ELSEVIER

Human Movement Science 21 (2002) 495–514

HUMAN MOVEMENT SCIENCE

www.elsevier.com/locate/humov

Automatic stride interval extraction from long, highly variable and noisy gait timing signals

Tom Chau [a,b,c,*], Sidra Rizvi [a,b]

[a] *Bloorview MacMillan Children's Centre, 350 Rumsey Road, Toronto, Ont., Canada M4G 1R8*

[b] *Institute of Biomaterials and Biomedical Engineering, 4 Taddle Creek Road, University of Toronto, Toronto, Canada M5S 3G9*

[c] *Graduate Department of Rehabilitation Science, 256 McCaul St., Toronto, Canada M5T 1W5*

Abstract

This paper presents a probabilistic algorithm for automatically extracting the stride interval time series from long, highly variable and noisy two-state timing signals. Long interstride temporal records are of particular interest in nonlinear dynamical analysis of gait. The proposed method consists of probabilistic estimation and extraction followed by post-extraction filtering. With noisy timing signals from 10 children with Spastic Diplegia, no statistical differences in the numbers of extracted strides ($p = 0.94$), the mean stride intervals ($p = 0.55$) and the scaling exponents ($p = 0.94$) (a measure of temporal heterogeneity) were found between series extracted by hand and by the probabilistic algorithm. The method is robust to noise and violations of normality. Results support the use of probabilistic extraction as an alternative to laborious manual extraction.

PsycINFO classification: 2240

Keywords: Gait dynamics; Quasi-periodic events; Gait timing signals; Automatic extraction

* Corresponding author. Present address: Bloorview MacMillan Children's Centre, 350 Rumsey Road, Toronto, Ont., Canada M4G 1R8. Tel.: +1-416-425-6220x3515; fax: +1-416-425-1634.

E-mail address: ttkchau@ieee.org (T. Chau).

0167-9457/02/$ - see front matter © 2002 Elsevier Science B.V. All rights reserved.
PII: S0167-9457(02)00125-2

1. Introduction

1.1. Significance of long gait time series

Recently, there has been a growing interest in the study of nonlinear dynamics in human movement such as postural sway during standing (Duarte & Zatisiorsky, 2001), eye movements (Schmeisser, McDonough, Bond, Hislop, & Epstein, 2001), targeted upper limb movements (Miyazaki, Kadota, Kudo, Masani, & Ohtsuki, 2001) and handwriting (Longstaff & Heath, 1999). In particular, much work has focussed on natural walking dynamics over long periods of time, i.e. in the order of hundreds and thousands of strides (see for example Hausdorff, Peng, Ladin, Wei, & Goldberger, 1995, 1996, 1997; West & Griffin, 1998, 1999). These experiments typically record a quasi-periodic event within the gait cycle such as heel contact (Hausdorff et al., 1995) or maximal knee extension (West & Griffin, 1998) over an extended period of walking. From these recordings, the stride interval time series is determined. Contrary to conventional belief that interstride fluctuations are strictly random, Hausdorff et al. (1995) found that the stride interval time series of young, healthy adults exhibits long-range correlations or long time memory. In fact, these long-range effects displayed statistical self-similarity, suggesting that movement dynamics are governed by an underlying fractal process. These surprising findings were independently confirmed by West and Griffin (1998, 1999). The potential clinical value of these discoveries lies in the fact that fractal dynamics are intrinsic to healthy gait and are compromised to varying degrees in the presence of neurological pathology (Hausdorff et al., 1997, 2000). Hausdorff et al. suggested that the measurement of this fractal property of human gait could play a role in the quantitative diagnosis of neurological pathologies and in the assessment of therapeutic interventions. Theoretically, fractal analyses may provide new insights into the mechanisms involved in the hierarchical control of walking (Hausdorff et al., 1995). The key point is that these important nonlinear dynamics can only be revealed by analyzing long records of walking, consisting of no less than several hundred strides. See Chau (2001) for a review of related literature.

1.2. The nature of the data

When studying the nonlinear dynamics of walking, the data in question are the quasi-periodic temporal signals typically collected in a quantitative gait analysis. The signals are not truly periodic as there is an approximately 3–4% natural variability in stride duration from stride to stride (Griffin, West, & West, 2000). Timing information is usually captured by force-sensitive resistor (FSRs) or footswitches placed strategically on the sole of the foot. The resultant information may be an analog force trace or if thresholded, a rectangular wave. In this paper, we focus exclusively on the thresholded timing signal. The top graph in Fig. 1 exemplifies an ideal (simulated) unilateral timing signal with 4% variability in the signal period. The temporal distance between the leading edges of two successive pulses can be used to define the stride interval time.

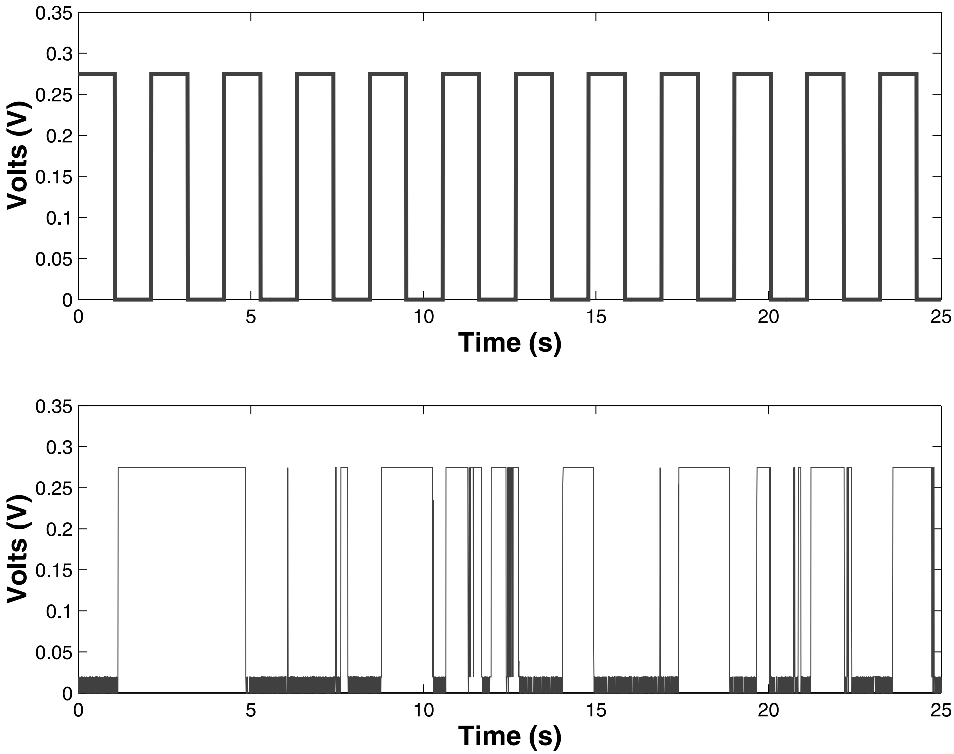

Fig. 1. Examples of an ideal timing signal (top) and a real signal obtained from a child with Spastic Diplegia (bottom).

In practice, the timing signals are not always clean and regularly spaced in time. There may be inadvertent foot contact due to pathological gait such as toe-dragging during the swing phase. Similarly, the signal may be lost momentarily as a result of inadequate foot contact during stance. The presence of electrical noise in the analog signal often introduces extraneous pulses in the thresholded signal. Further, compromised recovery of the footswitch due to blockage of the FSRs air passage, may keep the footswitch "on" beyond the actual contact duration. In contrast to the ideal rectangular wave, the bottom graph in Fig. 1 is an actual thresholded footswitch recording from a child with Spastic Diplegia. Note the irregularity in the spacing of the pulses and the presence of many short, but extraneous pulses.

1.3. Stride interval extraction – challenges

In the ideal case, the stride intervals can be easily extracted by identifying the successive location of quasi-periodic events. For example, the stride interval might be defined as the time between consecutive rising edges. Some specialized instrumentation, such as the ambulatory logging unit developed by Hausdorff, Forman, Pilgrim, Rigney, and Wei (1992), directly records the times corresponding to a quasi-periodic event (e.g. heel strike). In this case, the stride interval time series is obtained simply by taking the first difference of the raw data. In all cases, under ideal conditions, a series of stride interval times can be extracted from the temporal data.

The presence of noise as illustrated in Fig. 1 complicates the stride interval extraction process. Repeated events may become more difficult to identify and in the case where event times are directly available, outliers are likely to abound. Despite these additional challenges on top of the aforementioned signal contamination, in the worst cases, delineation of stride times can usually be done manually. Assisted by domain knowledge, extremely robust human pattern recognition can usually detect the sequence of most likely stride times. The domain-specific information may include the expected stride interval for the population under study and any known physical constraints of walking. As an example of the latter, a 5 ms pulse would not be considered a physiologically valid stride. Manual extraction is usually feasible because in a typical gait analysis, a small number of strides are recorded per trial. With very long time series, however, manual extraction of stride times can become a laborious and outright impractical process. Visual fatigue increases linearly with the time spent viewing a computer display terminal (Kaneko & Sakamoto, 2001) and in combination with mental fatigue, can increase the likelihood of human error (Schellekens, Sijtsma, Vegter, & Meijman, 2000). Further, executive functions of the brain can be compromised by durations of work in excess of eight consecutive hours (Proctor, White, Robins, & Echeverria, 1996).

At first, a seemingly easy solution to automatic extraction would be to filter the timing signal to eliminate noise. Then, one could reliably detect events by detecting state transitions in a timing signal such as that in Fig. 1. While this is acceptable when we are interested in global parameterizations such as peak amplitude or frequency, in the study of temporal dynamics, the precise times of individual events must be preserved. In a bi-state timing signal such as in Fig. 1, some of the short pulses may in fact indicate true event occurrences. Filtering them away would lead to an interevent series with significantly altered dynamics. Clearly, automatic processes that preserve the integrity of temporal information are needed for extracting stride intervals from very long time series. Ideally, these automatic processes would also be compatible with data collected with standard gait lab equipment, foregoing the need for specialized and often expensive hardware.

1.4. Roadmap

The remainder of the paper is organized as follows. In Section 2, we review some recent efforts in gait event detection. In Section 3, we present the probabilistic stride extraction method. Subsequently, in Section 4 we demonstrate its utility with long time series of thresholded (two-state) footswitch data collected from children with Spastic Diplegia.

2. Related work: Gait event detection

In the studies by Hausdorff et al., a custom ambulatory data logger provided the discrete series of heel strike times directly. Hence, extracting the stride interval time series simply involved taking the first difference and applying median filters to eliminate outliers (Hausdorff et al., 1995, 1996, 1997, 2000). In studies conducted by West and Griffin (1998, 1999) and Griffin et al. (2000), the analog trace of the knee extension angle was recorded. The stride interval series was defined as the time difference between successive local minima in the analog signal. It is presumed that these minima were numerically estimated.

Outside of studies by Hausdorff et al. and West and Griffin which focussed specifically on the fractal dynamics of gait, there has been very little work on extracting quasi-periodic events from long gait time series. However, automatic gait event detection has been a topic of much research in the area of functional electrical stimulation (FES). In the context of FES, the detection of multiple gait phases is a real-time task typically based on multiple sensor inputs from a few recent strides. In contrast, in the study of fractal gait dynamics, stride interval extraction is typically done off-line, taking into consideration the entire time record from a single sensor. Hence, gait event detection for FES and stride interval extraction for fractal analysis are not the same problem. Nonetheless, some insight into ways to automate the latter may be gained by surveying some recent detection methods for FES.

Recently proposed approaches to gait phase detection for FES have largely been variations of rule-based methods. Skelly and Chizeck (2001) developed a real-time fuzzy-rule base that predicted five gait phases based on the analog input from two insole FSRs per side. With three subjects, their system predicted heel strike within 8% (gait cycle duration) of the observer detection of heel strike. Pappas, Popovic, Keller, Dietz, and Morari (2001) used a heel-mounted gyroscope and three insole FSRs as two-state switches. They developed a rule-based algorithm consisting of seven hand-crafted rules to detect with extremely high accuracy, four gait phases in both able-bodied subjects and subjects with gait pathologies. Further, their system proved robust under various walking conditions including different terrains, stairs and varying speeds, and during standing up. Using rough sets and adaptive logic networks, Williamson and Andrews (2000) synthesized rule-based detectors to identify gait phases based upon a trio of accelerometer signals. An adaptive logic network was also used by Hansen, Haugland, Sinkjaer, and Donaldson (2002) to detect heel strike and foot lift and to discriminate between walking and standing. The input to the network was a 1 s window of the electroneurogram recorded from the sural nerve. Detection output was further refined by restriction rules based on fuzzy logic algorithms.

Despite the success of the above rule-based systems, it is difficult to derive rules either by hand or by inductive learning that will span all the possible ways that noise may contaminate a quasi-periodic timing signal. In fact, it is generally known that rule-based systems such as decision trees do not perform well with noisy data (Ripley, 1996). Given the high noise level in the signals of interest, rather than explicitly generate a rule-base, we elected to adopt a probabilistic approach. Probabilistic techniques have been used widely in the analysis of noisy temporal data (see, for example, Hochreiter & Mozer, 2001; Han & Gao, 2001; Jancovic & Ming, 2001). We will show that the proposed technique can be viewed as a probabilistic, single rule detection method.

3. The proposed method

Before proceeding, we define some general terminology. An event is an occurrence of the observed quasi-periodic phenomenon of interest, for example heel strike. The interevent time is the time between successive events. In the gait example, the interevent time is the time between consecutive heel strikes or the stride interval. The set of candidate event times, **T**, contains all potential event times. The implication is that only a subset of these times is valid. The set of probable event times, **U**, encompasses the most probable candidate event times. All time series are considered as finite, discrete sets. Relevant notation is summarized in Table 1. In general, bold upper case letters denote sets of times, while lower case subscripted letters denote individual times within the set. For example, the set of candidate event times can thus be written as $\mathbf{T} = \{t_1, \ldots, t_{n_T}\}$.

Although this paper proposes a method for stride interval extraction in the context of gait analysis, we will use generic terminology in describing the method to emphasize that it could be applied to other types of events and interevent times. Given an identified event time, t_e, the idea of probabilistic interevent time extraction is to choose the most probable time, t^*, of the next event. "Most probable" is based on an initial estimate of the mean interevent time from a subset of candidate event times.

3.1. Identification of candidate event times

The proposed probabilistic extraction algorithm operates on a set of candidate events. In this section, we describe a simple way in which this series might be obtained from typical gait measurements. Let $\{x_1, \ldots, x_N\}$ denote the raw, discrete time series of length N. If the data are thresholded two-state footswitch signals in the form of rectangular pulses, then without loss of generality, $x_i \in \{0, 1\}$. For

Table 1
Key notation used in paper

Symbol	Meaning
N	Length of raw discrete timing series
T	Set of n_T candidate event times
U	Set of n_U probable event times, $\mathbf{U} \subset \mathbf{T}$, $n_U \leqslant n_T$.
$\langle\cdot\rangle$	Mean value of a set
t_e	Event time under consideration
t^*	Most probable event time
τ	Crude estimate of mean interevent time

example, the high value of the wave typically indicates foot contact with the ground (stance phase) while the low value signifies the absence of contact (swing phase). A simple algorithm to extract event times from a two-state thresholded signal x_i, $i = 1, \ldots, N$, is as follows. Define $\mathbf{H}$ to be the set of indices such that x_i exceeds a threshold, θ. Similarly define $\mathbf{L}$ to be the set of indices such that x_i falls below the threshold θ. We thus have

$$\mathbf{H} = \{i | x_i \geqslant \theta\}, \tag{1}$$

$$\mathbf{L} = \{i | x_i < \theta\}, \tag{2}$$

where $i \in \{1, \ldots, N\}$ with N being the length of the discrete series. In the case of $x_i \in \{0, 1\}$, we might have $\theta = 0.5$. The sets $\mathbf{H}$ and $\mathbf{L}$ represent the indices corresponding to high and low values of x_i, respectively. The event times, for example heel strike times, are given by the transition from a low to a high value or in terms of indices, a difference of one between neighbouring low and high indices. To capture this, we increment every element of $\mathbf{L}$ by 1 to yield, $\mathbf{L}'$, such that $\mathbf{L}' = \{i + 1\}$, for all $i \in \mathbf{L}$. We are in effect shifting the elements of $\mathbf{L}$ to the right by one time step. Indices corresponding to low–high transitions should now occur both in $\mathbf{L}'$ and $\mathbf{H}$. Thus, the set of candidate event times $\mathbf{T} = \{t_i\}$ is simply determined by the intersection of the sets $\mathbf{H}$ and $\mathbf{L}'$, namely,

$$\mathbf{T} = \{t_i | i \in \mathbf{H} \cap \mathbf{L}'\}. \tag{3}$$

3.2. Probabilistic extraction of event times

Suppose that we have obtained a set of candidate event times $\mathbf{T} = \{t_1, t_2, \ldots, t_{n_T}\}$, possibly by the method described above or directly from hardware. We are now ready to formulate the probabilistic extraction technique. It involves two main procedures, estimation and extraction.

3.2.1. Estimation step

The first step entails the estimation of the mean interevent time from the raw series. We begin by computing a crude approximation to the interevent time series by simply taking the first difference of elements in $\mathbf{T}$. This yields a series $\Delta\mathbf{T}$, such that each element is given by

$$\Delta T_i = t_{i+1} - t_i \tag{4}$$

for $i = 1, \ldots, n_T - 1$. If the raw data resemble the ideal pulse train as in the top graph of Fig. 1, we could estimate the mean interevent time with $\langle \Delta\mathbf{T} \rangle$ where $\langle \cdot \rangle$ signifies the mean. However, recall that in practice, we are dealing with contaminated data (see Section 1.2), like the bottom graph of Fig. 1. Thus, $\Delta\mathbf{T}$ may contain many invalid interevent times. To obtain a better estimate, we propose to compute the mean using a subset $\Delta\mathbf{T}' \subset \Delta\mathbf{T}$, consisting only of elements which fall within the top 50th percentile of interevent times. Recognizing the potential asymmetry in the distribution of interevent times, we use the skewness of $\Delta\mathbf{T}$ to define the subset $\Delta\mathbf{T}'$.

The skewness of $\Delta\mathbf{T}$ is given by $\gamma = E[\Delta\mathbf{T} - \mu]^3/\sigma^3$ (Mendel, 1991), where $\gamma = 0$ for symmetrically distributed data while μ and σ are the mean and standard deviation of times in $\Delta\mathbf{T}$, respectively. Let p_{25} denote the 25th-percentile of the elements in $\Delta\mathbf{T}$. Define p_{50} and p_{75} similarly. The subset $\Delta\mathbf{T}'$ is then defined as follows:

$$\Delta\mathbf{T}' = \left\{ \Delta T_i \;\middle|\; \begin{array}{ll} \Delta T_i \geqslant p_{50}, & \gamma \leqslant -0.153, \\ p_{25} \geqslant \Delta T_i \geqslant p_{75}, & |\gamma| < 0.153, \\ \Delta T_i \leqslant p_{50}, & \gamma \geqslant 0.153, \end{array} \right. \tag{5}$$

where $i = 1, \ldots, n_T - 1$. To estimate the threshold skewness value, we first computed the skewness statistics for 10^4 simulated, normally distributed signals each of length 10^3. We then determined the values at which a hypothetical test γ would be considered statistically different from the 10^4 skewness values, at the 5% significance level. In other words, a γ value greater than or equal to $|0.153|$ is statistically different from that arising from normally distributed data.

With the subset $\Delta\mathbf{T}'$ we estimate the mean interevent time as $\langle\Delta\mathbf{T}'\rangle$ where $\langle\cdot\rangle$ denotes the average value. As outlined below, this crude estimate will serve as an initial guess for the next event time.

3.2.2. Extraction step

The next step in the procedure is the actual interevent time extraction. Given an event time, t_e, we want to find the most probable time of the next event in the series. Initially, $t_e = t_1$. We form a normal probability density function $f(t_j)$ to the right of t_e, centred around $\mu = t_e + \langle\Delta\mathbf{T}'\rangle$ and with variance $\sigma^2 = \langle\Delta\mathbf{T}'^2\rangle - \langle\Delta\mathbf{T}'\rangle^2$. This is simply the variance of the interevent times in the subset $\Delta\mathbf{T}'$. Specifically, the density is written as

$$f(t_j) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(\frac{-(t_j - \mu)^2}{2\sigma^2}\right) \tag{6}$$

for $t_j > t_e$, so that only times t_j to the right of t_e are considered. This is an intuitive restriction for, if we think of t_e as the present time, we cannot have the next event occur in the past, $t_j < t_e$.

Next, we determine the time t^* corresponding to the highest probability density, i.e.

$$f(t^*) = \max_j f(t_j) \tag{7}$$

again for $t_j > t_e$. This time t^* is added to the set of probable event times, $\mathbf{U}$. The above extraction process is then repeated with the new starting event time, $t_e = t^*$. Extraction is complete when $t_e + \langle\Delta\mathbf{T}'\rangle > t_{n_T}$ i.e. the probable location of the next step is beyond the end of the series. At the conclusion of the extraction process, we will have a series $\mathbf{U} \subset \mathbf{T}$ of the most probable successive candidate event times. The interevent time series $\Delta\mathbf{U}$ is obtained by simply taking the first difference of $\mathbf{U}$ such that $\Delta U_i = u_{i+1} - u_i$, $i = 1, \ldots, n_U - 1$, where n_U is the cardinality of the set $\mathbf{U}$. Fig. 2 portrays one iteration of the extraction process. Given a series of candidate event

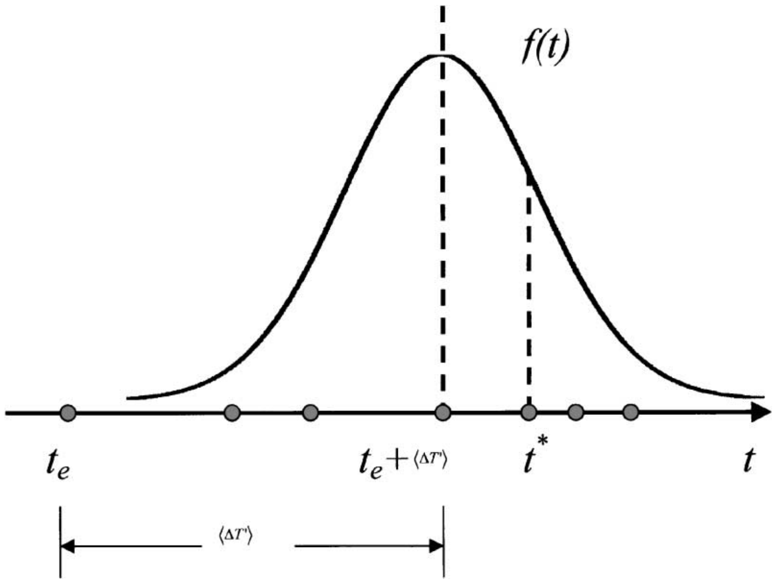

Fig. 2. One iteration of the extraction process. The horizontal axis represents time and the dots denote candidate event times. The current event time under consideration is t_e and $\langle \Delta \mathbf{T}' \rangle$ is the initial estimate of the mean interevent time. The probability density function $f(t)$, centred around $t_e + \langle \Delta \mathbf{T}' \rangle$, is superimposed on the timing data. The most probable location of the next event is at time t^*.

times $\mathbf{T} = \{t_1, t_2, \ldots, t_{n_T}\}$, the following is a summary of the probabilistic extraction algorithm.

Estimation step:

1. Difference the given series to obtain a series of candidate interevent times, $\Delta\mathbf{T}$, where $\Delta T_i = t_{i+1} - t_i$, $i = 1, \ldots, n_T - 1$.
2. Create a subset $\Delta\mathbf{T}' \subset \Delta\mathbf{T}$, according to the skewness of the distribution of interevent times $\Delta\mathbf{T}$.
3. Compute the initial estimate of the mean interevent time as $\langle \Delta\mathbf{T}' \rangle$.

Extraction step:

1. For a given event time t_e, initially with $t_e = t_1$,
   (a) Create a normal probability density function $f(t_j)$ with $\mu = t_e + \langle \Delta\mathbf{T}' \rangle$ and $\sigma^2 = \langle \Delta\mathbf{T}'^2 \rangle - \langle \Delta\mathbf{T}' \rangle^2$.
   (b) Choose the time of the next event t^*, such that $f(t^*) = \max_j f(t_j)$.
   (c) In the unlikely event of a tie, $f(t_k) = f(t_l) = \max_j f(t_j)$, $t_k < t_l$, break the tie by choosing the later time, i.e. $t^* = t_l$.
   (d) Add t^* to the set of probable times $\mathbf{U}$.
   (e) Set $t_e = t^*$ and repeat Step 1 until $t_e + \langle \Delta\mathbf{T}' \rangle > t_{n_T}$.
2. Compute the set of interevent times, $\Delta\mathbf{U}$.

3.3. Post-extraction filtering

The extracted set of probable interevent times, $\Delta\mathbf{U}$, may contain some outliers, i.e. interevent times which are unphysiologically long or short. However, we cannot blindly assume normally distributed event times when removing outliers. It is known that certain interevent time series, such as interneuronal spike intervals, are naturally positively skewed (Gerstein & Mandelbrot, 1964) due to underlying nonlinear dynamics. Hence, we must first check for violations of normality. One possible statistical test is the χ^2 test which compares the sample histogram frequencies against frequencies expected from a gaussian density (Bendat & Piersol, 2000). If normality is satisfied, fit a normal density to the values in $\Delta\mathbf{U}$. If normality is violated, check the skewness of the data in $\Delta\mathbf{U}$ and fit a skewed density. In either case, discard values outside of the 5th and 95th percentiles. Using the same notation as above, we thus have

$$\Delta\mathbf{U}' = \{\Delta U_i | p_5 < \Delta U_i < p_{95}\} \tag{8}$$

where p_5 and p_{95} are determined according to the fitting density. In practice, we have found that a shifted gamma density fits well with positively skewed stride interval time series. The filtered series $\Delta\mathbf{U}'$ represents the final extracted interevent time series.

Given a series of interevent times $\Delta\mathbf{U}$, the proposed post-extraction filtering procedure is as follows.

1. Using a standard statistical test, such as the χ^2 test, check that the values in $\Delta\mathbf{U}$ are normally distributed.
   (a) If $\Delta\mathbf{U}$ is normally distributed, fit a normal density to the data, i.e. determine the mean and standard deviation.
   (b) If $\Delta\mathbf{U}$ is not normally distributed, check the skewness of the distribution. Fit an appropriately skewed density, such as the γ density, i.e. estimate the necessary density parameters.
2. Based on the fitted density, determine p_5 and p_{95}.
3. Remove values of $\Delta\mathbf{U}$ outside of the 5th and 95th percentiles to obtain the subset $\Delta\mathbf{U}' \subset \Delta\mathbf{U}$.

3.4. Relationship to rule-based extraction

To classify gait events, rule-based methods essentially partition the sample space defined by the measurement variables. In the present case, we are concerned with a single timing signal and hence the one-dimensional sample space is made up of all possible time values. At each iteration of the probabilistic extraction algorithm, we can interpret the probability density function as determining a single probabilistic partition of the sample space. This partition can be viewed as a "rule" which favours the candidate event time with highest probability.

4. Example

We now illustrate the proposed interevent extraction method with a long gait time series. The particular event of interest is heel strike and thus the interevent time is the stride interval. As in the bottom graph of Fig. 1, the raw data are inherently noisy. The goal is to accurately extract the stride interval time series from the raw data, where "accurate" means agreement with manual extraction.

4.1. Methodology

The instrumentation included a FSR (Interlink Electronics), custom thresholding circuitry and a portable data logger (Valitec AD128). Ten children with Spastic Diplegia (age 6.0 ± 0.97 years) participated in the study in accordance with ethical guidelines prescribed by the University of Toronto and Bloorview MacMillan Children's Centre, Toronto, Canada. A unilateral two-state footswitch signal, with the high state signifying foot contact, was recorded over 10 min of walking in the 10 children with Spastic Diplegia. This yielded 10 individual signals. For each signal, manual and automatic extractions were carried out. Manual extraction involved the painstaking process of picking out heel strikes by visual inspection of the raw footswitch signal. Two flavours of automatic extraction were realized, namely probabilistic and generic or nonprobabilistic. The former refers to the proposed method while the latter simply involved the automatic extraction of all candidate events (low to high state transitions) without any post-processing of the time series.

As in Pappas et al. (2001), accuracy of automatic extraction was determined by comparing the results against the manually extracted time series. In particular, four quantities were determined to gauge the agreement with manual extraction. These were the number of extracted strides, the mean stride interval, the scaling exponent and the mean squared error between manually and automatically extracted series. Let $\mathbf{S} = \{s_1, \ldots, s_n\}$ represent an extracted stride interval series, with $\mathbf{S}^{\mathrm{A}}$ and $\mathbf{S}^{\mathrm{M}}$ denoting manually and automatically (by probabilistic or generic methods) extracted series, respectively. The number of extracted strides is $\mathrm{card}(\mathbf{S})$, where cardinality, $\mathrm{card}(\cdot)$, is the number of elements in a set. The absolute difference between the numbers of strides extracted automatically and manually is thus given by

$$|\mathrm{card}(\mathbf{S}^{\mathrm{M}}) - \mathrm{card}(\mathbf{S}^{\mathrm{A}})|. \tag{9}$$

The mean stride interval is the mean of the series $\mathbf{S}$, denoted as $\langle \mathbf{S} \rangle$. The absolute difference between mean stride intervals of manual and automatic extraction as a percentage of the former is

$$\frac{|\langle \mathbf{S}^{\mathrm{M}} \rangle - \langle \mathbf{S}^{\mathrm{A}} \rangle|}{\langle \mathbf{S}^{\mathrm{M}} \rangle} \times 100\%. \tag{10}$$

In Eqs. (9) and (10), absolute differences are taken to prevent the cancellation effect of oppositely signed differences when computing averages.

The scaling exponent is a measure of the strength of the nonlinear dynamics in a time series (Peng et al., 1994). See Hausdorff et al. (1995) for an implementable algorithm. That automatic extraction preserves these dynamics is critical to fractal analysis of gait. Lastly, the mean squared error between two stride interval series, $\mathbf{S}^{\mathrm{A}}$ and $\mathbf{S}^{\mathrm{M}}$, is

$$\frac{1}{n}\sum_{i=1}^{n}(s_i^{\mathrm{M}} - s_i^{\mathrm{A}})^2 \tag{11}$$

where n is the length of the shorter of the two series, i.e. $n = \min(\mathrm{card}(\mathbf{S}^{\mathrm{M}}), \mathrm{card}(\mathbf{S}^{\mathrm{A}}))$. In our calculations, the two series are shifted with respect to each other to minimize the mean squared error.

4.2. Results

Average performance measures for the different extraction procedures are summarized in Table 2. The averages are computed over the respective measures for each of the 10 signals. We find that for probabilistic and manual extractions, differences in the number of extracted strides ($p = 0.94$), the mean stride intervals ($p = 0.55$) and the scaling exponents ($p = 0.94$), are not statistically significant as determined by the nonparametric Kruskal–Wallis rank sum test. On the other hand, statistically comparing the results of generic versus manual extraction reveals significant differences in mean stride intervals and scaling exponents ($p = 0.034$), but no significant difference in the number of extracted strides ($p = 0.096$). This last result may be surprising, given that the mean numbers of extracted strides for manual (357) and generic extraction (621) appear drastically different. However, the enormous variability (± 330) in the number of strides extracted by the generic method leads to the conclusion of no statistical difference. The difference in mean squared error between probabilistic and generic extractions is significant ($p = 0.023$). Note that mean absolute differences are reported for the first two rows in Table 2 to capture both positive and negative deviations from manual extraction. The mean difference can be misleadingly small due to the cancellation effect of oppositely signed differences. The average length of the stride interval series, that is the product of the mean

Table 2
Automatic extraction compared against manual extraction-average performance over 10 signals

Measure	Manual	Automatic	
		Probabilistic	Generic
Mean no. strides	357 ± 161	361 ± 167	621 ± 330
Mean absolute difference between automatic and manual	–	20.3 ± 20.1	263.5 ± 210.0
Mean stride interval (s)	1.585 ± 0.394	1.463 ± 0.248	1.256 ± 0.624[a]
Absolute difference as percentage of manual (%)	–	6.7 ± 7.8	30.0 ± 18
Mean scaling exponent	0.714 ± 0.085	0.712 ± 0.116	0.635 ± 0.096[a]
Mean squared error with respect to manual (s^2)	–	0.69 ± 1.33	5.39 ± 13.24[b]

[a] Statistically different from manual extraction ($p = 0.034$).
[b] Statistically different from probabilistic extraction ($p = 0.023$).

stride interval and the mean number of strides in Table 2 is generally much less than 10 min, the designated length of the walking trial. This reduction is due to the removal of periods of nonwalking activity (e.g. resting) during the trial and to the elimination of certain strides by the extraction process.

Fig. 3 details a 20 s segment of probabilistic and generic extractions. It is evident from this graph that in the presence of noise, generic extraction erroneously detects too many heel strikes. Fig. 4 portrays a 35 stride segment of the extracted series from signal #9. Note the close agreement between the stride interval series due to probabilistic and manual extractions. In contrast, generic extraction yields a wildly deviant series. In Fig. 5, the number of strides extracted by each method are plotted for the 10 signals. On average, probabilistic extraction produced stride interval series within 20 strides of their manually derived counterparts. In contrast, generic extraction came within 264 strides of manual extraction, on average. Due to large variability in the number of extracted strides within each method, this difference is not statistically significant. Nonetheless, the absolute differences shed light upon how far generic extraction can potentially depart from the correct number of strides. Fig. 6 depicts the mean stride intervals for the stride interval series extracted from each

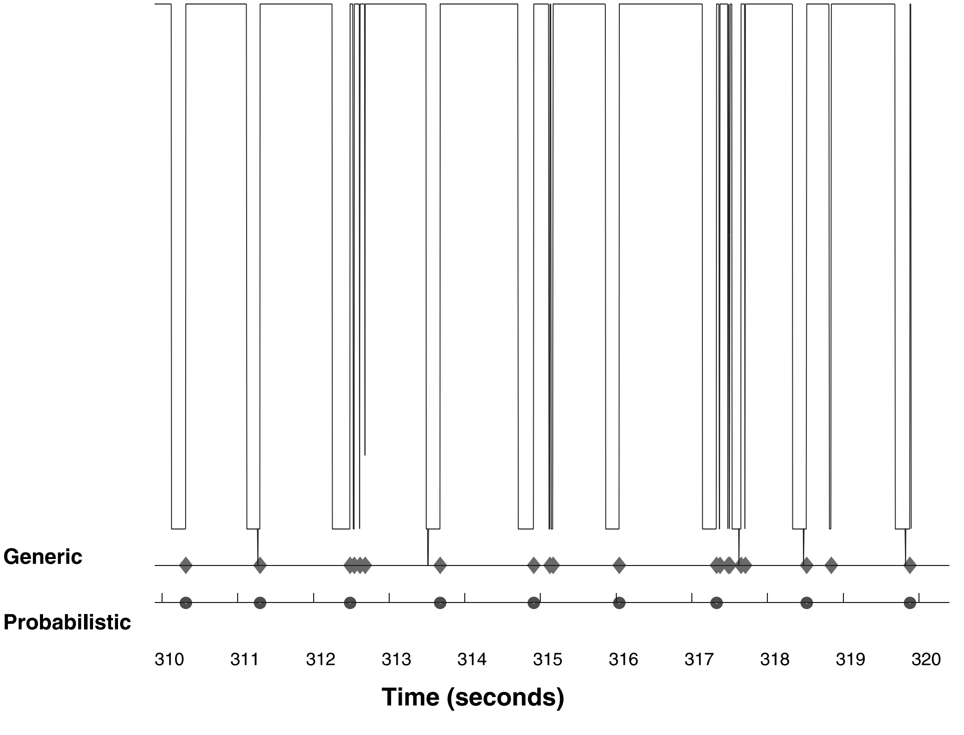

Fig. 3. Comparison of probabilistic and generic extraction of heel strike times. Both horizontal lines represent the time axis. The circles are the heel strike times identified through probabilistic extraction. The diamonds indicate heel strike times selected by generic extraction. One can visually verify that probabilistic extraction has isolated the most likely event times.

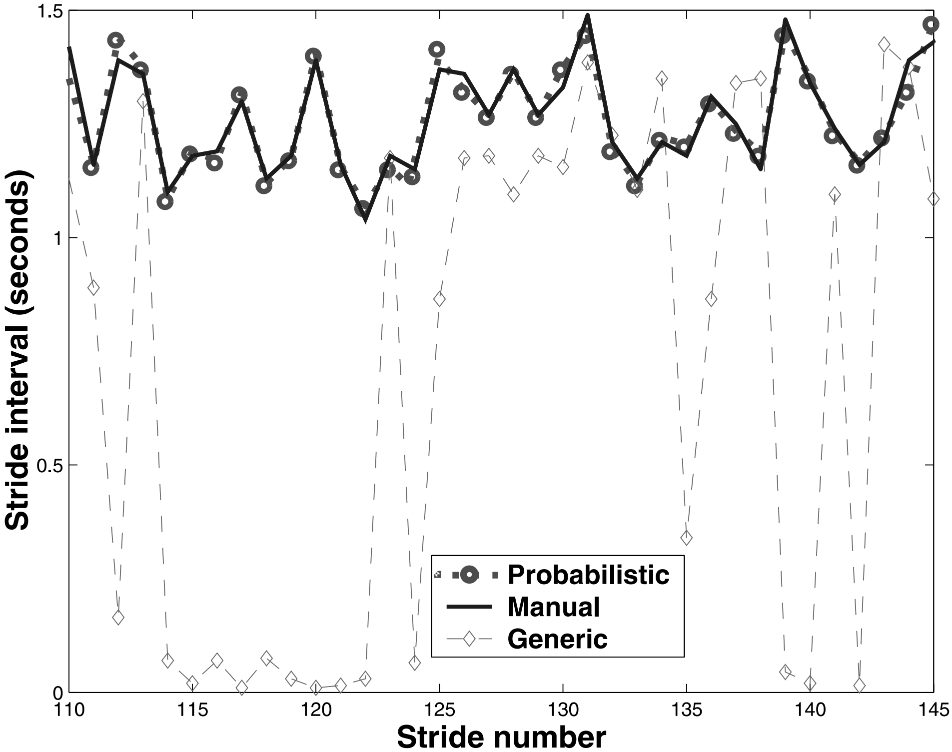

Fig. 4. A sample segment of extracted stride intervals.

of the 10 signals. Mean stride intervals of the probabilistic extracted series agreed within 6.7% of the mean values obtained from manual procedures, compared to an average 30% deviation in the generic extracted series.

In terms of extraction effort, the automatic probabilistic extraction algorithm required an average time of 1.1 ± 0.31 s on a standard desktop computer (Pentium III, 600 MHz) whereas manual extraction required an average of $1.44 \times 10^4 \pm 1.2 \times 10^3$ s (4 h $\pm$ 20 min). In other words, automatic extraction offered a time savings of four orders of magnitude.

4.3. Discussion

4.3.1. Preservation of dynamics

The close agreement between manual and probabilistic extraction in terms of the number of extracted strides, the mean stride interval and the mean scaling exponent indicates that the series arising from probabilistic extraction is a faithful reflection of the true interevent series. Of particular importance is the fact that the scaling exponents are statistically equivalent, implying that probabilistic extraction does indeed preserve the dynamics of the signal. The example also clearly illustrates that straight-

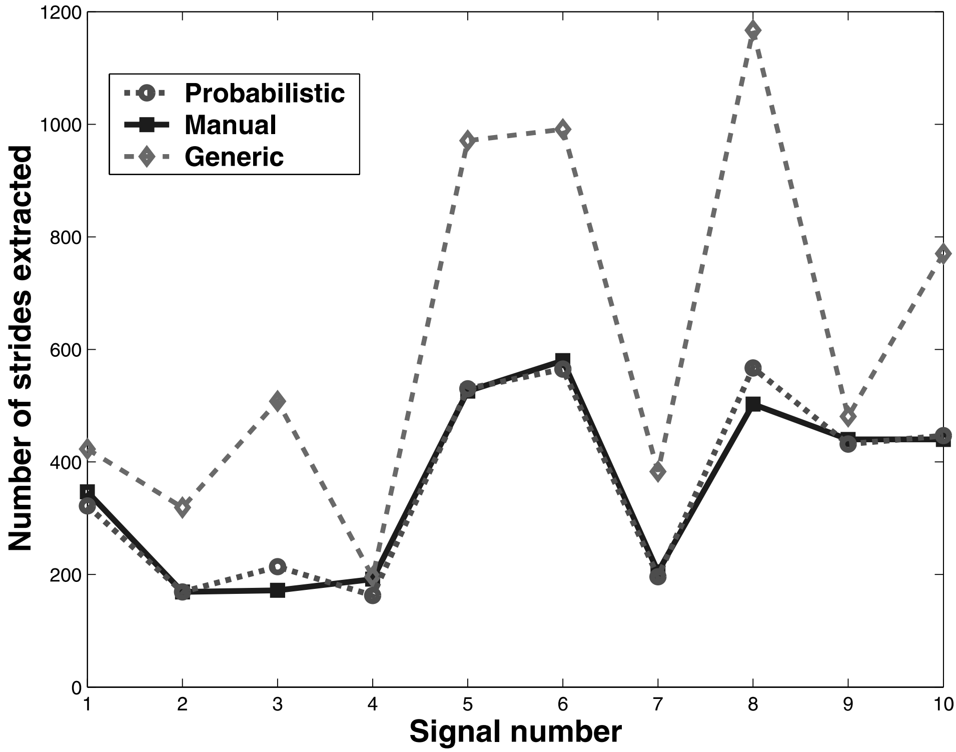

Fig. 5. Comparison of number of strides extracted by manual, probabilistic and generic methods.

forward generic extraction is inadequate. While the numbers of extracted strides, due to large variability, were not significantly different from manual extraction, the scaling exponents were diminished significantly. This means that the nonlinear dynamics of the original signal are lost by generic extraction and that the resulting signal is more random (scaling exponent closer to 0.5) than the true stride interval series (Hausdorff et al., 1996).

4.3.2. Robustness to long pauses

For long and noisy time series, manual extraction is not always 100% correct. The manual process is susceptible to human error, especially after four hours of exhausting visual analysis. On several occasions, the manually extracted series contained extremely long strides, in excess of 5 s. It is likely that the participant stopped walking or that the signal was lost and therefore this stride should be omitted from consideration. However, it was difficult for the human analyst to continuously gauge the absolute time scale and thus pauses in walking were erroneously added to the series. With post-extraction filtering, probabilistic extraction is advantageous in this regard as it automatically eliminates uncharacteristically long strides. For this reason, probabilistic extraction often yielded slightly shorter stride interval series in comparison with manually extracted series (Fig. 5).

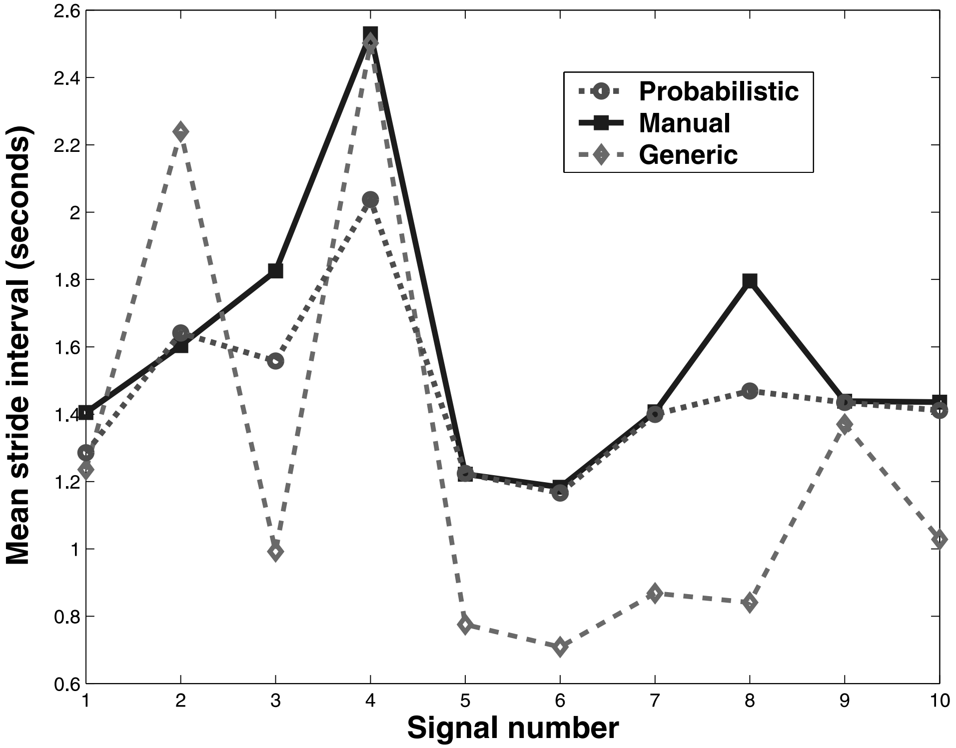

Fig. 6. Comparison of mean stride intervals.

4.3.3. Robustness to spurious short strides

Series extracted by the generic method were typically longer since they contained many spurious events between actual heel strikes. Due to the large number of false heel strike events and their proximity in time, the individual stride intervals and hence the mean stride interval were usually shortest with generic extraction (Fig. 6). An intuitive fix to generic extraction might be to simply remove outliers by discarding times outside of a specified number of standard deviations from the mean. However, as exemplified by the histogram of interevent times due to generic extraction for signal #6, the distribution of times may be bimodal (see Fig. 7). Hence, conventional outlier removal will generally not suffice. Probabilistic extraction automatically penalizes spuriously short strides by assigning low probabilities to the corresponding false heel strikes. In selecting the time of the next heel strike, more "probable" events, i.e. those in the neighbourhood of the expected stride interval are favoured. Probabilistic extraction is therefore robust to spuriously occurring false events.

4.3.4. Symmetry check justified

The example verifies that the post-extraction normality check is indeed necessary. Of the 10 signals, 9 were substantially positively skewed with $\gamma \geqslant 1.84$. Had we not checked for normality, we would have erroneously admitted some uncharacteristi-

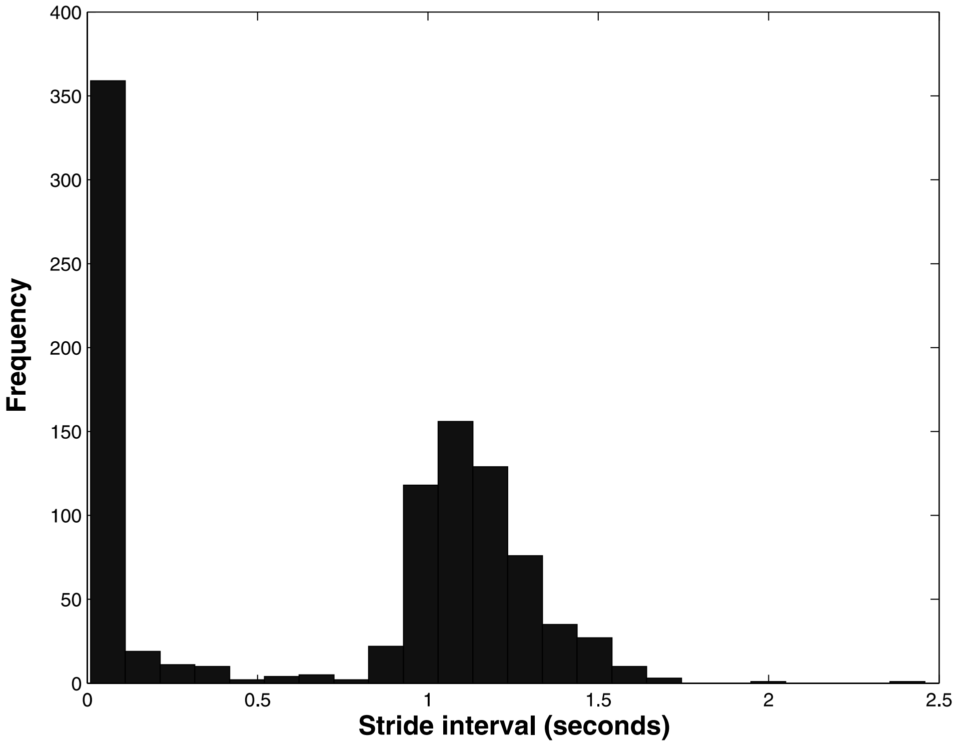

Fig. 7. The interevent times from generic extraction applied to signal #6 are bimodally distributed. Conventional outlier removal cannot be readily applied.

cally long stride intervals and possibly also eliminated some shorter, but valid stride intervals. Based on the experimental results, probabilistic extraction appears to be a viable alternative to manual extraction.

5. Concluding remarks

There are some limitations to the proposed method and some possible applications to other timing data.

5.1. Limitations

As extraction is based upon a single probability density function, the proposed method implicitly assumes a single speed of walking. While this is a limitation, the algorithm is robust to the inflated stride variability (21.2% ± 12 of the gait cycle on average) in the 10 signals used above.

The success of the extraction rests on the assumption that one can make a reasonable initial estimate of the interevent time. Our rather simple method of retaining data in the top 50 percentile may fail when the candidate event times are distributed in a multimodal fashion. In practice, we found the method could handle bimodal data provided the two modes were well separated. For more challenging event time distributions, the use of domain knowledge could greatly improve extraction accuracy. In other words, instead of estimating the mean interevent time from the raw data, one could use a known population value for the mean interevent time in Eq. (6).

5.2. Extensions

While we have only discussed two-state timing signals, the proposed technique could conceivably be applied to long analog traces, such as accelerations or joint angles. One would simply need to develop an appropriate method of extracting the candidate event times from the analog records. An example approach might be the isolation of times corresponding to successive signal peaks or valleys. Once these candidate event times are available, the probabilistic extraction algorithm could be applied verbatim.

The proposed technique could also be extended to detect other quasi-periodic gait events, particularly, in the study of pathological gait, such as cerebral palsy, where events may be occasionally missing or incomplete. Example events within the stance phase, might include heel-contact, foot-flat and toe-off. Hence, the method could be used to estimate other relevant temporal gait parameters such as single stance time and duration of swing phase (unilateral). If bilateral signals were recorded, double stance time could also be extracted.

Probabilistic extraction could conceivably be applied to other physiological data consisting of quasi-periodic events occurring over long time series. Examples include heart beat intervals (Goldberger et al., 1999), breathing rate time series (Szeto et al., 1992) and blood pressure fluctuations (Hughson, Maillet, Dureau, Yamamoto, & Gharib, 1995).

5.3. Conclusion

In this paper, we proposed a conceptually simple and easily implementable probabilistic stride interval extraction technique for long and noisy gait timing signals. The utility and robustness of the method was illustrated with 10 timing signals recorded from children with Spastic Diplegia. Close adherence to laborious manual extractions was demonstrated in terms of stride numbers, mean stride times and scaling exponents. The probabilistic extraction algorithm is compatible with data collected from standard gait lab equipment (generic footswitches and an off-the-shelf, entry level portable logger), foregoing the need for expensive custom hardware. In turn, this may facilitate more widespread study of fractal dynamics in human movement research, a pursuit that has already led to deeper understanding of movement coordination, control and pathology (see for example, Duarte & Zatisiorsky, 2001; Schmeisser et al., 2001; Miyazaki et al., 2001; Hausdorff et al., 1997).

Acknowledgements

The authors would like to acknowledge the support of the Natural Sciences and Engineering Research Council of Canada, the Hospital for Sick Children Foundation and the Bloorview Childrens Hospital Foundation.

References

Bendat, J., & Piersol, A. (2000). *Random data: analysis and measurement procedures* (third ed.). New York: John Wiley & Sons.

Chau, T. (2001). A review of analytical techniques for gait data. Part 1: fuzzy, statistical and fractal methods. *Gait and Posture, 13*, 49–66.

Duarte, M., & Zatisiorsky, V. (2001). Long-range correlations in human standing. *Physics Letters A, 283*, 124–128.

Gerstein, G., & Mandelbrot, B. (1964). Random walk models for spike activity of single neuron. *Biophysical Journal, 4*, 41–68.

Goldberger, A., Amaral, L., Hausdorff, J., Ivanov, P., Peng, C., & Stanley, H. (1999). Fractal dynamics in physiology: alterations with disease and aging. *Proceedings of the National Academy of Sciences of the United States of America* (Vol. 99, pp. 2466–2472).

Griffin, L., West, D., & West, B. (2000). Random stride intervals with memory. *Journal of Biological Physics, 26*, 185–202.

Han, J., & Gao, W. (2001). Robust speech recognition method based on discriminative environment feature extraction. *Journal of Computer Science and Technology, 16*, 458–464.

Hansen, M., Haugland, M., Sinkjaer, T., & Donaldson, N. (2002). Real time foot drop correction using machine learning and natural sensors. *Neuromodulation, 5*, 41–53.

Hausdorff, J., Forman, D., Pilgrim, D., Rigney, D., & Wei, J. (1992). A new technique for simultaneous monitoring of electrocardiogram and walking cadence. *American Journal of Cardiology, 70*, 1064–1071.

Hausdorff, J., Lertratanakul, A., Cudkowicz, M., Peterson, A., Kaliton, D., & Goldberger, A. (2000). Dynamic markers of altered gait rhythm in amyotrophic lateral sclerosis. *Journal of Applied Physiology, 88*, 2045–2053.

Hausdorff, J. M., Mitchell, S. L., Firton, R., Peng, C. K., Cudkowicz, M. E., Wei, J. Y., & Goldberger, A. L. (1997). Altered fractal dynamics of gait: reduced stride-interval correlations with aging and Huntington's disease. *Journal of Applied Physiology, 82*, 262–269.

Hausdorff, J. M., Peng, C. K., Ladin, Z., Wei, J. Y., & Goldberger, A. L. (1995). Is walking a random walk? Evidence for long-range correlations in the stride interval of human gait. *Journal of Applied Physiology, 78*, 349–358.

Hausdorff, J. M., Purdon, P. L., Peng, C. K., Ladin, Z., Wei, J. Y., & Goldberger, A. L. (1996). Fractal dynamics of human gait: stability of long-range correlation in stride interval fluctuations. *Journal of Applied Physiology, 80*, 1448–1457.

Hochreiter, S., & Mozer, M. (2001). A discrete probabilistic memory model for discovering dependencies in time. In *Artificial neural networks, ICANN 2001, Vol. 2130 of lecture notes in computer science* (pp. 661–668). Berlin: Springer-Verlag.

Hughson, R., Maillet, A., Dureau, G., Yamamoto, Y., & Gharib, C. (1995). Spectral analysis of blood-pressure variability in heart transplant patients. *Hypertension, 25*, 643–650.

Jancovic, P., & Ming, J. (2001). A probabilistic union model with automatic order selection for noisy speech recognition. *Journal of the Acoustical Society of America, 110*, 1641–1648.

Kaneko, K., & Sakamoto, K. (2001). Spontaneous blinks as a criterion of visual fatigue during prolonged work on visual display terminals. *Perceptual and Motor Skills, 92*, 234–250.

Longstaff, M., & Heath, R. (1999). A nonlinear analysis of the temporal characteristics of handwriting. *Human Movement Science, 18*, 485–524.

Mendel, J. (1991). Tutorial on high order statistics in signal processing and systems theory: theoretical results and some applications. *Proceedings of the IEEE* (79, pp. 278–305).

Miyazaki, M., Kadota, H., Kudo, K., Masani, K., & Ohtsuki, T. (2001). Fractal correlation of initial trajectory dynamics vanishes at the movement end point in human rapid goal-directed movements. *Neuroscience Letters, 304*, 173–176.

Pappas, I., Popovic, M., Keller, T., Dietz, V., & Morari, M. (2001). A reliable gait phase detection system. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 9*, 113–125.

Peng, C., Buldyrev, S., Havlin, S., Simons, M., Stanley, H., & Goldberger, A. (1994). Mosaic organization of DNA molecules. *Physical Review E, 49*, 1685–1689.

Proctor, S., White, R., Robins, T., & Echeverria, D. (1996). Effect of overtime work on cognitive function in automotive workers. *Scandinavian Journal of Work, Environment and Health, 22*, 124–132.

Ripley, B. (1996). *Pattern recognition and neural networks*. Cambridge: Cambridge University Press.

Schellekens, J., Sijtsma, G., Vegter, E., & Meijman, T. (2000). Immediate and delayed after-effects of long lasting mentally demanding work. *Biological Psychology, 53*, 37–56.

Schmeisser, E., McDonough, J., Bond, M., Hislop, P., & Epstein, A. (2001). Fractal analysis of eye movements during reading. *Optometry and Vision Science, 78*, 805–814.

Skelly, M., & Chizeck, H. (2001). Real-time gait event detection for paraplegic FES walking. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 9*, 59–68.

Szeto, H., Cheng, P., Decena, J., Yi, C., Wu, D., & Dwyer, G. (1992). Fractal properties in fetal breathing dynamics. *American Journal of Physiology, 263*, 141–147.

West, B. J., & Griffin, L. (1998). Allometric control of human gait. *Fractals, 6*, 101–108.

West, B. J., & Griffin, L. (1999). Allometric control, inverse power laws and human gait. *Chaos, Solitons and Fractals, 10*, 1519–1527.

Williamson, R., & Andrews, B. (2000). Gait event detection for FES using accelerometers and supervised machine learning. *IEEE Transactions on Rehabilitation Engineering, 8*, 312–319.