On using an adaptive neural network to predict lung tumor motion during respiration for radiotherapy applications

Marcus Isaksson and Joakim Jalden
Department of Electrical Engineering, Stanford University, Stanford, California 94036

Martin J. Murphy[a)]
Department of Radiation Oncology, Virginia Commonwealth University, Richmond, Virginia 23298

(Received 6 January 2005; revised 12 October 2005; accepted for publication 14 October 2005; published 29 November 2005)

In this study we address the problem of predicting the position of a moving lung tumor during respiration on the basis of external breathing signals—a technique used for beam gating, tracking, and other dynamic motion management techniques in radiation therapy. We demonstrate the use of neural network filters to correlate tumor position with external surrogate markers while simultaneously predicting the motion ahead in time, for situations in which neither the breathing pattern nor the correlation between moving anatomical elements is constant in time. One pancreatic cancer patient and two lung cancer patients with mid/upper lobe tumors were fluoroscopically imaged to observe tumor motion synchronously with the movement of external chest markers during free breathing. The external marker position was provided as input to a feed-forward neural network that correlated the marker and tumor movement to predict the tumor position up to 800 ms in advance. The predicted tumor position was compared to its observed position to establish the accuracy with which the filter could dynamically track tumor motion under nonstationary conditions. These results were compared to simplified linear versions of the filter. The two lung cancer patients exhibited complex respiratory behavior in which the correlation between surrogate marker and tumor position changed with each cycle of breathing. By automatically and continuously adjusting its parameters to the observations, the neural network achieved better tracking accuracy than the fixed and adaptive linear filters. Variability and instability in human respiration complicate the task of predicting tumor position from surrogate breathing signals. Our results show that adaptive signal-processing filters can provide more accurate tumor position estimates than simpler stationary filters when presented with nonstationary breathing motion. [DOI: 10.1118/1.2134958]

INTRODUCTION

This paper addresses the problem of predicting tumor motion from surrogate breathing signals for the purpose of synchronizing a radiotherapy beam to a moving target. In an earlier report[1] we demonstrated the use of adaptive signal-processing filters for this purpose while simultaneously accommodating the beam delivery system's response latency time. Here, we apply a nonlinear adaptive neural network to the problem and evaluate its effectiveness compared to simpler linear and nonadaptive filters.

Respiratory motion complicates the targeting of external radiation to tumors in the lung, pancreas, and other thoracic and abdominal sites. Historically, the response has been to encompass the expected range of tumor positions with a planning target volume, but this irradiates more normal tissue than one would like. Two new techniques are being developed to actively adapt beam delivery to breathing motion. The first approach synchronizes the radiation delivery cycle to the respiratory cycle so that the beam is on only when the tumor is believed to be within a certain target position window.[2–6] The second approach follows the moving tumor with the beam.[7–11]

Synchronized beam gating and tracking both require knowledge of the tumor's whereabouts throughout the breathing cycle. Continuous radiographic observation of the tumor during treatment has been employed in a beam-gating system,[5] but this approach often is not feasible. Instead, one can attempt to infer the internal tumor position indirectly from surrogate anatomical motion that can be monitored externally.[11–14] If the tumor motion and surrogate signal are spatially and temporally correlated, then the tumor position can be deduced. This indirect tracking strategy is the subject of our study. For the purposes of our discussion, we will use external chest markers as the surrogate motion indicator, although the concepts and results apply equally well to other breathing signals.

In the simplest implementation of indirect tumor tracking, it is assumed that the correlation between tumor and surrogate marker position is constant in amplitude and phase for time periods that are long compared to the duration of the treatment. The correlation is measured before beginning the fraction by observing tumor and marker motion simultaneously; then, during the treatment fraction only the marker position is observed, from which the tumor position is inferred. However, because normal breathing is variable[15–19] and tumor motion can follow complex trajectories (especially in the mid-to-upper lobes of the lung), there will be many patients for which this simple assumption is invalid.[20] To best treat these more difficult cases, the motion tracking

system should monitor and update the tumor and marker motion correlation during the actual delivery of radiation. Because the surrogate tracking strategy is intended specifically for situations that do not allow for continuous tumor localization, one must employ a hybrid tracking scheme in which periodic direct observations of the tumor position during treatment are combined with the continuous observation of the correlated marker motion.[11] The periodic tumor imaging data are used to update the tumor/marker correlation if and when it varies. This is the scenario that we consider.

Regardless of how one arrives at an estimate of the tumor position, there will be some delay before a dynamic system can make its corrective response. This delay can range from 50 ms for the beam to be gated to several hundred milliseconds for the beam to be physically realigned. To keep the beam on target the tumor position must be anticipated by the system delay (latency) time.

The problems of inferring tumor position from surrogate markers and predicting it ahead in time to compensate system lag have been addressed separately by a few other researchers, but in more limited approximations to the tumor tracking problem. Sharp *et al.*[21] have analyzed semiempirical filters for temporal prediction of respiration, but did not address tumor correlation with surrogate markers. Vedam *et al.* have analyzed stationary harmonic prediction models and linear predictive filters for the respiratory motion of the diaphragm,[22] and have also analyzed the correlation between the diaphragm and external chest motion,[12] but have not addressed the actual tumor correlation/prediction problem. Ahn *et al.*[23] have measured internal versus external motion correlations for lung and diaphragm sites and estimated the accuracy of inferring the internal tumor site from a simple stationary linear correlation with external marker motion. Their results showed a wide range in predictive accuracy using this simple approach, with the best case reducing the uncertainty of the tumor position by 75%, while a number of cases (especially in the mid/upper lung) had little or no predictive ability at all.

In the complete indirect tracking scheme the external marker position at time t is provided as input to the prediction/correlation control loop, which outputs the tumor position at the future time $t+\tau$. In our original study[1] we observed that semiempirical control loop filters have the inherent flexibility to deal with the complex, individualized, and variable character of normal breathing behavior, and proposed that adaptive filters would give better results than their stationary counterparts. Adaptive filters operate by using an initial set of observational data to initialize free parameters such that the filter output optimally matches the known target signal. The filter parameters are then progressively updated via later sampling of the input and target signals.

In this paper we demonstrate the use of an adaptive neural network filter to predict lung tumor motion ahead in time from surrogate marker measurements, for cases where both the breathing cycle and the tumor/marker correlation vary during treatment. For the demonstration we used three representative examples of respiratory motion data ranging from a very regular, stable tumor/surrogate position correlation to a very noisy, very irregular correlation. We show that, when presented with complex nonstationary signals, the adaptive filters can improve the predictability of tumor position compared to simpler linear and/or stationary algorithms.

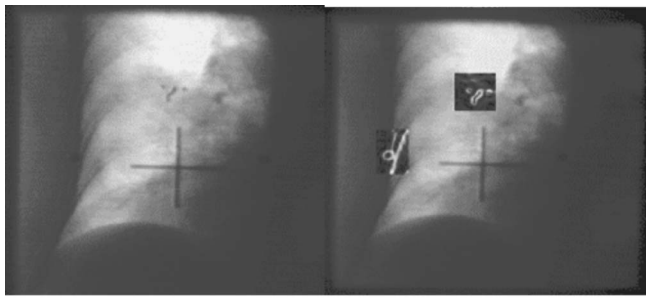

(a) (b)

FIG. 1. An AP fluoroscopic image frame for patient C showing the fiducials implanted in the lung tumor and the surrogate chest marker (highlighted by an edge filter in the square regions of interest).

METHOD AND MATERIALS

The respiratory motion data

For this study we chose three clinical examples of respiratory motion data (patients A, B, and C), from a cohort of eight patients undergoing image-guided radiosurgery.[10,24] These choices presented easy, medium, and hard-to-track tumor motions. One of the patients (A) had a pancreatic tumor; the other two (B and C) had midlobe lung tumors. Each patient had three fiducial markers implanted in his/her tumor for localization during therapy. In addition to the internal tumor fiducials, three small lead markers were affixed to the front external chest surface. Our data consist of fluoroscopic studies of the patients during free breathing. The fluoroscopic sequences were taken from AP and lateral viewpoints that were arranged so that both the external markers and internal fiducials were simultaneously visible in each image frame. This strategy ensured that the marker and tumor position data were exactly synchronized. Figure 1 shows an AP frame from patient C, illustrating the placement of the fiducials within the tumor and the external chest marker (visible within the enhanced regions of interest). The internal fiducials revealed the motion of the tumor itself, while the external markers acted as surrogates for the tumor motion. Patient A's pancreatic tumor motion followed the diaphragm/abdomen in a regular pattern, while patients B and C exhibited significant nonstationary correlations over time intervals of a few minutes.

It is well known that breathing patterns can change on a time scale of minutes, which means that a short observation interval of 30 s or less will not necessarily show the respiratory motion that occurs during the treatment fraction. To obtain data over realistic time scales, we observed the three patients in extended fluoroscopic sequences. Patient A was observed for two 45 s periods separated by 120 s without fluoroscopy. Patient B was observed in three 50 s sequences

spaced by 180 s intervals. For patient C there were two observations lasting 80 s that were separated by 180 s of no fluoroscopy.

The fluoroscopic sequences were videotaped and then analyzed offline. The analysis program used automatic contrast enhancement, edge filtering, and local thresholding to isolate each internal fiducial and external marker and extract its (x,y) coordinates in each frame. These data were compiled into files that provided a frame-by-frame record of the tumor coordinates correlated with the external marker coordinates.

We separated the analysis of tumor prediction into three parts. First, we measured the accuracy with which the internal tumor position could be inferred from the external marker position using three different filter configurations. Then, we measured the accuracy with which the external marker position could be predicted τ seconds ahead in time to allow for the delayed response of a breathing compensation technique. Finally, we combined the two processes by measuring the accuracy with which the internal tumor position could be inferred from the external marker position while predicting it τ seconds in advance. The results presented here are for tumor motion in the inferior/superior direction correlated to marker motion measured in the left/right and anterior/posterior directions.

The adaptive filtering process

Linear adaptive filters and neural networks are heuristic learning algorithms that configure themselves to replicate incoming signals without using any *a priori* models of the signal shape or generation process. In our application the filtering process consisted of supplying the measured trajectory $M(t)$ of the external chest marker to the filter input while taking from the filter output an estimate $P(t)$ of either the external marker position at a later time (for the temporal prediction tests) or the internal tumor position (for the correlation and prediction/correlation tests). The filter output was then compared to the actual tumor or marker coordinates $D(t)$ to get an error signal $\varepsilon(t)$ for adjusting the filter's free parameters. The error signal was also recorded as a running record of the predictive accuracy of the filter. We used one of the data sequences for each patient as training data to initialize the filter parameters, and then tested the filter performance against the other sets of data.

If the input signal $M(t)$ is stationary in time (i.e., its average over a fixed time interval does not change with time), then once the weights have been set by, e.g., a least-squares optimization during an initial training period, the filter will continue making an optimal estimate of the signal indefinitely. Furthermore, the temporal prediction accuracy will not change with latency. If the signal is not stationary then the filter will become increasingly inaccurate over time and the temporal prediction error will increase with latency. The solution to a nonstationary problem is to continually adjust the weights in response to the current signal characteristics so that the influence of the older signal samples diminishes as time passes. This defines an *adaptive* filter.

Our observations of tumor and/or surrogate organ positions were made at discrete time intervals. We will count observation intervals with the index i, and for compactness of notation will indicate the time of the current observation as t, the time of the previous observation as $t-1$, and the time of the ith preceding observation as $t-i$.

The filter design

For our tests we used a feed-forward neural network (NN) made of two input neurons with a sigmoid activation function feeding into a single-output neuron. The input signal for the external marker was first corrected for drifting dc offset by subtracting a running average of every marker sample to give zero mean and then rescaled to give variance of 1. The correlated tumor position was likewise centered on its mean and rescaled to unity variance. The rescaled input signal was divided into N sequential time samples using a tapped delay line, and the N samples were distributed in parallel to the two input neurons. This allowed the filter to remember the last N time samples. An additional input derived from the running time average of the input was also distributed to the two neurons. Each input to each neuron had an adjustable weight associated with it. Two additional weights were associated with the inputs from the first to the second layer, and a third weight on the output signal acted as a normalization. Therefore, there were altogether $2(N+1)+3$ adjustable weights to train. Figure 2 shows a schematic of the tapped delay line input, the neural network, and the complete filter configured for combined correlation/prediction in an actual real-time application.

A sigmoid activation function is nonlinear. If it is replaced by a linear function then the network reduces to a single layer perceptron, which is equivalent to a linear filter [Eq. (1)]

$$P(t) = \Sigma_i w_i M(t-i); \quad 0 < i <= N, \tag{1}$$

where w_i are the adjustable weights. Nonlinear neural networks are better than the linear perceptron in learning to recognize complex patterns in the input data, so we expect the NN to do better than a linear filter when presented with complex signals. At the same time, because the linear filter is a special case of the NN, we expect the NN to always do at least as well as the linear filter when operating under the same conditions. To investigate this we have tested both a linear and nonlinear filter design.

For combined prediction/correlation the linear and nonlinear filters were configured differently. The linear filter was set up in two stages—a prediction stage with N weights feeding its estimate of the surrogate marker position at $(t+\tau)$ into a correlation stage with another N weights. The nonlinear NN was set up as a single composite stage with $(2N+3)$ weights, as shown in Fig. 2. This allowed us to examine the relative merits of the two designs.

Choosing the signal history length for the input

The first step in filter optimization is to determine the signal history length for the input. This is particularly influ-

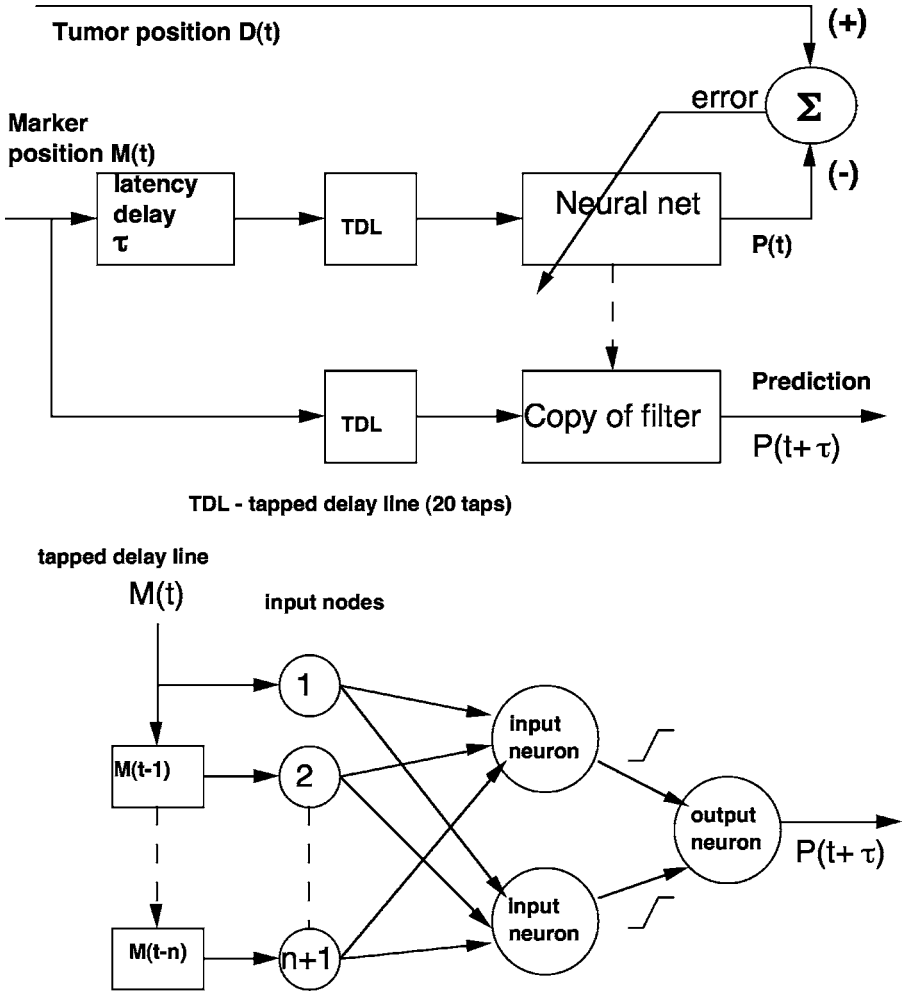

FIG. 2. Schematic of the neural network configured for real-time correlation and prediction. A tapped delay line distributes the incoming marker signal $M(t)$ to $2N$ input nodes of the neural network; the output of the first layer of neurons is transferred via a sigmoid activation function to the output neuron; the output neuron delivers the predicted tumor position $P(t)$ for comparison to the actual position $D(t)$.

ential for the temporal prediction function. The spacing of the samples, times the number of samples, determines the signal history length. If it is too short, the filter remembers only the most recent features and loses accuracy at longer latencies. If there are too few samples in the window, one loses temporal resolution; if there are too many samples in the window the training and adaptation rate decreases and the filter cannot react fast enough to change.

The sampling rate for our data was 10 Hz. We determined that 20–25 samples, corresponding to a history of 2–2.5 s, was sufficient. This (unsurprisingly) spanned approximately one average breathing cycle. A similar analysis for the correlation filter also indicated that 25 samples defined a good signal history window.

Optimizing the number of input weights

Each input requires a weight parameter. Too many weights will make the NN ill-conditioned and slow to adapt to changes in the input. With too few weights the filter becomes insensitive to detailed structure in the input. This is further complicated by the fact that the training process must work from a limited amount of training data. Too little training data or too many parameters can lead to overtraining of the network, resulting in instability. Therefore, one must find an optimal number of weights to balance sensitivity and robustness. As a general rule of thumb, the number of NN weights should be less than about 1/10 the number of training examples. Our data series were 400–600 samples long, which provided for about 500 training examples per patient. Therefore, we used $N=20$ input samples, which corresponded to 45 weights total.

If the data on the inputs are highly correlated, then the NN can use up many inputs (and weight parameters) on a small amount of useful information. Principal components analysis (PCA) can be used to concentrate the useful input information into fewer inputs. The data present on an initial number of N inputs are reconstructed on a set $M \leqslant N$ of principal component axes that are defined such that the data components reflected on the axes are maximally uncorrelated with one another. Each independent component of input is directed to its corresponding principal axis, which becomes a new input node. The PCA calculation produces a rank ordering of the principal axes. The axis that accounts for the largest part of the signal information is ranked first; the next most informative axis is ranked second, and so on. This allows the user to cut off the principal axis expansion at the number $M<N$ of components that account for most (if not all) of the useful information. This can reduce the dimension-

ality of the input and thus the number of weights. We took the initial 20 sequential input signal samples for the neural network and used PCA to reduce the number of input weights, then compared the accuracy of the filter to its performance without dimensionality reduction via PCA.

Training the filters

Because the individual sequences of data were comparatively short, the filters were trained on one sequence of data and then tested on a second sequence. The linear adaptive filter weights were trained and updated using the Least Mean Square (LMS) scheme[25]

$$w_i(t+1) = w_i(t) + 2\mu\varepsilon(t)\ D(t-i); \quad 0 < i < N, \tag{2}$$

where $\varepsilon(t) = P(t) - D(t)$ is the error in the current estimated signal sample $P(t)$. The parameter μ determines the speed of convergence. We have found from empirical trials that a good compromise between accuracy and speed of convergence is obtained when

$$\mu = 0.05/\text{trace}(R), \tag{3}$$

where $[R]$ is an estimate of the covariance matrix of the input signal samples. The nonlinear neural network was trained using the Levenberg-Marquardt algorithm.[25] The initial training used the full sampling rate of 10 Hz for the tumor and surrogate marker positions.

Training a nonlinear neural network is more difficult because the weights are initialized at random. Therefore, a single training pass can sometimes result in an exceptionally good filter, but sometimes result in a very bad one. The standard defense against this is to run the training process multiple times while averaging the resulting filters. However, too many training runs can overfit the filter, causing it to become too closely matched to the training data. To prevent this the training process can be monitored by using a third data set for each patient as a validation sequence. After each training epoch the filter is tested against the validation sequence. For the initial training epochs the performance on both the training and validation set improves, but if the filter begins to overfit the data then the validation performance will begin to decrease while the training performance continues to improve. This is a good time to stop the training.

Updating the filters during tests

Once the filters were initialized by the training process, they were applied to the test data sequences. The updating simulated clinical practice. For pure temporal prediction both filters were updated each time a new external surrogate position sample appeared at the input, i.e., at a rate of 10 Hz.

The process of adaptively updating the spatial correlation by periodically relocating the tumor position was simulated by supplying the filter with every time sample of the marker coordinates but only updating it with new tumor coordinates after every nth sample. The time interval between tumor position updates was tested at 1.0 s, 5.0 s, and infinity to observe the effect of update frequency on the accuracy of the inferred tumor position. (As the interval between updates increases, the adaptive filter approaches the behavior of a stationary filter.)

For prediction/correlation the two-stage linear filter had its temporal predictor stage updated at the external data sam-

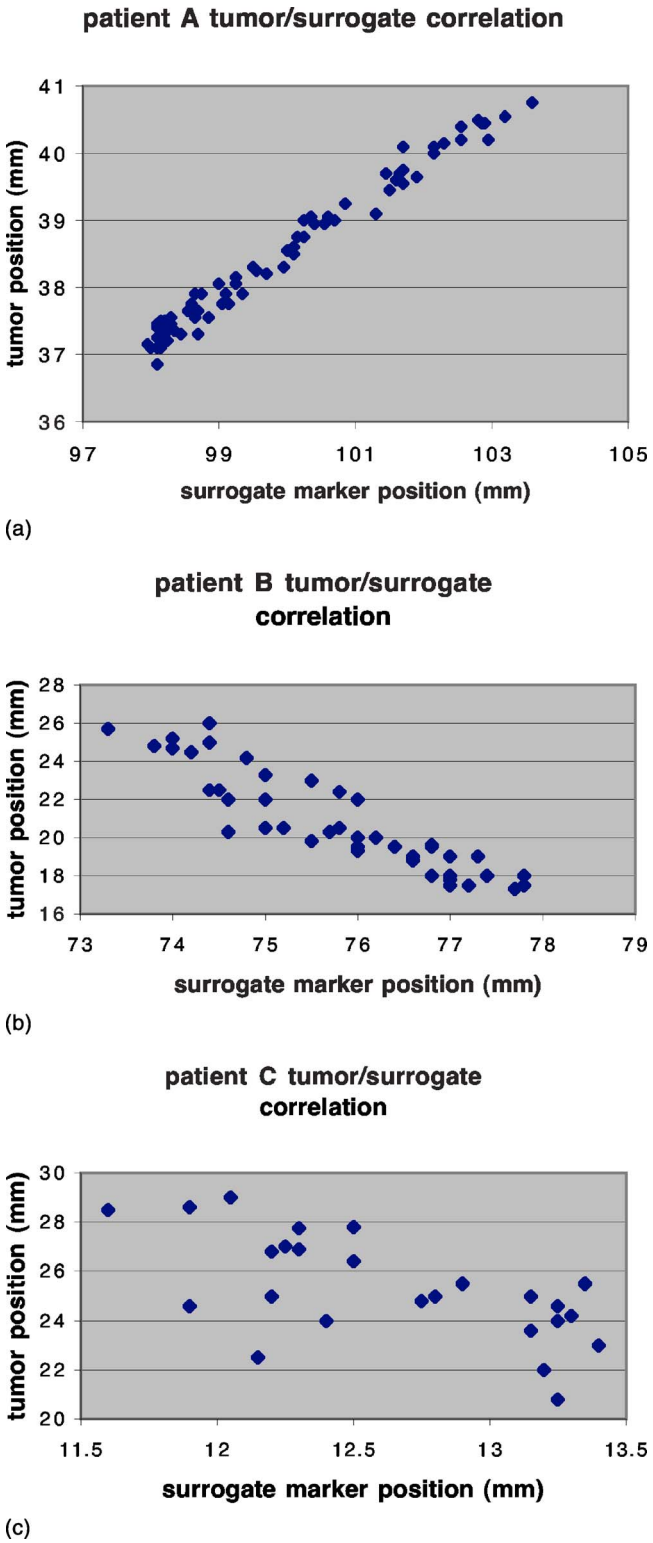

FIG. 3. The correlation plots of tumor versus marker position for (a) patient A; (b) patient B; and (c) patient C, derived by sampling a data point once every second during the test sequences.

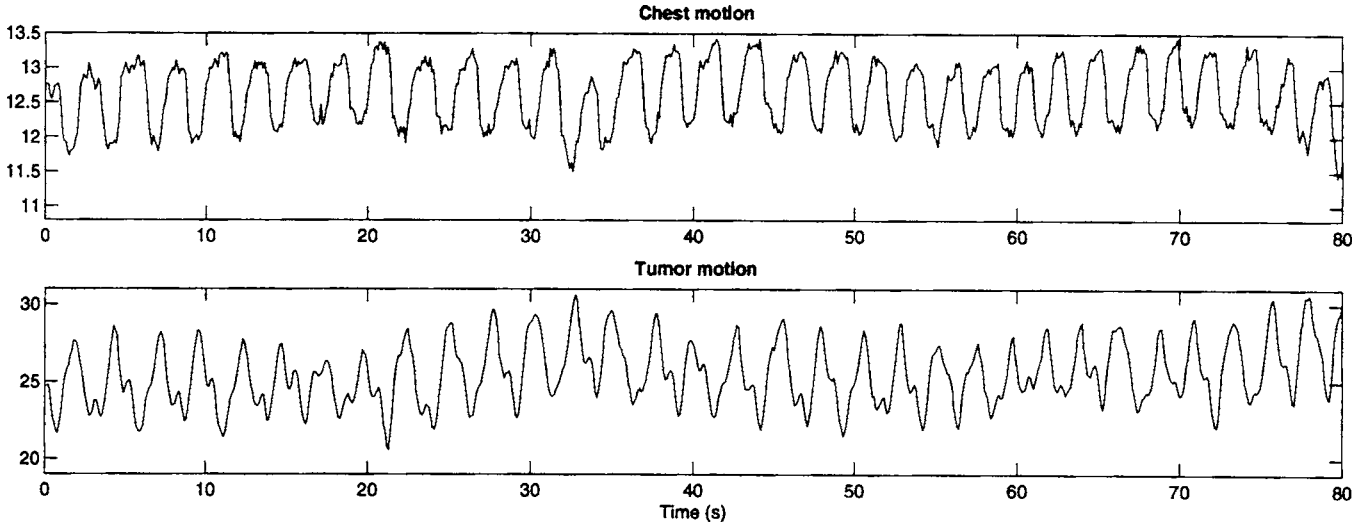

FIG. 4. The synchronized positions of patient C's tumor and external marker over a period of 60 s.

pling rate, while the correlation stage had its weights updated at the tumor sampling rate. The one-stage neural network combined both stages into a single filter that updated all of its weights at the tumor sampling rate.

Evaluation of the data

In a properly configured signal-processing filter algorithm the output should scale with the input; i.e., its performance should be independent of the absolute signal magnitude, thus allowing the signals to be normalized to, e.g., unit scale. Our performance measure for each of the three categories of test was the normalized root-mean-square error (nRMSE) between the predicted and actual signal over all the samples during the test interval

$$\mathrm{nRMSE} = [\Sigma_i (D_i - P_i)^2 / \Sigma_i (D_i - \mu_z)^2]^{1/2}, \tag{4}$$

where D_i is the ith observation, P_i is the estimate of the ith observation, and μ_z is the mean of all the observations. For an approximately sinusoidal signal the nRMSE is related to the peak-to-peak standard deviation of the predicted signal from the actual signal according to

$$\mathrm{nRMSE} = 2\sqrt{2}(S \,.\, D \,./ \delta S), \tag{5}$$

where δS is the peak-to-peak excursion of the signal (i.e., the full range of tumor displacement). A nRMSE of 28% therefore corresponds to a 10% peak-to-peak standard deviation. A nRMSE error of 100% indicates no correlation at all between the predicted and actual signal—one might as well take the average input signal as the predicted output. Equation (5) can be used to translate the filter's accuracy approximately into absolute distances of displacement.

RESULTS

The aim of this study was to demonstrate the effect of adapting the filter algorithms as new respiratory data were acquired. We preface the results with the following observations:

(1) a stationary filter is the limiting case of an adaptive filter when the interval between data samples used to update the filter weights increases to infinity;
(2) if the signals are spatially stationary (i.e., the tumor/surrogate correlation does not change in time) then adapting the filter will not improve results over a stationary filter; conversely, if the filter performance improves as the updating interval decreases, then the correlation is not stationary; and
(3) if the signal is temporally stationary, then once a filter has been trained it will predict future signal samples with constant accuracy regardless of the predict-ahead time (i.e., the latency); conversely, if the filter prediction loses accuracy as the latency increases, then the signal is not stationary.

If the signals are not stationary, then an adaptive filter should always do as well or better than its stationary counterpart.

Spatial correlation

The most important function of the filter in the indirect tracking application is prediction of the tumor position based on the surrogate position. Figure 3(a) shows the spatial correlation of tumor and external surrogate marker measured in one dimension for patient A, who had a pancreatic tumor. This motion correlation is an example of the kind of simple stationary pattern that can be predicted with equal accuracy using a stationary linear filter, an adaptive linear filter, and an adaptive neural network. Therefore, we will not elaborate further on the results for patient A.

Figure 3(b) shows the corresponding correlation for patient B—a lung patient with a moderately irregular respiration pattern. We would expect a stationary linear filter to give less accurate results for this patient than for patient A. Patient C presented the most irregular respiratory pattern among our three examples, resulting in the tumor/marker correlation in Fig. 3(c). Figure 4 shows the synchronized motion of patient C's tumor and external chest marker measured in one dimen-

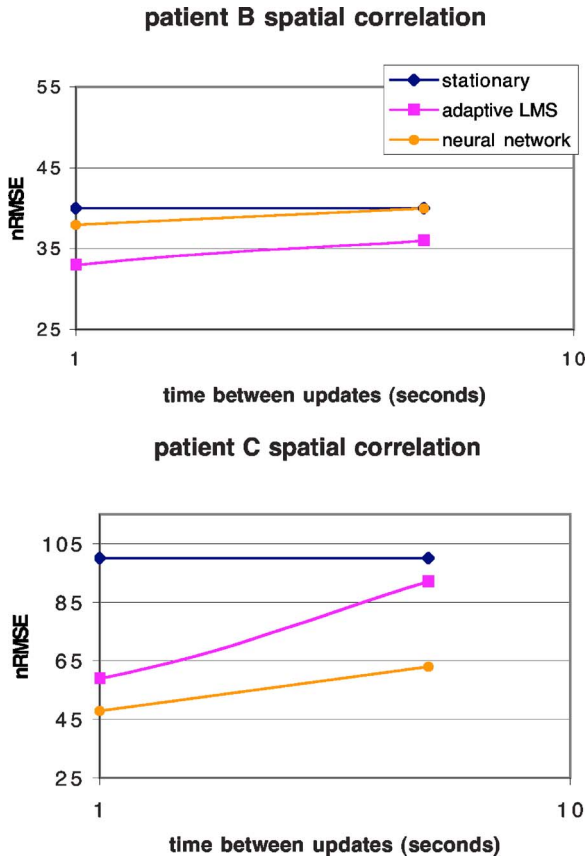

FIG. 5. The normalized root-mean-square error (nRMSE) in the spatial correlation of tumor and marker for patients B and C, using the stationary linear, adaptive linear, and adaptive neural network filters. The adaptive filters were updated at 1 and 5 s intervals.

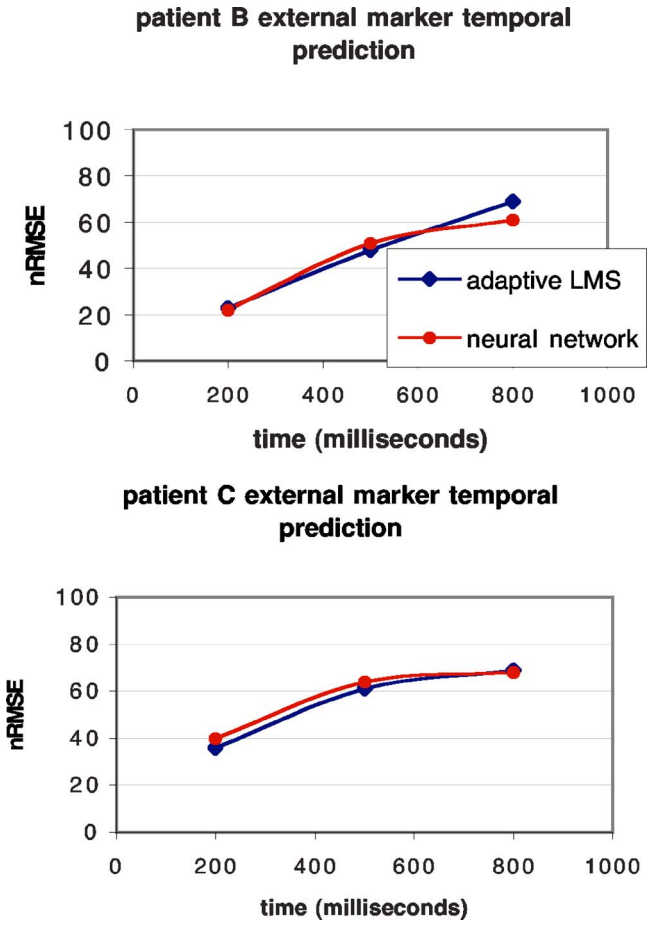

FIG. 6. The nRMSE in the temporal prediction of the motion of the external markers for patients B and C for lag times of 200, 500, and 800 ms, using the adaptive linear and NN filters.

sion over a 1-min interval. These data illustrate some of the effects that can confound attempts at temporally predicting the respiratory cycle and inferring tumor position from other surrogate signals. We note that the external marker's time sequence varies in amplitude and period by more than 30% from one cycle to the next, which makes it highly nonstationary. The tumor motion is clearly influenced not only be respiration but also by additional periodic rhythms that may include the heartbeat. The phase difference between the peak tumor and peak marker displacements varies over the observation interval, and the baseline fluctuations of the tumor and marker are not synchronized. We would expect the performance of a simple stationary filter to be worst of all for this patient.

Figure 3 can be used to estimate the accuracy of predicting tumor position from surrogate marker position using a simple linear stationary model. Such a model is equivalent to fitting a straight line to the correlation data and then using the fit to predict tumor position from the marker. For patient A the stationary model has nRMSE=14%. For patient B the stationary nRMSE=40%, while for patient C the stationary nRMSE=100%. Recall that when nRMSE=100% there is no predictive power at all—the average position of the tumor is the best estimate. These are the results against which one should compare the effectiveness of adapting the filters to new data as they are acquired.

Figure 5 shows the relative performance of the three correlation algorithms for patients B and C, for increasing intervals between tumor position updates. The improvement in performance as the update frequency increased shows that the correlation was nonstationary. Unsurprisingly, the stationary filter (infinite update interval) provided the poorest results in both cases. For patient C (the most irregular case) the stationary filter was completely unable to correlate the marker and tumor motions (as the simple linear correlation analysis already showed). For patient B (the moderately irregular case) the linear adaptive filter actually did a little better than the neural network (contrary to expectation), while the more complex and irregular patterns for patient C were handled better by the neural network.

Temporal prediction

This group of tests measured the predictability of the external chest marker position up to 800 ms in advance. Figure 6 summarizes the results for the linear and NN adaptive filter algorithms applied to lung tumor patients B and C. The de-

crease in accuracy as the latency increased shows that neither patient's temporal respiration pattern was stationary. The difference in performance between the linear filter and the neural network is not significant.

Combined spatial correlation and temporal prediction

Figure 7 summarizes the accuracy with which the filters could correlate the tumor motion with the external marker for both patients while looking ahead up to 800 ms to allow for system lag time. In each graph the prediction accuracy is plotted for 1 and 5 s tumor update intervals. For both patients the adaptive filters outperformed the stationary filter, and the adaptive neural network equaled or outperformed the adaptive linear filter.

DISCUSSION

The purpose of this study was to demonstrate the use and effectiveness of adaptive filters to predict tumor position from surrogate marker movement in a real-time control loop, using representative examples of clinical data. Our three case studies represented the range of breathing patterns observed in a small cohort of patients treated via image-guided radiosurgery. The pancreatic patient presented a straightforward tracking problem that could be handled by the least sophisticated filter. The two lung patients presented more challenging behaviors. The tumors were in the mid/upper lung, the correlations between tumor and chest motion were variable, and the tumor motion had additional nonrespiratory influences. Patient C was significantly more erratic than patient B. A comprehensive gating or tracking system should be prepared to deal with cases such as these, as well as even more difficult cases.

For lung patients B and C we found that, when compared to simple stationary algorithms, the adaptive filters gave an improvement in predicting the tumor position from the marker position. For patient C, where the simpler approaches could not be made to work at all, the more sophisticated neural network filter had a useful effect. This was not surprising, as the patient data were selected because they presented the type of problems that adaptive filters are designed to solve. There will be many clinical situations (such as patient A) where simpler algorithms will work effectively, but a control loop should be designed to handle the widest possible variety of actual patient behavior.

This study introduces many of the design issues that must be addressed to find the optimal filter setup. For example, the linear filter was designed as a two-stage predictor/correlator, while the neural network combined spatial correlation and temporal prediction in a single stage. The single-stage configuration allowed less frequent updating for the temporal prediction. Nevertheless, the NN outperformed the linear filter for the longer latency intervals. On the other hand, it appeared to reach a limiting accuracy as the latency period was reduced below 500 ms, which was most likely because it had less temporal resolution in the updating scheme.

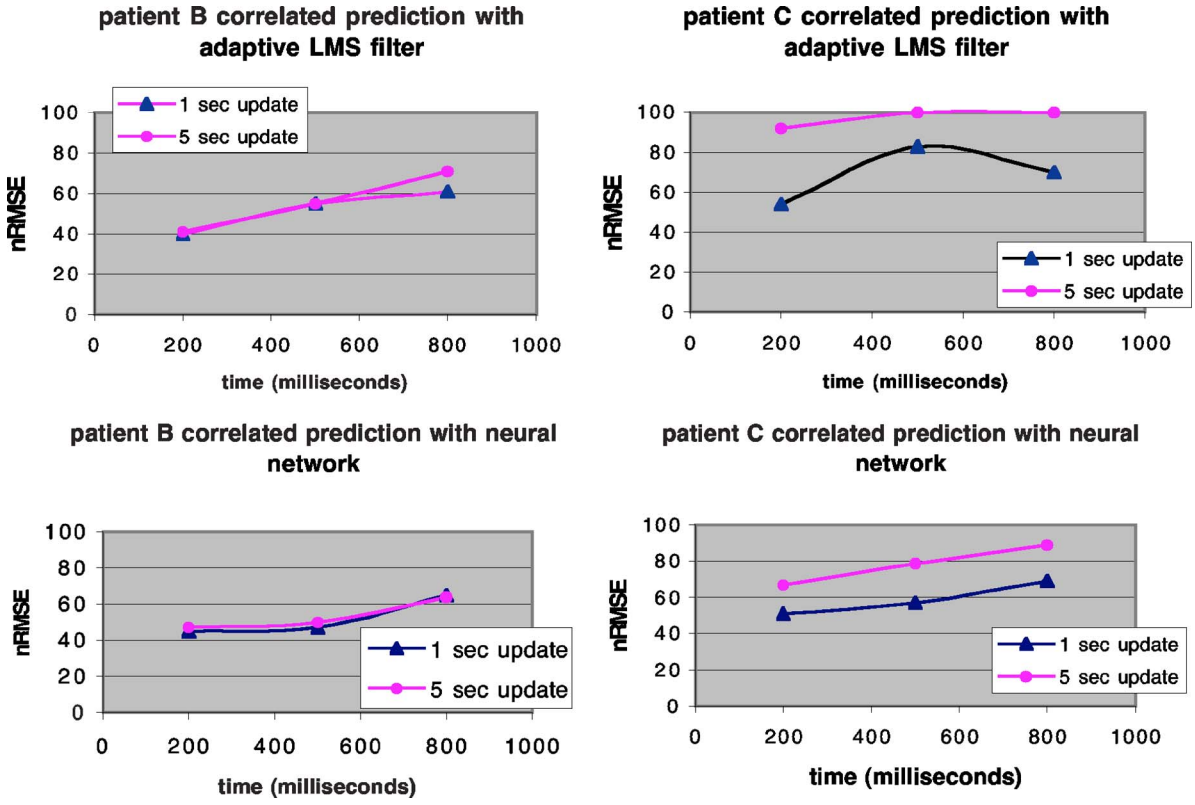

FIG. 7. For each filter, the nRMSE in the tumor position following spatial correlation and temporal prediction from the marker signal, for update intervals of 1 and 5 s and latencies of 200, 500, and 800 ms.

In our tests, dimensionality reduction of the NN inputs by principal component analysis (PCA) did not improve temporal prediction accuracy. This was not surprising—reducing the number of samples in the signal history window reduces time resolution. On the other hand, reduction from 20 to 5 inputs via PCA did improve the tumor spatial correlation results. This would indicate that PCA identifies those components of the surrogate marker position signal that provide the most useful information about the tumor position. This would be even more relevant to an obvious generalization of the NN filter for correlated tumor tracking. Some patients are chest breathers, some patients are abdominal breathers, and some patients switch back and forth. It is often not obvious where to put the surrogate marker(s) for the best results. A NN can combine data from more than one marker, allowing one to track multiple points on the chest and abdomen. Principal component analysis of the inputs would maximize the independent marker information on the inputs, and then the network weights would identify those inputs that best correlate with tumor position. We also note that cardiac influence on lung and pancreas tumors is readily observable.[20] A NN can combine both a surrogate breathing signal and a cardiac signal to improve the description of the overall tumor motion trajectory.

In summary, the usefulness of adaptive filtering was demonstrated by varying the adaptation interval between tumor position updates. As the interval increased, the adaptive filter's prediction/correlation accuracy diminished as it approached that of a stationary filter. Our results show that adapting the filter weights to new data improved the performance of the filter when presented with nonstationary signals. In the most extreme case the adaptive filter provided useful correlative predictability when the stationary filter was completely unable to do so. The results also show that a nonlinear neural network can give greater accuracy than a linear filter as the breathing patterns become more complex. We conclude that adaptability and nonlinearity can be useful attributes of control loops for the prediction and correlation of tumor motion and become more important as breathing motion becomes more complex and erratic.

ACKNOWLEDGMENTS

We would like to thank Professor Bernard Widrow of Stanford University's Department of Electrical Engineering for his expert advice and enthusiastic support for this study, which was supported in part by the Bio-X Research Initiative at Stanford University.

[a]Electronic mail: mjmurphy@vcu.edu

[1]M. J. Murphy, J. Jalden, and M. Isaksson, "Adaptive filtering to predict lung tumor breathing motion during image-guided radiation therapy," Proceedings of the 16th International Congress on Computer-assisted Radiology and Surgery (CARS 2002), 539–544 (2002).

[2]K. Ohara *et al.*, "Irradiation synchronized with respiration gate," Int. J. Radiat. Oncol., Biol., Phys. **17**, 853–857 (1989).

[3]H. D. Kubo and B. C. Hill, "Respiration gated radiotherapy treatment: A technical study," Phys. Med. Biol. **41**, 83–91 (1996).

[4]H. D. Kubo, P. M. Len, S. Minohara, and H. Mostafavi, "Breathing-synchronized radiotherapy program at the University of California Davis Cancer Center," Med. Phys. **27**, 346–353 (2000).

[5]H. Shirato, S. Shimizu, and T. Kunieda, "Physical aspects of a real-time tumor-tracking system for gated radiotherapy," Int. J. Radiat. Oncol., Biol., Phys. **48**, 1187–1195 (2000).

[6]Y. Seppenwoolde, H. Shirato, and K. Kitamura, "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**, 822–834 (2002).

[7]P. J. Keall, V. R. Kini, and S. Vedam, "Motion adaptive x-ray therapy: A feasibility study," Phys. Med. Biol. **46**, 1–10 (2001).

[8]S. Jiang, P. Zygmanski, and J. Kung, "Gated motion adaptive therapy (GMAT): Modification of IMRT MLC leaf sequence to compensate for tumor motion," Med. Phys. **29**, 1347–1348 (2002) (abstr).

[9]T. Neicu, H. Shirato, and Y. Seppenwoolde, "Synchronized moving aperture radiation therapy (SMART); Average tumor trajectory for lung patients," Phys. Med. Biol. **48**, 587–598 (2003).

[10]M. J. Murphy, "Tracking moving organs in real time," Semin. Radiat. Oncol. **14**, 91–100 (2004).

[11]A. Schweikard, H. Shiomi, and J. Adler, "Respiration tracking in radiosurgery," Med. Phys. **31**, 2738–2741 (2004).

[12]S. S. Vedam, V. R. Kini, P. J. Keall, V. Ramakrishnan, and R. Mohan, "Quantifying the predictability of diaphragm motion during respiration with a noninvasive external marker," Med. Phys. **30**, 505–513 (2003).

[13]Y. Tsunashima *et al.*, "Correlation between the respiratory waveform measured using a respiratory sensor and 3D tumor motion in gated radiotherapy," Int. J. Radiat. Oncol., Biol., Phys. **60**, 951–958 (2004).

[14]J. D. P. Hoisak, K. E. Sixel, R. Tirona, P. C. F. Cheung, and J. P. Pignol, "Correlation of lung tumor motion with external surrogate indicators of respiration," Int. J. Radiat. Oncol., Biol., Phys. **60**, 1298–1306 (2004).

[15]A. Hugelin and J. F. Vibert, "Is the respiratory rhythm multistable in man?," in *Respiratory Control: A Modeling Perspective*, Proc. Oxford Conference 1988, edited by G. D. Swanson, F. S. Grodins, and R. L. Hughson (Plenum, New York, 1989), pp. 353–359.

[16]C. P. Patil, K. B. Saunders, and B. M. Sayers, "Modeling the breath by breath variability in respiratory data," in *Respiratory Control: A Modeling Perspective*, Proc. Oxford Conference 1988, edited by G. D. Swanson, F. S. Grodins, and R. L. Hughson (Plenum, New York, 1988), pp. 343–352.

[17]G. C. Donaldson, "The chaotic behaviour of resting human respiration," Respir. Physiol. **88**, 313–321 (1992).

[18]P. Liang, J. J. Pandit, and P. A. Robbins, "Non-stationarity of breath-by-breath ventilation and approaches to modeling the phenomenon," in *Modeling and Control of Ventilation*, edited by S. J. G. Semple, L. Adams, and B. J. Whipp (Plenum, New York, 1995), pp. 117–121.

[19]G. Benchetrit, "Breathing pattern in humans: Diversity and individuality," Respir. Physiol. **122**, 123–129 (2000).

[20]C. Ozhasoglu and M. J. Murphy, "Issues in respiratory motion compensation during external-beam radiotherapy," Int. J. Radiat. Oncol., Biol., Phys. **52**, 1389–1399 (2002).

[21]G. C. Sharp, S. B. Jiang, S. Shimizu, and H. Shirato, "Prediction of respiratory tumour motion for real-time image-guided radiotherapy," Phys. Med. Biol. **49**, 425–440 (2004).

[22]S. S. Vedam, P. J. Keall, A. Docef, D. A. Todor, V. R. Kini, and R. Mohan, "Predicting respiratory motion for four-dimensional radiotherapy," Med. Phys. **31**, 2274–2283 (2004).

[23]S. Ahn *et al.*, "A feasibility study on the prediction of tumour location in the lung from skin motion," Br. J. Radiol. **77**, 588–596 (2004).

[24]R. I. Whyte, R. Crownover, M. J. Murphy, D. P. Martin, T. W. Rice, R. Rodebaugh, M. S. Weinhous, and Q.-T. Le, "Stereotactic radiosurgery for lung tumors: Preliminary report of a phase I trial," Ann. Thorac. Surg. **75**, 1097–1101 (2003).

[25]S. Haykin, *Kalman Filtering and Neural Networks* (Wiley Interscience, New York, 2001).