Gait Analysis with a Time-Of-Flight Camera

Rasmus Ramsbøl Jensen
Informatics and Mathematical Modelling
Technical University of Denmark
Email: raje@imm.dtu.dk

Introduction

In computer vision analysis of movement serves several purposes. It can be used in motion capture, which drives today's animation films. Studies have also been made on identifying people based on a unique pattern of movement. This kind of identification has the advantage that it can be done without interacting with the actor and also under varying quality and lighting conditions. Finally, the tracking of movement can be used in biomechanics. Either serving as a general study or as an assessment tool in the training of athletes or in the rehabilitation of patients.

The object of this project is to create a system that does gait analysis and requires minimal user interaction and only a simple setup. Such a system could bring the benefits of vision based gait analysis to the everyday life of the common physiotherapist.

The data

The image data are from a Time-Of-Flight (TOF) camera, which provides spatial data as well as intensity measures of the scene. The camera emits a modulated signal from an array of near infrared LEDs surrounding the lens. By calculating the phase shift between the emitted and the recorded signal a distance map of the scene is created. Knowing the lens properties these distances are converted into Euclidian coordinates of the scene. An intensity map is created by multiplying the amplitude by the square of the distances.

Tracking the subject

To track the subject throughout the sequence a Markov Random Field (MRF) is created for each frame based on several terms using both the depth and the intensity map. The algorithm used is a version of the Posecut algorithm adapted to work on both the spatial and intensity data from the TOF camera. The terms used are a likelihood function, a smoothness prior, a gradient term and a shape prior and the solution to the MRF is binary labelling only subject and background.

The likelihood function is based on a statistical depth model of the background and a likelihood that a pixel belongs to either subject or background. The statistical model is simply based on the frames before and after, the subject enters and exits the scene.

The smoothness prior is a simple assumption that data are grouped and not random. It penalizes ungrouped labelling.

Abstract

A system is presented that performs gait analysis based on spatial and intensity data from a Time-Of-Flight camera. The system requires no user interaction and there are no restrictions on the scene or the subject clothing. The project motivation is to create a system that brings algorithmic gait analysis to the everyday physiotherapist using a simple setup and with little manual time consumption.

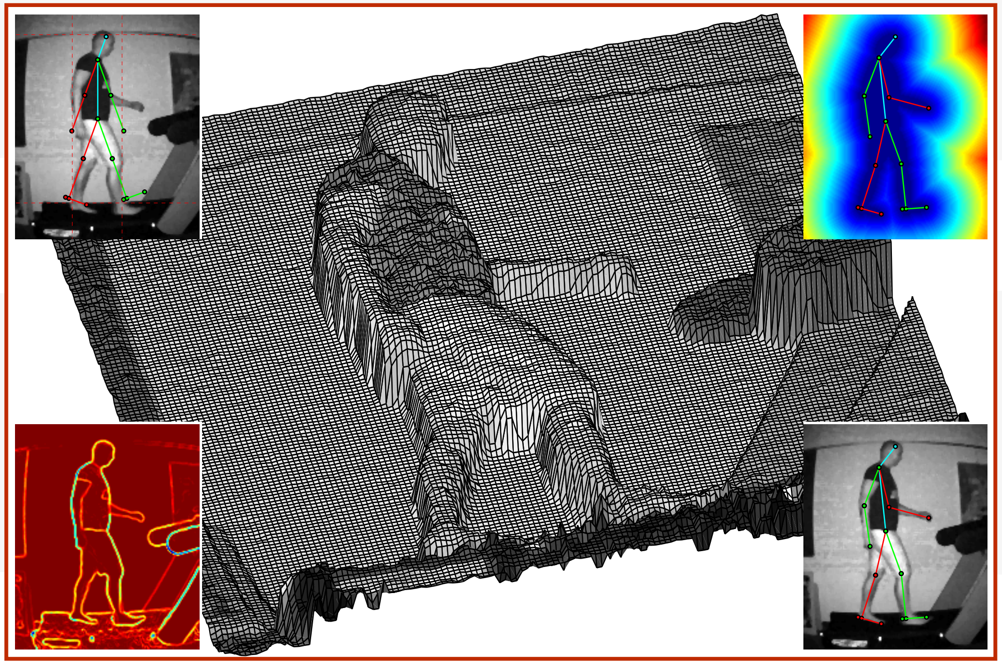

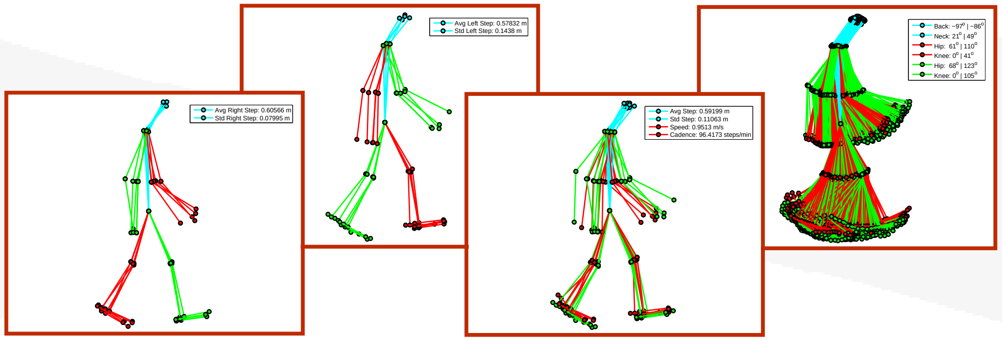

Tracking...continued

The gradient term works on the edges of the labelling and penalizes, if neighbours have different label in an area with very low gradient. This term is based on a mix of both depth and intensity.

Finally, a shape prior ensures humanlike segmentation using a simple human model. Labelling of subject has higher probability close to the model and high probability of background far from the model. The distance is calculated using Signed Euclidian Distance Transform (SEDT).

The algorithm is initialized by finding a frame and an area that differs a lot from the background model. As a rough guess on the initial pose the shape model is placed in this area using the vertical size as the model height. By optimizing the model pose the field energy is minimized. The MRF is solved using Graph Cuts. As an initial guess on the pose in the following frames the optimized pose is used and this process is repeated. The tracking of left and right is done in the depth map.

Analyzing the model

Having the pose estimated in every frame a sequential model is created that can be analyzed. To ensure smooth movement a smoothness prior is imposed on the movement. This is done by locally fitting low order polynomials to every point of the model as a function of time.

From the smoothened sequential model the steps are found by finding the local minima when the feet are furthest apart. This is used to calculate the step lengths, and from the timestamp the speed and cadence are also calculated.

Because the subject is walking on a treadmill the movement is a balanced cycle. The gait is cut into cycles, which are phase aligned. By fitting Fourier series to the new phase aligned data a cyclic model is created of the subject gait.

Conclusion

The system produces several gait parameters with no user interaction and with no restrictions on neither the background nor the subject clothing.

Because the camera provides a reference system there is no need for calibration or reference.

Studies against manual annotation in the sequence shows a very high correlation. The speed of the treadmill matches the output speed of the system perfectly.