

## A LOCATION-MIXTURE AUTOREGRESSIVE MODEL FOR ONLINE FORECASTING OF LUNG TUMOR MOTION

BY DANIEL CERVONE<sup>\*</sup>, NATESH S. PILLAI<sup>\*,1</sup>, DEBDEEP PATI<sup>†</sup>,  
ROSS BERBECO<sup>‡,2</sup> AND JOHN HENRY LEWIS<sup>‡,3</sup>

*Harvard University<sup>\*</sup>, Florida State University<sup>†</sup>, Brigham and Women's Hospital,  
Dana-Farber Cancer Institute and Harvard Medical School<sup>‡</sup>*

Lung tumor tracking for radiotherapy requires real-time, multiple-step ahead forecasting of a quasi-periodic time series recording instantaneous tumor locations. We introduce a location-mixture autoregressive (LMAR) process that admits multimodal conditional distributions, fast approximate inference using the EM algorithm and accurate multiple-step ahead predictive distributions. LMAR outperforms several commonly used methods in terms of out-of-sample prediction accuracy using clinical data from lung tumor patients. With its superior predictive performance and real-time computation, the LMAR model could be effectively implemented for use in current tumor tracking systems.

**1. Introduction.** Real-time tumor tracking is a promising recent development in External Beam Radiotherapy (XRT) for the treatment of lung tumors. In XRT, a compact linear accelerator is used to deliver photon radiation to the tumor locations in a narrow beam, minimizing exposure to nearby healthy tissue. As the location of the lung tumor is in constant motion due to respiration, some patients who undergo this treatment are implanted with a small metal marker (known as a fiducial) at the location of a tumor. During XRT, X-ray imaging reveals the location of the fiducial, thus providing the desired target of the radiation beam. Tumor tracking is an advanced technology that minimizes normal tissue exposure by moving the radiation beam to follow the tumor position [D'Souza, Naqvi and Yu (2005), Rottmann, Keall and Berbeco (2013), Schweikard et al. (2000)]. However, there is a system latency of 0.1–1.0 seconds (depending on the equipment used) that causes the aperture of the radiation beam to lag behind the real-time location of the tumor. This latency is estimated empirically by comparing the motion history

---

Received September 2013; revised April 2014.

<sup>1</sup>Supported in part by the NSF Grant DMS-11-07070.

<sup>2</sup>Supported in part by R21CA156068 from the National Cancer Institute.

<sup>3</sup>Supported in part by Award Numbers RSCH1206 from the Radiological Society of North America.

*Key words and phrases.* Lung tumor tracking, external beam radiotherapy, nonlinear time series, mixture autoregressive process, time series motifs, likelihood approximation, multiple-step prediction.

of the fiducial and radiation beam aperture. For tumor tracking XRT to be successful, hardware and software system latencies must be overcome by the [Introduction](#) of a predictive algorithm.

As accurate radiotherapy is essential for both minimizing radiation exposure to healthy tissue and ensuring the tumor itself is sufficiently irradiated, the subject of predicting tumor motion to overcome the system latency has received a good deal of attention in the medical community. Any possible forecasting approach must provide  $k$ -step ahead predictive distributions in real time, where  $k$  is approximately equal to the system latency multiplied by the sampling frequency of the tumor tracking imagery. Real-time forecasting requires that a ( $k$ -step ahead) prediction be made before any further data on the tumor's motion has been recorded.

Statistical methods for tumor prediction in the literature include penalized linear models [e.g., [Sharp et al. \(2004\)](#) and many others], the Kalman filter [[Murphy, Isaakson and Jalden \(2002\)](#)], state-space models [[Kalet et al. \(2010\)](#)] and wavelets [[Ernst, Schlaefer and Schweikard \(2007\)](#)]; machine learning methods include kernel density estimation [[Ruan and Keall \(2010\)](#)], support vector regression [[Ernst and Schweikard \(2009\)](#), [Riaz et al. \(2009\)](#)] and neural networks [[Murphy and Dieterich \(2006\)](#), [Murphy, Isaakson and Jalden \(2002\)](#)]. All of these examples include simulations of out-of-sample prediction using real patient data in order to assess forecasting accuracy. Because predictive performance varies considerably from patient to patient and across different equipment configurations, of particular importance to the literature are comparisons of different prediction methods for the same set of patients with the same conditions for data preprocessing [[Ernst et al. \(2013\)](#), [Krauss, Nill and Oelfke \(2011\)](#), [Sharp et al. \(2004\)](#)]. While standard "off-the-shelf" time series forecasting models can be applied to lung tumor tracking, better predictive performance can be achieved with a model that explicitly incorporates the dynamics of respiratory motion.

We propose a novel time series model which we call a location-mixture autoregressive process (LMAR). A future observation ( $Y_n$ ) given the observed history of the time series is assumed to follow a Gaussian mixture,

$$(1.1) \quad Y_n | Y_{n-1}, Y_{n-2}, \dots \sim \sum_{j=1}^{d_n} \alpha_{n,j} \mathbf{N}(\mu_{n,j}, \sigma^2),$$

where  $\sum_{j=1}^{d_n} \alpha_{n,j} = 1$  and  $\mu_{n,j}$  is of the form

$$(1.2) \quad \mu_{n,j} = \tilde{\mu}_{n,j} + \sum_{l=1}^p \gamma_l Y_{n-l}.$$

We refer to this as a location-mixture autoregressive model because the autoregressive part of the component means,  $\sum_{l=1}^p \gamma_l Y_{n-l}$ , is the same for all  $j$ , and only the location parameter,  $\tilde{\mu}_{n,j}$ , changes across the components in (1.1). Our model differs from other time series models that yield mixture-normal conditional

distributions (e.g., the class of threshold autoregressive models [Tong and Lim (1980)], including Markov-switching autoregressive models [Hamilton (1989)] and the mixture autoregressive models of Wong and Li (2000)) in that  $\tilde{\mu}_{n,j}$  in (1.2) depends on an unknown subseries of the time series, at least  $p$  observations in the past. The mixture weights,  $\{\alpha_{n,j}\}$ , also depend on the entire history of the observed time series, and the number of mixture components in our model,  $d_n$ , increases with  $n$ .

Another noteworthy characteristic of our model is that all parameters in (1.1) are obtained from a single, unknown  $(p+1) \times (p+1)$  positive definite matrix. This parsimonious parameterization is motivated in part by the need for real-time parameter estimation and forecasting. Compared with other mixture autoregressive models, LMAR is simpler to fit and admits accurate closed-form expressions for  $k$ -step ahead predictive distributions. While the data application we consider shows the promise and appeal of the LMAR model, we believe a thorough treatment of its theoretical properties (a future endeavor) is necessary before the LMAR model is a viable “off-the-shelf” method for diverse data sets.

We motivate our model in the context of time series motifs, which offer a geometric interpretation of the components in our model. In general terms, motifs catalog recurring patterns in time series and are commonly used in data mining tasks for which a symbolic representation of a time series is useful, such as event detection and time series clustering or classification [Fu (2011), Lin et al. (2002), Tanaka, Iwamoto and Uehara (2005), Ye and Keogh (2009)]. For the purposes of forecasting, predictive state representations [Boots and Gordon (2011), Littman, Sutton and Singh (2002), Shalizi (2003)] categorize time series motifs not as sub-series of the observed data, but as equivalence classes of conditional predictive distributions.

Section 2 of this paper discusses the important features of the data we use and graphically motivates our model. Section 3 formally introduces the LMAR model and describes parameter estimation and forecasting using principled methods that are feasible in real time. Section 4 describes the procedure for comparing out-of-sample prediction error under our model with competing forecasting methods for tumor tracking, including the selection of tuning parameters. The results of this comparison are discussed in Section 5, and Section 6 summarizes and points out future directions.

**2. Tumor tracking data.** We have data on 11 patients treated at the Radiation Oncology Clinic at the Nippon Telegraph and Telephone Corporation Hospital in Sapporo, Japan. A detailed discussion of the conditions and instruments involved in the data acquisition is available in Berbeco et al. (2005). The data is derived from observations of the position of gold fiducial markers implanted into the tumors of lung cancer patients. The marker position is determined via stereoscopic X-ray imaging conducted at 30 Hz. In each of the two stereoscopic images, the marker position is automatically detected using thresholding and edge detection.

The position of the marker in these two images is used to triangulate its position in 3D space relative to the radiation beam. Data consists of tumor positions measured over one or multiple days of radiotherapy treatment delivery (range 1–12), and for multiple sequences on each day, denoted *beams*. In our data set, there are a total of 171 such distinct sequences, with lengths varying from 637 observations (about 21 seconds at 30 observations per second) to 8935 observations (about 5 minutes).

Note that this paper focuses on within-beam forecasting—that is, each beam is treated independently and there is no information sharing between patients or within different beams from the same patient. Developing methodology for combining prediction models from distinct time series (both within and across patients) is an important area for further research.

**2.1. Features of the data.** Each observation in each sequence is a point in  $\mathbb{R}^3$ , representing the real-time 3D location of the lung tumor. The  $X$  axis is the lateral–medial (left–right) direction, the  $Y$  axis is superior–inferior, and the  $Z$  axis is anterior–posterior, with all measurements in millimeters.<sup>4</sup> Figure 1 shows the motion in each dimension during the first 100 seconds of a particular observation




FIG. 1. Sample time series of 3D locations of lung tumor. The  $X$  axis is the lateral–medial (left–right) direction,  $Y$  axis superior–inferior, and  $Z$  axis anterior–posterior.

<sup>4</sup>The origin is set to the isocenter, which is the center of rotation for the linear accelerator axis motions. During treatment, the patient is positioned so that this coincides with the centroid of the region being treated. However, there is uncertainty in determining this point, so the data is best thought of as relative tumor motion on each day.

TABLE 1

*Summary statistics for the first principal component of respiratory trace data, at the patient level*

Patient	Total beams	Total time (s)	Amplitude (mm)		Period(s)	
			Mean	SD	Mean	SD
1	4	212.27	14.57	6.98	3.66	1.16
2	2	136.87	13.74	1.84	3.89	1.06
3	2	80.93	9.84	3.16	3.97	0.56
4	38	2502.67	8.86	1.35	2.88	0.31
5	26	2769.33	7.90	1.66	3.61	0.68
6	28	2471.93	10.07	2.51	2.58	0.55
7	11	1661.37	9.66	2.41	5.05	1.09
8	8	832.80	14.38	4.02	3.15	1.18
9	15	2599.90	11.45	1.61	3.09	0.41
10	15	3497.67	14.88	3.65	3.77	0.64
11	22	3674.77	21.81	5.05	3.38	0.52

sequence. As expected with respiratory motion, the pattern is approximately periodic, with inhalation closely corresponding to decreasing values in the  $Y$  direction. However, the amplitude of each breath varies considerably (in Figure 1 the variation seems periodic, though this is not a typical feature of the data). The curves undergo gradual baseline location shifts and, while it may not be visually discerned from Figure 1, it is common for respiratory cycles to change periodicity, either sporadically or gradually over time. Table 1 shows the variability in period and amplitude of the respiratory traces, both within and between patients.

Due to the extremely high correlations between series of observations from different dimensions, it is useful to consider a lower-dimensional representation of the 3D process. Transforming each 3D sequence into orthogonal components using principal component analysis (PCA) loads the periodic respiratory dynamics onto the first component, representing about 99% of the total variance in the 3D data. The last two principal components still exhibit some periodic behavior (see Figure 2), but the signal is weak relative to the noise.<sup>5</sup> In addition to dimension reduction and useful interpretability, the PCA transformation prevents any loss of statistical efficiency if models are fit independently for each component. [Ruan and](#)

<sup>5</sup>A referee pointed out that while the first principal component gives the linear combination of the 3D data with maximum variance, it is not necessarily the most *forecastable* linear combination. Alternative linear transformations (e.g., forecastable components [[Goerg \(2013a\)](#)]) may load additional periodic features to the first component than we observe with PCA. In choosing an appropriate transformation, the goal is to find an orthogonal basis in which componentwise predictions have the smallest error when transformed back to the original basis. We do not explore this issue here; however, one advantage in using the first principal component is that the signal-to-noise ratio will be high, allowing for forecast procedures that are not well suited for measurement error in the observed data.

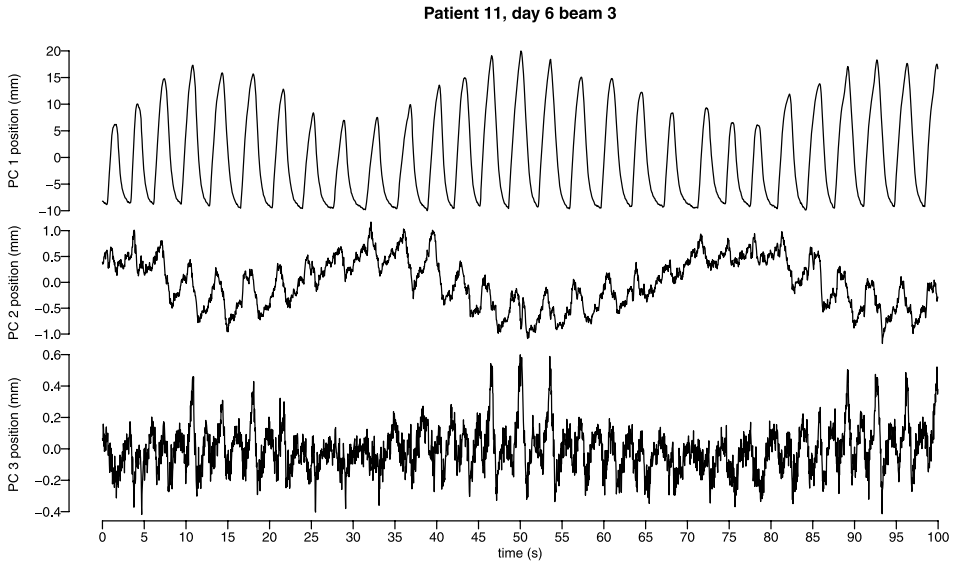


FIG. 2. *Time series of principal components. Components 2 and 3 exhibit periodic behavior, but with much smaller magnitude.*

Keall (2010) compared independent-component prediction before and after PCA using kernel density estimation, finding smaller 3D root mean squared prediction error when using the PCA-transformed data for prediction. When comparing several algorithms for predicting lung tumor motion, both Ernst et al. (2013) and Krauss, Nill and Oelfke (2011) used the principal components, then transformed their predictions to the original linear basis of the data.

For the remainder of this study, we focus on modeling the first principal component only, as it encodes such a large portion of the system dynamics. In clinical implementation, we would forecast independently on each orthogonal component and transform back to the original linear basis in order to inform the location of the radiation treatment beam.

*2.2. Time series motifs for forecasting: A graphical example.* Because the data are quasi-periodic, it is useful to look at short patterns that recur at possibly irregular intervals, which we call motifs (we provide a more rigorous definition of time series motifs in Section 3.2). Figure 3 highlights different motifs in the first principal component at the end of the exhale (start of the inhale) for a particular observation sequence. The highlighted areas appear to be heartbeats, which affect the location of the tumor differently depending on the real-time location of the tumor relative to the heart.

Observing repeated patterns within each time series in the data suggests a modeling/prediction framework that leverages this structure. In general, if the recent past of the time series resembles a motif we have observed previously in the data,




FIG. 3. Recurring patterns (coded by color and line type) in the first principal component of patient 10, day 1, beam 3. Areas boxed by lines of the same color/line type resemble one another. The behavior highlighted in these motifs is most likely caused by the patient's heartbeat.

then the shape of this motif should inform our predictions of future observations; this idea is formalized through predictive state representations [Littman, Sutton and Singh (2002), Shalizi (2003)]. For a graphical illustration, consider predicting 0.4 s (12 steps) ahead for the first principal component of the curve displayed in Figure 2. We have observed 100 seconds of the process, and it appears as though we have just observed the start of the exhale; the current observation at time  $t = 100$  seconds, as well as the previous 12 observations, are colored orange in Figure 4. Colored in black are segments earlier in the time series that resemble the current motif (specifically, we highlighted subseries of length 13 where the tenth point has the largest magnitude, and the 11th–13th points are decreasing).

To predict future observations, we can incorporate the points immediately succeeding the endpoints of black motifs. Figure 5 shows these trajectories (in gray), and the actual current trajectory of the process is shown in orange, with a point giving the value 0.4 s in the future. The gray curves provide reasonable forecasts for the future evolution of the time series and, indeed, the actual future value is close to where these trajectories predict.




FIG. 4. The most recent 0.43 s (13 observations) are in black. The thicker orange segments share similar local history.

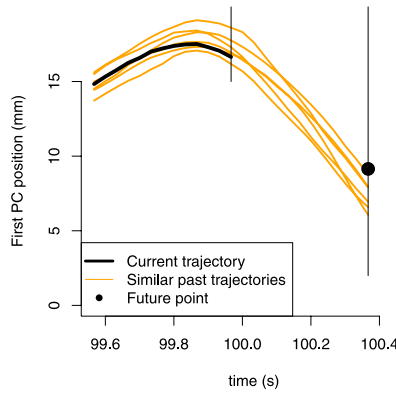


FIG. 5. The recent history of the process (thick black line) instantiates a motif. Previous instances of this motif, and their subsequent evolutions, are in orange and provide reasonable predictions for future points (black dot).

Our model, formally introduced in Section 3, implements the forecasting approach sketched in this subsection using an autoregressive model for the data-generating process.

**3. Location-mixture autoregressive processes.** Here, we define the LMAR process and provide computationally efficient algorithms for parameter estimation and  $k$ -step ahead forecasting. To establish terminology, we denote a *time series* as an ordered sequence of real numbers  $\{Y_i \in \mathbb{R}, i = 0, \pm 1, \pm 2, \dots\}$  measured at regular, equally spaced intervals. Also, a *subseries* of length  $p + 1$  is a subset of a time series  $\{Y_i, i = 0, \pm 1, \dots\}$  comprised of consecutive observations,  $Y_i, Y_{i+1}, \dots, Y_{i+p}$ . For notational ease, we will denote the subseries as  $Y_{i:(i+p)}$  or, equivalently,  $Y_{i+0:p}$ .

3.1. *A model for the data-generating process.* Let  $\{Y_i, i = -m, \dots, n\}$  be a time series. Also, assume  $\Sigma$  is a  $(p + 1) \times (p + 1)$  symmetric, nonnegative definite matrix, where  $\Sigma_{11}$  is the upper-left  $p \times p$  submatrix,  $\Sigma_{22}$  is the single bottom-right element, and  $\Sigma_{21}$  and  $\Sigma_{12}$  are the respective off-diagonal row and column vectors.  $p$  is assumed to be fixed and known. For notational ease, let  $\gamma = \Sigma_{11}^{-1} \Sigma_{12}$ ,  $\sigma^2 = \Sigma_{22} - \gamma' \Sigma_{12}$ , and  $\mathcal{J}_i = \{p + 1, \dots, i + m - p\}$ . Last, let

$$V_{ij} = \begin{pmatrix} Y_{i-p} - Y_{i-j-p} \\ \vdots \\ Y_{i-2} - Y_{i-j-2} \\ Y_{i-1} - Y_{i-j-1} \end{pmatrix}.$$



As in (1.1), we assume that the distribution of  $Y_i$  given  $Y_{-m}, \dots, Y_{i-1}$  is a normal mixture,

$$(3.1) \quad Y_i | Y_{(-m):(i-1)} \sim \sum_{j \in \mathcal{J}_i} \alpha_{i,j} \mathbf{N}(\mu_{i,j}, \sigma^2)$$

where  $\alpha_{i,j} = \frac{\exp(-(1/2) V'_{ij} \Sigma^{-1}_{11} V_{ij})}{\sum_{l \in \mathcal{J}_i} \exp(-(1/2) V'_{il} \Sigma^{-1}_{11} V_{il})}$  and  $\mu_{i,j} = Y_{i-j} + \gamma' V_{ij}$ .

The model in (3.1) defines the location-mixture autoregressive process with parameter  $\Sigma$  [abbreviated LMAR( $\Sigma$ )]. We can recognize the location-mixture form originally given in (1.1) by writing  $\mu_{i,j} = \tilde{\mu}_{i,j} + \sum_{l=1}^p \gamma_l Y_{i-l}$ , where

$$(3.2) \quad \tilde{\mu}_{i,j} = Y_{i-j} - \sum_{l=1}^p \gamma_l Y_{j-l}$$

and  $(\gamma_p \gamma_{p-1} \dots \gamma_1)' = \gamma$ . Thus, the distribution for  $Y_i | Y_{(-m):(i-1)}$  is a normal mixture with  $|\mathcal{J}_i|$  different mean components—each sharing a common autoregressive component but different location parameter—equal variance across components ( $\sigma^2$ ) and data-driven mixture weights ( $\alpha_{i,j}$ ). We assume (3.1) for all  $i \geq 0$ , but we do not make any distributional assumptions about  $Y_{(-m):(-1)}$ .

As  $\Sigma$  parameterizes the entire mixture distribution, the component means and mixture weights are linked through a common parameter which encourages self-similarity in the data-generating process. If two subseries  $Y_{(i-p):(i-1)}$  and  $Y_{(i-p-j):(i-1-j)}$  resemble one another in that  $V'_{ij} \Sigma^{-1}_{11} V_{ij}$  is small, then we have a large weight on the mixture component with mean  $Y_{i-j} + \gamma' V_{ij}$ . This means that the next observation of the process,  $Y_i$ , is centered near a previous value of the series  $Y_{i-j}$  inasmuch as the subseries of observations preceding  $Y_i$  and  $Y_{i-j}$  have a similar shape. Simply put, if  $Y_i$  and  $Y_{i-j}$  are preceded by similar values, then the components of  $V_{ij}$  will be close to 0. This drives up the mixture weight  $\alpha_{i,j}$ , implying the mean of  $Y_i$  will be close to  $\mu_{i,j}$  (which itself is close to  $Y_{i-j}$ ).

The dimension of  $\Sigma$ ,  $p + 1$ , can in principle be chosen using standard model selection methods (e.g., Bayes factors), though if the goal of fitting a LMAR model is prediction, we recommend cross-validation or hold-out testing for choosing  $p$ . For quasi-periodic time series, a reasonable choice for  $p$  might be anywhere between one-tenth and one-third of the average number of observations per period. Larger values of  $p$  increase the computational load in estimating  $\Sigma$  while favoring sparser component weights.

The model (3.1) specifies the role of time series motifs in the data-generating process, which was informally discussed in Section 2.2. To illustrate this, we introduce a latent variable  $M_i$  that takes values in  $\mathcal{J}_i$ , such that for all  $j \in \mathcal{J}_i$ ,

$$(3.3) \quad \mathbb{P}(M_i = j | Y_{(-m):(i-1)}) \propto \exp(-\frac{1}{2} V'_{ij} \Sigma^{-1}_{11} V_{ij}).$$

Then, given  $M_i = j$ , we induce the same distribution for  $Y_i$  as in (3.1) by assuming

$$(3.4) \quad Y_i | [M_i = j, Y_{(-m):(i-1)}] \sim N(Y_{i-j} + \gamma' V_{ij}, \sigma^2).$$

Expression (3.4) can be used to define a motif relation: each subseries of length  $(p + 1)$  is a *motif*, and  $Y_{(i-p):i}$  is an *instance* of motif  $Y_{(i-p-j):(i-j)}$  if  $M_i = j$  [thus yielding (3.4)]. We denote this by writing

$$(\text{motif}) Y_{(i-p-j):(i-j)} \rightarrow Y_{(i-p):i} (\text{instance}).$$

Note that our indexing set  $\mathcal{J}_i$  is defined in such a way that instances of a particular motif cannot overlap (share a common component  $Y_j$ ) with the motif itself.

Our definition of motifs is atypical of the literature for data mining tasks [Lin et al. (2002)] and predictive state representations of time series [Littman, Sutton and Singh (2002)]. For instance, the relationship that instantiates motifs (notated  $\rightarrow$ ) is not symmetric and is not an equivalence relation; for this reason we have defined a motif instance distinctly from a motif. Also, we define motifs as observed subseries of the data and motif instances as latent states (we do not observe  $M_i$ ). For most data mining tasks, time series motifs represent an equivalence class of observed subseries of the data (possibly transformed) [Fu (2011)], whereas predictive state representations of time series treat motifs as latent equivalence classes of predictive distributions [Shalizi (2003)].

However, our definition of motifs preserves the interpretation of geometric similarity we sketched in Section 2.2. From (3.3), we have  $M_i = j$  (meaning  $Y_{(i-p-j):(i-j)} \rightarrow Y_{(i-p):i}$ ) with high probability if  $V_{ij}$  is small with respect to the  $\Sigma_{11}^{-1}$  inner product norm. Our model thus expects a subseries that is an instance of a particular motif to be close to the motif, and  $\Sigma$  parameterizes this distance metric.

3.2. *Comparison with other mixture autoregressive processes.* We may compare the LMAR( $\Sigma$ ) to a general form of regime-switching autoregressive models, for which we can write the distribution function of  $Y_i$  conditional on all available history of the process  $Y_{(-m):(i-1)}$  as

$$(3.5) \quad F(y | Y_{(-m):(i-1)}) = \sum_{j=1}^d \alpha_{i,j} \Phi \left( \frac{y - (\beta_{0,j} + \sum_{l=1}^p \beta_{l,j} Y_{i-l})}{\sigma_j} \right),$$

where  $\sum_{j=1}^d \alpha_{i,j} = 1$  for all  $i$  and  $\Phi$  denotes the standard normal CDF. Models satisfying (3.5) can be represented in the framework of threshold autoregressive models [Tong (1978), Tong and Lim (1980); see Tong (1990) for a book-length treatment], which represent (3.5) using an indicator series  $\{M_i\}$  taking values on  $\{1, \dots, d\}$ , such that

$$(3.6) \quad Y_i = \beta_{0,M_i} + \sum_{l=1}^p \beta_{l,M_i} Y_{i-l} + \sigma_{M_i} \varepsilon_i,$$

where  $\{\varepsilon_i\}$  are i.i.d. standard normals. Generally,  $M$  is not observed, although there are notable exceptions such as the self-exciting threshold AR model of **Tong and Lim (1980)**.

A canonical model of this form is the mixture autoregressive model of **Le, Martin and Raftery (1996)** and **Wong and Li (2000)**, which assumes  $\{M_i\}$  are i.i.d. and independent of  $Y$ . Another special case of (3.6) is when  $M$  is a Markov chain, such as in the Markov-switching autoregressive models of **Hamilton (1989)** and **McCulloch and Tsay (1994)**. More general stochastic structure for  $M$  is considered by **Lau and So (2008)**, as well as in mixture-of-experts models in the machine learning literature [**Carvalho and Tanner (2005)**]. These models seem favorable over the mixture autoregressive models of **Wong and Li (2000)** when the data is seasonal or quasi-periodic, as is the case with the time series we consider.

The LMAR( $\Sigma$ ) process differs from (3.5) in that the mixture means, following (3.1)–(3.2), are given by

$$\mu_{i,j} = \tilde{\mu}_{i,j} + \sum_{l=1}^p \gamma_l Y_{i-l} = Y_{i-j} + \sum_{l=1}^p \gamma_l Y_{i-l} - \sum_{l=1}^p \gamma_l Y_{j-l},$$

instead of  $\mu_{i,j} = \beta_{0,j} + \sum_{l=1}^p \beta_{l,j} Y_{i-l}$  as in (3.5). Thus, for LMAR( $\Sigma$ ), the autoregressive coefficients ( $\gamma$ ) are fixed, and the normal-mixture form of the conditional distribution is induced by a location shift that is a function of a random subseries of past observations,  $\tilde{\mu}_{i,j}$ . The normal-mixture form of (3.5), however, is induced by a mixture distribution for autoregressive coefficients of the same lagged values of the time series. The mixture weights of the LMAR( $\Sigma$ ) process are also strongly data driven, depending on the entire history of the process. Unlike many forms of mixture autoregressive models, there is no prior distribution or conditional dependence structure assumed for  $M$ ; the distribution of  $M$  is supplied entirely by the data.

Another key difference is that LMAR( $\Sigma$ ) does not assume a fixed number of mixture components, as is clear from (3.1). But because the same autoregressive coefficient vector ( $\gamma$ ) parameterizes all mean components  $\mu_{i,j}$ , we actually have a much smaller parameter space than all the instances of (3.5) cited above, which include the parameters for the mixture components ( $d$  vectors of length  $p + 1$  for the means) as well as for the distribution of  $M$ . A small parameter space is advantageous in the context of our data application, as it facilitates rapid updating. Also, time constraints will not allow for any goodness-of-fit or model selection procedures for choosing structural parameters such as  $d$  or  $p$  in (3.5), or structural parameters for  $M$ . The only structural parameter in the LMAR( $\Sigma$ ) model is  $p$ , and in our analysis of this data set we found that predictive distributions were quite stable for different choices of  $p$ .

The most important distinction of the LMAR( $\Sigma$ ) model is the existence of good approximations for  $k$ -step ahead predictive distributions, for  $k \leq p$ , which are given in Section 3.4. Closed-form predictive distributions for  $k > 1$  are not available for many models of the form (3.5) [the exception is the Markov-switching

autoregressive models of Hamilton (1989); for a discussion see Krolzig (2000)]. Wong and Li (2000) recommended Monte Carlo estimates of  $k$ -step ahead predictive distributions, although Boshnakov (2009) found for them a closed-form representation as a normal mixture. Calculating the mixture component parameters for moderate  $k$ , however, is quite laborious. For the general model (3.5), De Gooijer and Kumar (1992) discussed the difficulty in  $k$ -step ahead forecasting and questioned whether predictive performance is improved over classes of linear time series models [also see Tong and Moenanaddin (1988) for a discussion of the robustness of medium-to-long range forecasts using threshold autoregressive models].

3.3. *Parameter estimation.* In order to be able to adjust radiotherapy treatments in real time to the patient’s breathing pattern, we seek estimation procedures that are fast enough to run online (in less than a few seconds). As a general rule, this favors approximate closed-form solutions to estimating equations over exact numerical or Monte Carlo methods. To estimate  $\Sigma$ , which is the only unknown parameter of this model, we take a conditional likelihood approach based on the conditional distribution  $Y_{0:n}|Y_{(-m):(-1)}$ . We assume the full-data likelihood can be written as

$$L(\psi, \Sigma) = L_1(\psi, \Sigma)L_2(\Sigma),$$

where  $L_1(\psi, \Sigma) \propto \mathbb{P}(Y_{(-m):(-1)}; \psi, \Sigma)$  and  $L_2(\Sigma) \propto \mathbb{P}(Y_{0:n}|Y_{(-m):(-1)}; \Sigma)$ . The distribution of the first  $m$  observations, and thus  $L_1$ , is left unspecified, and all information for  $\Sigma$  comes from  $L_2$ . If  $L_1$  depends on  $\Sigma$ , there will be some loss of efficiency when using only  $L_2$  for inference versus the complete-data likelihood, though under mild conditions the maximum conditional likelihood estimate is consistent and asymptotically efficient [Kalbfleisch and Sprott (1970)].

The conditional likelihood,  $L_2(\Sigma)$ , can be written as

$$(3.7) \quad L_2(\Sigma) = \prod_{i=0}^n \frac{1}{\sigma} \left[ \sum_{j \in \mathcal{J}_i} \exp\left(-\frac{1}{2\sigma^2}(Y_i - Y_{i-j} - \gamma'V_{ij})^2\right) \times \left( \frac{\exp(-V'_{ij}\Sigma^{-1}V_{ij}/2)}{\sum_{l \in \mathcal{J}_i} \exp(-V'_{il}\Sigma^{-1}V_{il}/2)} \right) \right].$$

To maximize (3.7), we augment the data to  $\{Y_{0:n}, M_{0:n}\}$ , with  $M_i$  as in (3.3). This invites the use of the Expectation–Maximization (EM) algorithm [Dempster, Laird and Rubin (1977)] to estimate  $\Sigma$ . The augmented-data (complete-data) conditional likelihood is

$$L_{2,\text{com}}(\Sigma) = \prod_{i=0}^n \frac{1}{\sigma} \prod_{j \in \mathcal{J}_i} \left[ \exp\left(-\frac{1}{2\sigma^2}(Y_i - Y_{i-j} - \gamma'V_{ij})^2\right) \times \left( \frac{\exp(-V'_{ij}\Sigma^{-1}V_{ij}/2)}{\sum_{l \in \mathcal{J}_i} \exp(-V'_{il}\Sigma^{-1}V_{il}/2)} \right) \right]^{1[M_i=j]}.$$

This can be simplified further. Let  $W'_{ij} = (V'_{ij}Y_i - Y_{i-j})$ , and recalling the notation for  $\sigma$  and  $\gamma$ , we have

$$(3.8) \quad L_{2,\text{com}}(\Sigma) = \prod_{i=0}^n \frac{\exp(-(1/2) \sum_{j \in \mathcal{J}_i} \mathbf{1}[M_i = j] W'_{ij} \Sigma^{-1} W_{ij})}{\sigma \sum_{l \in \mathcal{J}_i} \exp(-V'_{il} \Sigma_{11}^{-1} V_{il}/2)}.$$

The term  $\sum_{l \in \mathcal{J}_i} \exp(-V'_{il} \Sigma_{11}^{-1} V_{il}/2)$  can be viewed as an approximation of a Gaussian integral; if we assume that, for all  $i$ ,  $\{V_{il}, l \in \mathcal{J}_i\}$  resemble  $|\mathcal{J}_i|$  i.i.d. draws from some distribution  $V \sim N(0, \Omega)$ , then we have

$$(3.9) \quad \begin{aligned} & \sum_{l \in \mathcal{J}_i} \exp(-V'_{il} \Sigma_{11}^{-1} V_{il}/2) \\ & \approx |\mathcal{J}_i| \int \exp(-V' \Sigma_{11}^{-1} V/2) \frac{\exp(-V' \Omega^{-1} V/2)}{(2\pi)^{p/2} |\Omega|^{1/2}} dV \\ & = |\mathcal{J}_i| \left( \frac{|\Sigma_{11}^{-1} + \Omega^{-1}|}{|\Omega|} \right)^{1/2} \\ & = |\mathcal{J}_i| \left( \frac{|\Sigma_{11}|}{|\Sigma_{11} + \Omega|} \right)^{1/2}. \end{aligned}$$

Noting that  $\sigma |\Sigma_{11}|^{1/2} = |\Sigma|^{1/2}$ , and ignoring multiplicative constants, we arrive at an approximate augmented-data conditional likelihood:

$$L_{2,\text{com}}(\Sigma) \approx \left( \frac{|\Sigma_{11} + \Omega|}{|\Sigma|} \right)^{(n+1)/2} \exp\left(-\frac{1}{2} \sum_{i=0}^n \sum_{j \in \mathcal{J}_i} \mathbf{1}[M_i = j] W'_{ij} \Sigma^{-1} W_{ij}\right).$$

Typically  $\Sigma_{11} \ll \Omega$ , meaning

$$\begin{aligned} \partial(\log(|\Sigma_{11} + \Omega|) - \log(|\Sigma|)) &= \text{Tr}((\Sigma_{11} + \Omega)^{-1} \partial \Sigma_{11}) - \text{Tr}(\Sigma^{-1} \partial \Sigma) \\ &\approx -\text{Tr}(\Sigma^{-1} \partial \Sigma) \end{aligned}$$

as  $\partial \log(|\Sigma|)$  dominates  $\partial \log(|\Sigma_{11} + \Omega|)$ . This justifies the approximation  $\log(|\Sigma_{11} + \Omega|) - \log(|\Sigma|) \approx -\log(|\Sigma|)$  in the augmented-data conditional log-likelihood, as it will admit nearly the same maximizer. Thus, we have

$$(3.10) \quad \begin{aligned} \log(L_{2,\text{com}}(\Sigma)) &\approx -\frac{n+1}{2} \log(|\Sigma|) \\ &\quad - \frac{1}{2} \sum_{i=0}^n \sum_{j \in \mathcal{J}_i} \mathbf{1}[M_i = j] W'_{ij} \Sigma^{-1} W_{ij}. \end{aligned}$$

While (3.10) is much easier to work with than the logarithm of the exact conditional likelihood (3.8), the assumptions of this approximation are somewhat tenuous. Under this model (3.1), both conditional and marginal distributions of observations at each time point follow a normal mixture, meaning for  $l$  randomly

chosen from  $\mathcal{J}_i$ , we have a difference of normal mixtures (itself a normal mixture) for  $V_{il}$ , instead of i.i.d. normals as (3.9) suggests. We nevertheless proceed with approximation (3.10) in place of (3.8), noting that convergence of the EM algorithm needs to be more carefully monitored in this instance.

At each iteration of the EM algorithm, we maximize the so-called  $Q$  function:

$$\begin{aligned}
 (3.11) \quad Q^{(t)}(\Sigma) &= \mathbb{E}_{\Sigma^{(t)}}[\log(L_{2,\text{com}}(\Sigma))|Y] \\
 &\approx -\frac{n+1}{2} \log(|\Sigma|) - \frac{1}{2} \sum_{i=0}^n \sum_{j \in \mathcal{J}_i} \omega_{ij} W'_{ij} \Sigma^{-1} W_{ij},
 \end{aligned}$$

with  $\Sigma^{(t)} = \text{argmax}(Q^{(t-1)}(\Sigma))$  and  $\omega_{ij} = \mathbb{E}_{\Sigma^{(t)}}[\mathbf{1}[M_i = j]|Y]$ . Clearly,

$$\omega_{ij} = \frac{\exp(-W'_{ij}[\Sigma^{(t)}]^{-1} W_{ij}/2)}{\sum_{l \in \mathcal{J}_i} \exp(-W'_{lj}[\Sigma^{(t)}]^{-1} W_{lj}/2)}.$$

The maximizer of (3.11) can be found in closed form as a weighted sample covariance matrix,

$$(3.12) \quad \Sigma^{(t+1)} = \frac{1}{n+1} \sum_{i=0}^n \sum_{j \in \mathcal{J}_i} \omega_{ij} W_{ij} W'_{ij}.$$

Again, due to several different approximations used in maximizing the original conditional likelihood (3.7), it is necessary to monitor the convergence to a suitable (if slightly suboptimal) solution, as the log-likelihood is not guaranteed to increase at each iteration.

3.4. *A prediction model for fast implementation.* Exact closed-form expressions for  $k$ -step ahead predictive distributions are not available for the model (3.1). Because of the need for real-time forecasting of many steps ahead, we explore approximations to  $k$ -step ahead predictive distributions that are available in closed form. An immediate approach to doing so is to explore whether the approximate complete-data conditional log-likelihood used for inference (3.10) corresponds to a probabilistic model (perhaps misspecified) that admits closed-form predictive distributions. In other words, if the previous section derives an approximate log-likelihood (3.10) from an exact model (3.1), here we treat (3.10) as exact and explore corresponding approximate models.

Let  $Z_i = (Y_{i-p} \cdots Y_{i-1} Y_i)'$  for  $0 \leq i \leq n$ . Since  $W_{ij} = Z_i - Z_j$ , we may arrive at the likelihood expression (3.10) by assuming  $Z_i \sim N(Z_{i-M_i}, \Sigma)$  independently. This is obviously a misspecification, since for any  $k \leq p$ ,  $Z_i$  and  $Z_{i+k}$  contain duplicate entries and thus cannot be independent. But assuming the  $\{Z_i\}$  independent, and further assuming  $\mathbb{P}(M_i = j) = 1/|\mathcal{J}_i|$  independently for all  $i$ , we can

write the (conditional) likelihood for an independent multivariate normal mixture model, denoted  $L_a$  to distinguish from  $L_{2,\text{com}}$ :

$$(3.13) \quad L_a(\Sigma) = \prod_{i=0}^n \prod_{j \in \mathcal{J}_i} \left[ |\Sigma|^{-1/2} \exp\left(-\frac{1}{2} W'_{ij} \Sigma^{-1} W_{ij}\right) \right]^{\mathbf{1}_{\{M_i=j\}}}$$

Indeed, we see that  $L_a(\Sigma)$  is equal to the approximation of  $L_{2,\text{com}}(\Sigma)$  given in (3.10). Thus, the misspecified independent mixture model for  $Z_i$  yields the same likelihood ( $L_a$ ) as the approximation to  $L_2$ , the exact (conditional) likelihood corresponding to the data-generating process. Also, recall that  $M_i = j$  denotes  $Z_i$  as an instance of motif  $Z_j$ . The implied relation in (3.13) is that

$$(3.14) \quad Z_j \rightarrow Z_i \quad \text{if } Z_i|Z_j \sim N(Z_j, \Sigma)$$

and, indeed, this relation is closely connected to the one defined in (3.4). They appear equivalent, as (3.4) is recovered by assuming  $Z_i|Z_j \sim N(Z_j, \Sigma)$ , and then considering the conditional distribution  $Y_i|Y_{(-m):(i-1)}$ . However, for (3.14) to hold for all  $i$  requires the impossible assumption of  $Z_i$  being independent of  $Z_{i-1}$ , while the relation in (3.4) does not.

The corresponding  $Q$  function for this complete-data conditional likelihood (3.13) is

$$Q_a^{(t)}(\Sigma) = \sum_{i=0}^n -\frac{1}{2} \log(|\Sigma|) - \frac{1}{2} \sum_{j \in \mathcal{J}_i} \mathbb{E}_{\Sigma^{(t)}}[\mathbf{1}_{\{M_i = j\}}|Z] W'_{ij} \Sigma^{-1} W_{ij}$$

Working  $\mathbb{E}_{\Sigma^{(t)}}[\mathbf{1}_{\{M_i = j\}}|Z] = \omega_{ij}$ , we see that  $Q_a^{(t)}$  is identical to  $Q^{(t)}$  given in (3.11), confirming that the “same”  $\Sigma$  parametrizes both the original data-generating process assumed in (3.1) and its degenerate approximation that we will use to make predictions in (3.13). We may also think of maximizing  $Q$  as inferring motif instances given by the relation (3.14), that is, minimizing a distance metric.

The independent multivariate mixture distribution of  $\{Z_i\}$  considered here very easily provides  $k$ -step predictive distributions for  $k \leq p$ . If we have observed the process up to  $Y_n$  and wish to predict  $Y_{n+k}$ , this is equivalent to having observed  $Z$  up to  $Z_n$  and wishing to predict the last component of  $Z_{n+k}$ . Having observed  $Z_n$  completely, we have observed the first  $p - k + 1$  components of  $Z_{n+k}$ , and thus by the (misspecified) independence assumed for  $\{Z_i\}$ , the predictive distribution for  $Y_{n+k}$  depends only on these  $p - k + 1$  values. To write this, we denote  $\tilde{Z}_n^k$  as the first  $p - k + 1$  components of  $Z_{n+k}$  (or the last  $p - k + 1$  components of  $Z_n$ ); also, let  $\tilde{W}_{nj}^k = \tilde{Z}_n^k - \tilde{Z}_j^k$  and partition  $\Sigma$  into  $\Sigma_{11}^k$  as the upper-left  $(p - k + 1) \times (p - k + 1)$  submatrix,  $\Sigma_{22}^k$  as the single bottom-right element (thus identical to  $\Sigma_{22}$ ), and  $\Sigma_{12}^k, \Sigma_{21}^k$  accordingly. Then we have

$$(3.15) \quad Y_{n+k}|Y_{(-m):n} \sim \sum_{j \in \mathcal{J}_{n+k}} \alpha_j^k N(\mu_j^k, \sigma_k^2)$$

where:

- $\alpha_j^k = \mathbb{P}(M_{n+k} = j | \tilde{Z}_n^k) \propto \exp(-(\tilde{W}_{nj}^k)'[\Sigma_{11}^k]^{-1} \tilde{W}_{nj}^k/2),$
- $\mu_j^k = Y_{n+k-j} + \Sigma_{21}^k[\Sigma_{11}^k]^{-1} \tilde{W}_{nj}^k,$
- $\sigma_k^2 = \Sigma_{22}^k - \Sigma_{21}^k[\Sigma_{11}^k]^{-1} \Sigma_{12}^k.$

In terms of motifs, these predictive distributions result from considering the most recent subseries of the data of length  $p - k + 1$  as a partially observed motif instance,  $Z_{n+k}$ , which includes the future observation we wish to predict,  $Y_{n+k}$ . Using the implied motif relation in (3.14), we infer both the motif for which  $Z_{n+k}$  is an instance and derive predictive distributions using simple multivariate normal properties (3.15).

Of course, we use  $\hat{\Sigma}$ , the solution to (3.12), in place of  $\Sigma$  in the above expressions, acknowledging that the resulting predictive distributions fail to account for the uncertainty in our estimate of  $\Sigma$ .

3.5. *Interpreting  $\hat{\Sigma}$ .* Figure 6 shows estimates  $\hat{\Sigma}$  from two of the time series in our data. Interpreting these as covariance matrices, we see relatively high correlations across components, favoring instantiating motifs where the difference between the motif instance and the original motif is roughly linear with a slope near 0. Also, the diagonal terms are decreasing from top to bottom, implying that more weight is given to the most recent components of the observed time series when inferring the latent motif instance and making predictions.

**4. Evaluating out-of-sample prediction error with competing methods.**

We compare out-of-sample prediction performance for tumor tracking using the LMAR( $\Sigma$ ) model with three methods that are straightforward to implement and

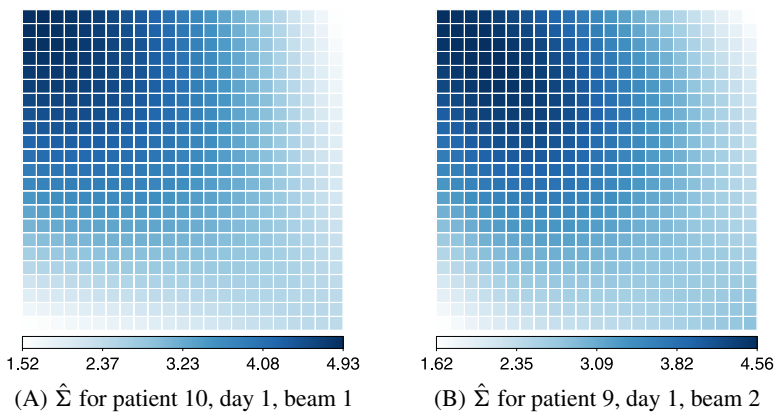


FIG. 6. Illustration of  $\hat{\Sigma}$  for two of the time series in our data, using  $p = 22$ . Note that the color scale differs slightly for each figure.



provide real-time forecasts. Neural networks (4.1) and ridge regression (4.2) both compare favorably to alternative methods with regards to prediction accuracy [Krauss, Nill and Oelfke (2011), Sharp et al. (2004)]. LICORS (4.3) is a non-parametric and nonregression forecasting method based on predictive state representations of the time series [Goerg and Shalizi (2012, 2013)]. For each method, Sections 4.4–4.6 discuss data preprocessing and computational considerations relevant for real-time tumor tracking.

4.1. *Feedforward neural networks.* Multilayer feedforward neural networks with at least one hidden layer have been used to forecast lung tumor motion by Murphy, Isaakson and Jalden (2002) and Murphy and Dieterich (2006), as well as in simultaneous comparisons of several methods [Ernst et al. (2013), Krauss, Nill and Oelfke (2011), Sharp et al. (2004)]. Using  $p \times h \times 1$  neural networks, we can predict  $Y_{i+k}$  as a function of  $Y_{(-m):i}$ . Let  $X_i = Y_{(i-p)+1:p}$ , then

$$(4.1) \quad \hat{Y}_{i+k} = \beta_0 + \beta' G(X_i),$$

where  $G(X_i) = (g(w_{01} + w'_1 X_i)g(w_{02} + w'_2 X_i) \cdots g(w_{0h} + w'_h X_i))'$  with activation function  $g$ ; here we assume  $g(x) = 1/(1 + \exp(-x))$ . Hyperparameters  $p$  and  $h$  are set by the user (as is the form of the activation function). Unknown parameters  $\beta_0, \beta, w_{01}, \dots, w_{0h}, w_1, \dots, w_h$  are estimated by minimizing the sum of squares using the R package `nnet` [Venables and Ripley (2002)]. Because the number of unknown parameters is large ( $w_1, \dots, w_h$  are  $p$ -vectors), to prevent overfitting, a regularization term is often used in the sum of squares minimization. Then, the model is fit by minimizing

$$(4.2) \quad C(Y, \theta) = \sum_{i=0}^{n-k} (\hat{Y}_{i+k} - Y_{i+k})^2 + \lambda \theta' \theta,$$

where  $\theta$  represents a vector of all unknown parameters stacked together and  $\lambda$  is a penalty hyperparameter that is supplied by the user, with higher values providing more shrinkage.

4.2. *Ridge regression.* The second competing method considered is a linear predictor of the form

$$(4.3) \quad \hat{Y}_{i+k} = \beta_0 + \beta' X_i,$$

with  $X_i = Y_{(i-p)+1:p}$  and where  $\beta_0, \beta$  are found by minimizing

$$(4.4) \quad C(Y, \beta_0, \beta) = \sum_{i=0}^{n-k} (\hat{Y}_{i+k} - Y_{i+k})^2 + \lambda(\beta_0^2 + \beta' \beta).$$

Nearly all studies involving forecasting lung tumor motion consider predictors of this form, usually referred to as ridge regression. However, since ridge regression

assumes  $\{Y_i\}$  to be independent [Hoerl and Kennard (1970)], the model implied by (4.3)–(4.4) is better described as fitting an autoregressive model of order  $p + k - 1$  (the first  $k - 1$  coefficients being 0) using conditional least squares, with an  $L_2$  penalty on the vector of autoregressive coefficients (yet we shall refer to this prediction method as ridge regression). Linear models lack many features that seem appropriate for this forecasting example, such as multimodal and/or heteroskedastic conditional distributions, yet still perform reasonably well and are commonly used as a baseline for comparing tumor prediction methods.

4.3. *Light cone reconstruction of states (LICORS)*. Mixed LICORS [Goerg and Shalizi (2013)] is a recent nonparametric forecasting method based on predictive state representations of spatiotemporal fields [Goerg and Shalizi (2012), Shalizi (2003)]. In the context of our forecasting example, mixed LICORS models  $Y_{i+k}|Y_{(-m):i}$  as depending only on the *past light cone* (with horizon  $p$ )  $X_i = Y_{(i-p)+1:p}$ ; furthermore,  $\varepsilon(X_i)$  is a minimal sufficient statistic for the predictive distribution of  $Y_{i+k}$ , so that

$$(4.5) \quad Y_{i+k}|Y_{(-m):i} \sim Y_{i+k}|X_i \sim Y_{i+k}|\varepsilon(X_i),$$

and if  $\varepsilon(X_i) = \varepsilon(X_j)$ , then  $Y_{i+k}|\varepsilon(X_i) \sim Y_{j+k}|\varepsilon(X_j)$ . Without loss of generality, we may assume  $\varepsilon$  takes values in  $\mathcal{S} = \{s_1, \dots, s_K\}$ , and for simpler notation let  $S_i = \varepsilon(X_i)$  and denote  $\mathbb{P}_j(Y_{i+k}) = \mathbb{P}(Y_{i+k}|S_i = s_j)$ . The unknown parameters of this model are the mapping  $\varepsilon$ , the number of predictive states  $K$  and the predictive distributions of the predictive states  $\{\mathbb{P}_j, 1 \leq j \leq K\}$ . For fixed  $K$ , the remaining parameters are estimated by maximizing

$$(4.6) \quad C(Y, \varepsilon, \mathbb{P}_1, \dots, \mathbb{P}_K) = \prod_{i=0}^{n-k} \sum_{j=1}^K \mathbb{P}_j(Y_{i+k}) \mathbb{P}(S_i = j|X_i),$$

which acts as a likelihood, except for  $\mathbb{P}_j$  being unknown. Goerg and Shalizi (2013) maximized (4.6) with a nonparametric variant of the EM algorithm using weighted kernel density estimators to approximate the unknown densities of the predictive distributions  $\{\mathbb{P}_j, 1 \leq j \leq K\}$ ; they also advocated data-driven procedures for choosing the number of predictive states  $K$ .

It is possible to embed the LMAR model in a parametric (Gaussian) mixed LICORS framework, treating  $\{V_{ij}, j \in \mathcal{J}_i\}$  as the past light cone  $\ell_i^-$  and  $\{V_{ij} \text{ where } M_i = j\}$  as the predictive state  $S_i = \varepsilon(\ell_i)$ . While this choice of  $\varepsilon$  does provide a minimal sufficient statistic for the predictive distribution of  $Y_i$  (or  $L_i^+$ ) under the LMAR model, it will not provide any dimension reduction or parsimony since  $\varepsilon(\ell_i)$  will almost surely be unique for each  $i$  under our model assumptions.

Fitting the mixed LICORS model to the time series in our data and using it for forecasting was accomplished using the R package LICORS [Goerg (2013b)]. Note that point forecasts using the inferred model (4.5) will be a weighted average of the means of the predictive states  $s_i \in \mathcal{S}$ .

4.4. *Data preprocessing.* Similar to Krauss, Nill and Oelfke (2011), we use a total of 80 seconds of data (2400 observations) from each time series, 40 seconds for model fitting and 40 seconds for out-of-sample prediction given the model fit to the first 40 seconds of data. This necessitates removing time series for which we have fewer than  $2400 + k$  observations, where  $k$  is the forecast window. This eliminates 61 of the 171 time series in our data base, unfortunately including all time series from patients 1, 2 and 3. An additional 15 time series were eliminated because there were several gaps in the observation sequence. This leaves us with 95 total time series; patient 8 has only one time series and patient 6 has the next fewest series with 9. Patient 11 has the most time series with 21. While each time series is three dimensional, we predict using only the first principal component (the principal component transformation is estimated from the initial 40 s of training data) as discussed in Section 2.1.

4.5. *Tuning hyperparameters.* Because of the need for real-time model fitting and prediction, all tuning and hyperparameters for the methods we consider must be specified prior to the administration of radiotherapy—before any data is observed. This suggests finding specifications for each model that perform reasonably well for all patients, though perhaps sub-optimally for each patient individually. Indeed, this is the approach usually taken in the literature [Ernst et al. (2013), Krauss, Nill and Oelfke (2011), Sharp et al. (2004)]. Because patients are typically given several or many instances of radiotherapy during different sessions, there seems to be potential for more patient-specific tuning of hyperparameters, though this is left as a separate problem for now.

Table 2 lists the hyperparameters and/or tuning parameters for each of the prediction methods we consider. As described in Section 4.4, since the first 40 seconds of each time series will not be used to evaluate out-of-sample prediction, we may use these subseries to find sensible, patient-independent values for all hyperparameters. Each 40 second subseries is further divided, where for a given set of hyperparameters each prediction method is fit to the first 30 seconds of data

TABLE 2  
*List of global, patient-independent hyperparameters to be tuned for each prediction method*

Method	Hyperparameter	Description
LMAR	$p$	Motif length (3.14)
Neural networks	$p$	Length of input vector $X_i$ (4.1)
	$h$	Number of neurons in hidden layer (4.1)
	$\lambda$	Shrinkage; L2 penalty (4.2)
Ridge regression	$p$	Length of input vector $X_i$ (4.3)
	$\lambda$	Shrinkage; L2 penalty (4.4)
Mixed LICORS	$p$	Length of input vector $X_i$ (4.5)

(900 observations), and then the remaining 10 seconds are used to generate out-of-sample predictions, for which we store the vector of errors.

Using a course grid search over the parameter space given in Table 2, predictive error [both root mean squared error (RMSE) as well as median absolute error (MAE), which is more robust to heavy-tailed error distributions] is averaged across patients, allowing us to choose the best set of patient-independent hyperparameter values [Krauss, Nill and Oelfke (2011)]. Note that different hyperparameter values are chosen for different forecast windows.

*4.6. Computational considerations.* In addition to providing real-time forecasts, tumor tracking models require parameters that can be estimated very quickly so that accurate (forecast-assisted) radiotherapy can begin as soon as possible after observing a short window of training data.

Ridge regression yields almost instantaneous estimates of parameters necessary for prediction [ $\beta$  in (4.3)], since (4.4) can be minimized in closed form. Fitting neural networks (4.1), however, requires numerical optimization of (4.2). This was carried out using the `nnet` package in R, which implements the BFGS algorithm [Venables and Ripley (2002)]. Because (4.2) is not convex, we recommend several random starting points for initiating the optimization, inasmuch as time allows; the dimension of the parameter space and the convergence criteria for the numerical optimization are both extremely important considerations in addition to the length of the time series being fit. For example, on a Lenovo X220 laptop with an Intel Core i5-2520 M 2.50 GHz processor, a  $45 \times 6 \times 1$  neural network required about 10 seconds to fit on 1200 observations when using `nnet`'s default convergence criteria, with 10 randomly initialized starting points.

The computation time in fitting the LMAR( $\Sigma$ ) depends critically on both the convergence criteria for the EM algorithm as well as the initial value of  $\Sigma$  used. Typically, the likelihood (3.7) or log-likelihood is used, however, the EM updates given in (3.12) are only approximate, meaning the likelihood is not guaranteed to increase at every iteration. We found that using the approximate log-likelihood (3.10) to check convergence yielded convergence in the exact log-likelihood. This being the case, other metrics could possibly be used to check convergence that are quicker to calculate than (3.10), such as the Frobenius norm of differences in the updates of  $\hat{\Sigma}$ . To obtain good starting values, the algorithm can be run before having observed the entire training sequence using a simple starting value of a diagonal matrix. Using a relative tolerance of 0.0001 for the approximate log-likelihood, we were able to compute  $\hat{\Sigma}$  in no more than four seconds for each of the time series considered. R code for fitting the LMAR model is included in this paper's supplementary materials [Cervone et al. (2014)].

The value of  $m$  for the LMAR model may also trade off estimation speed and accuracy; we used  $m = 400$ , though found essentially identical results for  $m = 200$  and  $m = 300$  (higher values of  $m$  favor faster, but less precise, estimation of  $\Sigma$ ).

Parameter estimation for mixed LICORS took several minutes on our machine. However, much of this computational cost is accrued in inferring  $K$ , the number of predictive states. The procedure described in Goerg and Shalizi (2013) and implemented in the LICORS R package is to start at an upper bound for the number of predictive states, optimize the likelihood approximation (4.6) and then merge the two states whose predictive distributions are closest (measured by some distance or a hypothesis test). The optimizing and merging steps are repeated until we either have 1 state remaining or, alternatively, all pairwise tests for equality among predictive distributions are rejected. Then, cross-validation is used to choose among these candidate models indexed by different values of  $K$ .

While there may be some loss in prediction accuracy, estimation speed can be improved by fixing  $K$  (perhaps tuning it as in Section 4.5). Furthermore, initializing the nonparametric EM algorithm with informative starting values (learned from previously observed respiratory trace curves) and relaxing the convergence criteria may substantially increase estimation speed with little loss in predictive performance.

**5. Prediction results for tumor tracking data.** The results of out-of-sample predictions using the LMAR model, as well as the methods discussed in Section 4, are provided in this section. Point forecasts are discussed in Sections 5.1–5.3 and interval/distributional forecasts in Section 5.4.

*5.1. Results for point forecasts.* The measures of predictive performance we consider are root mean squared error (RMSE) and median absolute error (MAE), as well as the fraction of time each forecasting method obtains the minimum prediction error among the methods compared. We report these quantities for each of the 8 patients, at forecast windows of 0.2 s (6 observations), 0.4 s (12 observations) and 0.8 s (18 observations) in Table 3.

We stress that RMSE may not be the most useful summary of predictive performance since the error distributions are heavy tailed, and in the application of radiotherapy, we are more concerned with whether or not the treatment beam was localized to the tumor than with the squared distance of the treatment beam to the tumor.<sup>6</sup> For this reason, we feel that the median (more generally, quantiles of the distribution function for absolute errors) is the best summary of predictive performance for this data context. Ultimately, the dosimetric effects of these errors are of most interest, but their determination is complicated and beyond the scope of this work.

Two further points of emphasis regarding the accuracy summaries are that while we eliminated time series with unevenly spaced observations from consideration,

---

<sup>6</sup>However, the loss function implied in the model fitting and point prediction is squared error loss, which is the simplest for many computation reasons.

TABLE 3

Summary of errors in point forecasts for all four methods and all three forecast windows considered. RMSE is root mean squared error, MAE is median absolute error, and Best refers to the proportion of time for which the absolute prediction error is smallest among the methods considered. For each metric, the most desirable value among the four methods for each patient/forecast window combination is in **bold**

Patient	Method	0.2 s forecast			0.4 s forecast			0.6 s forecast		
		RMSE	MAE	Best	RMSE	MAE	Best	RMSE	MAE	Best
4	LMAR	0.52	0.24	0.27	0.99	0.39	0.27	<b>1.18</b>	<b>0.44</b>	<b>0.31</b>
	NNs	<b>0.46</b>	<b>0.22</b>	<b>0.28</b>	<b>0.90</b>	0.39	<b>0.28</b>	1.20	0.48	0.27
	Ridge	0.53	0.31	0.20	1.08	0.62	0.17	1.50	0.86	0.18
	LICORS	0.58	0.25	0.25	1.05	<b>0.37</b>	<b>0.28</b>	1.43	0.52	0.24
5	LMAR	0.56	<b>0.25</b>	<b>0.30</b>	0.96	0.42	0.29	<b>1.15</b>	<b>0.51</b>	<b>0.30</b>
	NNs	<b>0.55</b>	0.27	0.27	<b>0.89</b>	<b>0.40</b>	<b>0.30</b>	<b>1.15</b>	<b>0.51</b>	<b>0.30</b>
	Ridge	0.58	0.31	0.25	1.01	0.56	0.23	1.39	0.78	0.23
	LICORS	0.79	0.35	0.19	1.33	0.63	0.18	1.79	0.89	0.17
6	LMAR	<b>0.77</b>	<b>0.40</b>	<b>0.29</b>	<b>1.54</b>	<b>0.82</b>	<b>0.30</b>	<b>2.00</b>	<b>1.06</b>	<b>0.34</b>
	NNs	1.01	0.46	0.24	1.74	0.93	0.24	2.43	1.38	0.22
	Ridge	0.83	0.42	0.28	1.59	0.88	0.28	2.14	1.28	0.28
	LICORS	1.37	0.57	0.19	2.17	1.19	0.18	2.92	1.75	0.15
7	LMAR	<b>0.40</b>	<b>0.15</b>	<b>0.35</b>	<b>0.85</b>	<b>0.27</b>	<b>0.37</b>	<b>1.23</b>	<b>0.41</b>	<b>0.36</b>
	NNs	0.43	0.19	0.26	0.88	0.36	0.25	1.35	0.51	0.25
	Ridge	0.44	0.26	0.20	1.00	0.59	0.16	1.56	0.96	0.17
	LICORS	0.62	0.25	0.20	1.05	0.41	0.21	1.56	0.56	0.23
8	LMAR	1.27	<b>0.62</b>	<b>0.27</b>	<b>2.63</b>	1.46	0.26	3.57	2.00	0.24
	NNs	<b>1.26</b>	0.68	<b>0.27</b>	2.71	<b>1.27</b>	<b>0.28</b>	<b>3.46</b>	<b>1.76</b>	<b>0.29</b>
	Ridge	1.44	0.69	0.20	2.86	1.54	0.19	4.11	2.26	0.19
	LICORS	1.50	0.64	0.26	2.89	1.33	<b>0.28</b>	3.70	<b>1.76</b>	0.28
9	LMAR	<b>0.58</b>	<b>0.22</b>	<b>0.39</b>	<b>1.29</b>	<b>0.52</b>	<b>0.35</b>	<b>2.03</b>	<b>0.90</b>	<b>0.30</b>
	NNs	0.73	0.32	0.24	1.69	0.64	0.26	2.45	0.92	0.24
	Ridge	0.81	0.34	0.22	1.68	0.73	0.22	2.42	0.98	0.25
	LICORS	1.35	0.53	0.15	2.20	0.98	0.17	2.64	1.19	0.20
10	LMAR	<b>0.88</b>	<b>0.36</b>	<b>0.34</b>	<b>1.73</b>	<b>0.77</b>	<b>0.33</b>	<b>2.55</b>	<b>1.19</b>	<b>0.30</b>
	NNs	1.09	0.44	0.25	2.16	0.93	0.24	2.98	1.35	0.24
	Ridge	0.95	0.45	0.24	1.84	0.94	0.24	2.67	1.41	0.26
	LICORS	1.62	0.61	0.17	2.20	1.10	0.19	3.25	1.56	0.20
11	LMAR	<b>1.13</b>	<b>0.44</b>	<b>0.32</b>	<b>2.59</b>	<b>1.06</b>	<b>0.29</b>	<b>3.70</b>	<b>1.49</b>	<b>0.31</b>
	NNs	1.24	0.50	0.25	2.95	1.19	0.24	3.99	1.70	0.23
	Ridge	1.19	0.63	0.22	2.69	1.51	0.21	3.99	2.40	0.21
	LICORS	1.64	0.57	0.21	3.04	1.09	0.26	4.21	1.65	0.25

we still have quite a few time series with unusual motion in our data base. Without actually observing the patient, we are not sure whether observed deviations from normal breathing are caused by exogenous factors or are instances of rel-

evant components of the data-generating process, such as coughs, yawns, deep breaths, etc. The other point is that there is a lot of disparity in the measures of predictive performance within the literature on this subject; in addition to working with different data sets, obtained from differing equipment, some authors account for the between-patient variation in respiratory dynamics by scaling or normalizing all curves or by comparing errors from a prediction method against errors from making no prediction and just using the lagged value of the series. When using evaluation procedures of Krauss, Nill and Oelfke (2011) and Murphy and Dieterich (2006), we produced very similar results with ridge regression and linear models. However, the error summaries we present here, in comparison with the LMAR model, are not directly comparable to these results.

5.2. *Quantitative summaries of point forecasts.* Summarizing Table 3, we see that ridge regression is actually suboptimal in all accuracy measures for all patients and forecast windows. The LMAR model strongly outperforms the other three methods for all forecast windows for patients 6, 7, 9, 10 and 11; neither neural networks nor LICORS appear to be optimal for any patient across all forecast windows, although neural networks perform well for patients 4, 5 and 8, while LICORS predicts well for patients 4, 8 and 11. Between-patient differences prevent any particular forecasting method from dominating other methods across patients, but the LMAR model seems to offer the most accurate overall point forecasts given these results.

5.3. *Qualitative summaries of point forecasts.* When looking at the predicted time series for each method used, the general pattern we observe is that LMAR outperforms the other three methods when the data undergo changes in shape, period or amplitude—or, more generally, when the test data do not resemble the training data. Figure 7 shows one (atypically dramatic) instance of such behavior. The top curve is the first 40 seconds of the time series, on which all prediction methods were trained. The next four curves give the predicted time series at a window of 0.2 s for LMAR (red), NN (blue), ridge regression (green) and LICORS (purple). It is clear from the figure that the end of the training period for this time series coincided with a dramatic change in the patient's respiration.

Both neural networks and LICORS suffer from the range of the curve being larger (dropping below  $-5$  mm and exceeding 10 mm) after the training period; for both methods, the training data bounds the range of point forecasts, regardless of the input vector for future test cases. For LICORS, when the test data is below the minimum of the training data ( $-5$  mm), the single predictive state associated with the minimal values of the training data will dominate, leading to brief periods of static forecasts. With this time series, this particular predictive state represents an abrupt transition between sharp exhale and sharp inhale. Thus, the forecasts for the test data are dramatic overestimates throughout the “U” shaped motifs starting around  $t = 47$ , where the patient does not actually fully inhale.

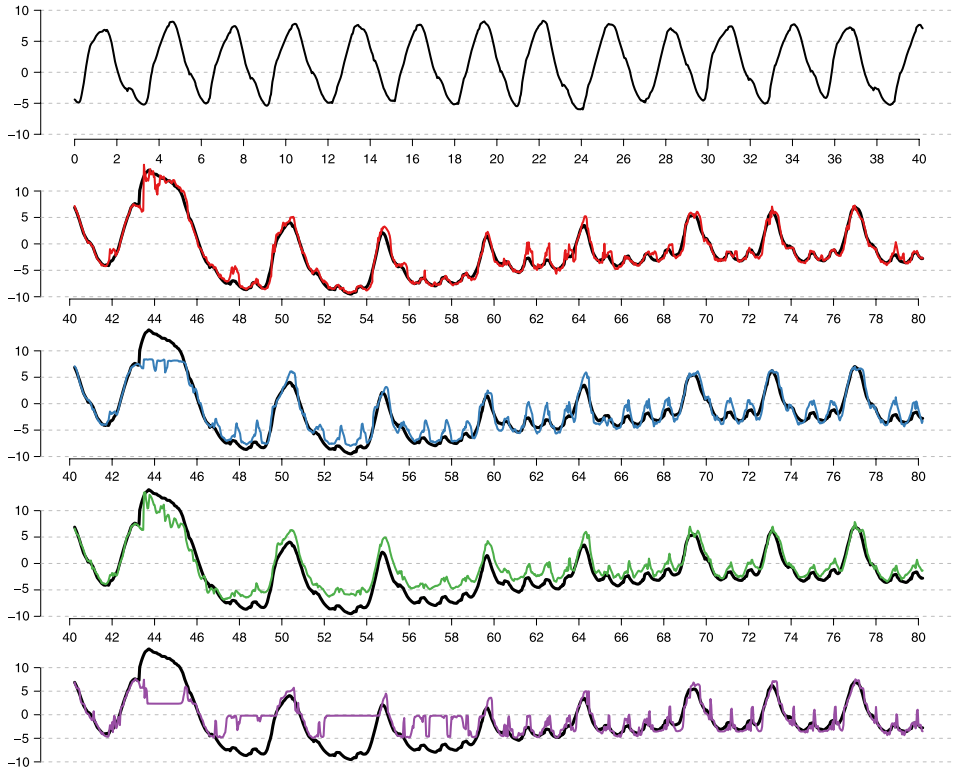


FIG. 7. Predictions for patient 9, day 3, beam 6 with a forecast window of 0.2 s. Location (mm) is the y axis and time (s) the x axis. The 40 s training sequence is top, with predictions for the next 40 s from LMAR in red, NN in blue, ridge regression in green and LICORS in purple.

Ridge regression seems to accurately predict the magnitudes of increases and decreases, yet the predictions are off by a nearly constant factor for  $t \in (48, 68)$ . In the context of the ridge regression model (4.3), this suggests that  $\beta$  is correctly specified, but perhaps  $\beta_0$  is time varying. The LMAR model includes an autoregressive term for the most recent  $p$  observations in its forecast, and thus, like ridge regression, accurately predicts rates of change in the time series. Moreover, the stochastic location-mixture component in the LMAR prediction adjusts predictions for gradual magnitude shifts in the data.

Another reason why the LMAR model works relatively well when the test data differ from the training data is that the form of the dependence of forecasts on the most recent  $p$  observations evolves, whereas it remains static for the other three methods. While the parameters of the model are not re-estimated during real-time prediction, LMAR uses the entire history of the time series in making forecasts, not just the first 40 seconds alongside the most recent  $p$  observations, as is the case with the other three methods. With appropriate parallel computing resources, all methods could theoretically update parameters continuously (or periodically)






FIG. 8. Predictions for patient 4, day 6, beam 1 with a forecast window of 0.2 s. Location (mm) is the y axis and time (s) the x axis. The 40 s training sequence is top, with predictions for the next 40 s from LMAR in red, NN in blue, ridge regression in green and LICORS in purple.

throughout treatment. [Murphy and Dieterich \(2006\)](#) continuously retrained neural networks using the updated history of the respiratory trace. While they did not compare this to the alternative of not actively updating the forecast model, [Krauss, Nill and Oelfke \(2011\)](#) did so and found a small improvement in RMSE of about 1–3%.

When the time series are more well behaved, all four methods perform quite well; in fact, neural networks tend to have the lowest errors when all four curves are accurate. Figure 8 shows the training and prediction test series for a strongly periodic respiratory trace. We should expect the performance of neural networks to be superior when the dynamics of the tumor motion are stable, as the parameter space for neural networks is far larger; in theory, feedforward neural networks with at least one hidden layer can approximate any continuous function arbitrarily well [[Hornik, Stinchcombe and White \(1989\)](#)], including time series prediction.

5.4. *Interval and distributional forecasts.* Unlike commonly used time series models in the tumor-tracking literature, the LMAR model provides multimodal,

heteroskedastic predictive distributions, which are theoretically appropriate for forecasting respiratory motion. Despite this, our analysis of predictive performance has focused exclusively on the accuracy of point forecasts because in current implementations of tumor-tracking systems, there is no clinical value in obtaining interval or distributional forecasts. The treatment beam has a fixed width and is always on, meaning an interval or distributional forecast does not alter the optimal course of action of a tumor-tracking system already supplied with a point forecast. However, interval/distributional forecasts would prove valuable if we could, for instance, suspend the treatment beam instantaneously if the predicted probability of the tumor location being enclosed by the treatment beam fell below a certain threshold.

Table 4 gives a summary of the performance of out-of-sample interval and distributional forecasts to complement the summaries of point forecasts. The LMAR model, by specifying a data-generating process, naturally provides full predictive distributions as a by-product of point prediction. The same is true for ridge regression (assuming the typical homoskedastic Gaussian structure for the residuals) and LICORS. Neural networks do not naturally provide predictive distributions; following Tibshirani (1996), we obtain them by bootstrapping, while assuming prediction errors are (heteroskedastic) independent Gaussians, with mean 0 and variance estimated by bootstrapping.

We expect LMAR prediction intervals to undercover, since uncertainty in the estimation of  $\Sigma$  is omitted from our forecasts. While this is indeed the case, for all patients and forecast windows, 90% prediction intervals have between 84% and 94% coverage—a more appropriate range than any other method can claim.

The logarithmic score in Table 4 refers to the negative logarithm of the predictive density evaluated at the true observation, averaged over each out-of-sample prediction (the result in Table 4 then averages each of these scores over all beams from the same patient). The logarithmic score is a *proper* scoring rule—its expected value is minimized by the oracle (or true) predictive distribution—thus, lower values indicate a better fit between the predictive distributions and realized values of a patient's time series [Gneiting, Balabdaoui and Raftery (2007)].

Generalizing across patients and forecast windows, in comparison to the other methods considered, the LMAR model seems to most accurately characterize prediction uncertainty.

**6. Discussion.** The location-mixture autoregressive (LMAR) model introduced in this paper provides accurate, real-time forecasts of lung tumor motion. Our method achieves better performance on out-of-sample prediction for forecast windows of 0.2 s, 0.4 s and 0.6 s for the majority of the patients considered than existing methods such as neural networks [which performed best in a prediction comparison study of Krauss, Nill and Oelfke (2011)] and penalized linear models (a common baseline for judging predictive performance). We also note that uncertainty quantification is quite straightforward using our model, whereas it is hard to do using neural networks.

TABLE 4

Summary of interval and distributional forecasts for all four methods at all three forecast windows. The interval coverage considered is 90% confidence intervals. Log PS refers to the log probability score of the predictive distribution. For each metric, the most desirable value among the four methods for each patient/forecast window combination is in **bold**

Patient	Method	0.2 s forecast		0.4 s forecast		0.6 s forecast	
		Coverage	Log PS	Coverage	Log PS	Coverage	Log PS
4	LMAR	0.84	0.72	<b>0.86</b>	1.30	<b>0.93</b>	1.37
	NNs	0.88	<b>0.57</b>	0.83	1.34	0.85	1.58
	Ridge	0.85	0.80	0.84	1.53	0.84	1.86
	LICORS	<b>0.89</b>	0.70	0.84	<b>1.03</b>	0.84	<b>1.32</b>
5	LMAR	<b>0.87</b>	<b>0.71</b>	<b>0.88</b>	<b>1.20</b>	<b>0.93</b>	<b>1.30</b>
	NNs	0.85	0.72	0.78	1.52	0.80	1.75
	Ridge	0.85	0.91	0.84	1.53	0.82	1.91
	LICORS	0.84	1.04	0.82	1.46	0.79	1.78
6	LMAR	<b>0.87</b>	1.25	<b>0.88</b>	<b>1.85</b>	<b>0.93</b>	<b>2.07</b>
	NNs	0.79	1.31	0.74	2.16	0.76	2.53
	Ridge	<b>0.87</b>	<b>1.22</b>	0.85	1.91	0.83	2.26
	LICORS	0.79	1.58	0.70	2.57	0.66	2.82
7	LMAR	0.85	<b>0.30</b>	<b>0.85</b>	<b>0.87</b>	<b>0.89</b>	<b>1.09</b>
	NNs	<b>0.88</b>	0.48	0.84	1.35	0.84	1.82
	Ridge	0.86	0.63	0.83	1.49	0.82	1.95
	LICORS	0.84	0.78	0.77	1.16	0.76	1.59
8	LMAR	<b>0.89</b>	1.67	0.91	2.30	0.94	2.60
	NNs	0.94	<b>1.53</b>	0.82	2.36	<b>0.90</b>	2.59
	Ridge	0.88	1.82	0.85	2.51	0.82	2.90
	LICORS	0.94	1.71	<b>0.90</b>	<b>2.11</b>	0.88	<b>2.39</b>
9	LMAR	<b>0.89</b>	<b>0.87</b>	<b>0.90</b>	<b>1.65</b>	<b>0.92</b>	<b>2.07</b>
	NNs	0.86	1.02	0.78	2.20	0.80	2.77
	Ridge	0.81	1.54	0.81	2.21	0.81	2.64
	LICORS	0.86	1.62	0.81	1.98	0.79	2.31
10	LMAR	<b>0.86</b>	<b>1.18</b>	<b>0.88</b>	<b>1.94</b>	<b>0.91</b>	2.33
	NNs	0.84	1.23	0.76	2.25	0.79	2.65
	Ridge	0.83	1.35	0.84	2.03	0.84	2.44
	LICORS	<b>0.86</b>	1.61	0.82	2.02	0.81	<b>2.31</b>
11	LMAR	0.85	<b>1.38</b>	<b>0.87</b>	2.13	<b>0.91</b>	2.36
	NNs	0.87	1.50	0.80	2.70	0.83	2.91
	Ridge	0.86	1.63	0.85	2.44	0.85	2.84
	LICORS	<b>0.88</b>	1.56	0.83	<b>1.99</b>	0.82	<b>2.25</b>

The LMAR model is similar to other autoregressive models that yield multimodal conditional distributions, such as the class of threshold autoregressive models [Tong (1978)], yet the parameter space consists of just a single, low-dimensional covariance matrix, and the model admits accurate closed-form ap-

proximations of multiple-step ahead predictive distributions. The LMAR model also has a useful interpretation in the context of time series motifs, which can describe the data-generating process and the form of forecasts.

While the predictive performance of our method on this data set is very encouraging, the parameter inference for the LMAR model presented here is approximate, and the assumptions of both the model and its inference may not be appropriate for some other nonlinear time series. Formalizing and generalizing the LMAR model is thus a fruitful area for future work.

Real-time prediction of lung tumor motion presents additional challenges to those presented in this work. It is preferable to have as short a training window as possible, since during this time the patient may be irradiated without actually receiving the benefit of tumor tracking. While some training is actually necessary to estimate the system latency in some cases (we have treated it as fixed throughout this work), the 40 seconds used for training in this paper (while typical in the literature on the subject) could ideally be reduced.

Also, one can consider patient-specific hyperparameter values and/or tuning parameters or modify the model to borrow information across the patients. Due to the need for real-time model fitting before we can forecast, it is most likely infeasible to apply any model selection criteria (either within-model, such as for hyperparameters, or between-model) after having begun to observe data. More study of between-patient and within-patient variability in model fits could help researchers use more patient-optimal prediction methods (as well as begin prediction after a shorter training sequence, as they would not need to rely solely on the observed data for parameter estimation).

The parametric simplicity of the LMAR model, as well as its formalization as a statistical model as opposed to a prediction algorithm, enable generalizations of our procedure to include hierarchical models and other statistical structures that address the challenges of delivering accurate external beam radiotherapy. Combined with its excellent predictive performance on real data, the LMAR model represents a promising new contribution to this area of research.

**Acknowledgments.** The authors would like to thank Dr. Seiko Nishioka of the Department of Radiology, NTT Hospital, Sapporo, Japan, and Dr. Hiroki Shirato of the Department of Radiation Medicine, Hokkaido University School of Medicine, Sapporo, Japan, for sharing the patient tumor motion data set with us. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Cancer Institute, National Science Foundation or the National Institutes of Health.

#### SUPPLEMENTARY MATERIAL

**Supplement: Code** (DOI: [10.1214/14-AOAS744SUPP](https://doi.org/10.1214/14-AOAS744SUPP); .zip). R Code used for fitting and forecasting with the LMAR model.

## REFERENCES

- BERBECO, R. I., NISHIOKA, S., SHIRATO, H., CHEN, G. T. and JIANG, S. B. (2005). Residual motion of lung tumours in gated radiotherapy with external respiratory surrogates. *Phys. Med. Biol.* **50** 3655–3667.
- BOOTS, B. and GORDON, G. J. (2011). An online spectral learning algorithm for partially observable nonlinear dynamical systems. In *Proceedings of the 25th International Conference on Artificial Intelligence* 293–300. AAAI Press, San Francisco, CA.
- BOSHPAKOV, G. N. (2009). Analytic expressions for predictive distributions in mixture autoregressive models. *Statist. Probab. Lett.* **79** 1704–1709. [MR2731114](#)
- CARVALHO, A. X. and TANNER, M. A. (2005). Mixtures-of-experts of autoregressive time series: Asymptotic normality and model specification. *IEEE Trans. Neural Netw.* **16** 39–56.
- CERVONE, D., PILLAI, N., PATI, D., BERBECO, R. and LEWIS, J. H. (2014). Supplement to “A location-mixture autoregressive model for online forecasting of lung tumor motion.” DOI:10.1214/14-AOAS744SUPP.
- D’SOUZA, W. D., NAQVI, S. A. and YU, C. X. (2005). Real-time intra-fraction-motion tracking using the treatment couch: A feasibility study. *Phys. Med. Biol.* **50** 4021–4033.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B* **39** 1–38. [MR0501537](#)
- DE GOOIJER, J. G. and KUMAR, K. (1992). Some recent developments in non-linear time series modelling, testing, and forecasting. *Int. J. Forecast.* **8** 135–156.
- ERNST, F., SCHLAEFER, A. and SCHWEIKARD, A. (2007). Prediction of respiratory motion with wavelet-based multiscale autoregression. In *Medical Image Computing and Computer-Assisted Intervention* (N. Ayache, S. Ourselin and A. Maeder, eds.) 668–675. Springer, New York.
- ERNST, F. and SCHWEIKARD, A. (2009). Forecasting respiratory motion with accurate online support vector regression (SVRpred). *Int. J. Comput. Assisted Radiol. Surg.* **4** 439–447.
- ERNST, F., DÜRICHEN, R., SCHLAEFER, A. and SCHWEIKARD, A. (2013). Evaluating and comparing algorithms for respiratory motion prediction. *Phys. Med. Biol.* **58** 3911–3929.
- FU, T.-C. (2011). A review on time series data mining. *Engineering Applications of Artificial Intelligence* **24** 164–181.
- GNEITING, T., BALABDAOUI, F. and RAFTERY, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **69** 243–268. [MR2325275](#)
- GOERG, G. M. (2013a). Forecastable component analysis. In *Proceedings of the 30th International Conference on Machine Learning* 64–72. JMLR W&CP.
- GOERG, G. M. (2013b). LICORS: Light cone reconstruction of states—predictive state estimation from spatio-temporal data. R package version 0.2.0.
- GOERG, G. M. and SHALIZI, C. R. (2012). LICORS: Light cone reconstruction of states for non-parametric forecasting of spatio-temporal systems. Technical report, Statistics Dept., Carnegie Mellon Univ., Pittsburgh, PA.
- GOERG, G. M. and SHALIZI, C. R. (2013). Mixed LICORS: A nonparametric algorithm for predictive state reconstruction. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics* 289–297. JMLR W&CP.
- HAMILTON, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* **57** 357–384. [MR0996941](#)
- HOERL, A. E. and KENNARD, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **12** 55–67.
- HORNIK, K., STINCHCOMBE, M. and WHITE, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Netw.* **2** 359–366.
- KALBFLEISCH, J. D. and SPOTT, D. A. (1970). Application of likelihood methods to models involving large numbers of parameters. *J. Roy. Statist. Soc. Ser. B* **32** 175–208. [MR0270474](#)

- KALET, A., SANDISON, G., WU, H. and SCHMITZ, R. (2010). A state-based probabilistic model for tumor respiratory motion prediction. *Phys. Med. Biol.* **55** 7615–7631.
- KRAUSS, A., NILL, S. and OELFKE, U. (2011). The comparative performance of four respiratory motion predictors for real-time tumour tracking. *Phys. Med. Biol.* **56** 5303–5317.
- KROLZIG, H.-M. (2000). Predicting Markov-switching vector autoregressive processes. Technical report, Economics Dept., Oxford Univ.
- LAU, J. W. and SO, M. K. P. (2008). Bayesian mixture of autoregressive models. *Comput. Statist. Data Anal.* **53** 38–60. [MR2528591](#)
- LE, N. D., MARTIN, R. D. and RAFTERY, A. E. (1996). Modeling flat stretches, bursts, and outliers in time series using mixture transition distribution models. *J. Amer. Statist. Assoc.* **91** 1504–1515. [MR1439090](#)
- LIN, J., LONARDI, S., KEOGH, E. and PATEL, P. (2002). Finding motifs in time series. In *Proceedings of the 2nd Workshop on Temporal Data Mining* 53–68. ACM, New York.
- LITTMAN, M. L., SUTTON, R. S. and SINGH, S. P. (2002). Predictive representations of state. In *Proceedings of Advances in Neural Information Processing Systems* **14** 1555–1561. MIT Press, Cambridge, MA.
- MCCULLOCH, R. E. and TSAY, R. S. (1994). Statistical analysis of economic time series via Markov switching models. *J. Time Series Anal.* **15** 523–539.
- MURPHY, M. J. and DIETERICH, S. (2006). Comparative performance of linear and nonlinear neural networks to predict irregular breathing. *Phys. Med. Biol.* **51** 5903–5914.
- MURPHY, M. J., ISAAKSON, M. and JALDEN, J. (2002). Adaptive filtering to predict lung tumor motion during free breathing. In *Computer Assisted Radiology and Surgery* (H. U. Lemke, M. W. Vannier, K. Inamura, A. G. Farman, K. Doi and J. H. C. Reiber, eds.) 539–544. Springer, New York.
- RIAZ, N., SHANKER, P., WIERSMA, R., GUDMUNDSSON, O., MAO, W., WIDROW, B. and XING, L. (2009). Predicting respiratory tumor motion with multi-dimensional adaptive filters and support vector regression. *Phys. Med. Biol.* **54** 5735–5748.
- ROTTMANN, J., KEALL, P. and BERBECO, R. (2013). Markerless EPID image guided dynamic multi-leaf collimator tracking for lung tumors. *Phys. Med. Biol.* **58** 4195–4204.
- RUAN, D. and KEALL, P. (2010). Online prediction of respiratory motion: Multidimensional processing with low-dimensional feature learning. *Phys. Med. Biol.* **55** 3011–3025.
- SCHWEIKARD, A., GLOSSER, G., BODDULURI, M., MURPHY, M. J. and ADLER, J. R. (2000). Robotic motion compensation for respiratory movement during radiosurgery. *Comput. Aided Surg.* **5** 263–277.
- SHALIZI, C. R. (2003). Optimal nonlinear prediction of random fields on networks. In *Discrete Mathematics and Theoretical Computer Science Proceedings (volume AB)* 11–30. DMTCS, Nancy, France.
- 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** 425–440.
- TANAKA, Y., IWAMOTO, K. and UEHARA, K. (2005). Discovery of time-series motif from multi-dimensional data based on MDL principle. *Mach. Learn.* **58** 269–300.
- TIBSHIRANI, R. (1996). A comparison of some error estimates for neural network models. *Neural Comput.* **8** 152–163.
- TONG, H. (1978). On a threshold model. In *Pattern Recognition and Signal Processing* (C. H. Chen, ed.) 575–586. Sijthoff & Noordhoff, Amsterdam.
- TONG, H. (1990). *Nonlinear Time Series: A Dynamical System Approach*. Oxford Univ. Press, New York. [MR1079320](#)
- TONG, H. and LIM, K. S. (1980). Threshold autoregression, limit cycles and cyclical data. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **42** 245–292.
- TONG, H. and MOEANADDIN, R. (1988). On multi-step non-linear least squares prediction. *J. R. Stat. Soc. Ser. D (Statistician)* **37** 101–110.

- VENABLES, W. N. and RIPLEY, B. D. (2002). *Modern Applied Statistics with S*, 4th ed. Springer, New York.
- WONG, C. S. and LI, W. K. (2000). On a mixture autoregressive model. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **62** 95–115. [MR1747398](#)
- YE, L. and KEOGH, E. (2009). Time series shapelets: A new primitive for data mining. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 947–956. ACM, New York.

D. CERVONE  
N. S. PILLAI  
DEPARTMENT OF STATISTICS  
HARVARD UNIVERSITY  
1 OXFORD ST.  
CAMBRIDGE, MASSACHUSETTS 02138  
USA  
E-MAIL: [dcervone@fas.harvard.edu](mailto:dcervone@fas.harvard.edu)  
[pillai@stat.harvard.edu](mailto:pillai@stat.harvard.edu)

D. PATI  
DEPARTMENT OF STATISTICS  
FLORIDA STATE UNIVERSITY  
117 N WOODWARD AVE  
P.O. BOX 3064330  
TALLAHASSEE, FLORIDA 32306  
USA  
E-MAIL: [debdeep@stat.fsu.edu](mailto:debdeep@stat.fsu.edu)

R. BERBECO  
J. H. LEWIS  
RADIATION ONCOLOGY  
BRIGHAM AND WOMEN'S HOSPITAL,  
DANA-FARBER CANCER INSTITUTE  
AND  
HARVARD MEDICAL SCHOOL  
75 FRANCIS ST., ASBI-L2  
BOSTON, MASSACHUSETTS 02115  
USA  
E-MAIL: [rberbeco@lroc.harvard.edu](mailto:rberbeco@lroc.harvard.edu)  
[jhlewis@lroc.harvard.edu](mailto:jhlewis@lroc.harvard.edu)