Determining Height and Gender of a Subject Using Gait

Michael Kuhlman, Clarkson University*

Abstract

The aim of this study is to determine the height and gender of a subject from a video feed using computer vision algorithms. Psychological studies have shown that most sex-specific traits of gait are variations in lateral motion in the pelvic and shoulder regions. We compared the variance of movement from the traditional side view to the front view for determining gender. The front view captures the most lateral motion and therefore is the best view for determining gender.

We used Poser 7, a design and animation software package, to model walking subjects and to determine the accuracy of current height-estimation algorithms. We compared results from computer generated video to real world video using the same height-estimation algorithm.

1. Introduction

Gait as a biometric has many advantages over other forms of human identification. One can analyze a subject's gait at distance and it is noninvasive, which makes it suitable for security applications. Facial recognition, on the other hand, requires very high resolution pictures, whereas a silhouette for gait analysis consists of a few hundred pixels. Most studies of gait analysis are for human identification. However, for security applications, it would be useful to be able to determine physical attributes of a subject, such as gender, height, weight, etc to decrease the number of suspects.

Historicallly, the field of gait analysis has been dominated by males, and naturally, the majority of subjects are also men (as opposed to psychophysics experiments, where there is less skew in the gender ratio). Typically, in most datasets, less than one quarter of the subjects on average are women. Given that the goal of this project is gender classification, the limiting factor of a database's usefulness is the number of women in it. The experiments were performed on the USF, UMD and CMU databases. We also looked into generating our own silhouettes using the animation software Poser 7.

---

*This was work that was done under the supervision of Prof. Rama Chellappa, as a part of the MERIT program at the University of Maryland.

2. Gender Discrimination

2.1 Psychophysics and Gait Analysis

Studies in psychophysics have shown that humans can determine the gender of a walker from a point light display. A point light display only shows specific parts of the body; no other cues are given other than motion/location of these body parts as seen in Figure 1. From the side view, humans have a recognition rate of merely 63%. Later studies have shown that humans can identify gender properly with a recognition rate of 79% from the front [2]. [2] argued that the major gender-specific component of gait is the ratio of the lateral velocity of the shoulders to the lateral velocity of the hips. The majority of gender-specific information in gait is therefore lost when looking from the side view. Only subtle cues from the ankles can be seen from the side for gender identification.

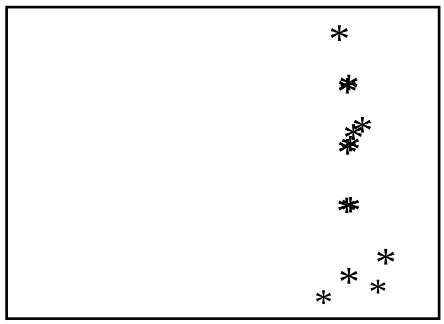

Figure 1. An example of a point-light-display representation of a subject from the side view.

There are many methods currently employed for gait analysis [5]. Such methods that analyze side-view silhouettes, however, are for human identification purposes, not for gender discrimination. Methods such as hidden Markov models (HMM) and dynamic time warping (DTW) have not been used to distinguish the gender of the subject.

2.2 Gender Discrimination from the Side View

There are few gender-specific traits that humans can pick up from the side view. There are subtle differences in how men and women move their ankles while walking [2]. Studies that utilized Dynamic Time Warping and partial Procrustes distance [3, 4] could not classify subjects by gender when using gait sequences from the USF dataset.

2.3 Gender Discrimination from the Front View–Variance Images

Since localized lateral velocity contains gender-specific information (as opposed to shape features), we decided to analyze the variance of front-view silhouette changes since feature extraction is more difficult. We calculated the variance of each pixel’s value across an entire gait cycle to make a variance image which represented the movement of the subject in a single picture. The variance of each row was summed, converting the variance image into a single column vector. We then compared column vectors of women to column vectors of men to try to find gender-specific trends in variance.

2.4 UMD Database

Figure 2. Silhouettes of two subjects from UMD database.

The UMD dataset consists of gait sequences from both the front and side view. Eight men and women were selected at random from the UMD database. We analyzed one complete gait cycle from the front and side.

2.5 CMU Database

Figure 3. Silhouettes of two subjects from CMU database.

The CMU database consists of 25 subjects and two women walking on a treadmill. We randomly chose six males and took the two female and generated variance images of three gait cycles from each subject.

3. Height Estimation

There exist automated methods for measuring heights of moving objects viewed by stationary cameras [6]. Using a reference cube as shown in Figure 4, we can calculate the vanishing points of a scene using metrology. We then apply single-view mensuration algorithm to each of the video frames to estimate the height of the subject. This height-estimation algorithm, however, has a positive bias for unknown reasons. In addition, no comparisons using ground-truth data have been done. Such biases are similar for both real and computer generated gait sequences. Using Poser 7, we attempted to track down what types of movement cause the small bias in height estimation. We constructed several computer generated sequences, locking various types of motion and ran them through the algorithm. So far, we have not found a specific form of movement which causes the positive bias. We will have to conduct more research regarding ground truth in Poser to better understand why there is positive bias in this algorithm.

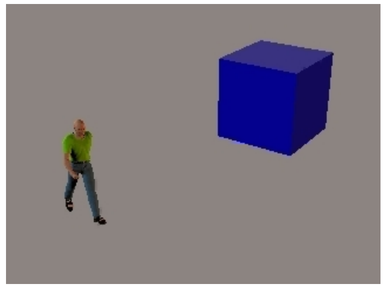

Figure 4. Computer generated image from Poser 7 to better understand positive bias in height estimation algorithm.

4. Poser 7

Poser 7 is a human modeling and animation software package that is worth noting. Although originally intended for amateur use in computer animation and photorealism, Poser has many built in features that make it useful for scientific human modeling. As a quick pilot study for this project, for example, we were able to generate synthetic data of subjects walking from the front view.

Poser 7's key feature relevant to gait analysis is it's built in walk designer. The walk designer has several "blends" and "tweaks" that lets you change certain parameters of the gait. This tweaked walk can then be applied to a specific model for a desired length of time (in frames). The model will then walk along a predefined path in space or in place. A rendered woman walking is shown in Figure 5.

Figure 5. ¾ view of a woman walking in Poser 7.

Although Poser only includes a few models, each character has several body "morphs" that can change body types (ex: stocky, pear-shaped, etc) or feature proportions (arm length, leg width, etc). Combining body morphs with walk tweaks will result in enough samples of unique characters. Third party models and walk designers can be purchased from DAZ 3D and other websites.

Depending on the preview settings, one can generate silhouettes of walking subjects. These computer generated silhouette sequences, however, are far more accurate than the product of any background subtraction algorithm (individual strands of hair can be seen in the silhouette)

Figure 6. Computer generated silhouettes from Poser.

Other potential uses for computer generated sequences would include calibration of algorithms, effects of frame rate on human identification, controls for experiments, and the ability to produce androgynous figures for gender classification.

5. Results

5.1 Results on the UMD dataset

Gender could not be discriminated from variance images of side-view sequence from the USF and UMD databases. Gender could also not be discriminated from front-view sequences in the UMD database. Results were inconclusive with the UMD database due to poor background subtraction. There is noticeable noise in Figure 7.

Figure 7. Variance Images from UMD database.

5.2 Results on the CMU dataset

Given the small number of women in the CMU database, clustering was our only option for analysis. Results for gender classification using variance look promising but is still inconclusive. The analysis of the variance images as seen in Figure 8 is rather simple using the CMU dataset. We first sum up the variance of each row producing a variance vector. Then, we normalize the vectors to accommodate for subject height. Finally, we take the variance of the variance vectors and the results are clustered rather nicely as seen in Figure 9. This same relationship of variance values holds true for the UMD dataset, yet values are not nearly as well clustered as the CMU dataset.

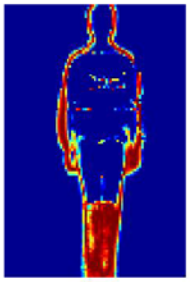 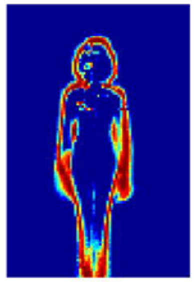

Figure 8. Variance Images from CMU database (not scaled for height).

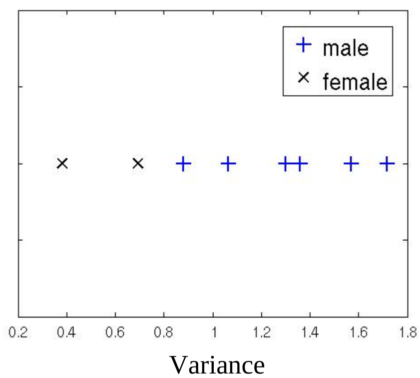

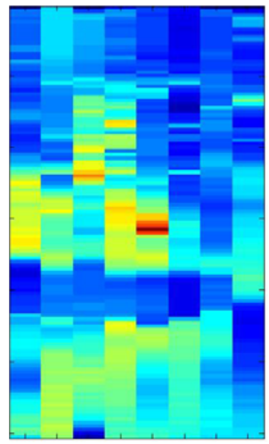

Figure 9. Results of the CMU dataset. (Left) Variance of variance vectors of subjects. (Right) Variance vectors of variances images (right most two are female subjects; scaled for height).

6. Discussion

The psychological experiments previously mentioned took place on treadmills, giving an absolute coordinate system that greatly aids both human and computer vision. Results from the UMD database were invalid partly due to the subjects moving. The subjects were also not walking directly towards the camera.

Variance also does not distinguish lateral motion from vertical motion. We observed that in both computer generated and UMD front-view gait sequences, the centroid of the silhouette has greater vertical variance than lateral variance by a factor of 2. In order to have properly scaled images when tracking a moving object, we put the X value of the centroid of the head of the silhouette in the center of the frame, possibly tampering with lateral motion.

6.1 CMU Dataset

Why is there a global variance difference in the CMU dataset as opposed to localized differences around the hips and shoulders? This method does not differentiate changes between the arms and hips, which move along the same rows of the image (this information is lost when we generate the variance vector). Arm swing is dependent on shoulder swing [2], and the arms temporarily block the view of the hip for a portion of the gait cycle. It is highly probable that the variance due to the arm swing is greater than the variance due to hip swing by several magnitudes, making hip variance measurements insignificant. This explains why there is greater total variance among men over all.

Taking the variance of the variance vectors is not the only way to filter the data. Similar trends are found with summing all variances in the variance vector (note that height is normalized which does not affect summations). Better clustering for this small set of data was found by taking the variance of the variance vectors. Taking the variance of the entire variance image was not considered since the variance images themselves were not normalized for height.

6.2 Further work:

Many human identification algorithms focus more on the legs than the arms. This is because carrying objects (commonplace on the field) changes arm motion [5]. Since variance image analysis focuses greater on shoulder and arm movement than hip and leg movement, it would be interesting to see how well variance can determine gender when subjects are carrying objects. Better tracking methods would also allow gender discrimination of moving subjects. Finally, more data of women walking on treadmills will be required to further the study of gender discrimination from the front.

Poser 7's capabilities of computer simulations for gait analysis should be researched further. Ground truth is not currently built into Poser and we wish to develop a method for establishing the ability to determine a figure's exact height in Poser. This will be helpful in calibrations of various video-fed algorithms.

Acknowledgements

I would like to thank Dr. Rama Chellappa for his mentoring and Ashok Veeraraghavan and Aswin C. Sankaranarayanan for their valuable insights.

References

[1] L.D. Kozlowski and J. E. Cutting, "Recognizing the Sex of a Walker from a Dynamic Point Light Display," *Perception and Psychophysics*, **21,** pp. 575-580, 1977

[2] G. Mather and Linda Murdoch, "Gender Discrimination in Biological Motion Displays Based on Dynamic Cues," *Proceedings*: *Biological Sciences*, **258**, pp 273-279, 1994

[3] A Veeraraghavan, R. Chellappa, and A.K. Roy-Chowdhury, "The Function Space of an Activity," *CVPR, 2006*

[4] A Veeraraghavan, A.K. Roy-Chowdhury, and R. Chellappa, "Matching Shape Sequences in Video with Applications in Human Movement Analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **27**, pp 1896-1909, 2005.

[5] M. S. Nixon, T. N. Tan, and R. Chellappa, "Human Identification Based on Gait," *Springer-Verlag New York, Inc. Secaucus, NJ, USA*, 2006.

[6] Shao, J. and Zhou, S. and Chellappa, R., **"**Video Metrology**",** under review, 2006.