articulated chain, the objective end-effector, and a squared object. In this case, we use the image constraint to avoid the object. The experiments show that adding the image constraint outperforms inverse kinematics, in order to achieve the desired articulated chain configuration.


4.2. Using Real Images. In this test we apply the inverse kinematics with image constraint approach on a real stereoscopic sequence of human motions. Besides, the 3D joints' positions of the sequence are manually annotated for a quantitative comparison. The sequence has 450 frames corresponding to 15 seconds in real-time. The main objective of this test is to show that the proposed approach performs well with real




(a)



(b)



(c)




(d)

FIGURE 10: HumanEva *box* sequence for the subject 1, using Inverse kinematics.

images. In addition, this experiment shows that using this approach it is possible to solve the problem using only one image.

In Figure 9 is shown a frame of the stereoscopic sequence while applying the inverse kinematics approach with and without image constraint. We define an articulated chain in 3D space, with 2 rotational joints of 3 DOF each one. In the case of using image constraint, we apply

the approach firstly by using the left camera ($c = 1$) and then, by using both cameras ($c = 2$). The results show that when the inverse kinematics approach loses the elbow position, the inverse kinematics image constraint approach estimates the position of the elbow inside the silhouette. Using this stereoscopic sequence, there are not any significant differences between using one or both cameras.




(a)



(b)



(c)



(d)

FIGURE 11: HumanEva *box* sequence for the subject 1, applying Inverse kinematics using image-based constraints.

In order to perform a quantitative evaluation the mean squared error is used. Formally, the error between an estimated 3D joint X^e and the truly performed one X^{GT} from ground truth data is computed as

$$D(\mathbf{X}^e, \mathbf{X}^{GT}) = \frac{1}{i} \sum \|\mathbf{X}^e - \mathbf{X}^{GT}\| \quad (9)$$

where i is the number of frames. Specifically, we compare the manually annotated elbows positions and the estimated elbows positions using the inverse kinematics with image constraints in the case of one view and in the case of two views. In addition, we also can compare with the results of the priority inverse kinematics (PIK) approach [16] because the same test sequence is used to evaluate its performance.

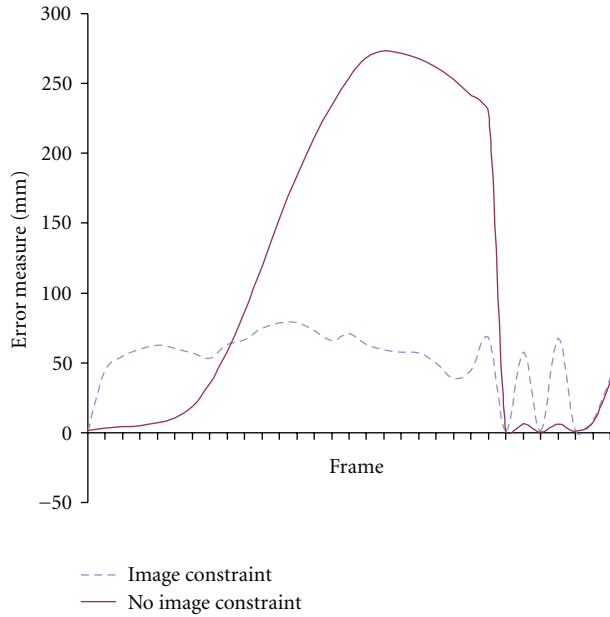


FIGURE 12: Elbow's estimation error by frame of the *box* sequence.




FIGURE 13: Knee's estimation error by frame of the *box* sequence.

Table 1 summarizes the results that shows how the inverse kinematics with image constraints has less error. It can also be observed how using the second view does not significantly improve the results.

TABLE 1: Comparison by using the manually annotated sequence.

	PIK (mm)	IBIK-1 view (mm)	IBIK-2 views (mm)
Left elbow	46.54	20.05	19.81
Right elbow	42.40	19.86	19.07

TABLE 2: Overall error of the estimation of the 3D positions of the internal joints for the two sequences (the elbow in the case of the arm and the knee in the case of the leg) for the two sequences.

	IK (mm)	IBIK (mm)
<i>Box</i> sequence	47.72	21.39
<i>Walking</i> sequence	40.69	16.35

4.3. HumanEva Test. This test evaluates our system using two views of two sequences of real motions, *walking* and *box*, of the subject 1 of HumanEva dataset. These sequences have a total of 3050 frames per view. Due to the fact that this database also contains the 3D positions of the joints by using markers, the objective of these experiments is to make a quantitative evaluation of our approach. We define an articulated chain in 3D space, with 2 rotational joints of 3 DOF each one, in order to estimate the configurations of the arm (for the *box* sequence) and the leg (for the *walking* sequence).

By using the mean square distance of (9), Table 2 shows the obtained results applying inverse kinematics with and without the image-based constraints.

Visual results for the *box* sequence are shown in Figures 10 and 11, where it is possible to see how the inverse kinematics approach loses the elbow position, and how, by adding the image-based constraints, the elbow's estimation lies inside the silhouette. This fact can also be observed in the graphic of Figure 12 where the elbow's estimation error by frame of the two approaches is displayed.


Figures 13, 14 and 15 show the results of the *walking* sequence, where a similar performance than previous sequence can also be observed. In this case, by adding the image-based constraint the recovered motion of the leg is more natural, avoiding the artifacts caused by the inverse kinematics approach in the knee's estimation.

5. Conclusion


The invasiveness of the sensor system, the high dimension of the posture space, and the modeling approximations in the mechanical model of the human body are sources of errors that accumulate and result in an approximate posture that could not be sufficient in biomechanical applications that study human motion precisely. Specifically, applications based on computer vision and inverse kinematics approaches presents the problem that no information was available to locate the internal joints and this forced the IK approach to make a somewhat arbitrary decision about what was the optimal angle for these joints.




(a)



(b)



(c)



(d)

FIGURE 14: HumanEva *walking* sequence for the subject 1, applying Inverse kinematics estimation.

In this paper, we present how to add image constraints to the inverse kinematics formulation in order to solve this problem. We have proposed a criterion that tries to guide the articulated chain to the body projection into the image. In this way, impossible chain configurations are avoided. Experiments using synthetic images show how this approximation performs correctly and, how to solve difficult situations that occur when there are motions that do not

imply to the end-effectors. Besides, we have evaluated our approach using real images, including sequences of a known human motion database in order to compute quantitative results. The computed error, about 2 centimeters, can be considered as sufficiently small to permit its use in motion capture applications. Moreover, adding the image constraint implies that the solution of the kinematic chain is more independent on initial configuration.

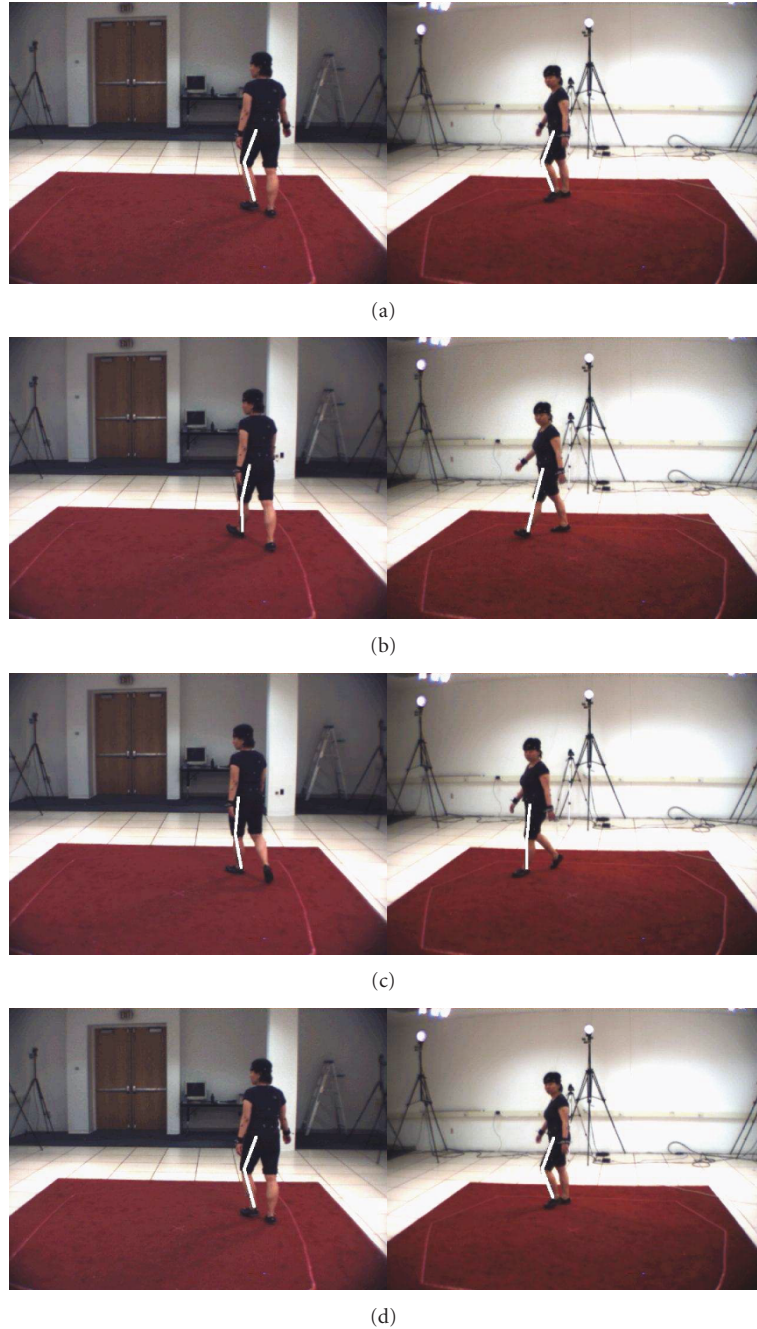


FIGURE 15: HumanEva walking sequence for the subject 1, applying Inverse kinematics using image-based constraints.

As future work, we plan to generalize this approach to include more tasks by using the priority strategy. In this way, it would be possible to use more complex models of the human body in order to achieve better estimations.

Acknowledgment

This work has been supported by the Spanish MEC under projects TIN2007-67993 and TIN2007-67896. Dr. J. Varona also acknowledge the support of a Ramon y Cajal (cofunded by the European Social Fund) Postdoctoral fellowship from the Spanish MEC.

References

- [1] L. Mündermann, S. Corazza, and T. P. Andriacchi, "The evolution of methods for the capture of human movement leading to markerless motion capture for biomechanical applications," *Journal of NeuroEngineering and Rehabilitation*, vol. 3, article 6, 2006.
- [2] T. B. Moeslund, A. Hilton, and V. Krüger, "A survey of advances in vision-based human motion capture and analysis," *Computer Vision and Image Understanding*, vol. 104, no. 2-3, pp. 90–126, 2006.
- [3] B. Zou, S. Chen, C. Shi, and U. M. Providence, "Automatic reconstruction of 3D human motion pose from uncalibrated

- monocular video sequences based on markerless human motion tracking,” *Pattern Recognition*, vol. 42, no. 7, pp. 1559–1571, 2009.
- [4] J. Varona, A. Jaume-i-Capó, J. González, and F. J. Perales, “Toward natural interaction through visual recognition of body gestures in real-time,” *Interacting with Computers*, vol. 21, no. 1-2, pp. 3–10, 2009.
 - [5] D. Tolani, A. Goswami, and N. I. Badler, “Real-time inverse kinematics techniques for anthropomorphic limbs,” *Graphical Models*, vol. 62, no. 5, pp. 353–388, 2000.
 - [6] J. Zhao and N. I. Badler, “Inverse kinematics positioning using nonlinear programming for highly articulated figures,” *ACM Transactions on Graphics*, vol. 13, no. 4, pp. 313–336, 1994.
 - [7] A. Liégeois, “Automatic supervisory control of the configuration and behavior of multibody mechanisms,” *IEEE Transactions on Systems, Man and Cybernetics*, vol. 7, no. 12, pp. 868–871, 1977.
 - [8] F. Chaumette and S. Hutchinson, “Visual servo control. II. Advanced approaches [Tutorial],” *IEEE Robotics and Automation Magazine*, vol. 14, no. 1, pp. 109–118, 2007.
 - [9] E. Marchand and N. Courty, “Controlling a camera in a virtual environment,” *Visual Computer*, vol. 18, no. 1, pp. 1–19, 2002.
 - [10] P. Baerlocher and R. Boulic, “An inverse kinematics architecture enforcing an arbitrary number of strict priority levels,” *Visual Computer*, vol. 20, no. 6, pp. 402–417, 2004.
 - [11] C. A. Klein, C. Chu-Jenq, and S. Ahmed, “A new formulation of the extended Jacobian method and its use in mapping algorithmic singularities for kinematically redundant manipulators,” *IEEE Transactions on Robotics and Automation*, vol. 11, no. 1, pp. 50–55, 1995.
 - [12] T. Horprasert, D. Harwood, and L. S. Davis, “A statistical approach for real-time robust background subtraction and shadow detection,” in *Proceedings of the 7th IEEE International Conference on Computer Vision, Frame Rate Workshop (ICCV ’99)*, pp. 1–19, Kerkyra, Greece, September 1999.
 - [13] D. G. Bailey, “An efficient euclidean distance transform,” in *Proceedings of the 10th International Workshop Combinatorial Image Analysis (IWCLA ’04)*, pp. 394–408, Auckland, New Zealand, December 2004.
 - [14] L. Sigal and M. J. Black, “Humaneva: synchronized video and motion capture dataset for evaluation of articulated human motion,” Tech. Rep. CS-06-08, Brown University, Providence, RI, USA, 2006.
 - [15] G. R. Bradski and V. Pisarevsky, “Intel’s computer vision library: applications in calibration, stereo, segmentation, tracking, gesture, face and object recognition,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR ’00)*, vol. 2, pp. 796–797, Hilton Head Island, SC, USA, 2000.
 - [16] R. Boulic, J. Varona, L. Unzueta, M. Peinado, A. Suescun, and F. Perales, “Evaluation of on-line analytic and numeric inverse kinematics approaches driven by partial vision input,” *Virtual Reality*, vol. 10, no. 1, pp. 48–61, 2006.