Physiol. Meas. **25** (2004) 1425–1436 PII: S0967-3334(04)78902-9

Models for estimating the change-point in gas exchange data

G E Kelly[1], **J K Lindsey**[2] **and A G Thin**[3]

[1] Department of Statistics, University College Dublin, Belfield, Dublin 4, Republic of Ireland
[2] Faculty of Economics, University of Liège, Liège, Belgium
[3] School of Life Sciences, Heriot-Watt University, Edinburgh, UK

E-mail: gabrielle.kelly@ucd.ie

Received 6 April 2004, accepted for publication 9 September 2004
Published 4 October 2004
Online at stacks.iop.org/PM/25/1425
doi:10.1088/0967-3334/25/6/008

Abstract

In subjects undertaking an incremental exercise test to exhaustion, the onset of metabolic acidosis can be detected by an increased rate of carbon dioxide output ($\dot{V}CO_2$) relative to the rate of increase of oxygen uptake ($\dot{V}O_2$). To locate the change-point (the gas exchange threshold) in such subjects, a two-line regression model relating these two quantities has been used, where the location of the change-point is unknown. We argue that statistical models where the change-point is set on time (rather than $\dot{V}O_2$) are more appropriate. This is because $\dot{V}O_2$ is not monotone in time. We use novel statistical methodology of hidden Markov models to demonstrate the existence of the change-point. We use time series models, to estimate the position of the change-point. In these models distributions other than the multivariate normal are considered. For some subjects, the variance of $\dot{V}CO_2$ increases with time because of increasing ventilation and this is also modelled. The results are illustrated using gas exchange data on three healthy subjects who performed a 20 W min^{-1} workrate ramp test.

Keywords: gas exchange threshold, change-point, statistical modelling, hidden Markov models

Introduction

Incremental work rate tests on a cycle ergometer are frequently used to assess the aerobic fitness of subjects. The onset of an exercise-induced metabolic acidosis can be detected by an increase in the rate of carbon dioxide output ($\dot{V}CO_2$) relative to the rate of increase in oxygen uptake ($\dot{V}O_2$) (Beaver *et al* 1986, Sue *et al* 1988). The breath-by-breath values of carbon dioxide output ($\dot{V}CO_2$) are plotted against oxygen uptake ($\dot{V}O_2$) and related by a two-segment linear regression. Thus, the point of increase in ratio is modelled by fitting two regression lines with an unknown join point to the data. The $\dot{V}O_2$ value of the intersection of the two regression lines is known as the gas exchange threshold (GET). There is extensive physiological and statistical literature on this problem that is known as the change-point problem. An overview can be found in Carlstein *et al* (1994) and Krisnaiah and Miao (1998).

Nonparametric approaches to the change-point problem have been proposed by Hall and Titterington (1994) and Loader (1996). These do not require the mean function to be specified *a priori*, i.e. do not require the regression to be linear. However they assume repeated observations on an individual are independent. These approaches have been applied in the physiological literature by Sherrill *et al* (1990) and Magalang and Grant (1995).

It would be more realistic if the change-point was set on time instead of $\dot{V}O_2$ because $\dot{V}O_2$ is not monotone in time. Setting the change-point on $\dot{V}O_2$ implies that for a number of breaths (i.e. a period of time) around the onset of metabolic acidosis, the person is switching between two metabolic states (i.e. aerobic to aerobic plus supplemental anaerobiosis metabolism), (Koike *et al* 1990), which is not biologically reasonable. However, it is recognized, that the time of occurrence of the change-point in a test, will in part depend on the size of the work rate ramp (i.e. for a given subject it will occur after a shorter time period on a high work rate ramp compared to a low one). But since individuals typically undergo repeat tests to monitor their fitness, it is the difference in repeat estimates of the change-point under the same test conditions which is important. An increase in the time of occurrence of the GET would therefore be indicative of an increase in fitness. However, to make comparisons between individuals or to norms, the time values can be converted to $\dot{V}O_2$ as described below. Kelly *et al* (2001 and 2002) extended the classical model by addressing the situation where observations are correlated and by using time series techniques. They introduced autoregressive moving average (ARMA) models for this type of data. The distribution in the model was assumed to be normal. Moreover, because standard time series software assumes the time points are equally spaced, the analysis proceeded on a breath-by-breath basis.

We shall extend this work, firstly by setting the change-point on time and then by considering different models for the response $\dot{V}CO_2$ as a function of $\dot{V}O_2$. We shall first look at continuous-time hidden Markov chain models for these data. We shall again consider time series models but the software we use does not require the time points to be equally spaced. We also look at distributions other than the multivariate normal. We consider a subclass of elliptically-contoured distributions which are a multivariate generalization of the power-exponential family introduced by Gómez *et al* (1998) and discussed by Lindsey (1999). They include the multivariate normal as a special case. They cover both thick- and thin-tailed distributions and so can be useful in providing robustness against outliers. We shall also consider the multivariate Student-*t* and the asymmetric multivariate Laplace, the latter discussed in Kotz *et al* (2001). The choice is dictated by the fact that these are the only known ones with covariance matrices. The covariance matrix can be structured for time series dependence. For some subjects the variance changed with the location regression function; this is also incorporated in these models. Different models will be compared using Akaike's information criterion (Akaike 1973).

Methods

Experimental design

Data from three healthy subjects during exercise are examined in detail. This is part of a larger study described in Kelly *et al* (2001, 2002). Each subject exercised to exhaustion on an electromagnetically braked cycle ergometer (Excalibur, Lode BV, Groningen, the Netherlands). The test began with 1 min of unloaded cycling, a 4 min warm-up at 20 watts (W) and then a 20 W min^{-1} workrate ramp. Subjects wore a nose clip and mouthpiece connected to a heated wire flow meter (SensorMedics, CA, USA). The flow signal was digitally integrated to give tidal volume and respired gases were continuously analysed by rapidly responding O_2 (paramagnetic) and CO_2 (infrared) analysers (response times 110 and 90 ms respectively) (SensorMedics Vmax 229, SensorMedics, CA, USA). The analysers were calibrated and the transients checked prior to each test using a calibration syringe, precision O_2 and CO_2 gas mixes and a rapidly switching solenoid valve (SensorMedics, CA, USA). Expiratory flow and fractional concentrations of expired O_2 and CO_2 were measured and digitally integrated on a breath-by-breath basis to provide estimates of minute ventilation ($\dot{V}E$), $\dot{V}O_2$ and $\dot{V}CO_2$.

During the initial part of an incremental exercise test, the changes in gas exchange measured breath-by-breath at the mouth are non-linear with respect to the work rate ramp as a result of both a non-linear lag in O_2 uptake kinetics and a lag in CO_2 output due to an initial adjustment of CO_2 stores in the body. Therefore data points prior to the completion of the first minute of work rate ramp, and any points beyond where it was evident that CO_2 stores were still filling, were excluded from the analysis (Beaver *et al* 1986). Furthermore towards the end of an incremental exercise test, significant hyperventilation is manifest in an attempt to compensate for the exercise-induced metabolic acidosis. Since this occurs beyond the region of interest, data points above 75% of the maximum $\dot{V}O_2$ value for a given test were excluded as they could potentially influence the determination of the GET. In no case was there evidence of marked hyperventilation below this point. The data also contain occasional breath values that are clearly artifactual; these typically result from swallowing, coughing or premature ending of the breath for some other reason. Using the criteria of Lamarra *et al* (1987), such breaths were deleted from the data before performing the analysis.

Statistical analysis

The $\dot{V}CO_2$ versus time profile is different for different subjects and thus is modelled differently in each case. We consider hidden Markov models and time series models.

Hidden Markov models

We denote the $\dot{V}CO_2$ values at time t by y_t and the corresponding $\dot{V}O_2$ values by x_t. We introduce a discrete random variable s_t in each time period, referred to as the state of the system at time t, that takes values on the integers 0 and 1 and indicates the regime from which a particular observation y_t has been drawn. State 0 refers to aerobic metabolism and state 1 to aerobic and anaerobic metabolism. State 1 is absorbing. Once the transition to some anaerobic metabolism has taken place it cannot be reversed. We note that s_t is hidden. s_t is Markov i.e. s_t depends only on the most previous value s_{t-1}, but this induces longer term dependence in the observed sequence y_t. Now, $s_t = k$ indicates that the observation y_t is drawn from a density $f(y_t;\theta_k)$ where θ_k represents the unknown parameters in the model in regime k, $k = 0$ or 1. We consider the cases where $f(y_t;\theta_k)$ is given by the normal, the power-exponential and the asymmetric Laplace distributions, all with location parameter given by $\alpha_k + \beta_k x_t$. Thus the mean function for the two states is given by two adjoining lines. The power-exponential and asymmetric Laplace distributions are described in the appendix. When the probability of being in state 0 first reaches 0 at some time t i.e. $P(s_t = 0) = 0$ then this t estimates the change-point. Maximum likelihood is used to estimate the parameters of the models, similar to that described in MacDonald and Zucchini (1997), although they work in discrete time with discrete distributions.

Time series models

We consider autoregressive (AR) processes of order 1, AR(1). Letting Z_t be a purely random process with a mean of zero and constant variance then a process W_t is said to be AR(1) if $W_t = \beta W_{t-1} + Z_t$, where β is a constant. In the change-point model, the location function is modelled as $\alpha_1 + \beta_1 x_t$ for $t \leqslant \tau$ and $\alpha_2 + \beta_2 x_t$ for $t > \tau$ where τ represents the unknown change-point. The observations in a series are assumed (1) to be independent; (2) to follow an AR(1) process or (3) to follow an AR(1) process with the variance an increasing function of the location parameter. In other words the models are given by two adjoining lines:

$$y_t = \alpha_1 + \beta_1 x_t + u_t \qquad \text{for} \quad t \leqslant \tau$$

and

$$y_t = \alpha_2 + \beta_2 x_t + u_t \qquad \text{for} \quad t > \tau.$$

The noise u_t is assumed to be independent or follow an AR(1) process. As for the hidden Markov models, different distributions (for u_t) will be considered: the multivariate Student-*t*, the multivariate power-exponential, and the asymmetric Laplace distribution. These models can all be fitted using the second author's `elliptic` function in *R* (Ihaka and Gentleman 1996). We compute the log likelihood, $\log(\hat{L})$, which is the logarithm of the 'probability' or likelihood of the observed data, under the different models. We then compute Akaike's information criterion (AIC) (Akaike 1973) which is given by the formula $-\log(\hat{L}) + p$ where p is the number of unknown parameters in the model. We use AIC to select the best model. In our examples, this amounts to choosing the model which minimizes the estimated negative log likelihood, $-\log(\hat{L})$, with an extra penalty for the distribution functions containing an unknown parameter, multivariate power-exponential, Student-*t*, and asymmetric Laplace and models where the variance increases with the location parameter. The penalty is necessary because typically the greater the number of parameters in the model the better the fit i.e. the greater the estimated log likelihood.

Fitting is done for each subject separately because typically estimates of the change-point need to be constructed on a per subject basis. Due to the discontinuity in the location function, the maximum likelihood estimate (mle) of the change-point is found by taking a grid of values for time and selecting the time point which maximizes the likelihood. At each step, we use the estimate from the previous as an initial estimate, which speeds up convergence.

Finally the time change-point is converted to a $\dot{V}O_2$ value by interpolation. This is given by the average of the five $\dot{V}O_2$ values corresponding to breaths at times adjacent to the time change-point. Three values were also used but there was little difference in the results.

Results

We take breath-by-breath values of $\dot{V}CO_2$ and $\dot{V}O_2$ as our data.

Subject 1

The plot of $\dot{V}O_2$ versus time is shown in figure 1. It can be seen that $\dot{V}O_2$ is not monotone in time as mentioned in the introduction.

The first question to address is whether or not a change-point exists. The two-state hidden Markov chain model with a normal distribution gives a log likelihood of 244.2. Figure 2 shows the probability of being in state 0 for these data. There is a clear transition between the two states going from probability 1 of being in state 0 to probability 0. The change-point in time is approximately 253 and in terms of $\dot{V}O_2$ is 1.688 l min^{-1}. Using

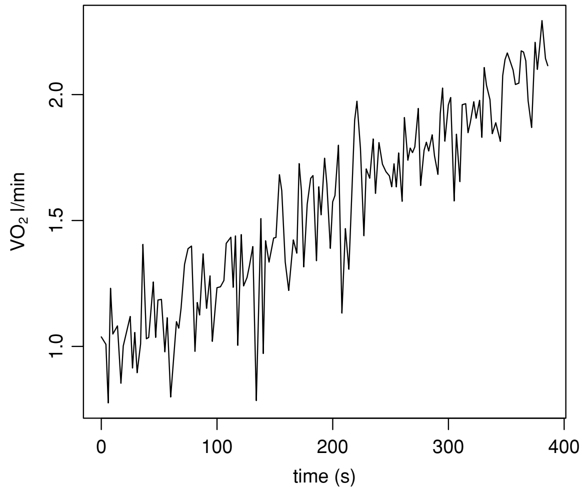

Figure 1. A plot of VO_2 in litres per min versus time for subject 1.

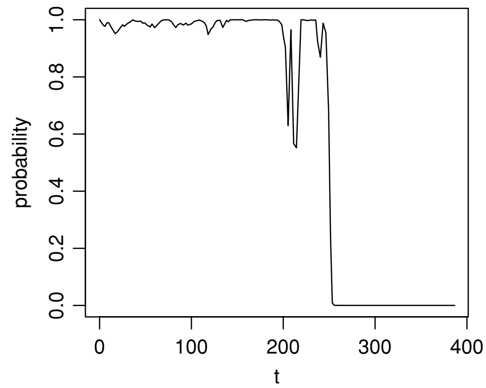

Figure 2. Probability of state 0 versus time for the two-state hidden Markov chain model for subject 1.

the power-exponential distribution gives a log likelihood of 249.3, and since it has only one extra parameter over the normal, by AIC this is a significant improvement. The estimate of the power parameter is $\hat{\kappa} = 0.51$, not far from the multivariate Laplace. The estimate of the change-point is again approximately 253. The Student-*t* gave a log likelihood of 246.8 with the associated degrees of freedom 5. The change-point in time is again about 253. The asymmetric Laplace model gives a log likelihood of 245.6, not as good a fit as the others. The asymmetry parameter $\hat{\kappa} = 1.137$ indicating some asymmetry. The estimate of the change-point in time is 249 and in terms of $\dot{V}O_2$ is again $1.688\ \mathrm{l\ min^{-1}}$.

We then consider the time series models. Firstly the two-line regression model, assuming independent normal distributions, gives the mle of the change-point at time 236 ($\dot{V}O_2 = 1.723\ \mathrm{l\ min^{-1}}$) corresponding to a log likelihood of 249.0. A local maximum occurs at times 243–245 where the log likelihood is 248.7. Because the change-point is estimated, the ratio of the residuals sums of squares from the one- and two-line models does not have an F-distribution under the null hypothesis of no change in slope (Davies 1987). However, the log

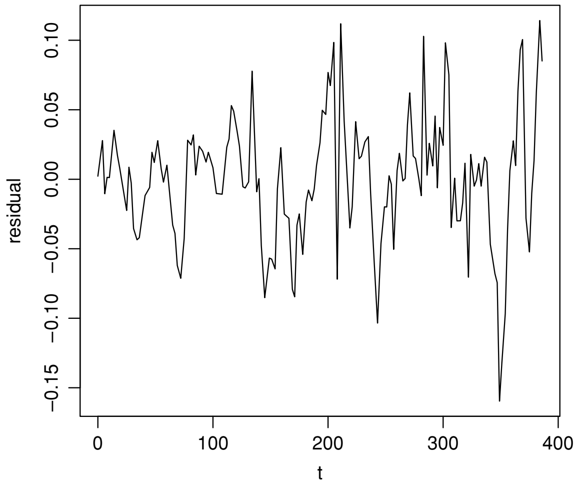

Figure 3. Residuals versus breath index from the two-line regression model of VCO_2 versus VO_2 with independent normal distributions for subject 1.

likelihood for a single line to these data is 187.6. Thus, using Akaike's information criterion, the change-point model is clearly an improvement.

The residuals from the two-segment regression model assuming independence were examined and are shown in figure 3. Positive residuals tend to be followed by negative ones, indicating autocorrelation. The Durbin–Watson statistic testing for serial correlation is significant with a *P*-value of 0.0002. The multivariate normal distribution was then fitted with continuous-time AR(1) correlation. This has a maximum at times 235–236 ($\dot{V}O_2 = 1.723$ l min^{-1}) with a log likelihood of 281.6. There are local maxima at times 208–210 and at times 243–245. The autoregressive parameter is estimated to be 0.81. Figure 4 shows the correlogram of the recursive residuals from this model, indicating little correlation. We note that this model provides a better fit than the hidden Markov model by AIC.

Next, the multivariate power-exponential distribution was fitted to the data. The mle of the change-point in time is $\hat{\tau} = 235$–236 with $\dot{V}O_2 = 1.723$ l min^{-1}, with a log likelihood of 283.4 and the estimate of the power parameter $\widehat{\kappa} = 40.4$, strongly indicating non-normality with short tails. The asymmetric multivariate Laplace model also estimates the change-point to be at 235–236 with a log likelihood of 279.3. The asymmetry parameter is 0.00 with standard error 0.077. Again the multivariate Student-*t* model has a maximum at time 235–236 with a log likelihood of 280.8. The estimated degrees of freedom is 67.9, indicating a distribution close to normal. A plot of the residuals for the Laplace model showed that the residuals increase with the location parameter. Indeed, this pattern can also be seen in the fit of the power-exponential and multivariate Student-*t* models. This indicates heteroscedasticity.

Therefore for each of the distributions considered above, the variance was modelled as a linear function of the location parameter. This allows for heteroscedasticity and also permits the variance function to change at the change-point. For the normal model, the maximum is at time 209–210, corresponding $\dot{V}O_2 = 1.461$ l min^{-1}, with a log likelihood 290.3, a considerable improvement on the previous fit. For the power-exponential model, the mle is 209–210 and the power parameter approached $\hat{\kappa} = \infty$, indicating a short-tailed distribution. However there is no improvement in the log likelihood compared to the normal by AIC. For the multivariate Student-*t* model, the estimated degrees of freedom at the mle approached ∞,

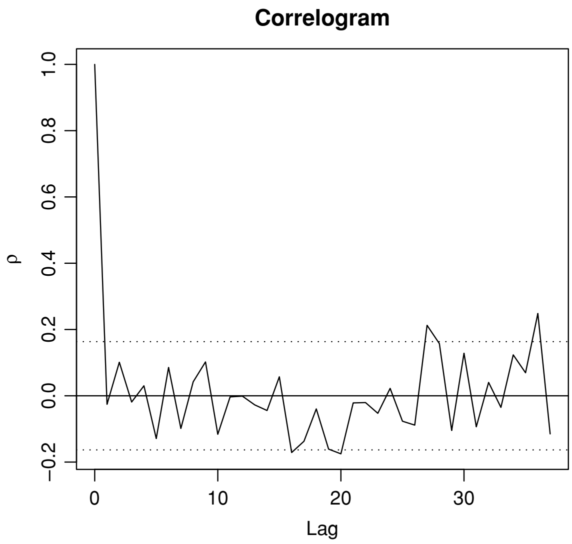

Figure 4. Correlogram of the residuals from the two-line regression model of VCO_2 versus VO_2 with normal distribution and AR(1) for subject 1.

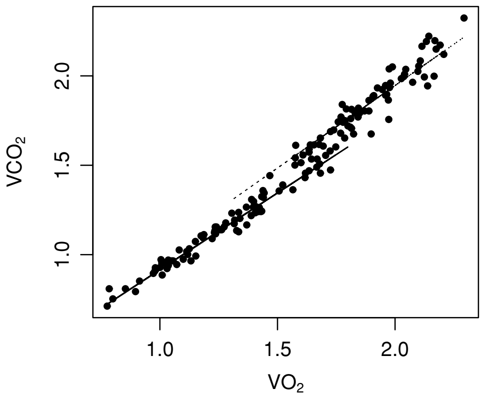

Figure 5. The data and estimated regression lines from the model with the power exponential distribution, AR(1), and variance increasing with the location parameter for subject 1.

indicating a normal model. The fit given by the asymmetric multivariate Laplace model, is a log likelihood of 288.0 at the change-point of 208–210 which is worse than the normal. The asymmetry parameter is 0.00. Therefore the power exponential gives the best fit. The estimated regression lines are shown in figure 5 for the power exponential with AR(1) and variance a linear function of the location parameter. We see that the traditional approach of taking the change-point on $\dot{V}O_2$ at the intersection of the two lines is not a good estimate of the change-point because the two lines are almost parallel. The lines overlap because the plot is in terms of $\dot{V}O_2$, which is not monotone in time, and not time itself; if the plot was in terms of time, there would be no overlap.

Standard likelihood theory for finding confidence intervals does not apply here to the change-point. The likelihood ratio statistic does not have a chi-squared distribution in the simplest of models, i.e. that of estimating the change-point in an i.i.d. normal mean-shift setup

Table 1. Results from hidden Markov models and time series models for subject 1.

Distribution	Change-point in time	Change-point in $\dot{V}O_2$	$-\log(\widehat{L})$	AIC
	Hidden Markov models			
Normal	253	1.688	−244.2	−238.2
Power exponential	253	1.688	−249.3	−242.3
Student-*t*	253	1.688	−246.8	−239.8
Asymmetric Laplace	249	1.688	−245.6	−238.6
	Time series models			
Normal independent	236	1.723	−249.0	−244.0
Normal+AR(1)	235–236	1.723	−281.6	−275.6
Power exponential+AR(1)	235–236	1.723	−283.4	−276.4
Student-*t*+AR(1)	235–236	1.723	−280.8	−273.8
Asymmetric Laplace+AR(1)	235–236	1.723	−279.3	−272.3
Normal+AR(1)+varfn	209–210	1.461	−290.2	−283.3
Power exp+AR(1)+varfn	209–210	1.461	−292.1	−284.1
Student-*t*+AR(1)+varfn	209–210	1.461	Same as normal+AR(1)+varfn	
a.Laplace+AR(1)+varfn	208–210	1.461	−288.0	−280.0

AR(1) refers to an autoregressive process of order 1 and varfn refers to variance function an increasing function of the mean was fitted.

(Hinkley 1970). The bootstrap method is also not readily available as i.i.d. residuals cannot be formed due to the heteroscedastic and the time series nature of the models.

The results for subject 1 are summarized in table 1.

Subject 2

A similar analysis was carried out for a second subject. First, the two-state hidden Markov chain model with a normal distribution was fitted to the data. As for subject 1, there is a clear transition between the two states going from probability 1 of being in state 0 to probability 0. The log likelihood under this model is 266.3, with time change-point approximately 325, corresponding to $\dot{V}O_2 = 2.082\ \mathrm{l\ min^{-1}}$. For the power-exponential model the log likelihood and change-point were the same with estimated power parameter 0.966, close to 1.0, the value for the normal. The results for the Student-*t* were also the same as the normal. With the asymmetric Laplace distribution the model did not converge. $\dot{V}CO_2$ was then plotted versus $\dot{V}O_2$ with the estimated regression lines from the normal model. The change-point is taken to be at time 325. We saw again that taking the change-point on $\dot{V}O_2$ at the intersection of the two lines was not a good estimate of the change-point because the two lines were almost parallel. This result is also true for the other models fitted below.

The results for the time series models were as follows. The independent normal model estimates the change-point to be between 319 and 320, corresponding $\dot{V}O_2 = 2.068\ \mathrm{l\ min^{-1}}$, with a log likelihood of 271.1. The AR(1) model, again with the normal distribution, gives a maximum also at 319–320 with a log likelihood of 300.7. However the likelihood is far from quadratic with several local maxima. The power exponential model with AR(1) gave a change-point of 319–320 with a log likelihood of 302.5, not a great improvement over the normal. The mle of the power parameter is 0.006 indicating a long-tailed distribution. The Student-*t* with AR(1) also estimated the change-point to be at 319–320. The estimated degrees of freedom are 43.6 and the fit is identical to the normal. This is in contradiction with the power exponential model. A plot of the recursive residuals indexed by breath number

Table 2. Results from hidden Markov models and time series models for subject 2.

Distribution	Change-point in time	Change-point in $\dot{V}O_2$	$-\log(\widehat{L})$	AIC
	Hidden Markov models			
Normal	325	2.082	−266.3	−260.3
Power exponential	325	2.082	Same as normal	
Student-*t*	325	2.082	Same as normal	
Asymmetric Laplace	–	–	No convergence	
	Time series models			
Normal independent	319–320	2.068	−271.1	−266.1
Normal+AR(1)	319–320	2.068	−300.7	−294.7
Power exponential+AR(1)	319–320	2.068	−302.5	−295.5
Student-*t*+AR(1)	319–320	2.068	Same as normal+AR(1)	
Asymmetric Laplace+ AR(1)	–	–	No convergence	
Normal+AR(1)+varfn	296–297	2.071	−321.9	−314.9
Power exp+AR(1)+varfn	297	2.071	−323.7	−315.7
Student-*t*+AR(1)+varfn	296–297	2.071	Same as normal+AR(1)+varfn	
a.Laplace+AR(1)+varfn	–	–	No convergence	

AR(1) refers to an autoregressive process of order 1 and varfn refers to variance function an increasing function of the mean was fitted.

for the normal model showed that the variance increases with time. Therefore, the model with the normal distribution was fitted, where the variance increases as a function of the location parameter and the correlation structure is AR(1). The improvement in log likelihood is considerable, now being 321.9, with the mle of the change-point occurring at time 296–297, corresponding $\dot{V}O_2 = 2.071\ \mathrm{l\,min^{-1}}$. For the same model with the power exponential the mle of the change-point is 297 with a log likelihood of 323.7. This is a slight improvement over the normal by AIC. The mle of the power parameter now tended to ∞ indicating a short-tailed distribution. This model with the multivariate Student-*t* distribution gives estimated degrees of freedom tending to ∞. Thus it is the same as the normal. Here we see, the Student-*t* and power-exponential models give contradictory results on the nature of the distribution while giving similar answers for the change-point. It is possible that the data cannot distinguish between the shape of the tail of the distribution and an increasing variance with time. This model with the asymmetric Laplace distribution did not converge indicating an inappropriate model for these data.

The results for subject 2 are summarized in table 2.

Subject 3

In a third data set, the estimate of the change-point is about the same, 368 with corresponding $\dot{V}O_2 = 2.767\ \mathrm{l\,min^{-1}}$ for all Markov models. By AIC there was little difference between the models, apart from the skew Laplace which had a slightly worse fit. The analogue of figure 5 for these data where the change-point is estimated to be 368 shows a jump at the change-point rather than a smooth transition.

The normal independence model estimates the time change-point to be at 362, corresponding $\dot{V}O_2 = 2.847\ \mathrm{l\ min^{-1}}$ with a log likelihood of 268.5. When the AR(1) is introduced, the change-point is again estimated to be at 361–362 and the log likelihood is 332.6. The Student-*t* gave a fit similar to the normal with estimated degrees of freedom tending to ∞. The fit of the skew Laplace distribution was a little worse than these two.

The estimated change-point was 361–362 with estimated log likelihood 330.3 and estimated asymmetry parameter 0.0. When the model includes a variance function which increases with the location parameter as described above, there is little improvement in the fit of either the normal or the skew Laplace model. Thus the normal model with AR(1) correlation adequately describes the data.

Discussion

The v-slope method of Beaver *et al* (1986) for estimating the GET has found much prominence in the physiology literature. They required that the change in slope from the lower line segment to the upper be greater than 0.1, which is *ad hoc*. Using a more complex gas measurement than considered here, smoothing of $\dot{V}CO_2$ based on changes in the fractional concentration of end-tidal CO_2, interpolation and a 9s moving-average filter were used to smooth the breath-to-breath irregularities. They argued that filtering would not obscure the existence of a change-point. However, it may change the estimate of its location and thus induce bias. Moreover, for these 'corrected' data, the noise may still not be independent so that a model where there is dependence between successive observations is more appropriate. For comparison the Beaver estimates (unsmoothed data) for subjects 1, 2 and 3 were 1.515, 1.904 and 1.879 $\mathrm{l\,min^{-1}}$ respectively. The corresponding estimates from Kelly *et al* (2001 and 2002) were 1.706, 2.353 and 2.168 $\mathrm{l\,min^{-1}}$. There is a discrepancy for subject 3 with the results here, as a result of the breaths being treated as equally spaced (an increase between two breaths was not as sharp as appeared), or in the Beaver case because the data were unsmoothed.

We have found that models which have an abrupt transition at the threshold fit our data well. Our time series models fitted better than the hidden Markov models. However the Markov models provide a strong indication of the existence of a change-point. The time series models assume the change-point exists at the outset. By fixing the change-point on time, we also avoid finding the intersection point of almost parallel lines and thus a change-point perhaps outside the range of the data. We do note that the slope of the work rate ramp does have an influence on the slope of the line above the change-point, i.e. the greater the work rate ramp, the steeper the slope. This means that it could be argued that selecting a steeper ramp would solve the problem of almost parallel lines. However this is not so simple in practice—the 'optimum' ramp is not known until after the test.

Our findings are different to those of Lamarra *et al* (1987), who found breath-by-breath noise from their models (which did not have a time series component) approximated an uncorrelated Gaussian stochastic process, with a standard deviation that was largely independent of metabolic rate. The relative rates of both pulmonary ventilation and blood flow affect gas exchange. The uptake of O_2 in the lungs is more dependent on pulmonary blood flow than ventilation. In contrast, CO_2 output is closely related to the level of pulmonary ventilation (West 1969). Therefore ventilatory fluctuations between breaths would be expected to result in greater breath-by-breath variability in the measured $\dot{V}CO_2$ than $\dot{V}O_2$. Indeed, we found this to be the case and for two of our subjects, we found that the variability of the residuals from a two-line model increased with time. The models we have developed allow for this heteroscedasticity. For all three subjects, these residuals exhibited autocorrelation. This has also been incorporated in our models. In conclusion, we argue that, by considering more complex models than those traditionally used, the GET may be estimated with greater accuracy. In addition, they are suited to simpler gas measurement systems and thus have more widespread applicability than traditional methods.

Acknowledgments

We wish to thank R Gentleman and R Ihaka, and the R core group for developing the R software which is freely available under the GNU license and with which the data were analysed. The models were fitted with the function `elliptic` and the hidden Markov chain model with the function `cphidden` from the second author's R public libraries, respectively called `growth` and `repeated`, available at `http://www.luc.ac.be/~jlindsey/rcode.html`.

Appendix

Multivariate power-exponential distribution

The multivariate power-exponential distribution is given by

$$f(\mathbf{y}; \boldsymbol{\mu}, \Sigma, \kappa) = \frac{n\Gamma(n/2)}{\pi^{n/2}\sqrt{|\Sigma|}\Gamma(1+n/2\kappa)2^{1+n/2\kappa}} \exp\{-[(\mathbf{y}-\boldsymbol{\mu})^T\Sigma^{-1}(\mathbf{y}-\boldsymbol{\mu})]^\kappa/2\}$$

where the mean and covariance matrix are

$$\mathrm{E}(\mathbf{Y}) = \boldsymbol{\mu},$$
$$\mathrm{cov}(\mathbf{Y}) = \frac{2^{1/\kappa}\Gamma((n+2)/2\kappa)}{n\Gamma(n/2\kappa)} \times \Sigma.$$

This distribution includes the multivariate normal ($\kappa = 1$), the multivariate Laplace ($\kappa = 0.5$) and the multivariate uniform ($\kappa = \infty$) distributions as special cases.

Asymmetric multivariate laplace distribution

The multivariate power-exponential distribution contains, as a special case, one possible definition of a multivariate Laplace distribution. Kotz *et al* (2001) give another possibility:

$$f(\mathbf{y}; \Sigma, \boldsymbol{\kappa}) = \frac{2\mathrm{e}^{\mathbf{y}^\mathrm{T}\Sigma^{-1}\boldsymbol{\kappa}}}{(2\pi)^{n/2}|\Sigma|^{1/2}} \left(\frac{\mathbf{y}^\mathrm{T}\Sigma^{-1}\mathbf{y}}{2+\boldsymbol{\kappa}^\mathrm{T}\Sigma^{-1}\boldsymbol{\kappa}}\right)^{(2-n)/4} K_{(2-n)/2}\left(\sqrt{(2+\boldsymbol{\kappa}^\mathrm{T}\Sigma^{-1}\boldsymbol{\kappa})\mathbf{y}^\mathrm{T}\Sigma^{-1}\mathbf{y}}\right)$$

where $K_u(\cdot)$ is the modified Bessel function of the third kind and $\boldsymbol{\kappa}$ is the vector of skew parameters, one for each observation. For $\boldsymbol{\kappa} = \mathbf{0}$, this is a symmetric multivariate Laplace distribution which is elliptically-contoured; otherwise, it is not a member of that family. Below, we set all values of the vector $\boldsymbol{\kappa}$ to be identical so that only one new parameter is introduced; this implies that the distribution has the same skew over all time points.

Here, the mean and covariance matrix are given by

$$E(\mathbf{Y}) = \boldsymbol{\kappa}$$
$$\mathrm{cov}(\mathbf{Y}) = \Sigma + \boldsymbol{\kappa}\boldsymbol{\kappa}^\mathrm{T}.$$

As with the Student-*t* and power-exponential distributions, zero correlation does not imply independence.

References

Akaike H 1973 Information theory and an extension of the maximum likelihood principle *Second Int. Symp. on Inference Theory* ed B N Petrov and F Cs'aki (Budapest: Akademiai Kiado) pp 267–81

Beaver K L, Wasserman K and Whipp B J 1986 A new method for detecting anaerobic threshold by gas exchange *J. Appl. Physiol.* **60** 2020–7

Carlstein E, Műller H-G and Siegmund D (ed) 1994 *Change-Point Problems (Lecture Notes* vol 23*)* (Hayward, CA: Institute of Mathematical Statistics)

Davies R B 1987 Hypothesis testing when a nuisance parameter is present only under the alternative *Biometrika* **74** 33–43

Gómez E, Gómez-Villegas M A and Marin J M 1998 A multivariate generalization of the power exponential family of distributions *Commun. Stat.* A **27** 589–600

Hall P and Titterington D 1994 Edge-preserving and peak-preserving smoothing *Technometrics* **34** 429–40

Hinkley D 1970 Inference about the change-point in a sequence of random variables *Biometrika* **57** 1–17

Ihaka R and Gentleman R 1996 R: a language for data analysis and graphics *J. Comput. Graph. Stat.* **5** 299–314

Kelly G E, Thin A, Daly L and McLoughlin P 2001 Estimation of the gas exchange threshold in humans: a time series approach *Eur. J. Appl. Physiol.* **85** 586–92

Kelly G E, Thin A, Daly L and McLoughlin P 2002 Erratum: Estimation of the gas exchange threshold in humans: a time series approach *Eur. J. Appl. Physiol.* **87** 588

Koike A, Weiler-Ravell D, McKenzie D K, Zanconato S and Wasserman K 1990 Evidence that the metabolic acidosis threshold is the anaerobic threshold *J. Appl. Physiol.* **68** 2521–6

Kotz S, Kozubowski T J and Podgorski K 2001 *The Laplace distribution and generalizations. A Revisit with Applications to Communications, Economics, Engineering and Finance* (Basel: Birkhauser) chapter 6

Krisnaiah P K and Miao B Q 1998 Review about estimation of change points *Handbook of Statistics* vol 7 (Amsterdam: North-Holland) pp 375–402

Lamarra N, Whipp B, Ward S and Wasserman K 1987 Effect of interbreath fluctuations on characterizing exercise gas exchange kinetics *J. Appl. Physiol.* **62** 2003–12

Lindsey J K 1999 Multivariate elliptically contoured distributions for repeated measurements *Biometrics* **55** 1277–80

Loader C 1996 Change point estimation using nonparametric regression *Ann. Stat.* **24** 1667–78

MacDonald I L and Zucchini W 1997 *Hidden Markov and Other Models for Discrete-Valued Time Series* (London: Chapman and Hall)

Magalang U G and Grant B J B 1995 Determination of gas exchange threshold by nonparametric regression *Am. J. Respir. Crit. Care Med.* **151** 98–106

Sherrill D L, Anderson S J and Swanson G 1990 Using smoothing splines for detecting ventilatory thresholds *Med. Sci. Sports Exerc.* **22** 684–9

Sue D Y, Wasserman K, Moricca R B and Casaburi R 1988 Metabolic acidosis during exercise in patients with chronic obstructive pulmonary disease. Use of the V-slope method for anaerobic threshold determination *Chest* **94** 931–8

West J B 1969 Ventilation-perfusion inequality and overall gas exchange in computer models of the lung *Respir. Physiol.* **7** 88–110