

Analyzing Gait Using a Time-of-Flight Camera

Rasmus R. Jensen, Rasmus R. Paulsen, and Rasmus Larsen

Informatics and Mathematical Modelling, Technical University of Denmark
Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark
{raje,rrp,r1}@imm.dtu.dk
www.imm.dtu.dk

Abstract. An algorithm is created, which performs human gait analysis using spatial data and amplitude images from a *Time-of-flight* camera. For each frame in a sequence the camera supplies cartesian coordinates in space for every pixel. By using an articulated model the subject pose is estimated in the depth map in each frame. The pose estimation is based on likelihood, contrast in the amplitude image, smoothness and a shape prior used to solve a Markov random field. Based on the pose estimates, and the prior that movement is locally smooth, a sequential model is created, and a gait analysis is done on this model. The output data are: Speed, Cadence (steps per minute), Step length, Stride length (stride being two consecutive steps also known as a gait cycle), and Range of motion (angles of joints). The created system produces good output data of the described output parameters and requires no user interaction.

Keywords: Time-of-flight camera, Markov random fields, gait analysis, computer vision.

1 Introduction

Recognizing and analyzing human movement in computer vision can be used for different purposes such as biomechanics, biometrics and motion capture. In biomechanics it helps us understand how the human body functions, and if something is not right it can be used to correct this.

Top athletes have used high speed cameras to analyze their movement either to improve on technique or to help recover from an injury. Using several high speed cameras, bluescreens and marker suits an advanced model of movement can be created, which can then be analyzed. This optimal setup is however complex and expensive, a luxury which is not widely available. Several approaches aim to simplify tracking of movement.

Using several cameras but without bluescreens nor markers [11] creates a visual hull in space from silhouettes by solving a spacial Markov random field using graph cuts and then fitting a model to this hull.

Based on a large database [9] is able to find a pose estimate in sublinear time relative to the database size. This algorithm uses subsets of features to find the nearest match in parameter space.

An earlier study uses the *Time-of-flight (TOF)* camera to estimate pose using key feature points in combination with an articulated model to solve problems with ambiguous feature detection, self penetration and joint constraints [13].

To minimize expenses and time spent on multi camera setups, bluescreens, markersuits, initializing algorithms, annotating etc. this article aims to deliver a simple alternative that analyzes gait.

In this paper we propose an adaptation of the *Posecut* algorithm for fitting articulated human models to grayscale image sequences by Torr et al. [5] to fitting such models to *TOF* depth camera image sequences. In particular, we will investigate the use of this *TOF* data adapted *Posecut* algorithm to quantitative gait analysis.

Using this approach with no restrictions on neither background nor clothing a system is presented that can deliver a gait analysis with a simple setup and no user interaction. The project object is to broaden the range of patients benefiting from an algorithmic gait analysis.

2 Introduction to the Algorithm Finding the Pose

This section will give a brief overview of the algorithm used to solve the problem of finding the pose of the subject. To do a gait analysis the pose has to be estimated in a sequence of frames. This is done using the adapted *Posecut* algorithm on the depth and amplitude stream provided by a *TOF* camera [2] (Fig. 1 shows a depth map with amplitude coloring). The algorithm uses 4 terms to define an energy minimization problem and find the pose of the subject as well as segmenting between subject and background:

Likelihood Term: This term is based on statistics of the background. It is based on a probability function of a given pixel being labeled background.

Smoothness Prior: This is a prior based on the general assumption that data is smooth. Neighbouring pixels are expected to have the same label with higher probability than having different labels.

Contrast Term: Neighbouring pixels with different labels are expected to have values in the amplitude map that differs from one another. If the values are very similar but the labels different, this is penalized by this term.

Shape Prior: Trying to find the pose of a human, a human shape is used as a prior.

2.1 Random Fields

A frame in the sequence is considered to be a random field. A random field consists of a set of discrete random variables $\{X_1, X_2, \dots, X_n\}$ defined on the index set I . In this set each variable X_i takes a value x_i from the label set $L = \{L_1, L_2, \dots, L_k\}$ presenting all possible labels. All values of $x_i, \forall i \in I$ are represented by the vector \mathbf{x} which is the configuration of the random field and takes values from the label set L^n . In the following the labeling is a binary problem, where $L = \{\text{subject}, \text{background}\}$.




Fig. 1. Depth image with amplitude coloring of the scene. The image is rotated to emphasize the spatial properties.

A neighbourhood system to X_i is defined as $N = \{N_i | i \in I\}$ for which it holds that $i \notin N_i$ and $i \in N_j \Leftrightarrow j \in N_i$. A random field is said to be a Markov field, if it satisfies the positivity property:

$$P(\mathbf{x}) > 0 \quad \forall \mathbf{x} \in L^n \quad (1)$$

And the Markovian Property:

$$P(x_i | \{x_j : j \in I - \{i\}\}) = P(x_i | \{x_j : j \in N_i\}) \quad (2)$$

Or in other words any configuration of \mathbf{x} has higher probability than 0 and the probability of x_i given the index set $I - \{i\}$ is the same as the probability given the neighbourhood of i .

2.2 The Likelihood Function

The likelihood energy is based on the negative log likelihood and for the background distribution defined as:

$$\Phi(\mathbf{D} | x_i = \text{background}) = -\log p(\mathbf{D} | x_i) \quad (3)$$

Using the Gibbs measure without the normalization constant this energy becomes:

$$\Phi(\mathbf{D} | x_i = \text{background}) = \frac{(\mathbf{D} - \mu_{\text{background},i})^2}{\sigma_{\text{background},i}^2} \quad (4)$$

With no distribution defined for pixels belonging to the subject, the subject likelihood function is set to the mean of the background likelihood function. To estimate a stable background a variety of methods are available. A well known method, models each pixel as a mixture of Gaussians and is also able to update these estimates on the fly [10]. In our method a simpler approach proved sufficient. The background estimation is done by computing the median value at each pixel over a number of frames.

2.3 The Smoothness Prior

This term states that generally neighbours have the same label with higher probability, or in other words that data are not totally random. The generalized Potts model where $j \in N_i$ is given by:

$$\psi(x_i, x_j) = \begin{cases} K_{ij} & x_i \neq x_j \\ 0 & x_i = x_j \end{cases} \quad (5)$$

This term penalizes neighbours having different labels. In the case of segmenting between background and subject, the problem is binary and referred to as the Ising model [4]. The parameter K_{ij} determines the smoothness in the resulting labeling.

2.4 The Contrast Term

In some areas such as where the feet touches the ground, the subject and background differs very little in distance. Therefore a contrast term is added, which uses the amplitude image (grayscale) provided by the *TOF* camera. It is expected that two adjacent pixels with the same label have similar intensities, which implies that adjacent pixels with different labels have different intensities. By decreasing the cost of neighbouring pixels with different labels exponentially with an increase in difference in intensity, this term favours neighbouring pixels with similar intensities to have the same label. This function is defined as:

$$\gamma(i, j) = \lambda \exp\left(\frac{-g^2(i, j)}{2\sigma_{background, i}^2}\right) \quad (6)$$

Where $g^2(i, j)$ is the gradient in the amplitude map and approximated using convolution with gradient filters. The parameter λ controls the cost of the contrast term, and the contribution to the energy minimization problem becomes:

$$\Phi(\mathbf{D}|x_i, x_j) = \begin{cases} \gamma(i, j) & x_i \neq x_j \\ 0 & x_i = x_j \end{cases} \quad (7)$$

2.5 The Shape Prior

To ensure that the segmentation is human like and wanting to estimate a human pose, a human shape model consisting of ellipses is used as a prior. The model is based on measures from a large Bulgarian population study [8], and the model is simplified such that it has no arms, and the only restriction to the model is that it cannot overstretch the knee joints. The hip joint is simplified such that the hip is connected in one point as studies shows that a 2D model can produce good results in gait analysis [3]. Pixels near the shape model in a frame are more likely to be labeled subject, while pixels far from the shape are more likely to be background.

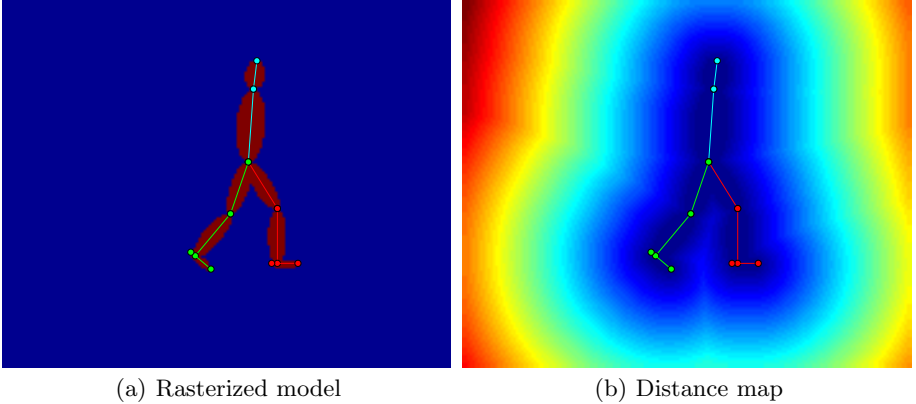


Fig. 2. Raster model and the corresponding distance map

The cost function for the shape prior is defined as:

$$\Phi(x_i|\Theta) = -\log(p(x_i|\Theta)) \quad (8)$$

Where Θ contains the pose parameters of the shape model being position, height and joint angles. The probability $p(x_i|\Theta)$ of labeling subject or background is defined as follows:

$$p(x_i = \text{subject}|\Theta) = 1 - p(x_i = \text{background}|\Theta) = \frac{1}{1 + \exp(\mu * (\text{dist}(i, \Theta) - d_r))} \quad (9)$$

The function $\text{dist}(i, \Theta)$ is the distance from pixel i to the shape defined by Θ , d_r is the width of the shape, and μ is the magnitude of the penalty given to points outside the shape. To calculate the distance for all pixels to the model, the shape model is rasterized and the distance found using the *Signed Euclidian Distance Transform (SEDT)* [12]. Figure 2 shows the rasterized model and the distances calculated using the *SEDT*.

2.6 Energy Minimization

Combining the four energy terms a cost function for the pose and segmentation becomes:

$$\Psi(\mathbf{x}, \Theta) = \sum_{i \in V} \left(\Phi(\mathbf{D}|x_i) + \Phi(x_i|\Theta) + \sum_{j \in N_i} (\psi(x_i, x_j) + \Phi(\mathbf{D}|x_i, x_j)) \right) \quad (10)$$

This Markov random field is solved using *Graph Cuts* [6], and the pose is optimized in each frame using the pose from the previous frame as initialization.

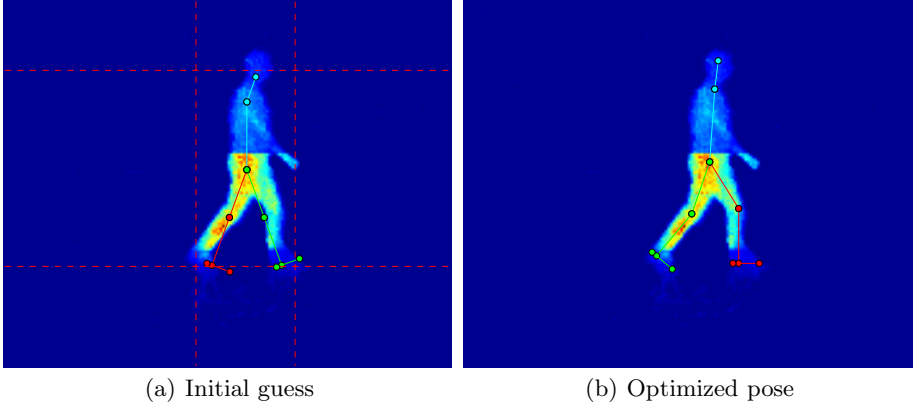


Fig. 3. Initialization of the algorithm

2.7 Initialization

To find an initial frame and a pose, the frame that differs the most from the background is chosen based on the background log likelihood function. As a rough guess on where the subject is in this frame, the log likelihood is summed first along the rows and then along the columns. These two sum vectors are used to guess the first and last rows and columns that contains the subject (Fig 3(a)). From the initial guess the pose is optimized according to the energy problem by searching locally. Figure 3(b) shows the optimized pose. Notice that the legs change place during the optimization. This is done based on the depth image such that the closest leg is also closest in the depth image (green is the right side in the model) and solves an ambiguity problem in silhouettes.

The pose in the remaining frames is found using the previous frame as an initial guess and then optimizing on this. This generally works very well, but problems sometimes arise when the legs pass each other as feet or knees of one leg tend to get stuck on the wrong side of the other leg. This entanglement is avoided by not allowing crossed legs as an initial guess and instead using straight legs close together.

3 Analyzing the Gait

From the markerless tracking a sequential model is created. To ensure local smoothness in the movement before the analysis is carried out a little postprocessing is done.

3.1 Post Processing

The movement of the model is expected to be locally smooth, and the influence of a few outliers is minimized by using a local median filter on the sequences of

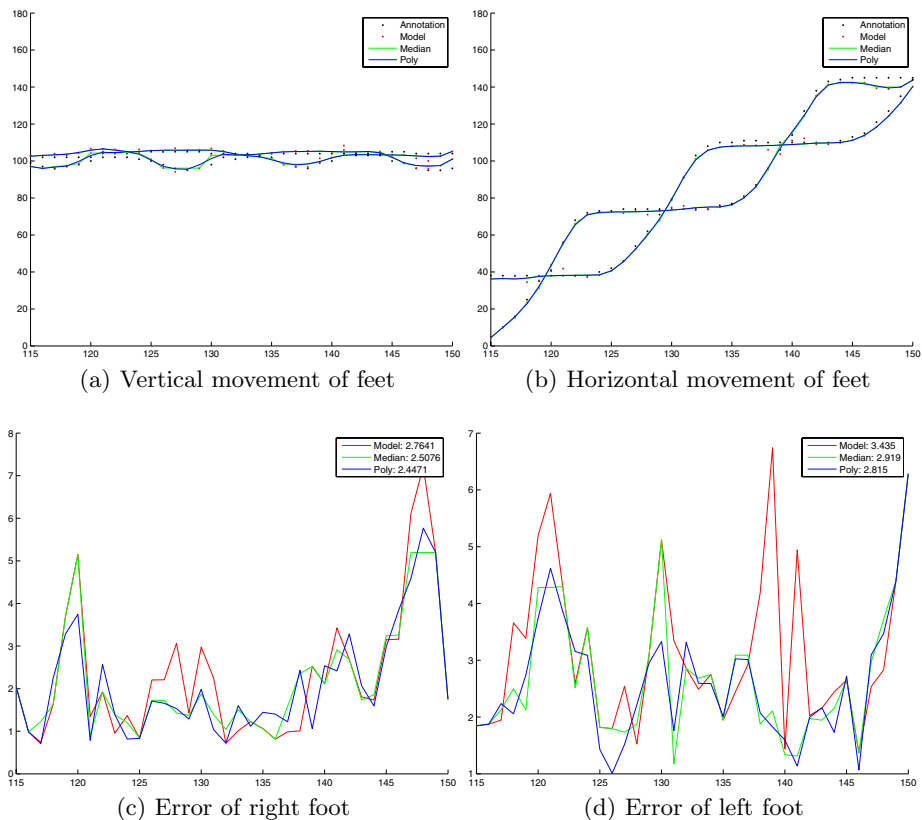


Fig. 4. (a) shows the vertical movement of the feet for annotated points, points from the pose estimate, and for curve fittings (image notation is used, where rows are increased downwards). (b) shows the points for the horizontal movement. (c) shows the pixelwise error for the right foot for each frame and the standard deviation for each fitting. (d) shows the same but for the left foot.

point and then locally fitting polynomials to the filtered points. As a measure of ground truth the foot joints of the subject has been annotated in the sequence to give a standard deviation in pixels of the foot joint movement. Figure 4 shows the movement of the feet compared to the annotated points and the resulting error. The figure shows that the curve fitting of the points gives an improvement on the accuracy of the model, resulting in a standard deviation of only a few pixels. If the depth detection used to decide which leg is left and which is right fails in a frame, comparing the body points to the fitted curve can be used to detect and correct the incorrect left right detection.

3.2 Output Parameters

With the pose estimated in every frame the gait can now be analyzed. To find the steps during gait, the frames where the distance between the feet has a

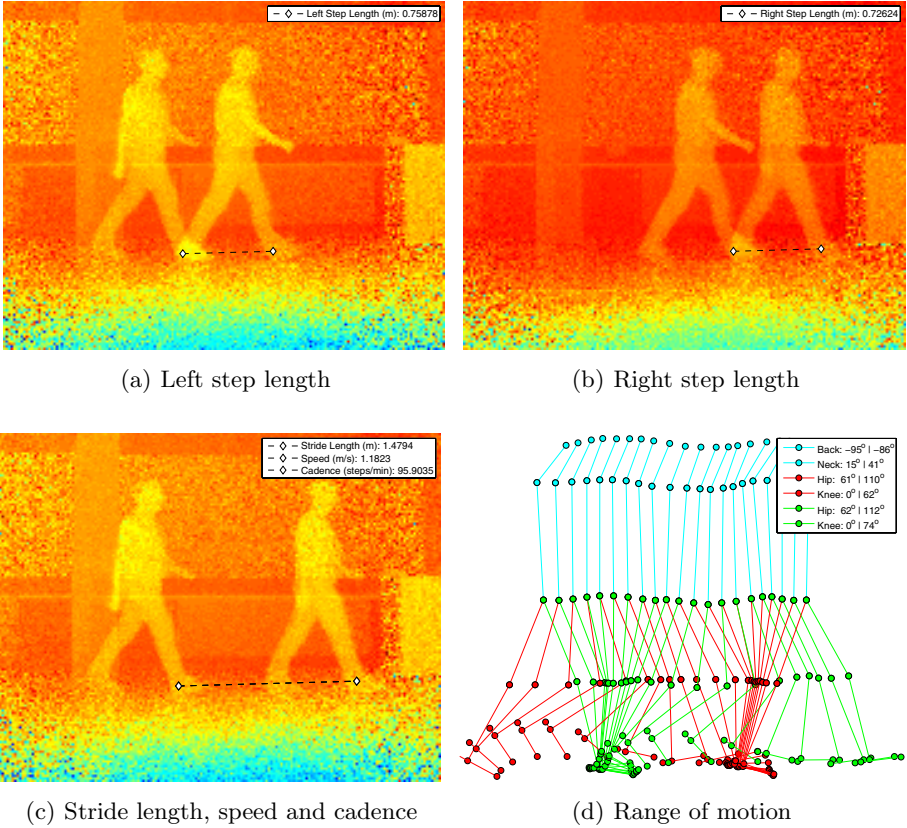


Fig. 5. Analysis output

local maximum are used. Combining this with information about which foot is leading, the foot that is taking a step can be found. From the provided Cartesian coordinates in space and a timestamp for each frame the step length (Fig. 5(a) and 5(b)), stride length, speed and cadence (Fig. 5(c)) are found. The found parameters are close to the average found in a small group of subjects aging 17 to 31 [7], even though based only on very few steps and therefore expected to have some variance, this is an indication of correctness. The range of motion is found as the clockwise angle from the x-axis in positive direction for the inner limbs (femurs and torso) and the clockwise change compared to the inner limbs for the outer joints (ankles and head). Figure 5(d) shows the angles and the model pose throughout the sequence.

4 Conclusion

A system is created that autonomously produces a simple gait analysis. Because a depth map is used to perform the tracking rather than an intensity map,

there are no requirements to the background nor to the subject clothing. No reference system is needed as the camera provides a such. Compared to manual annotation in each frame the error is very little. For further analysis on gait the system could easily be adapted to work on a subject walking on a treadmill. The adaption would be that there is no longer a general movement in space (it is the treadmill conveyor belt moving) hence speed and stride lengths should be calculated using step lengths. With the treadmill adaption, averages could be found of the different outputs as well as standard deviations.

Currently the system uses a 2-dimensional model and to optimize precision in the joint angles the subject should move in an angle perpendicular to the camera. While the distances calculated depends little on the angle of movement the joint angles have a higher dependency. This dependency could be minimized using a 3-dimensional model. It does however still seem reasonable that the best results would come from movement perpendicular to the camera, whether using a 3-dimensional model or not.

The camera used is the SwissRanger™ SR3000 [2] at a framerate of about 18 Fps, which is on the low end in tracking movement. A better precision could be obtained with a higher framerate. This would not augment processing time greatly, due to the fact that movement from one frame to the next will be relatively shorter, bearing in mind that the pose from the previous frame is used as an initialization for the next.

Acknowledgements

This work was in part financed by the ARTTS [1] project (Action Recognition and Tracking based on Time-of-Flight Sensors) which is funded by the European Commission (contract no. IST-34107) within the Information Society Technologies (IST) priority of the 6th framework Programme. This publication reflects only the views of the authors, and the Commission cannot be held responsible for any use of the information contained herein.

References

1. Artts (2009), <http://www.artts.eu>
2. Mesa (2009), <http://www.mesa-imaging.ch>
3. Alkjaer, E.B., Simonsen, T., Dygre-Poulsen, P.: Comparison of inverse dynamics calculated by two- and three-dimensional models during walking. In: 2001 Gait and Posture, pp. 73–77 (2001)
4. Besag, J.: On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)* 48(3), 259–302 (1986)
5. Bray, M., Kohli, P., Torr, P.H.S.: Posecut: simultaneous segmentation and 3D pose estimation of humans using dynamic graph-cuts. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3952, pp. 642–655. Springer, Heidelberg (2006)
6. Kolmogorov, V., Zabini, R.: What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26(2), 147–159 (2004)

7. Latt, M.D., Menz, H.B., Fung, V.S., Lord, S.R.: Walking speed, cadence and step length are selected to optimize the stability of head and pelvis accelerations. *Experimental Brain Research* 184(2), 201–209 (2008)
8. Nikolova, G.S., Toshev, Y.E.: Estimation of male and female body segment parameters of the bulgarian population using a 16-segmental mathematical model. *Journal of Biomechanics* 40(16), 3700–3707 (2007)
9. Shakhnarovich, G., Viola, P., Darrell, T.: Fast pose estimation with parameter-sensitive hashing. In: *Proceedings Ninth IEEE International Conference on Computer Vision*, vol. 2, pp. 750–757 (2003)
10. Stauffer, C., Grimson, W.E.L.: Adaptive background mixture models for real-time tracking. In: *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, vol. 2, pp. 246–252 (1999)
11. Wan, C., Yuan, B., Miao, Z.: Markerless human body motion capture using Markov random field and dynamic graph cuts. *Visual Computer* 24(5), 373–380 (2008)
12. Ye, Q.-Z.: The signed Euclidean distance transform and its applications. In: *1988 Proceedings of 9th International Conference on Pattern Recognition*, vol. 1, pp. 495–499 (1988)
13. Zhu, Y., Dariush, B., Fujimura, K.: Controlled human pose estimation from depth image streams. In: *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, pp. 1–8 (2008)