

Real-time profiling of respiratory motion: baseline drift, frequency variation and fundamental pattern change

D Ruan¹, J A Fessler², J M Balter³ and P J Keall¹

¹ Department of Radiation Oncology, Stanford University, Stanford, CA 94305-5304, USA

² Department of Electrical Engineering & Computer Science, The University of Michigan, Ann Arbor, MI 48109-2122, USA

³ Department of Radiation Oncology, The University of Michigan, Ann Arbor, MI 48109-2122, USA

E-mail: druan@stanford.edu

Received 12 February 2009, in final form 22 June 2009

Published 22 July 2009

Online at stacks.iop.org/PMB/54/4777

Abstract

To precisely ablate tumor in radiation therapy, it is important to locate the tumor position in real time during treatment. However, respiration-induced tumor motions are difficult to track. They are semi-periodic and exhibit variations in baseline, frequency and fundamental pattern (oscillatory amplitude and shape). In this study, we try to decompose the above-mentioned components from discrete observations in real time. Baseline drift, frequency (equivalently phase) variation and fundamental pattern change characterize different aspects of respiratory motion and have distinctive clinical indications. Furthermore, smoothness is a valid assumption for each one of these components in their own spaces, and facilitates effective extrapolation for the purpose of estimation and prediction. We call this process ‘profiling’ to reflect the integration of information extraction, decomposition, processing and recovery. The proposed method has three major ingredients: (1) real-time baseline and phase estimation based on elliptical shape tracking in augmented state space and Poincaré sectioning principle; (2) estimation of the fundamental pattern by unwarping the observation with phase estimate from the previous step; (3) filtering of individual components and assembly in the original temporal-displacement signal space. We tested the proposed method with both simulated and clinical data. For the purpose of prediction, the results are comparable to what one would expect from a human operator. The proposed approach is fully unsupervised and data driven, making it ideal for applications requiring economy, efficiency and flexibility.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Current radiotherapy techniques enable highly targeted dose delivery. However, internal organ motion, particularly that caused by respiration, makes it challenging to track or predict the location of the target volume. Multiple studies have shown that target motion results in observable degradation in the treatment effectiveness and accuracy (Engelsman *et al* 2001, Ozhasoglu and Murphy 2002, Gierga *et al* 2004, Jiang *et al* 2003).

Treatment margins are often used to account for the effect of target motion (ICRU 1993, 1999, Langen and Jones 2001). To further reduce margin in order to spare normal tissue radiation, people have applied passive or active breath holding (Wong *et al* 1999, Mah *et al* 2000, Rosenzweig *et al* 2000, Nelson *et al* 2005, Venkat *et al* 2008). As an alternative, respiratory gating has been applied to reduce the tumor motion during beam-on time by limiting radiation exposure to only a portion of the breathing cycle (Kubo *et al* 2000, Berbeco *et al* 2005, Cui *et al* 2008). The effectiveness of margin reduction by breath hold or gating has been studied (Engelsman *et al* 2005). Dynamically tracking and compensating for the target motion have the potential to improve both delivery efficiency and accuracy.

For both gating and dynamic compensation, it is crucial to precisely locate the target position in real time. Observation noise and system delay require estimation schemes that are robust and capable of prediction. Linear prediction models and its many variants (Sharp *et al* 2004, Wu *et al* 2004) often have difficulty in incorporating the semi-periodic structure. On the other hand, a purely periodic model (Lujan *et al* 2003, Neicu *et al* 2003) does not sufficiently account for variations: deviation from periodicity is considered ‘abnormal’ or ‘irregular’ (Ruan *et al* 2006).

In this study, we take a holistic perspective, and consider the observed breathing trajectory as a periodic fundamental pattern that is frequency modulated by a phase warping function and displacement modulated by a baseline drift function. The observation is further corrupted by independent additive noise. Note that it is reasonable to assume the phase warping functions and the baseline drifts to be piecewise smooth: the smoothness in the phase warping function corresponds to the semi-periodic structure. For tracking and prediction, we propose the following procedure: first, the variations in baseline, frequency and fundamental pattern are decoupled: second, each component is filtered to achieve denoising or prediction, based on a smoothness assumption; finally, the components are assembled back into the original signal space to obtain the ultimate tracking or prediction results.

The decomposition module produces real-time estimates for baseline drift, frequency variation and fundamental pattern, which are of clinical significance in their own aspects. The baseline drift estimate serves as an indicator for abrupt patient movement, and may be used to trigger interruption of treatment and/or adjustment of patient placement. For treatment with multileaf collimator (MLC), the physical configuration of the leaves limits the system’s tracking ability perpendicular to the leaf track, and it may be more reasonable to track the baseline position along such a direction. Large abrupt changes in breathing frequency typically indicate abnormality and require clinical attention. The fundamental pattern is of its own merit. For lung tumors, radiotherapy planning is often performed based on a pre-treatment 4D CT scan (Keall *et al* 2006, D’Souza *et al* 2007, Colgan *et al* 2008), and it is implicitly assumed that the target motion during the treatment closely resembles that observed in the planning data set. Keeping track of the fundamental pattern provides a way to validate this assumption in real time: when the patient’s breathing pattern deviates significantly from the planning trace, it is necessary to adjust the treatment delivery plan accordingly, or even replan in some cases.

We present the basic theory of the proposed methodology in section 2. Section 3 describes the test data. Section 4 provides experimental results and discusses their implications. Section 5 summarizes the study and proposes future work.




Figure 1. Illustration of the decomposition process.

2. Methods

We model the variation in respiratory motion as the composite result of four sources: baseline drift, frequency variation, fundamental pattern change and additive random observation noise. Each of these factors characterizes a different aspect of respiratory motion and has distinct physiological interpretation and clinical consequences. This work aims to develop an online profiling system that provides real-time estimating of baseline drift, frequency variation and pattern change.

This section lays the theoretical foundation of this study. For simplicity, we present the derivation for the case of scalar observations, but all analyses generalize to high-dimensional situations naturally. The different types of variations in breathing are reflected in the discrete observations according to the generative model in (1): a fundamental pattern that is modulated both in frequency and in amplitude, and corrupted by additive noise before discrete samples are obtained.

Mathematically, the generative model for the position observation $s(t)$ is given by

$$s_i = [f(\phi(t)) + b(t) + w(t)]|_{t=t_i}, \quad i = 1, 2, \dots, \quad (1)$$

where

- f : 2π -periodic fundamental pattern.
- ϕ : the phase warping function that maps the time tag to the corresponding phase value. As a warping function, it is monotonically nonincreasing.
- b : temporally piecewise slow varying baseline drift function.
- w : independent zero mean additive noise.

For profiling, we need to characterize the instantaneous breathing pattern by estimating (f, ϕ, b) simultaneously. This is an ill-posed problem. In Ruan *et al* (2008a), we have introduced a real-time approach to systematically estimate the baseline drift b from the observations. Yet the decomposition of fundamental pattern f and warping map ϕ still remains to be solved and is particularly challenging. Intuitively, given the phase warp ϕ , an estimate for the fundamental pattern f can be obtained by first ‘inverse-warping’ the observations with ϕ^{-1} , and then performing minimum mean-squared error (MMSE) estimation on f , based on the periodicity assumption, as illustrated in figure 1.

Now we focus on acquiring a continuous real-time phase estimate from noisy discrete observations. The two most prominent candidates are the Hilbert phase (Gabor 1946) and the linear interpolating phase. The Hilbert phase, though theoretically interesting, suffers from sensitivity to short-lived small fluctuations (Vainstein and Vakman 1983), thus unsuitable for

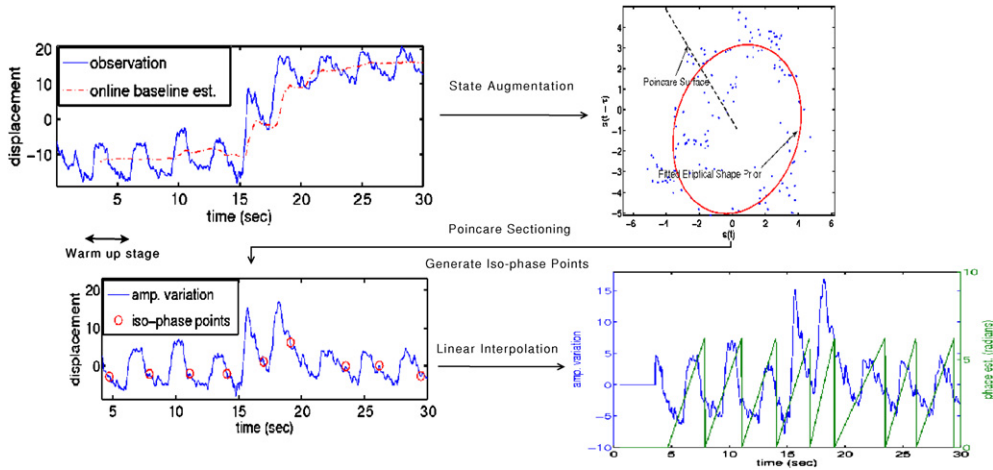


Figure 2. An illustration of the online phase estimation procedure: augment the observations \Rightarrow ellipse fitting \Rightarrow Poincaré sectioning \Rightarrow map Poincaré intersections back to time–displacement relation as iso-phase points \Rightarrow linear interpolate/extrapolate to generate continuous phase.

real-time applications. We will briefly summarize the linear phase method and its applicability in respiratory motion analysis hereafter.

Given marker events that identify the completion of a cycle (e.g., the R-peaks of an electrocardiogram), the linear phase is defined via interpolation between the neighboring marker events:

$$\phi^L(t) = 2\pi \frac{t - t_n}{t_{n+1} - t_n} + 2\pi n \quad (t_n \leq t < t_{n+1}), \quad (2)$$

where t_n represents the time tag for the n th marker event. In addition to its simplicity, an advantage of the linear interpolating phase is that it is locally defined, so perturbations and errors do not propagate over time. The major challenge in applying the linear phase method to a breathing trajectory is to efficiently identify the marker events. Historically, Neicu *et al* (2003) identified the peaks and troughs from a low-pass-filtered breathing trajectory, and subsequently obtained average tumor trajectory (ATT). There exists, however, an intrinsic dilemma with such an approach: the identification of extreme values is very sensitive to noise, due to diminished values of local gradients; while noise suppression via low-pass filtering introduces uncertainty to the detected extreme locations.

For the purpose of warping, it suffices to identify iso-phase points. We propose to combine an elliptical shape prior in an augmented state space and Poincaré sectioning idea to automatically produce the iso-phase events in real time. Figure 2 illustrates the basic procedure for the online phase estimation.

2.1. State augmentation and an elliptical shape prior

As in Ruan *et al* (2007, 2008b), we first form a state vector using the current s_i and a past observation of a fixed temporal delay s_{i-k} to capture the first-order system dynamics. Assuming uniform sampling with interval Δ (s), the discrete time lag k corresponds to a delay of $\tau = k\Delta$ s. This augmentation naturally distinguishes different stages of hysteresis. Insufficient as a target position tracker per se, the elliptical shape prior provides enough

low-order information in characterizing the semi-periodicity intrinsic to the time–displacement observations (Ruan *et al* 2008a). We briefly summarize the procedure for estimating the ellipse parameters below. Further details about adaptivity and recursive learning can be found in Ruan *et al* (2008a).

Let (x, y) denote the coordinates of a point in the two-dimensional augmented state space, and define vector $z = (x^2 \ xy \ y^2 \ x \ y \ 1)$. The point (x, y) falls on the ellipse parameterized by $\mathbf{a} = (a \ b \ c \ d \ e \ f)$ if and only if it satisfies the following quadratic curve equation:

$$F(\mathbf{a}, \mathbf{z}) = \mathbf{a}^T \mathbf{z} = ax^2 + bxy + cy^2 + dx + ey + f = 0 \tag{3}$$

with negative discriminant, i.e., $b^2 - 4ac < 0$.

The optimal coefficient \mathbf{a} is the eigenvector that corresponds to the only positive eigenvalue for the generalized eigendecomposition of the pair (\mathbf{S}, \mathbf{C}) , where

$$\mathbf{S} \triangleq \sum_i \mathbf{z}_i \mathbf{z}_i^T, \quad \text{where } \mathbf{z}_i \text{ is generated with } (x = s_i, y = s_{i-k}),$$

and

$$\mathbf{C} = \begin{bmatrix} \tilde{\mathbf{C}} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{bmatrix}, \quad \text{with} \quad \tilde{\mathbf{C}} \triangleq \begin{bmatrix} 0 & 0 & 2 \\ 0 & -1 & 0 \\ 2 & 0 & 0 \end{bmatrix}.$$

The conversion from the quadratic parameterization with \mathbf{a} to the standard Euclidean representation of the ellipse is well known. For simplicity, we call the converted parameters $(x_0, y_0, r_1, r_2, \theta)$, where (x_0, y_0) is the estimated center of the ellipse, r_1 and r_2 are the lengths of the major and minor axes respectively, and θ is the angular direction of the major axis. In other words, any point on the ellipse satisfies

$$\left\{ \frac{\cos \theta (x - x_0) + \sin \theta (y - y_0)}{r_1} \right\}^2 + \left\{ \frac{-\sin \theta (x - x_0) + \cos \theta (y - y_0)}{r_2} \right\}^2 = 1. \tag{4}$$

The projection of the center of the ellipse onto the $s(t)$ axis, i.e., x_0 , provides a reasonably good estimate for baseline drift in real time.

2.2. Phase tracking by Poincaré sectioning

The bottleneck for applying the linear phase estimator is to identify marker events and picking peaks and troughs of the breathing pattern is difficult. It is sufficient to obtain marker events for a *definite but arbitrary* phase, which is the focus of our discussion hereafter.

To put the proposed Poincaré sectioning method into context, we first recall the well-known Poincaré–Bendixson theorem for planar dynamic systems.

Theorem 1 Bendixson (1901). *Given a differentiable real dynamic system defined on an open and simply connected subset of the plane, then every non-empty compact ω -limit set of an orbit, which contains no fixed points, is a periodic orbit.*

Based on this theorem, it is obvious that the trajectory in the augmented state space is a periodic orbit, as there exists no fixed point in the state evolution process. It is important to distinguish between the ‘periodic orbit’ in the state space, referred to as in the Poincaré–Bendixson theorem, and the more intuitive ‘semi-periodicity’ with respect to the time evolution. The former only concerns the pattern of the augmented samples in the state

space, and is oblivious of the exact time scale. To illustrate this fundamental difference, let us consider a trajectory of samples according to

$$\begin{aligned} x(t) &= \sin(2\pi t/T_n) & \text{for } \sum_{i=1}^{n-1} T_i \leq t < \sum_{i=1}^n T_i; \\ y(t) &= \cos(2\pi t/T_n) & \text{for } \sum_{i=1}^{n-1} T_i \leq t < \sum_{i=1}^n T_i, \end{aligned} \quad (5)$$

where $\{T_i\}$ are positive numbers that control the local speed of evolution. Note that the trajectory defined above is *NOT* periodic in time for nonidentical T_i 's. On the other hand, this produces a periodic orbit in the state space, which can be expressed as $x^2 + y^2 = 1$, and discrete samples from this stipulated generative process (5) fall on this orbit. Therefore, we expect noisy observations from semi-periodic signal to generate scatter plots about an orbit in the state space, oblivious to its own frequency variation. For the purpose of identifying the iso-phase events, we take advantage of the invisibility of pure frequency variation, and regard it as an implicit phase warping procedure. Then we could generate a Poincaré section by choosing a Poincaré plane⁴ and recording the intersections of such plane with the given trajectory. If a system has a natural period associated with it, then the Poincaré plane corresponds to a definite (but arbitrarily) chosen phase (Hirsch *et al* 2003). For breathing motion, there is often no definite natural period with respect to time, since the frequency is subject to small variation, yet there exists one in the state space we have defined, illustrated by the orbit. The intersecting points between the trajectory and the Poincaré plane are subsequently identified as iso-phase events.

The placement of the Poincaré surface is of high relevance for the usefulness of the extracted intersections. An optimal surface maximizes the number of intersections, or equivalently, minimizes the time interval between them. We have already argued that our trajectory is a periodic orbit according to the Poincaré–Bendixson theorem, so it suffices to find a plane that intersects each period exactly once. Fortunately, we have already constructed an elliptical shape prior which reasonably approximates the structure of the trajectory in the state space. This leads to a simple construction of the Poincaré surface: a ray starting from the center of the estimated ellipse center. To maintain consistency of such a choice, we need to fix the relative direction of the ray to the orientation of the ellipse θ in (4), which is constantly changing with new observations. Note that the requirement for a legitimate Poincaré section is much less stringent than that for peaks/trough: it only requires consistency, but not extrema. It is also feasible to use multiple Poincaré surfaces if desired.

Once we have identified the intersections between the Poincaré surface and the trajectory as iso-phase events, the phases before the last marker event can be obtained by linear interpolation (2). We use linear extrapolation to extend phase estimates for time tags that are beyond the last marker event as

$$\phi(t) = \begin{cases} 2\pi \frac{t - t_n}{t_{n+1} - t_n} + 2\pi n & \text{for } t_n \leq t < t_{n+1}; \\ 2\pi \frac{t - t_N}{t_N - t_{N-1}} + 2\pi N & \text{for } t_N < t. \end{cases} \quad (6)$$

⁴ The Poincaré plane is of codimension one, so it is indeed a line in our case. We stick with the general terminology of ‘plane’ or ‘hyperplane’.

2.3. Fundamental pattern estimation

Once we have estimated the baseline drift b and the phase ϕ , the estimation of fundamental pattern f is quite straightforward based on minimum mean-squared error (MMSE) estimation, with the following steps:

- Obtain a continuous time–displacement function from discrete samples $\{s_i\}$. With finite noisy observations s_i , we model the continuous trajectory \tilde{s} with linear combinations of synthesis functions:

$$\tilde{s}(t) = \sum_k c_k \eta(t - k\Delta).$$

With finite noisy number of observations s_i , it is advantageous to use a non-interpolant synthesis function η , rather than the infinitely supported sinc or cardinal splines. In our implementation, we adopt the third-order B-spline for the purpose of interpolation using recursive filtering (Unser *et al* 1993a, 1993b).

- Estimate the fundamental pattern f by inverse warping. Because f is a 2π periodic function, we have an equivalent representation by restricting it to a single period: $\tilde{f} = f|_{[0, 2\pi)}$, and the MMSE estimate is simply the average:

$$\tilde{f}(\theta) = \text{mean}\{\tilde{s}(t) : t = \phi^{-1}(2k\pi + \theta) \text{ for some } k \in \mathcal{Z}\}. \quad (7)$$

The above estimation is for a globally static fundamental pattern. In some cases, the oscillatory magnitude of the breathing is also temporally varying, and it may be desirable to apply adaptive learning. It is typical to discount previous observations with exponentially fading memory. Mathematically, a discount factor $\lambda \in (0, 1]$ is introduced, and the goal is to find the \tilde{f} such that

$$\tilde{f}^*(\theta) = \arg \min_{\tilde{f}(\theta)} \int_{\theta=0}^{2\pi} \sum_{i=1}^{K(\theta)} \lambda^{K(\theta)-i} \|s(\phi^{-1}(2\pi i + \theta)) - \tilde{f}(\theta)\|^2, \quad (8)$$

where $K(\theta)$ is the biggest integer such that $K(\theta)2\pi + \theta \leq \phi(t)$. The adaptive weight λ controls the discounting ratio between past samples and the more current ones. It reflects the tradeoff between system responsiveness and noise sensitivity, and thus is also called a ‘forgetting factor’. In particular, $\lambda = 0$ indicates a myopic estimation using only the most recent samples, and $\lambda = 1$ corresponds to weighting all the samples equally, ignoring their relative temporal distance to the current instant.

The solution to (8) is given by the weighted average of historical data corresponding to the same phase:

$$\tilde{f}_t^*(\theta) = \frac{1}{C_t(\theta)} \sum_{i=1}^{K(\theta)} \lambda^{K(\theta)-i} s(\phi^{-1}(2\pi i + \theta)), \quad (9)$$

where $C_t(\theta) = \sum_{i=1}^{K(\theta)} \lambda^{K(\theta)-i}$ is the normalization factor. Note that the $\lambda = 1$ case corresponds to real-time averaging, which at the end of the horizon, reduces to (7).

3. Material

For proof of the principle, we simulated three sets of data so that we could have access to ‘ground-truth’ for verification purposes. In the first test, we simulated noise-free, strictly periodic data with both ideal sinusoidal and modified cosine models as described in Lujan

Table 1. RPM data set information.

ID vs Para.	1	2	3	4	5	6	7	8	9	10	11	12
STD	2.9	6.5	13.1	2.8	4.9	2.8	4.3	7.6	2.1	7.7	13.0	6.6
P-P	10.9	25.0	48.9	9.0	13.1	11.5	17.8	26.9	13.1	37.4	38.9	32.5
Period (s)	4.5	4.6	7.2	5.6	4.4	5.4	4.7	9.7	4.7	4.1	3.1	5.2

The data are adjusted to have globally zero mean; average periods were estimated with the subspace projection method (Ruan *et al* 2006).

et al (2003). The discrete sinusoidal and modified cosine waveforms were generated, respectively, according to

$$\begin{aligned} s_i^{\sin} &= s(i \Delta t) = s_0 + a \sin(\pi i \Delta t / T - \phi) \\ s_i^{\text{modified cos}} &= s(i \Delta t) = s_0 - a \cos^{2n}(\pi i \Delta t / T - \phi), \end{aligned} \quad (10)$$

where we used the typical value $n = 2$. The parameters in (10) were set to be period $T = 5$ s, sampling interval $\Delta t = 1/30$ s, corresponding to a sampling frequency of 30 Hz, baseline position $s_0 = 0$, magnitude $a = 5$ cm, and systematic phase offset $\phi = 0$.

In the second test, we introduced slow frequency modulation to generate semiperiodic observations.

$$\begin{aligned} s_i^{\sin} &= s_0 + a \sin \{ \pi i \Delta t / T_n - \phi \}, & \text{for } \sum_{j=1}^{n-1} T_j / \Delta < i \leq \sum_{j=1}^n T_j / \Delta; \\ x_i^{\text{modified cos}} &= s_0 - a \cos^{2n} \{ \pi i \Delta t / T_n - \phi \}, & \text{for } \sum_{j=1}^{n-1} T_j / \Delta < i \leq \sum_{j=1}^n T_j / \Delta, \end{aligned}$$

where T_i are distinctive values that are NOT integer multiples of each other, so that the resulting waveform is not strictly periodic.

In the third simulation, we tested the ability of the proposed method in identifying and decoupling the variation in the fundamental pattern. The data were generated according to

$$\begin{aligned} s_i^{\sin} &= s_0 + a_n \sin \{ \pi i \Delta t / T - \phi \}, & \text{for } \sum_{j=1}^{n-1} T_j / \Delta < i \leq \sum_{j=1}^n T_j / \Delta; \\ x_i^{\text{modified cos}} &= s_0 - a_n \cos^{2n} \{ \pi i \Delta t / T - \phi \}, & \text{for } \sum_{j=1}^{n-1} T_j / \Delta < i \leq \sum_{j=1}^n T_j / \Delta, \end{aligned}$$

where a_n takes on slowly varying values.

For real clinical data, we used the RPM system to obtain the trajectories of external fiducials placed on the chests of 12 patients. The time–displacement relationship was recorded at 30 Hz and assumed to be highly correlated with superior–inferior diaphragm motion (Vedam *et al* 2004). We centered and scaled the unit-less RPM data so that their dynamic range mimicked the SI motion for chest and lung tumors (Seppenwoolde *et al* 2002, Sharp *et al* 2004). Characteristic parameters for the RPM data used in our experiment are reported in table 1.




Figure 3. Online profiling results for simulated data with no variation in baseline, frequency or fundamental pattern oscillatory amplitude. Top row (a(#)) presents sinusoidal waveform; bottom row (b(#)) presents modified cosine waveform. Leftmost column presents discrete observations (blue dots), the fitted ellipse (red line) in the augmented state space and the Poincaré section (black dashed line); middle column presents simulated trajectory (blue), the iso-phase markers (red circle) and the corresponding linear phase modular 2π (green line); rightmost column presents anti-warped observation (blue dots) and estimated fundamental pattern (red line).

4. Results and discussions

4.1. Test results on simulated data

We first report the results from the simulation study. Note that all simulated data have an artificial static baseline, so a ‘good’ baseline estimation algorithm should provide a stable estimate. The performance of ellipse based baseline estimation has been justified in Ruan *et al* (2008a), so we skip that part, and focus on the validation of the warping function and the fundamental pattern. Figure 3 reports the results from the first test, where there is no variation in any of the components. The phase estimation is stable as expected, resulting in an accurate fundamental pattern estimate for profiling. Due to the stability of each component, it suffices to illustrate the result for an arbitrary time instant.

In the second simulated data set, frequency variation was introduced. We were pleasantly surprised to observe that both the Poincaré structure and the estimated fundamental pattern exhibited minimal (almost unnoticeable) variation during the course of real-time profiling. This is the ideal. The generated variation lies only in frequency (or phase). Upon correctly resolving such variation, the unwarping procedure decouples this from the estimated fundamental patterns, making them consistent over time. Thus, it is sufficient to show the Poincaré sectioning and the profile for an arbitrary time instant in figure 4. Note that since the modified cosine model does not ‘exactly’ correspond to an ellipse in the augmented space, as most real breathing trajectories, the corresponding phase estimate is not perfect. However, the

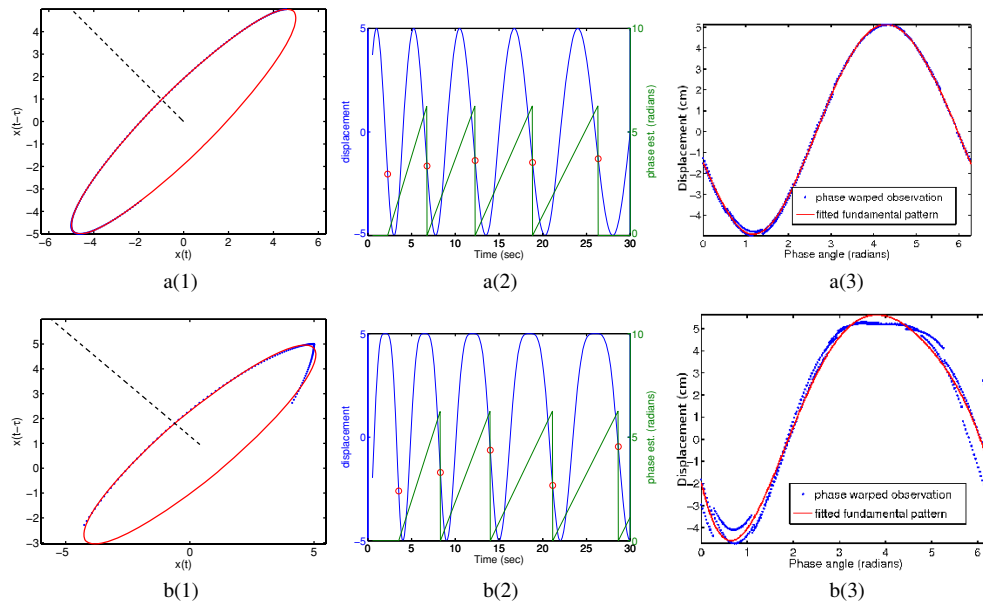


Figure 4. Online profiling results for simulated data with variation in frequency. Top row (a#) presents sinusoidal waveform; bottom row (b#) presents modified cosine waveform. Leftmost column presents discrete observations (blue dots), the fitted ellipse (red line) in the augmented state space and the Poincaré section (black dashed line) for identifying the iso-phase points; middle column presents simulated trajectory (blue), the iso-phase markers (red circle) and the corresponding linear phase modular 2π (green line); rightmost column presents anti-warped observation (blue dots) and estimated fundamental pattern (red line).

seemingly obvious fluctuations in the projected-back displacement values actually correspond to very small errors in the estimated warping function. This is because the error in estimated phase warping results from the fluctuation in displacement scaled by the inverse of local gradient. The significant local gradient increases the robustness of phase estimation with respect to displacement fluctuation, in contrast to the case of using peaks and troughs.

In the third simulation, changes in fundamental pattern were introduced by amplitude variation. The optimal decomposition should be able to provide uniformly spaced iso-phase points and capture the variations in fundamental pattern. Figure 5 illustrates the evolution of the estimated ellipse, Poincaré section and fundamental patterns. Note that despite the changes in oscillatory amplitude, the proposed method correctly identified the iso-phase points, subsequently attributed the variation to the fundamental amplitude change.

4.2. Test results on clinical data

For clinical data, the validation process becomes more difficult, since there is no ground-truth available. As we commented before, the quality of the fundamental pattern estimate relies almost solely on the performance of the phase estimator. As a starting point, figure 6 reports the online phase estimate results for each of the patient data. Visually, the estimated phase correctly captures the frequency variation over time.

To obtain a quantitative measure of performance, we adopt a cross-validation type idea. The observed trajectory is partitioned into two sections ($s_{\text{train}} = [0, T_1]$, $s_{\text{test}} = (T_1, T_{\text{end}}]$).

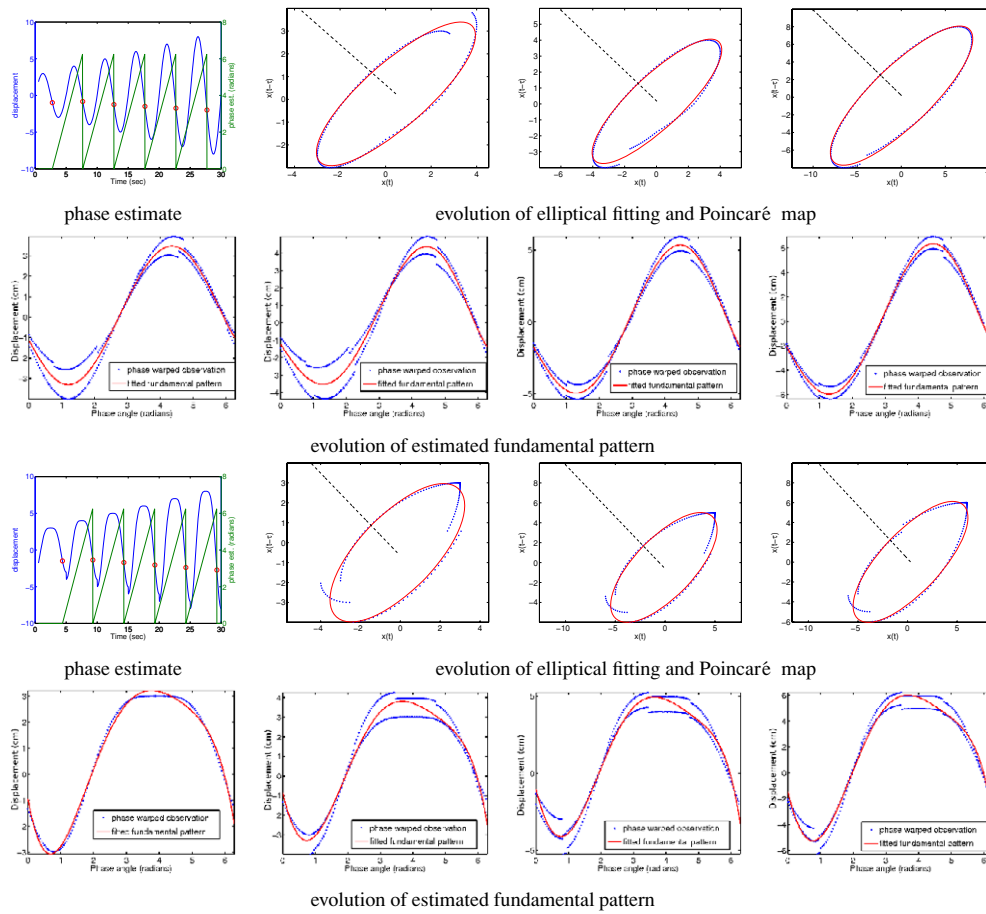


Figure 5. Online profiling results for simulated data with variation in fundamental pattern. Phase estimation, evolution of Poincaré map and fundamental pattern for amplitude modulated sinusoid are in rows 1 and 2; estimates for amplitude modulated modified cosine waveform are in rows 3 and 4.

We apply the proposed decomposition method to the former section, and use a simple zeroth order extrapolation for baseline, frequency and fundamental pattern, respectively. The error is defined as the discrepancy between the predicted displacement and the *a priori* known values for the latter section. Figure 7 illustrates the prediction performance and table 2 summarizes the prediction performance in terms of mean absolute value (MAE), root mean squared error (RMSE) and maximum error (ME). For MAE and RMSE, the error quantities were averaged over the continuous prediction horizon between 1 and 3 s.

$$\begin{aligned}
 \text{MAE} &= \frac{1}{2} \int_{\tau=1}^3 |\hat{s}(T_1 + \tau) - s(T_1 + \tau)|; \\
 \text{RMSE} &= \left\{ \frac{1}{2} \int_{\tau=1}^3 (\hat{s}(T_1 + \tau) - s(T_1 + \tau))^2 \right\}^{1/2}; \\
 \text{ME} &= \sup_{\tau \in (1,3)} |\hat{s}(T_1 + \tau) - s(T_1 + \tau)|.
 \end{aligned}
 \tag{11}$$




Figure 6. Real-time phase estimate results from clinical data in table 1. Patient IDs are ordered left to right, top to bottom.

Table 2. Prediction error.

ID vs Para.	1	2	3	4	5	6	7	8	9	10	11	12
MAE	0.9	2.3	1.7	0.6	7.5	2.5	1.4	11.7	2.4	0.6	2.4	13.7
RMSE	1.2	2.5	1.9	0.7	10.0	3.9	1.8	15.5	3.5	1.1	1.7	15.0
ME	2.0	4.3	3.8	1.7	22.0	10.3	3.5	28.5	8.4	4.8	4.5	24.4

The unit for displacement is consistent with that presented in table 1. The error metrics were defined in (11).




Figure 7. Illustration of prediction (for one individualized-period ahead) based on the simultaneous real-time estimate of mean position as in Ruan *et al* (2008a) and phase estimate in figure 6. Based on available observations (blue), forward prediction (red) is made, and compared to the unseen 'truth' (black).

For trajectories that exhibit relatively regular trends, the proposed method worked well in predicting the target positions, even for quite large prediction lengths. Our method cannot predict prospective target positions that differ dramatically from the historical knowledge, as in the case of patients 5, 8 and 12. These data sets violate at least one of the consistency assumptions: they either have very different baseline from the historical trend (patient 12), experience sudden frequency variations (patient 8), or demonstrate dramatic pattern changes

(patient 5). These scenarios are impossible to predict, even for human observers. Such challenge points to the necessity of monitoring real-time prediction/estimation errors closely, and rigorous studies of abnormality detection mechanisms.

4.3. General comments

- Section 2 provides theoretical development for scalar observations, which applies to real position management system (RPM, Variation Medical Systems, Palo Alto, CA) and other types of low-dimensional external surrogates. When higher dimensional data are available, one could track and predict each dimension independently and process the signal along each dimension independently and simultaneously. It is feasible to utilize the correlation among various physical dimensions, particularly in phase estimation, to improve the performance further.
- We adopted a uniform sampling setup for the sake of simplicity, but all analyses generalize naturally to nonuniformly acquired observations. The low-dimensional nature of the elliptical shape model in the augmented state space also enables handling of sparse samples with robustness.
- We have used linear phase (6) for interpolation and extrapolation, because it is simple and preserves monotonicity of the warping function. Essentially, it is an interpolation method with piecewise linear basis functions. If smoother phase transition is desired, one could adopt more sophisticated formulations such as monotone splines (Ramsey 1988).
- We have demonstrated the use of a single Poincaré surface in identifying the iso-phase marker events in the experiments. Generalization to the use of multiple Poincaré sectioning is straightforward. The tradeoff resulting from such a choice is also clear: as the number of Poincaré surfaces increases, the linearly interpolated/extrapolated phase map has more flexibility in approximating any monotonic warping function, and potentially better represents the real phase; on the other hand, the more Poincaré surfaces, the more susceptible the extracted iso-phase points are to observation noise. Proper separation of the Poincaré surfaces is advised. From our practical experience, a single sectioning is often sufficient for general estimation/tracking and prediction tasks. We suggest using no more than 4 sections in the presence of non-negligible observation noise.

5. Summary and future work

We have proposed an integrated method in performing real-time profiling for respiratory motion. By estimating the parameters for a low-dimensional elliptical shape model in an augmented state space, the proposed method decouples baseline drifts, frequency variation and amplitude change in fundamental oscillatory pattern. Poincaré sectioning is utilized to provide a set of iso-phase maker events, which guides the assignment of phase values to each temporal point, using a simple linear interpolation/extrapolation filter. The observations are subsequently warped to obtain a fundamental pattern defined on the basic interval $[0, 2\pi)$. This online decomposition procedure is efficient, and robust to observation noise and/or sparse samples. The decoupling of various contributing components to the variation in respiratory motion helps characterize the instantaneous change, and has the potential to facilitate the corresponding clinical decisions. Furthermore, smoothness assumptions are reasonable with respect to each component in their own spaces. By filtering baseline drift, phase warp function and fundamental pattern respectively and then assembling them in the original signal space, we have provided a predictor that incorporates most physical priors for respiratory motion.

In the future, we will investigate the fusion of estimations for higher dimensional data, since frequency variation and baseline drifting are correlated among different physical dimensions. We will also study the very relevant problem of change detection, and the strategies for alarming when abnormalities occur.

From validation point of view, we will further test the proposed method on other external surrogate signals and internal tumor trajectories, and analyze the performance in the presence of different noise statistics.

Acknowledgments

The authors gratefully acknowledge the anonymous reviewer for his/her helpful suggestions and the support from Calypso medical.

References

- Bendixson I 1901 Sur les courbes définies par des équations différentielles *Acta. Math.* **24** 1–88
- Berbeco R I, Mostafavi H, Sharp G C and Jiang S B 2005 Towards fluoroscopic respiratory gating for lung tumours without radiopaque markers *Phys. Med. Biol.* **50** 4481–90
- Colgan R, McClelland J, McQuaid D, Evans P M, Hawkes D, Brock J, Landau D and Webb S 2008 Planning lung radiotherapy using 4d ct data and a motion model *Phys. Med. Biol.* **53** 5815–30
- Cui Y, Dy J G, Alexander B and Jiang S B 2008 Fluoroscopic gating without implanted fiducial markers for lung cancer radiotherapy based on support vector machines *Phys. Med. Biol.* **53** N315–27
- D'Souza W D, Nazareth D P, Zhang B, Deyoung C, Suntharalingam M, Kwok Y, Yu C X and Regine W F 2007 The use of gated and 4d ct imaging in planning for stereotactic body radiation therapy *Med. Dosim.* **32** 92–101
- Engelsman M, Damen E M, De Jaeger K, van Ingen K M and Mijnheer B J 2001 The effect of breathing and set-up errors on the cumulative dose to a lung tumor *Radiother. Oncol.* **60** 95–105
- Engelsman M, Sharp G C, Bortfeld T, Onimaru R and Shirato H 2005 How much margin reduction is possible through gating or breath hold? *Phys. Med. Biol.* **50** 477–90
- Gabor D 1946 Theory of communication *J. IEE* **93** 426–57
- Gierga D P, Chen G T Y, Kung J H, Betke M, Lombardi J and Willett C G 2004 Quantification of respiration-induced abdominal tumor motion and its impact on IMRT dose distributions *Int. J. Radiat. Oncol. Biol. Phys.* **58** 1584–95
- Hirsch M, Smale S and Devaney R 2003 *Differential Equations, Dynamical Systems, and An Introduction to Chaos* (New York: Academic)
- ICRU 1993 Prescribing, recording, and reporting photon beam therapy *ICRU Technical Report 50* (Bethesda, MD: ICRU)
- ICRU 1999 Prescribing, recording and reporting photon beam therapy (supplement to ICRU Report 50) *ICRU Technical Report 62* (Bethesda, MD: ICRU)
- Jiang S B, Pope C, Al Jarrah K M, Kung J H, Bortfeld T and Chen G T 2003 An experimental investigation on intra-fractional organ motion effects in lung IMRT treatments *Phys. Med. Biol.* **48** 1773–84
- Keall P J *et al* 2006 The management of respiratory motion in radiation oncology report of AAPM task group 76 *Med. Phys.* **33** 3874–900
- Kubo H D, Len P M, Minohara S and Mostafavi H 2000 Breathing-synchronized radiotherapy program at the university of California Davis cancer center *Med. Phys.* **27** 346–53
- Langen K M and Jones D T L 2001 Organ motion and its management *Int. J. Radiat. Oncol. Biol. Phys.* **50** 265–78
- Lujan A E, Balter J M and Haken R K T 2003 A method for incorporating organ motion due to breathing into 3D dose calculations in the liver: sensitivity to variations in motion *Med. Phys.* **30** 2643–9
- Mah D, Hanley J, Rosenzweig K E, Yorke E, Braban L, Ling C C and Magras G 2000 Technical aspects of the deep inspiration breath hold technique in the treatment of thoracic cancer *Int. J. Radiat. Oncol. Biol. Phys.* **48** 1175–85
- Neicu T, Shirato H, Seppenwoolde Y and Jiang S B 2003 Synchronized moving aperture radiation therapy (smart): average tumour trajectory for lung patients *Phys. Med. Biol.* **48** 587–98
- Nelson C, Starkschall G, Balter P, Fitzpatrick M J, Antolak J A, Tolani N and Prado K 2005 Respiration-correlated treatment delivery using feedback-guided breath hold: a technical study *Med. Phys.* **32** 175–81
- Ozhasoglu C and Murphy M J 2002 Issues in respiratory motion compensation during external-beam radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **52** 1389–99
- Ramsey J O 1988 Monotone regression splines in action *Stat. Sci.* **3** 425–61

- Rosenzweig K E *et al* 2000 The deep inspiration breath-hold technique in the treatment of inoperable non-small-cell lung cancer *Int. J. Radiat. Oncol. Biol. Phys.* **48** 81–7
- Ruan D, Fessler J A, Balter J M and Sonke J-J 2006 Exploring breathing pattern irregularity with projection-based method *Med. Phys.* **33** 2491–9
- Ruan D, Fessler J A and Balter J M 2007 Real-time prediction of respiratory motion based on nonparametric local regression methods *Phys. Med. Biol.* **52** 7137–52
- Ruan D, Fessler J A and Balter J M 2008a Mean position tracking of respiratory motion *Med. Phys.* **35** 782–92
- Ruan D, Fessler J A, Balter J M, Berbeco R I, Nishioka S and Shirato H 2008b Inference of hysteretic respiratory tumor motion from external surrogates: a state augmentation approach *Phys. Med. Biol.* **53** 2923–36
- Seppenwoolde Y, Shirato H, Kitamura K, Shimizu S, Herk M, Lebesque J V and Miyasaka K 2002 Precise and real-time measurement of 3D tumor motion in lung due to breathing and heartbeat, measured during radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **53** 822–34
- Sharp G C, Jiang S B, Shimizu S and Shirato H 2004 Prediction of respiratory tumour motion for real-time image-guided radiotherapy *Phys. Med. Biol.* **49** 425–40
- Unser M, Aldroubi A and Eden M 1993a B-spline signal processing: part I. Theory *IEEE Trans. Signal Proc.* **41** 821–33
- Unser M, Aldroubi A and Eden M 1993b B-spline signal processing: part II. Efficient design and applications *IEEE Trans. Signal Proc.* **41** 834–48
- Vainstein L A and Vakman D E 1983 *Frequency Analysis in the Theory of Oscillations and Waves* (Moscow: Nauka)
- Vedam S S, Keall P J, Docef A, Todor D A, Kini V R and Mohan R 2004 Predicting respiratory motion for four-dimensional radiotherapy *Med. Phys.* **31** 2274–83
- Venkat R B, Sawant A, Suh Y, George R and Keall P J 2008 Development and preliminary evaluation of a prototype audiovisual biofeedback device incorporating a patient-specific guiding waveform *Phys. Med. Biol.* **53** N197–208
- Wong J W, Sharpe M B, Jaffray D A, Kini V R, Robertson J M, Stromberg J S and Martinez A A 1999 The use of active breathing control (ABC) to reduce margin for breathing motion *Int. J. Radiat. Oncol. Biol. Phys.* **44** 911–9
- Wu H, Sharp G C, Salzberg B, Kaeli D, Shirato H and Jiang S B 2004 A finite state model for respiratory motion analysis in image guided radiation therapy *Phys. Med. Biol.* **53** 57–72