33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

Subspace Detection of the Impulse Response Function from Intra-Partum Cardiotocography

Philip A. Warrick[1], *Member, IEEE*, Emily F. Hamilton[1,2]

***Abstract*— Recording of maternal uterine pressure (UP) and fetal heart rate (FHR) during labour and delivery is a procedure referred to as cardiotocography (CTG). We model this as an input-output system to estimate its dynamics in terms of an impulse response function (IRF). CTG data is very noisy and missing data are common. In this paper, we identify the models using subspace methods, which incorporate noise-suppression and permit the use of non-contiguous data. Using contiguous data, the subspace method performed better than linear regression; more of the 57 CTG pathological records in our database were modelled (30 vs. 26). Allowing non-contiguous data, even more pathological records were modelled with this approach (49). Furthermore, the models were discriminating; compared to linear regression, the IRF gain showed statistically significant differences more often between normal and pathological records (in 15/18 vs. 10/18 epochs) over the final three hours of labour.**

I. INTRODUCTION

Labour and delivery is routinely monitored electronically with sensors that measure maternal uterine pressure (UP) and fetal heart rate (FHR), a procedure referred to as cardiotocography (CTG). Temporary decreases in FHR are known as decelerations and reflect events such as compression of the umbilical cord by uterine contractions, malfunction of the fetal heart muscle, or premature separation of the placenta. Generally, larger insults are indicated by recurring episodes of deep, long decelerations whose onsets occur late with respect to the uterine contractions. We extract information from UP-FHR by treating the pair as an input-output system using system identification to estimate system dynamics in terms of an impulse response function (IRF).

A significant challenge to such modelling is that CTG often contains intervals of missing data on the UP or FHR signals. Signals with non-contiguous intervals were ignored in previous studies using non-parametric linear regression, which is inapplicable to such data [1], [2]. The main objective of this study was to extract more information from the measured signals by processing this imperfect data.

We achieved these goals by applying subspace methods to this system identification problem. Modelling of contiguous data this way was more successful than with linear regression; when applied to non-contiguous data, both the epochs modelled and the CTG recordings having models increased further. Using a database of CTG recordings labelled by outcome data available after birth, we compared the models over the final three hours of labour and with respect to their outcome classes (either normal or pathological) and found that the IRF gain was more consistently discriminating, compared to models from linear regression.

These findings have direct clinical significance. Increasing the number and amount of recordings amenable to analysis with a discriminating technique has the potential to warn clinicians on more cases and earlier on when intervention could potentially avert the pathological outcome.

Corresponding author: Philip Warrick (e-mail: philip.warrick@perigen.com). The authors acknowledge the financial support of this work by Perigen Canada.

[1]Perigen Canada, Montreal, Quebec, Canada. [2]Department of Obstetrics and Gynecology, McGill University, Montreal, Quebec, Canada

II. METHODS

A. Data

We used a database consisting of 111 intrapartum CTG recordings for pregnancies having a birth gestational age >36 weeks and having no known genetic malformations and having a minimum of 3 hrs. of recording.

Each recording was labelled by outcome according to its arterial umbilical-cord base deficit and neonatal indications of neurological impairment. An elevated base deficit (BD) measurement is an important indicator for metabolic acidosis of sufficient degree to cause neurological injury [3], [4]. The recordings consisted of approximately equal numbers of cases labelled normal (54 cases: BD < 8 mmol/L) and severely pathological (57 cases: BD ≥ 12 mmol/L, accompanied by either death or evidence of hypoxic ischemic encephalopathy). This proportion of pathological cases is much higher than their natural incidence.

The CTG was acquired from fetal monitors used clinically. The FHR was reported at a uniform sampling rate of 4 Hz while the UP was acquired at 1 Hz and up-sampled to 4 Hz. In the majority of cases, the UP or FHR sensors were external (attached to the maternal abdomen); the UP was acquired by tocography and the FHR was acquired from an ultrasound probe. Otherwise they were acquired internally via an intrauterine (IU) probe and/or a fetal scalp electrode.

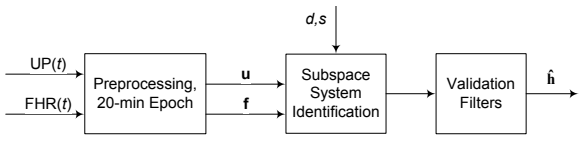

Fig. 1. Overall block diagram. Preprocessing cleans and segments the UP and FHR into epochs of input **u** and output **f**. Subspace system identification computes an IRF and estimates the best delay d and scale s. The resulting models $\hat{\mathbf{h}}$ must pass the validation filters to be considered valid.

978-1-4244-4122-8/11/$26.00 ©2011 IEEE

B. Overall processing

We modelled UP-FHR system dynamics by linear system identification, as shown schematically in Fig. 1. A preprocessing step cleaned and segmented the UP and FHR into 20 min epochs of input and output data (**u** and **f**). Next, using subspace system-identification methods, we estimated the IRF $\hat{\mathbf{h}}$ and determined the best values for the IRF delay d and scale s. Finally, the models were validated using a continuous second-order parametric fit of the IRF and surrogate testing.

C. Preprocessing

The FHR and UP signals may be temporarily interrupted by loss of sensor contact. Because of these gaps in the signal, detrending by standard FIR filtering was not used. Instead for each 20 minute epoch, the very low frequency content of each signal (offset and long-term trend) was estimated and subtracted using a low-order (p=3) Chebyshev polynomial approximation. Epochs were advanced by 10 min increments.

D. Linear model

Let the input, UP, and output, FHR, at time sample k ($k = 1 \ldots N$) be denoted by $u(k)$ and $f(k)$, respectively. The linear response $f(k)$ of a discrete-time system to an arbitrary input signal $u(k)$ is given by the convolution sum:

$$f(k) = \sum_{i=d}^{d+M-1} (h_i \Delta t) u(k-i) = \mathbf{h} * \mathbf{u}(k) \qquad (1)$$

where Δt is the sampling period, and **h** is the IRF beginning at delay sample d, and of length M. $\mathbf{u}(k)$ is the length-M vector of input samples $[u_{k-d-M+1} \ \ldots \ u_{k-d-1} \ u_{k-d}]$ used to compute $f(k)$ at sample k. For causal (physically realizable) systems, $d \geq 0$, but under certain conditions, such as input measurement delay, d may be negative [5].

E. Subspace method

A previous study directly estimated **h** above using linear regression and the pseudo-inverse of the input autocorrelation matrix [1]. This approach assumed contiguous input and output data (i.e., with gaps no greater than $15s$ in length). However, we wanted to relax this requirement in order to permit processing more of the data, which often included temporary gaps of longer duration.

Subspace methods are well-suited to this problem because they permit such non-contiguous data to be included in the estimation. In addition, they are very applicable to noisy data, such as CTG, because of their general noise model and inherent use of singular value decomposition (SVD) within a non-iterative, regression-based estimation. As well, by incorporating an estimate of initial state, all epoch data are used in subspace model estimates; data need not be discarded due to initial filter-length effects. Finally, they require very few tuning parameters apart from a scaling factor s.

The *PO-MOESP* subspace method [6] is based on a state-space input-output model that incorporates process and measurement noise, written in innovation form as:

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + Ke(k) \\ f(k) &= Cx(k) + Du(k) + e(k) \end{aligned} \qquad (2)$$

where $x(k)$ is the state, the innovation $e(k)$ is a white-noise sequence and K is the Kalman gain. Whereas direct use of Kalman filter methods require matrices A, B, C and D to be specified, these parameters and the initial state $x(0)$ are all estimated by subspace methods. We used the LTI-Toolbox [7] implementation of *PO-MOESP* for our single-input, single-output (SISO) state-space model estimates.

This avoids direct estimation of **h** and the need for length-M contiguous data as in (1), but **h** is available indirectly as

$$h(k) = \begin{cases} 0 & k < 0 \\ D & k = 0 \\ CA^{k-1}B & k > 0 \end{cases} \qquad (3)$$

We found that using no direct feedthrough term (i.e., constraining D=0) produced the most consistent models with the best fidelity on the output prediction, as measured by the minimum description length (MDL).

The order r (i.e., the state-vector dimensionality) of the state-space model was selected by inspecting the most significant eigenvalues of the system matrix A. By experimentation, we found that regardless of the value of the maximum order s we provided, a system order $r = 2$ generally produced models with the best fidelity. This finding was consistent with a previous study that fitted a second-order continuous low-pass system to the IRF [8].

By varying s over selected values 2^j, $j = 3, 4, \ldots 7$, we produced multi-resolutional models that adapted to the time constants of the system. If model estimation at $s = 2^j$ failed, we continued the procedure with $s = 2^{j+1}$, up to j_{max}.

F. Delay detection

As described in [1], input measurement delay associated with the UP sensor may result in a negative IRF delay. In contrast, the physiological response is expected to have a positive delay. The combination of these two delays can produce an FHR response that occurs before (negative d) or after (positive d) the measured UP contraction onset. Therefore, we developed an algorithm to determine the best d for each epoch. We set the bounds on d to ± 80 s.

We searched for delays corresponding to the zero-crossings of a Chebyshev approximation of the first IRF lag $h_1(d)$. Due to the periodicity of $h_1(d)$ there were often multiple such zero-crossings. We ranked these candidates according to their MDL and used them to smooth the epoch-to-epoch delay estimates using a median filter and nearest-neighbour approach, as described in [1].

G. Fitting a low-pass second-order parametric model

To reduce the parametrization even further, we fitted the resulting impulse-response function to a continuous second-order low-pass system, as described in [8]:

$$\hat{\mathbf{h}}_p = f(G, \zeta, \omega_n, d), \qquad d = \{d_{min} \ldots d_{max}\} \qquad (4)$$

where ζ = damping factor, ω_n = undamped natural frequency, G = gain and d = delay. We accepted models at this stage (Fit2) if the VAF of the IRF fit was > 90%.

	Normal	Pathological
UP continuity	94.9%	84.8%
FHR continuity	89.7%	78.9%

TABLE I

CONTINUITY LEVELS FOR THE INPUT UP AND OUTPUT FHR SIGNALS

H. Validating models using surrogate data

In order to be confident that the resulting model captured system dynamics rather than noise, we also estimated models on a set of surrogate FHR signals. In [1] we used the amplitude adjusted Fourier transform (AAFT) method to generate surrogates from contiguous data. For the non-contiguous data of this study, however, we used the more general approach of simulated annealing.

In simulated annealing [9], samples are randomly swapped and a cost function E determines whether the interchange is accepted. E was set to the sum of squared differences between the auto-correlation functions of the original FHR and the surrogates. If the surrogate FHR more closely matched the original ($\Delta E < 0$), the swap was accepted. Otherwise it was accepted with probability $p = e^{-\Delta E/T}$. The samples were first fully randomized and T was set to a high value. In subsequent iterations, T was gradually lowered so that p reduced towards zero until E converged.

The gap locations in the surrogates were fixed to those in the original; gaps were excluded from swapping but included in the auto-correlation function on the assumption that they had equal influence on the cost function. In order to allow more variability in the surrogates, rather than fixing the gap data to a constant as suggested in [9], we randomized gap values over the range of the signal just before and after the gap. To speed up calculations, we limited the length of the autocorrelation function to 256 samples (a conservative upper bound on the first zero crossing, as determined by preliminary experimentation) and used the approach of [10] to re-calculate only those terms of the auto-correlation that were changed by the swap.

Despite these optimizations, the processing load was high so we limited generation to 3 surrogate FHRs and accepted the model if the fidelity of the model with the original FHR was greater than all of the models generated from the surrogate data. This gave a confidence level of 75% that a model capturing spurious dynamics would be rejected.

III. RESULTS

A. Continuity levels

Table I shows the overall average levels of UP and FHR continuity in our CTG database. While both UP and FHR are affected by sensor contact problems, continuity was lower for FHR, likely because Doppler FHR detection is prone to maternal heart rate interference. Continuity was lower in pathological compared to normal records, consistent with the clinical reality that these cases undergo more intervention.

	Regression	Subspace	Subspace
Continuous data	yes	yes	no
Validation method	AAFT	AAFT	Fit2/Annealing
Epochs with models	455	464	874
Epochs validated	188 (41.3%)	231 (49.8%)	392 (44.9%)
Records with models	43	43	56
Records validated	26 (60.5%)	30 (70.0%)	49 (86.0%)

TABLE II

COMPARISON OF OVERALL PROCESSING USING LINEAR REGRESSION AND SUBSPACE ALGORITHMS FOR SYSTEM IDENTIFICATION (PATHOLOGICAL CASES)

B. Processing summary

The results of the processing shown in Table II are limited to pathological records because better processing of these cases have more potential to improve detection of pathology.

We first processed contiguous data with the subspace approach, using AAFT surrogates for validation, and compared the overall results to those obtained using linear regression [1]. For the 43 records having models in at least 3 epochs, 30/43=70.0% of these records had models that passed the validation, compared to 26/43=60.5% records with the linear regression approach. The subspace approach produced models for a slightly higher number of epochs overall, but a greater proportion of these passed the validation step (231/464=49.8% vs. 188/455=41.3%). These results indicate that for identical signals, the subspace approach created useful models more often than the previous approach, permitting 4 more pathological cases to be considered.

Next we included non-contiguous data from the records and applied the subspace approach using simulated annealing surrogates for validation. Compared to the subspace results with contiguous data, 13 more records had a least three models (56 vs. 43). Of these records, a higher number (49) had models that passed the validation step and the percent of records validated increased to 49/56=87.5%. The overall number of epochs with models increased to 874, indicating that modelling had occurred in the non-contiguous data. With the reduced signal quality presented to the subspace algorithm and the use of the simulated annealing validation, the proportion of validated models also decreased slightly to 392/874=44.9%. Results from normal records were similar: 407 epochs had validated models and 52/54 records had at least 3 epochs with validated models.

C. IRFs of typical epochs

Fig. 2 shows the preprocessed UP and FHR, the IRF, and the model-predicted FHR for an epoch from a pathological case with non-contiguous signal using original and surrogate FHR. With the original FHR, the IRF delay d was 11.0 s and the gain G was -0.651 bpm/mmHg. The variance accounted for (VAF) [1] of this model was 59.7%. The VAF of the model using surrogate data was lower (37.2%), as were models generated from the other surrogate FHR, so this model passed validation. The predicted FHR using interpolated contiguous UP showed a predicted (and plausible) deceleration in missing FHR.

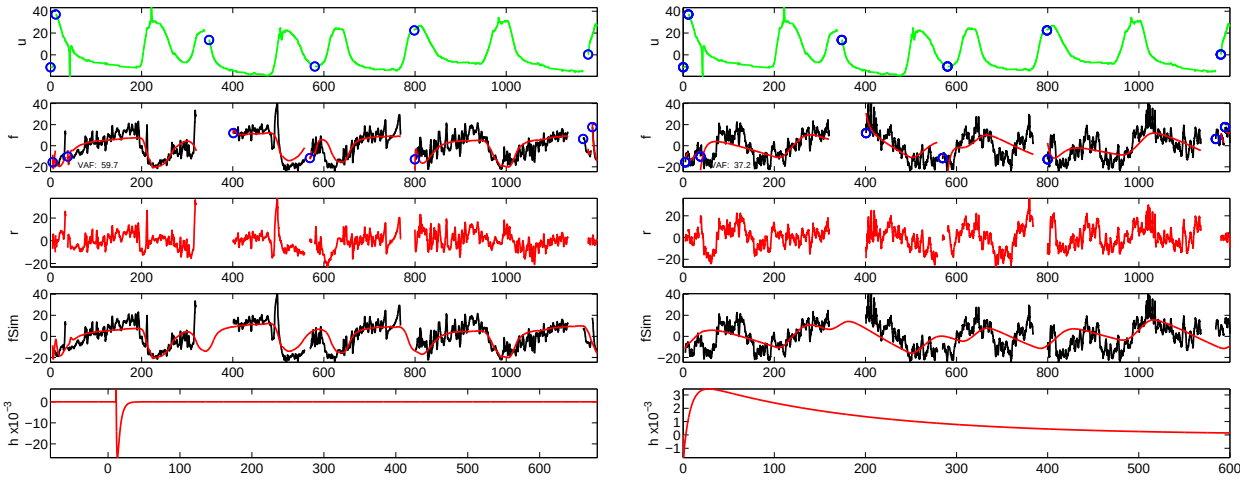

(a) Measured FHR (b) Surrogate FHR

Fig. 2. a) Example UP (green) and FHR with numerous discontinuities (blue circles) for a pathological case. The predicted FHR is shown in red and the (preprocessed) measured FHR in black. The residual error r is shown in the third panel. The fourth panel shows the predicted FHR using interpolated contiguous UP showing a predicted deceleration in missing FHR. The IRF model is shown in the fifth panel. The model parameters were: VAF=59.7%, delay d=11.0 s, and gain G=-0.651 bpm/mmHg. b) The same epoch with surrogate FHR. A model was estimated, but it had much lower VAF (37.2%). The units of FHR and UP are beats per minute (bpm) and mmHg, respectively. All horizontal time axes have units of seconds.

D. IRF parameters by class

Figure 3 shows class average results for gain G. Overall both classes had trends towards larger negative G over time. Pathological cases had larger negative G. There were statistically significant class differences in 15/18 epochs (blue asterisks). We used the Kolmogorov-Smirnov distribution test, rejecting the null hypothesis at $p < 0.05$. These differences were more consistent than with linear regression, where 10/18 epochs were statistically significant [2]. Delay d and the damping factor ζ also had differences (not shown), but only in the last few epochs (d) or intermittently (ζ).

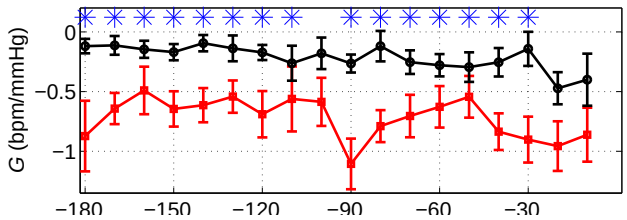

(a)

Fig. 3. IRF gain G over the last 180 min of labour for normal (black) and pathological (red) cases. The blue asterisks indicate statistically significant differences ($p < 0.05$) between the two classes for that epoch.

IV. DISCUSSION

From the previous work [2], we expect to see more discrimination from the IRF delay parameter. In future studies we will prefer models with higher values of the scale s because they capture longer (and clinically more significant) FHR responses. This should reduce the estimate variance that may have obscured class differences in this parameter. It is also possible that with successfully modelling of many more pathological records than before, the delays of these cases fall into more than one sub-population of pathology.

These results demonstrate that subspace system identification extracted more CTG information, allowing analysis of more recordings, especially of the important pathological cases. Since these recordings tend to be noisier and have lower continuity than normal CTG (see Table I) the subspace approach is very suitable. As well, given the very promising discrimination results of the IRF gain, in future studies we expect to use model parameters in the detection of fetal pathology (hypoxia) during labour and delivery.

REFERENCES

[1] P. A. Warrick, E. F. Hamilton, D. Precup, and R. E. Kearney, "Identification of the dynamic relationship between intra-partum uterine pressure and fetal heart rate for normal and hypoxic fetuses," *IEEE Transactions on Biomedical Engineering*, vol. 56, no. 6, pp. 1587–1597, June 2009.

[2] P. Warrick, E. Hamilton, D. Precup, and R. Kearney, "Classification of normal and hypoxic fetuses from systems modeling of intrapartum cardiotocography," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 4, pp. 771–779, 2010.

[3] J. Low, R. Victory, and E. Derrick, "Predictive value of electronic fetal monitoring for intrapartum fetal asphyxia with metabolic acidosis," *Obstet Gynecol*, vol. 93, pp. 285–291, 1999.

[4] J. T. Parer, T. King, S. Flanders, M. Fox, and S. J. Kilpatrick, "Fetal acidemia and electronic fetal heart rate patterns: Is there evidence of an association?" *Journal of Maternal-Fetal & Neonatal Medicine*, vol. 19, no. 5, pp. 289–294, May 2006.

[5] I. W. Hunter and R. E. Kearney, "Two-sided linear filter identification," *Med Biol Eng Comput*, vol. 21, pp. 203–209, 1983.

[6] M. Verhaegen and V. Verdult, *Filtering and system identification: a least squares approach*. Cambridge, UK.: CUP, 2007.

[7] M. Verhaegen, V. Verdult, and N. Bergboer. (2007) The LTI-Toolbox. Delft University of Technology. [Online]. Available: http://www.dcsc.tudelft.nl/ datadriven/lti/ltitoolbox_product_page.html

[8] P. A. Warrick, R. E. Kearney, D. Precup, and E. F. Hamilton, "Low-order parametric system identification for intrapartum uterine pressure-fetal heart rate interaction," in *The 2007 IEEE Engineering in Medicine and Biology 29th Annual Conference*, 2007, pp. 5043–5046.

[9] T. Schreiber and A. Schmitz, "Surrogate time series," *Physica D: Nonlinear Phenomena*, vol. 142, no. 3-4, pp. 346–382, Aug. 2000.

[10] I. W. Hunter and R. E. Kearney, "Generation of random sequences with jointly specified probability density and autocorrelation functions," *Biological Cybernetics*, vol. 47, no. 2, pp. 141–146, June 1983.