2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

Human Motion Patterns from Single Camera Cues for Medical Applications

Evan Ribnick, Vassilios Morellas, and Nikolaos Papanikolopoulos
{ribnick, morellas, npapas}@cs.umn.edu
University of Minnesota, Minneapolis, MN, USA

Abstract— **Physical constraints that underly the formation of periodic motions can be effectively used to accurately reconstruct the periodic motion from even single camera views. As shown in our earlier work, this reduces to a problem of geometric inference. In this paper, we focus on periodic motions exhibited by humans, which are generally not perfectly periodic, and explore the suitability of the reconstruction techniques in these scenarios. We examine the degree of periodicity of human gait empirically, including the applicability of our motion model. Importantly, we illustrate the usefulness of these techniques by applying them to the task of clinical gait analysis. A computational tool to analyze periodic human motion can prove to be invaluable in medical applications either in terms of assessing deviations from normal patterns or evaluating changes resulting from therapy or other clinical procedures.**

I. INTRODUCTION

Humans exhibit a variety of periodic motion patterns in everyday activities and as such periodicity has been recognized as an important cue when analyzing human motion. There exists a dominant trend in the literature of methods performing analysis of periodic motion in image coordinates. This trend however prevents the use of these techniques in tasks such as motion analysis, recognition, and classification, since the appearance of a motion in image coordinates varies greatly with the viewing angle.

A great deal of work related to periodic motion has focussed on detection and analysis in image coordinates. Several techniques use Fourier analysis of pixel-coordinate movements to detect, segment, or classify motions [1]–[7]. Recent work has shown the use of Fourier analysis to segment and extract multiple periodic motions in video sequences simultaneously [5]. Other types of image-based analysis for detection, representation, and classification of repetitive motions, including human gestures and facial expressions have been focused in the literature [8]–[10].

Some have recognized the important connection between periodic motion and 3D inference. A method for estimating the structure of an articulated body undergoing repetitive motion by considering snapshots separated by exactly one period in time using geometric constraints is reported in [14]. In another method, training data is used to learn Fourier-based representations of periodic human motions which can be used to infer their structure in 3D [13]. More recently, work on computing the 3D trajectory of periodic motions based on their image-coordinate trajectories and geometric constraints can be found in [11], [12]. Regarding the applicability of the method appearing in [13] compared to the one in [12], [13] is limited to the trajectories seen in training data and an orthographic projection is assumed, indicating reconstruction only for motions without a translational component.

Vision-based human motion analysis and the closely related task of activity classification have drawn a great deal of interest from both a medical and a security and surveillance perspective. There is a large body of research in this field. The reader is referred to the survey papers [15]–[17], [19]. Research in this area can be generally divided into two major categories: (i) those that perform analysis from a single camera view (monocular systems), and (ii) those that require multiple views (e.g., binocular and trinocular stereo). Typically they either operate directly in image coordinates [20], or learn mappings from image-coordinate appearances to 3D state vectors [21]. These techniques are affected by the viewing angle of the camera and one can only hope to overcome this constraint by learning classifiers based on motions viewed from many different angles.

Gait analysis in a clinical setting which is the main focus of this paper, is another very important application domain. The classical work in [22] gives a thorough characterization of the gait of adult men, and clearly illustrates the connection between a person's gait and the presence/absence of physical debilitations. Technology-assisted gait analysis for medical diagnosis is rather sparse. Gait is a useful cue for diagnosis of pathological disorders, specifically cerebral palsy and poliomyelitis and attention deficit disorders [23]. Signals of interest include motion patterns of lower extremities. Fourier analysis is performed on the resulting signals, and self-organizing maps are used to cluster the data. Data from normal and pathological subjects are then well separated. Similarly, [24] develops a system for diagnosing Parkinson's disease. The system analyzes static images of the walking person wearing a conveniently colored track suit, and an artificial neural network is trained in order to distinguish between diseased and normal subjects based on several features extracted from the images. Other works use commercial motion capture systems to analyze gait [25]–[27] .

Many motions of interest, including walking and running, are inherently repetitive in nature. Nonetheless, it is clear

This work is supported in part by the Department of Homeland Security, the Center for Transportation Studies and the ITS Institute at the University of Minnesota, the Minnesota Department of Transportation, the U.S. Army Research Laboratory and the U.S. Army Research Office, and the National Science Foundation.

that human motion is quite complex and and does not follow any rigid mathematical model. In order to effectively analyze these types of movements, a more complete understanding of the periodicity exhibited by natural human motion is required. In light of these challenges, this paper aims to examine the suitability of techniques which are predicated on periodicity for real human motion, and illustrates their usefulness by considering the application of clinical gait analysis.

II. Periodic Motion Reconstruction

In this section, we briefly review the basic formulation and technique for reconstructing periodic motions in 3D. For more detailed discussion, the reader is referred to [11], [12].

A. Definitions and Basic Equations

In this article, we define periodic motion as any movement that is periodic in terms of its velocity (in 3D world coordinates). In terms of the 3D position of the point, we have:

$$\mathbf{p}(t+nT) = \mathbf{p}(t) + n\boldsymbol{\Delta}_{\mathbf{p}_T}, \tag{1}$$

where $\boldsymbol{\Delta}_{\mathbf{p}_T} \triangleq (\Delta_{X_T}, \Delta_{Y_T}, \Delta_{Z_T})$ is the displacement per period of the point, which is constant over any period of length T. For example, if the point being tracked is on the foot of a walking person, then the stride length is equal to $\|\boldsymbol{\Delta}_{\mathbf{p}_T}\|_2$.

Since samples are taken at discrete times determined by the video frame rate, we represent times using discrete indices of the form t_k^i. This represents the time of the k-th sample in the i-th period. We can then arrive at the following expression for the position at time t_k^i:

$$\begin{pmatrix} X_k^i \\ Y_k^i \\ Z_k^i \end{pmatrix} = \begin{pmatrix} X_k^0 \\ Y_k^0 \\ Z_k^0 \end{pmatrix} + i \begin{pmatrix} \Delta_{X_T} \\ \Delta_{Y_T} \\ \Delta_{Z_T} \end{pmatrix}, \tag{2}$$

where the sample at time t_k^i is expressed relative to t_k^0, the k-th sample in the $0th$-period.

When a periodic motion is projected into the image using the pinhole camera model, we arrive at the pixel-coordinate trajectory described by the following equation:

$$\begin{pmatrix} u_k^i \\ v_k^i \end{pmatrix} = \frac{1}{Z_k^0 + i\Delta_{Z_T}} \begin{pmatrix} f_x & 0 \\ 0 & f_y \end{pmatrix} \begin{pmatrix} X_k^0 + i\Delta_{X_T} \\ Y_k^0 + i\Delta_{Y_T} \end{pmatrix} + \begin{pmatrix} c_x \\ c_y \end{pmatrix}. \tag{3}$$

Note that we have placed the origin of the world coordinate system at the camera center, with the Z-axis parallel to the camera's optical axis. The quantities f_x, f_y, c_x, and c_y are intrinsic parameters of the camera representing the focal length and image plane center in pixel units. Thus, we see that it is possible to express the projection of any sample projected into image coordinates, (u_k^i, u_k^i), as a function of the corresponding 3D sample at time t_k^0 and the inter-period displacement $\boldsymbol{\Delta}_{\mathbf{p}_T}$ (3).

B. Minimizing the 3D Geometric Error

Rearranging the terms in (3), we can obtain expressions for X_k^0 and Y_k^0 in terms of estimates of Z_k^0 and $(\Delta_{X_T}, \Delta_{Y_T}, \Delta_{Z_T})$:

$${}^i\hat{X}_k^0 = \frac{u_k^i - c_x}{f_x}(\hat{Z}_k^0 + i\hat{\Delta}_{Z_T}) - i\hat{\Delta}_{X_T} \tag{4}$$

$${}^i\hat{Y}_k^0 = \frac{v_k^i - c_y}{f_y}(\hat{Z}_k^0 + i\hat{\Delta}_{Z_T}) - i\hat{\Delta}_{Y_T}, \tag{5}$$

where " $\hat{}$ " denotes that a quantity is an estimate, and ${}^i\hat{X}_k^0$ and ${}^i\hat{Y}_k^0$ are approximations of X_k^0 and Y_k^0 based on the estimates and the image-coordinate samples of period i. Such equations can be formed for each sample $k = 0, 1, ..., N-1$ and each period $i = 0, 1, ..., M-1$.

Ideally ${}^{i_1}\hat{X}_k^0 = {}^{i_2}\hat{X}_k^0$ and ${}^{i_1}\hat{Y}_k^0 = {}^{i_2}\hat{Y}_k^0$ for any sample k and any pair of periods i_1 and i_2. Therefore, making use of (4) and (5), we can obtain a pair of equations as follows:

$$\begin{aligned} {}^{i_1}\hat{X}_k^0 - {}^{i_2}\hat{X}_k^0 = {} & \frac{u_k^{i_1} - u_k^{i_2}}{f_x}\hat{Z}_k^0 + (i_2 - i_1)\hat{\Delta}_{X_T} \\ & + \frac{i_1 u_k^{i_1} - i_2 u_k^{i_2} + c_x(i_2 - i_1)}{f_x}\hat{\Delta}_{Z_T} = 0 \end{aligned} \tag{6}$$

$$\begin{aligned} {}^{i_1}\hat{Y}_k^0 - {}^{i_2}\hat{Y}_k^0 = {} & \frac{v_k^{i_1} - v_k^{i_2}}{f_y}\hat{Z}_k^0 + (i_2 - i_1)\hat{\Delta}_{Y_T} \\ & + \frac{i_1 v_k^{i_1} - i_2 v_k^{i_2} + c_y(i_2 - i_1)}{f_y}\hat{\Delta}_{Z_T} = 0. \end{aligned} \tag{7}$$

Two equations of the form (6) and (7) can be obtained for every sample k, for every pair of periods i_1 and i_2. This results in a total of $2N\binom{M}{2}$, where M is the number of periods, and N is the number of samples from each period. If we stack all these equations together in matrix form, the result is an over-constrained homogeneous linear system, which can be solved by solving the following optimization:

$$\begin{aligned} &\text{minimize } \|A\mathbf{X}\|_2 \\ &\text{subject to } \|\mathbf{X}\|_2 = 1, \end{aligned} \tag{8}$$

where A is the coefficient matrix, and:

$$\mathbf{X} \triangleq \begin{pmatrix} \hat{Z}_0^0 & \hat{Z}_1^0 & \hat{Z}_2^0 & ... & \hat{Z}_{N-1}^0 & \hat{\Delta}_{X_T} & \hat{\Delta}_{Y_T} & \hat{\Delta}_{Z_T} \end{pmatrix}^T. \tag{9}$$

This optimization (8) can be performed efficiently using Singular Value Decomposition (SVD), where $A = U\Sigma V^T$, and the minimizer $\mathbf{X}^*$ is the last column of V [28].

C. Constant Period Estimation

The period of motion T must also be known in order to perform reconstruction. It can be estimated using Fourier analysis of the image-coordinate velocity signals, $\dot{u}(t)$ and $\dot{v}(t)$. Specifically, different linear combinations of the signals of the form:

$$v_\phi(t) = \begin{pmatrix} \cos(\phi) \\ \sin(\phi) \end{pmatrix}^T \begin{pmatrix} \dot{u}(t) \\ \dot{v}(t) \end{pmatrix}, \tag{10}$$

parameterized by the angle ϕ are computed. The Power Spectral Density (PSD) of each signal is calculated and all

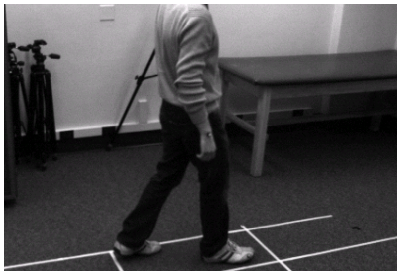

Fig. 1. An example from the gait sequences used to analyze the periodicity of real human motion.

the PSDs are superimposed with weights proportional to their sparsity (as given by the ℓ-0.5 norm). Finally, the period of motion is taken as the inverse of the largest spectral component in this weighted sum of PSDs.

III. The Periodicity of Real Human Motion

Theoretical results have shown that we can reconstruct a periodic trajectory in 3D by using only its appearance from a single camera-view. However, it is important to bear in mind that the primary interest in this work is in the analysis of human motion in video. Given this intended application area, it is then natural to ask the following question: *How periodic is real human motion?* Reconstruction techniques associated with natural human motion may not strictly adhere to any parametric motion model. We explore this question in detail in the form of quantitative analysis and provide some interesting results. The data analyzed in these experiments consists of gait sequences collected using a commercial motion-capture system. In each sequence, an infra-red marker was placed on the subject's foot (in this case near the toe), and the motion-capture system returned accurate 3D trajectories for this point of interest. The data set contains 16 gait sequences from a single subject, with each gait sequence comprised of approximately $3-4$ full strides. An example of one image from this data set is shown in Figure 1.

It is important to note that this data illustrates the case of a cooperative subject. In real scenarios, data may be collected in an unstructured way and may involve uncooperative or even adversarial subjects (for example, video from a surveillance camera) resulting in significant fluctuations. However, even in such situations relatively constant periodicity is often exhibited over short lengths of time.

A. Temporal Periodicity

We examine different ways to evaluate deviations from periodicity. The first one is temporal periodicity, i.e., the constancy of the temporal duration of each stride of the gait. By manually demarcating the beginning of each stride in every sequence we find that over a set of 16 gait sequences, the mean period is $2.46sec$ and the standard deviation is $0.11sec$ (4.6% of the mean). Additionally, the average absolute change in period length between two successive strides is $0.06sec$. That is to say, the deviation in period between two successive strides is on average $0.06sec$.

B. Spatial Periodicity

Spatial periodicity is defined as the constancy of both the stride length and direction. As before, we make use of manual demarcations of individual strides. Over the set of 16 sequences, we find that the mean stride length is $106.1cm$, with a standard deviation of $7.53cm$ (7.1% of the mean). The average absolute change in stride length between two successive strides was found to be $2.69cm$. In addition, the average absolute change in stride *direction* between two successive strides is 2.17^o.

C. Inter-Period Prediction Error

According to our definition, periodic motion can be described as $\mathbf{p}(t+T) = \mathbf{p}(t)+\mathbf{\Delta}_{\mathbf{p}_T}$. The displacement between any two samples separated by exactly one period in time is a constant factor $\mathbf{\Delta}_{\mathbf{p}_T}$, referred to as the inter-period displacement. In this experiment, we examine the suitability of the motion model, which we have used in the derivation of our reconstruction techniques, to natural human motion. For each sequence, $\mathbf{\Delta}_{\mathbf{p}_T}$ is first estimated by computing the mean displacement between each pair of samples separated by one period in time. Then, the position of each sample $\mathbf{p}(t+T)$ is predicted as $\hat{\mathbf{p}}(t+T) = \mathbf{p}(t) + \mathbf{\Delta}_{\mathbf{p}_T}$, and the error between the prediction $\hat{\mathbf{p}}(t+T)$ and the actual $\mathbf{p}(t+T)$ is measured. This is a form of a composite measure, which captures the overall periodicity of each gait sequence and measures the applicability of our motion model. Over the set of 16 gait sequences the average prediction error computed is $9.63cm$.

IV. Application: Clinical Gait Analysis

In the following set of experiments, we explore the applicability of the proposed reconstruction techniques to clinical gait analysis. Traditionally, gait analysis is performed manually by an expert physician who simply observes the gait of a patient in order to asses the degree of his/her debilitation. This process is subjective and the results of a manual analysis may vary from one physician to another for the same patient. In light of the current state-of-the-practice, it may be both possible and desirable to standardize this type of gait analysis through the use of technology. This may also make the results more accurate, precise, and repeatable. One possibility is to use commercial motion capture systems in order to collect accurate data about gait [25]. However, such systems are composed of multiple infra-red cameras (a minimum of six), are quite expensive, and require some technical expertise to operate, rendering them impractical in clinical settings. More recently, there have also been attempts to introduce further automation into the diagnostic procedure in the form of single-camera systems [23], [24].

Specifically, we explore the specific case of people whose gaits have been affected by strokes, and we compare them to the gaits of healthy people. In particular, we concern ourselves here with patients who exhibit *paretic* gait, in which one side of the body has been partially paralyzed by a stroke. We will show that it is possible to distinguish between the gaits of these patients and those of healthy individuals.

Furthermore, our results indicate that it may be possible to infer the degree of debilitation exhibited from only a small number of points on the body.

A. Data Collection

The data used for these experiments consisted of gait information from both healthy individuals and people who have suffered strokes. In all cases, the points on the body that were tracked included the subjects' ankle and knee joints – this is the only information that was used from each subject. The data was collected from the natural gaits of these subjects.

1) Healthy Gait: The set of healthy gait data consisted of samples from 6 different people without any known medical conditions affecting their movement. Several sequences were filmed of each person, from both left- and right-side viewing angles, using a standard resolution video camcorder. Small colored markers were attached to the ankle and knee joints over the subjects' clothing, and they were tracked using color classification. In total, there were 27 left-side and 32 right-side track sequences of healthy gait.

2) Paretic Gait: Samples of debilitated gait came from 8 different individuals exhibiting partial paralysis, consisting of 3 who are left-side and 5 who are right-side paretic. Due to constraints of anonymity, data was collected using a commercial motion capture system, which provides accurate 3D trajectories of infra-red markers attached to the subject's body. Since many of the sample trajectories consisted of just over 1 full stride of gait, longer sequences were synthesized by concatenating 3 periods of motion from the raw data. Then, 2 virtual camera positions and orientations were chosen (such that the left and right legs, respectively, could be seen), and the 3D trajectories were projected into the image coordinates of these virtual cameras. This comprised the data which we used as input for these experiments. Note, that this is similar to the tracking approach taken in related work such as [13], except that here we use the more realistic perspective projection instead of the orthographic assumption. As before, only tracks of the subjects' ankle and knee joints were used. In total, there were 7 tracks of left-side paretic gait, and 13 of right-side paretic gait.

B. Trajectory Reconstruction

Given the image-coordinate tracks described in the previous section, the constant period of motion was estimated using the technique described earlier. Since the period of motion is a property of the body as a whole, only the estimate from the ankle joint of each sequence was used. Next, each image-coordinate trajectory was resampled, so that the period of each sequence became 100 samples/stride. This was a design choice, which serves to make all subsequent analysis invariant to both walking speed and camera frame rate. These resampled image-coordinate trajectories were then reconstructed using the 3D geometric error cost function. Reconstruction was performed using the constant period formulation, since the gaits collected here maintained perfect periodicity relatively reliably. Each sequence in these gait samples consisted of tracks from two rigidly-connected body points (i.e., ankle and knee), each reconstructed knee-joint trajectory was scaled relative to the corresponding reconstructed ankle trajectory in order to enforce the rigidity constraint. Computation of the scale factors are omitted from this article due to space constraints.

C. Gait Distance Computations

To carry out high-level analysis, we must first compute the distances between all pairs of reconstructed ankle/knee sequences in a consistent fashion. Recall that, in each sequence, we have already rescaled the reconstructed knee trajectories with respect to those of the ankles in order to enforce the rigid body constraints. However, before inter-sequence distances can be computed, we must appropriately scale each entire ankle/knee pair (in 3D) in a way that enables us to compute scale-invariant distances. This is accomplished by choosing one period of one ankle trajectory as a reference period. Then, for each sequence, every one-period segment of ankle trajectory is scaled, rotated, and translated so that it is best-aligned with the reference period. For the one-period segment that most closely matches the reference period, the scale-factor computed for that alignment is then used as the overall scaling of that sequence. Once each sequence has been appropriately scaled, we can compute the distance between each pair of reconstructed ankle/knee trajectories (i.e., the inter-sequence distances). Distances are computed using the Orthogonal Procrustes Distance, which calculates the translation and rotation that minimizes the MSE between the two sets of points, including both the ankle and knee reconstructions.

D. Sequence Classification

In this particular experiment we begin to examine our ability to classify between healthy and paretic gait using the distances computed above. Here we consider each ankle/knee gait sequence independently, ignoring (for now) our knowledge of which subject it was generated by. All healthy sequences are used, but only the paretic-side gait of the stroke victims are included in this experiment. A simple nearest-neighbor classification is used, in which a test sequence is assigned the label of its nearest-neighbor in the training data. We choose 10-fold cross-validation in order to test the classification performance. Note that our goal here is not to develop a sophisticated new classification scheme, but only to explore the informativeness of these reconstructions as a proof-of-concept pilot study. The classification results are summarized in Figure 2(a). Note that overall classification accuracy is close to 95% for individual gait sequences. Additionally, the mean and standard deviations of the inter- and intra-class distances were computed. These results are summarized in Figure 2(b). As can be seen, the healthy and paretic gait classes are relatively tightly clustered, since the mean of the inter-class distances is much larger than the means of the intra-class distances. This explains our ability to classify accurately, and further supports the assertion that

Healthy Gait	96.6%
Paretic Gait	90.0%
Total	94.9%

(a)

Distance	mean	std
Intra-class: Healthy Gait	0.56	0.23
Intra-class: Paretic Gait	0.71	0.69
Inter-class	0.95	0.64

(b)

Fig. 2. (a) 10-fold cross-validation classification accuracy for individual sequence classification. (b) Means and standard deviations of inter- and intra-class distances.

these reconstructions contain significant information regarding the presence of gait debilitation, even if they are only from two tracked points on the body. Another observation is that the standard deviation of the paretic gait intra-class distances is relatively large. This can be attributed to the fact that the set of paretic gaits contained subjects with both left- and right-side paralysis, and each person's paralysis resulted in a different gait appearance, leading to large variability within this class.

E. Subject Classification

So far we have treated each gait sequence individually, regardless of which subject it was from. However, taking the identity of the subject into account may aid in performing a more robust analysis. In this experiment, we consider a more realistic case, which is perhaps more similar to the way analysis would be performed in a clinical setting. Instead of measuring only inter-sequence distances, here we consider *inter-person* distances, taking into account all available gait sequences for each subject[1]. Here the goal is to classify each *person* or subject as belonging to either the healthy or paretic gait category. Note that this is more similar to the type of automated analysis that could be performed in a doctor's office: first several gait samples would be collected from the current subject in question, then this person's gaits would be compared to both the healthy and debilitated subjects in the database in order to make a diagnosis. To compute the inter-person distance between two subjects, we simply take the median of all inter-sequence distances between these two people. Inter-sequence distances were computed as described earlier. The median was used here due to its robustness to outliers. In this experiment each subject was classified as either healthy or paretic based on all gait sequences from one side of his/her body. As before, classification was performed here using a simple nearest-neighbor scheme, where in this case a subject was assigned the same label as the subject in the training set with which the minimum inter-person distance was computed. In each case, one subject was used as the test set, while the training set consisted of all remaining subjects (i.e., the database). Using this classification scheme, we achieved a *classification accuracy of 100%* for both the healthy and paretic subjects. As can be seen from this result, classification accuracy is improved by treating each subject as a whole, instead of considering only individual gait sequences separately. In addition, the use of the median for computing the inter-person distances made our classification more robust to outliers, in which one inter-sequence distance might be erroneous due to factors such as inaccurate reconstruction or poor tracking performance. Finally, some statistics of the inter- and intra-class distances were again computed, except that in this case we use the inter-person distances rather than just comparing individual gait sequences. The results are summarized in Table 3. As we expect, it appears that the class of healthy subjects is relatively tightly clustered with respect to the inter-class distances, as is the case with the paretic subjects. This helps explain our ability to classify subjects correctly, and again supports our assertion that gait reconstructions using only two points on the body can still contain significant information about the subject.

[1]We take the left- and right-side gaits of healthy individuals separately, and use only the paretic side of debilitated subjects.

Inter-Person Distance	mean	std
Intra-class: Healthy Subject	0.52	0.11
Intra-class: Paretic Subject	0.88	0.70
Inter-class	1.06	0.67

Fig. 3. Means and standard deviations of inter- and intra-class *inter-person distances*.

F. Manifold Embedding

We further explore the information content of these gait reconstructions by examining their embedding in a low-dimensional manifold space. Each gait sequence was taken as a point on the manifold in some high-dimensional space, and embeddings were computed using ISOMAP. Only the inter-sequence distance matrix was required as input, so it was not necessary to explicitly represent each gait sequence as a point in this space.

A plot of the two-dimensional embedding of the healthy and paretic gait sequences is shown in Figure 4, where healthy gaits are indicated by blue points, and paretic gaits by red points. One fact that is quite striking about this embedding is that, barring three outlier points, the healthy and paretic gaits are linearly separable in this low-dimensional space.

We can draw two interesting conclusions from this observation. First, this result seems to indicate that most of the variation in this set of data can be explained by only two degrees of freedom. So even though the gaits represented by these high-dimensional points are quite complex motions, most of the difference between them can be parameterized with two variables. Second, this embedding explains our previous results for both the sequence classification and the subject classification tasks. It is clear that, at least in this particular embedding, it would be difficult to develop any classifier which could achieve perfect classification accuracy based only on the individual sequences. This is evidenced by the few outliers depicted in Figure 4 – the three healthy gaits embedded in the midst of all the paretic data points. However, when we switch to subject classification using inter-person distances, this problem is alleviated, since we take into account all gait sequences from each subject, and small numbers of outliers no longer affect the classification.

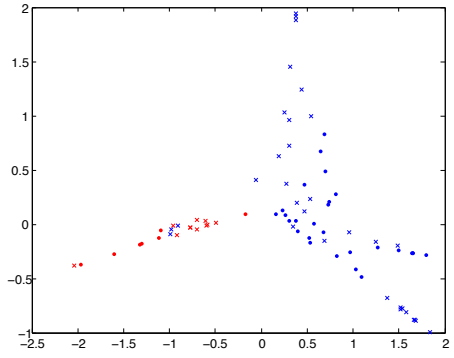

Fig. 4. Two-dimensional ISOMAP embedding of the healthy and paretic gait sequences. Blue/red points indicate healthy/paretic, and circle/cross points represent left/right-side, respectively.

Interestingly, the manifold embedding indicates not only linear separability in 2D but also contains information about the degree or severity of the subject's paralysis. Informally, we can observe that paretic gaits closer to the healthy gaits in the manifold embedding (i.e., red points farther to the right in Figure 4) show less paralysis than those farther from the healthy gaits. Unfortunately, it is difficult to illustrate the differences in gait in a static image. As such, video clips showing examples of three different gait sequences with different levels of paralysis can be found online at *http://www.ece.umn.edu/users/ribn0003/strokedata/*. Included are paretic gait sequences from three different parts of the manifold – moving from right to left in Figure 4. Each sequence shows the motion-capture tracks of the subject's knees and ankles, where each sequence was synthesized from the real motion-capture as described earlier.

V. Conclusions and Future Work

In this paper we have built on our previous work on reconstructing periodic motions in 3D in order to explore its effectiveness in human motion analysis. In order to justify the use of these reconstruction techniques, we have performed an empirical study regarding the periodicity of real human motion. Using motion-capture data from several instances of normal gait, measures of both temporal and spatial periodicity were considered, as well as the suitability of the proposed motion model for periodic trajectories. We have found that human motion can indeed adhere closely to strict periodicity over a duration of 3-4 strides.

Finally, the applicability of the proposed algorithms to the domain of human motion analysis was demonstrated in clinical gait analysis. Reconstructions of two points on the body could be used to accurately classify gait sequences as belonging to either healthy individuals, or to subjects whose movement has been affected by stroke. A classification accuracy of 100% was achieved when classifying between sets of trajectories consisting of multiple instances of each subject's gait. Furthermore, it was observed that a low-dimensional manifold embedding of the reconstructed trajectories contained information regarding the severity of the debilitation.

References

[1] R. Polana and R. Nelson, Detection and Recognition of Periodic, Nonrigid Motion, *Int'l. J. Comp. Vis.*, Vol. 23, No., 3, 1997.
[2] F. Liu and R.W. Picard, Finding Periodicity in Space and Time, *Proc. IEEE Conf. Comp. Vis. and Patt. Rec.*, 1998.
[3] R. Cutler and L.S. Davis, Robust Real-Time Periodic Motion Detection, Analysis, and Applications, *IEEE Trans. Patt. Anal. and Mach. Intel.*, Vol. 22, No. 8, 2000.
[4] D. Ormoneit et al., Representing Cyclic Human Motion Using Functional Analysis, *Img. and Vis. Comp.*, Vol. 23, 2005.
[5] A. Briassouli and N. Ahuja, Extraction and Analysis of Multiple Periodic Motions in Video Sequences, *IEEE Trans. Patt. Anal. and Mach. Intel.*, Vol. 29, No. 7, 2007.
[6] P.-S. Tsai et al., Cyclic Motion Detection for Motion Based Recognition, *Patt. Rec.*, Vol. 27, No. 12, 1994.
[7] H. Fujiyoshi et al., Real-Time Human Motion Analysis by Image Skeletonization, *IEICE Trans. Inf. and Sys.*, Vol. E87-D, No. 1, 2004.
[8] S.M. Seitz and C.R. Dyer, View-Invariant Analysis of Cyclic Motion, *Int'l. J. Comp. Vis.*, Vol. 25, 1997.
[9] C.J. Cohen et al., Dynamical System Representation, Generation, and Recognition of Basic Oscillatory Motion Gestures, *Proc. Int'l. Conf. Aut. Face and Gesture Rec.*, 1996.
[10] J. Davis et al., Categorical Representation and Recognition of Oscillatory Motion Patterns, *Proc. IEEE Conf. Comp. Vis. and Patt. Rec.*, Vol. 1, 2000.
[11] E. Ribnick and N. Papanikolopoulos, Estimating 3D Trajectories of Periodic Motions from Stationary Monocular Views, *Proc. European Conf. Comp. Vis.*, 2008.
[12] E. Ribnick and N. Papanikolopoulos, View-Invariant Analysis of Periodic Motion, *Proc. IEEE/RSJ Int'l. Conf. Intel. Robots and Sys.*, 2009.
[13] Z. Zhang and N.F. Troje, 3D Periodic Human Motion Reconstruction from 2D Motion Sequences, *Neural Computation*, Vol. 19, 2007.
[14] S. Belongie and J. Wills, Structure from Periodic Motion, *Proc. Int'l. Wkshp. on Spatial Coherence for Visual Motion Anal.*, 2004.
[15] T.B. Moeslund and E. Granum, A Survey of Computer Vision–Based Human Motion Capture, *Comp. Vis. and Img. Und.*, Vol. 81, 2001.
[16] T.B. Moeslund et al., A Survey of Advances in Vision-Based Human Motion Capture and Analysis, *Comp. Vis. and Img. Und.*, Vol. 104, 2006.
[17] J.K. Aggarwal and Q. Cai, Human Motion Analysis: A Review, *Proc. IEEE Nonrigid and Articulated Motion Wkshp.*, 1997.
[18] J.M. Wang et al., Gaussian Process Dynamical Models, in *Advances in Neural Information Processing Systems 18*, 2006.
[19] L. Wang et al., Recent Developments in Human Motion Analysis, *Patt. Rec.*, Vol. 36, 2003.
[20] E. Shechtman and M. Irani, Matching Local Self-Similarities across Images and Videos, *Proc. IEEE Conf. Comp. Vis. and Patt. Rec.*, 2007.
[21] A. Elgammal and C.-S. Lee, Inferring 3D Body Pose from Silhouettes using Activity Manifold Learning, *Proc. IEEE Conf. Comp. Vis. and Patt. Rec.*, Vol. 2, 2004.
[22] M.P. Murray, Gait as a Total Pattern of Movement, *American J. Physical Medicine*, Vol. 46, No. 1, 1967.
[23] H.M. Lakany et al., Human Walking: Tracking and Analysis, *IEE Colloquium on Motion Anal. and Tracking*, 1999.
[24] H. Lee et al., Human Gait and Posture Analysis for Diagnosing Neurological Disorders, *Proc. IEEE Int'l. Conf. Img. Proc.*, Vol. 2, 2000.
[25] K.R. Kaufman et al., Gait Characteristics of Patients with Knee Osteoarthritis, *J. Biomech.*, Vol. 34, No. 7, 2001.
[26] J. Michalak et al., The Embodiment of Sadness and Depression – Gait Patterns Associated with Dysphoric Mood, *Psychosomatic Medicine*, Vol. 71, 2009.
[27] N.F. Troje, Decomposing Biological Motion: A Framework for Analysis and Synthesis of Human Gait Patterns, *J. Vis.*, Vol. 2, 2002.
[28] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, 2003.