

PAPER

Intra-fraction motion prediction in MRI-guided radiation therapy using Markov processes

To cite this article: Seyed Ali Mirzapour *et al* 2019 *Phys. Med. Biol.* **64** 195006

View the [article online](#) for updates and enhancements.



MR Safe
4D Phantom
for MRgRT

QUASAR MRI^{4D}
BY MODUS QA

LEARN MORE ➔



PAPER

RECEIVED
1 February 2019REVISED
25 July 2019ACCEPTED FOR PUBLICATION
1 August 2019PUBLISHED
23 September 2019

Intra-fraction motion prediction in MRI-guided radiation therapy using Markov processes

Seyed Ali Mirzapour¹, Thomas Mazur², Gregory Sharp³ and Ehsan Salari^{1,4}¹ Department of Industrial, Systems, and Manufacturing Engineering, Wichita State University, Wichita, KS, United States of America² Department of Radiation Oncology, Washington University in St. Louis, St. Louis, MO, United States of America³ Department of Radiation Oncology, Massachusetts General Hospital, Boston, MA, United States of America⁴ Author to whom correspondence should be addressed.E-mail: ehsan.salari@wichita.edu**Keywords:** intra-fraction motion, motion predictive model, Markov processes, MRI-guided radiotherapy

Abstract

Internal organ motion during radiation delivery may lead to underdosing of cancer cells or overdosing of normal tissue, potentially causing treatment failure or normal-tissue toxicity. Organ motion is of particular concern in the treatment of lung and abdominal cancers, where breathing induces large tumor displacement and organ deformation. A new generation of radiotherapy devices is equipped with on-board MRI scanners to acquire a real-time movie of the patient's anatomy during radiation delivery. The goal of this research is to develop, calibrate, and test motion predictive models that employ real-time MRI images to provide the short-term trajectory of respiration-induced anatomical motion during radiation delivery. A semi-Markov model predicts transitions between the phases of a respiratory cycle, and a Markov model predicts transitions to future respiratory cycles, leading to accurate motion forecasting over longer-term horizons. The intended application for this work is real-time tracking and re-optimization of intensity-modulated radiation delivery.

1. Introduction

MRI-guided radiation therapy (MRgRT) systems have enabled exciting, novel applications in radiation oncology (Lagendijk *et al* 2008, Fallone *et al* 2009, Mutic and Dempsey 2014). These platforms integrate MRI scanners with radiation delivery to provide high contrast soft tissue imaging without added exposure to ionizing radiation. MRI has conclusively demonstrated its clinical value in radiation treatment of many treatment sites (Metcalf *et al* 2013). With MRI on-board the treatment system, additional imaging degrees of freedom such as diverse pulse sequences and customizable imaging geometries provide added flexibility in the treatment of many sites (Noel *et al* 2015, McPartlin *et al* 2016, Yang *et al* 2016, Menten *et al* 2017).

MRgRT systems can acquire MR images before every treatment to quantify day-to-day variation in positioning and anatomy. This imaging has enabled online adaptation in which treatment-specific plans are produced to optimize tumor targeting and organ-at-risk sparing (Acharya *et al* 2016, Lamb *et al* 2017). Early clinical results demonstrate the promise of this day-by-day feedback applied to treatment of abdominal cancers (Henke *et al* 2018, Luterstein *et al* 2018). On-board intra-fraction MRI can also be acquired to monitor anatomy changes over the duration of radiation delivery (Crijns *et al* 2012). For example, on the MRIdian system (Viewray, Inc.), planar, sagittal images at a location intersecting the target are acquired at four frames per second during treatment.

Beam gating can be realized in MRgRT by tracking the treatment target and interrupting treatment whenever the target moves too far beyond a specified boundary (Lamb *et al* 2017). While beam gating is not a new technique, gating derived from MRI provides added information for decision making compared to previously used surrogates for monitoring motion such as implanted markers, pressure-actuated bellows, and external optical monitoring (Keall *et al* 2002, Heerkens *et al* 2014). With MRI, both target position and surrounding normal tissue are directly visualized. Additionally, because the coordinate systems for imaging and radiation delivery are co-registered, overlap between fluence and anatomy can be quantified.

Synchronizing motion monitoring and delivery suggests opportunities for more sophisticated intra-fraction motion management. An example of a feed-forward technique that has recently been demonstrated on a non-clinical MRgRT system is target tracking with *multi-leaf collimator* (MLC) systems (Crijns *et al* 2012). With MLC tracking, planned leaf motion is modulated by a signal that follows measured motion (Keall *et al* 2001, Tacke *et al* 2010, Keall *et al* 2014). Another recently-demonstrated feedback technique in MRgRT is tumor trailing whereby the beam aperture is continuously updated according to time-averaged drift in target position (Fast *et al* 2018). Expanding further on these control techniques, a closed-loop feedback scheme derived from intra-fraction data could potentially be used to update volumetric dose and associated dose-volume-histogram criteria at a fixed frequency through treatment and compare instantaneous and planned results. With an appropriate choice of feedback scheme, treatment margins and normal-tissue dose can potentially be reduced without compromising delivery efficiency or target dose.

For any feedback derived from intra-fraction imaging, performance can deteriorate due to lag between signal acquisition, decision to act based on that signal, and actuation of that decision. On the MRIdian system, a lower bound on this latency is pre-determined by the imaging frequency (4 Hz). Additional sources of latency include image registration, signal calculation, and beam actuation. On a first-generation MRIdian system incorporating three ^{60}Co sources, latency in beam gating was measured to be less than 400 milliseconds (ms) (Green *et al* 2018). For more sophisticated feedback methods, a higher latency is likely due to the added complexity of operations such as image registration, dose calculation, dose accumulation, and leaf motion.

Latency can be mitigated to some extent by signal prediction at a future instance, ideally at or beyond the anticipated latency. Many strategies for predicting respiratory motion derived from a variety of signals in RT applications have been investigated, each with advantages and disadvantages (Sharp *et al* 2004, Wu *et al* 2004, Kakar *et al* 2005, Ren *et al* 2007, Putra *et al* 2008, Kalet *et al* 2010, Krauss *et al* 2011, Yun *et al* 2012, Ernst *et al* 2013, Park *et al* 2016, Song *et al* 2017). The linear prediction method (Sharp *et al* 2004, Vedam *et al* 2004, Kalet *et al* 2010) and its generalization *auto-regressive moving average* (ARMA) are extensively used to predict tumor motion (McCall and Jeraj 2007, Ren *et al* 2007). Despite widespread applications in literature, performance of these methods typically suffers at long latencies and also in the presence of highly non-linear motions (Park *et al* 2016). More advanced techniques include signal processing applications such as Kalman filters (Sharp *et al* 2004, Kalet *et al* 2010, Hong *et al* 2011) and sinusoidal models (Vedam *et al* 2004, Verma *et al* 2011). While more sophisticated than linear models, these approaches can be prone to over-fitting and suffer from high computation complexity (Verma *et al* 2011). Machine learning techniques such as support vector machines (D’Souza *et al* 2009, Ernst and Schweikard 2009, Riaz *et al* 2009) and neural networks (Sharp *et al* 2004, Ernst and Schweikard 2009, Murphy and Pokhrel 2009, Cheong *et al* 2018) have shown promise; however, these methods require substantial computing time in model training and often can be difficult to tune due to challenges in interpreting model results (Park *et al* 2016).

Another class of motion predictive model is the *finite state model* (FSM) (Wu *et al* 2004, Kalet *et al* 2010) where a respiratory motion is modeled using a finite number of states. Wu *et al* developed an FSM to describe the transitions between different stages of breathing and investigated the performance of their proposed prediction model for lung tumor motion tracked using the real-time tumor tracking (RTTT) system (Wu *et al* 2004). They reported a *root mean squared error* (RMSE) value of 1.5 mm for a 1D motion signal. Their FSM was later extended to a 3D motion signal using hidden Markov models (HMM) by Kalet *et al* (2010). In their studies, tumor motion was broken down into linear breathing states and sequences of states. Each state is determined by the previous state and is distinguished by velocity. They reported an RMSE value of 1.88 mm for their HMM compared to 2.27 mm for the linear model using a prediction horizon of 200 ms. However, their approach was limited to short-term prediction (less than 1000 ms) of the target trajectory, and lag between predicted and actual signals can be qualitatively observed.

In this paper, we develop and test the performance of Markov and semi-Markov models to predict intra-fraction motion. In particular, respiration-induced motion is modeled as cyclic, probabilistic transitions between a finite set of discrete states. Markov processes are powerful tools to model the sequence of probabilistic transitions between discrete states. Semi-Markov processes in particular provide additional flexibility in modeling the time spent at each state before transitioning to the next state. Our objective is to predict the short-term trajectory of the motion signal in real time while allowing for quantifying the uncertainty associated with observing a given motion trajectory.

The prediction performance is evaluated using motion signals inferred from cine MRI acquired on the MRIdian platform over nine patients. A significant challenge for any prediction approach applied to this data is the limited temporal resolution of the acquisitions (4 frames per second) relative to previously considered motion signals. The minimum prediction horizon of interest—corresponding to the time between image measurements—is thus 250 ms. Weighing different prediction techniques, we hypothesize that Markov processes will enable highly accurate prediction at this horizon and beyond. With a goal of ultimately realizing sophisticated feedback techniques, prediction is tested at integer intervals of the current imaging frequency from 250 ms up to

3000 ms. Performance variation with selected motion signal, training data and prediction horizon are quantified. In section 2, we develop a motion predictive model using Markov and semi-Markov models. We then propose a mechanism for anomaly detection to continuously monitor changes in breathing pattern. In section 3, we present the results obtained from applying the proposed prediction model to nine patient cases. Section 4 summarizes our findings and highlights potential future directions. Finally, section 5 concludes the paper.

2. Developing a motion predictive model

The cine MRI provided by the MRIdian system consists of a sequence of 2D sagittal images acquired at a rate of 4 frames per second. To track anatomical motion from the cine-MR images, local or global image features such as the GTV centroid (local) or aggregate changes in the 2D *deformation vector field* (DVF) over the anatomical region of interest (global) can be used to generate a surrogate motion signal (Cai *et al* 2015, Mazur *et al* 2016), which may be used in motion mitigation strategies. Here, we focus on the anatomical motion induced due to breathing which demonstrates a cyclic pattern. Each cycle can be divided into a sequence of motion phases, but there are significant variations in the amplitude and duration of the cycles as well as the time spent at each motion phase. Thus, we employ a probabilistic model to describe the transitions between different cycle types and motion phases within a cycle using Markov and semi-Markov processes, respectively.

In this section, we first present the proposed probabilistic model to describe a generic respiratory motion signal. We next discuss how this model will be used to predict the short-term trajectory of the motion signal in real time. Lastly, we propose an anomaly detection approach to identify irregular motion behavior observed in the motion signal.

2.1. Motion phase labeling

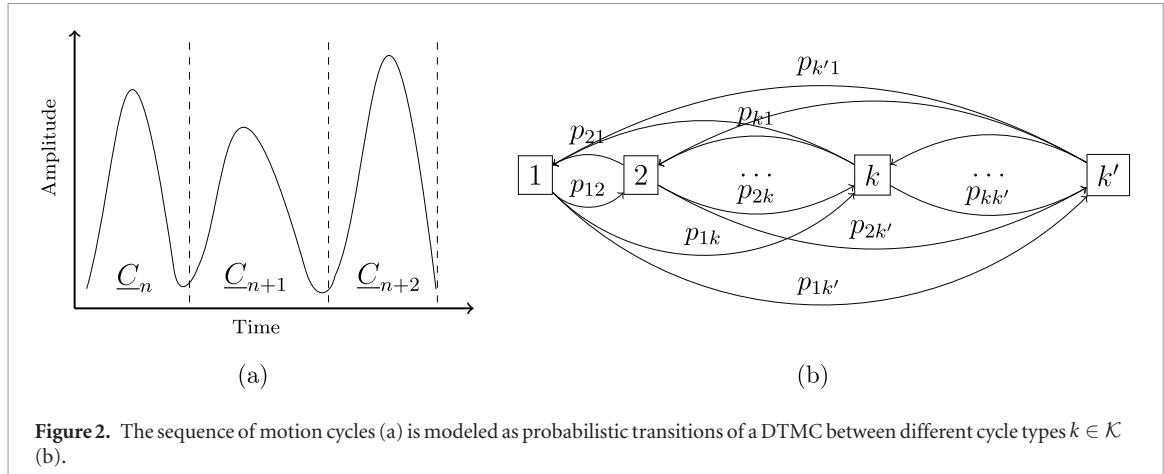
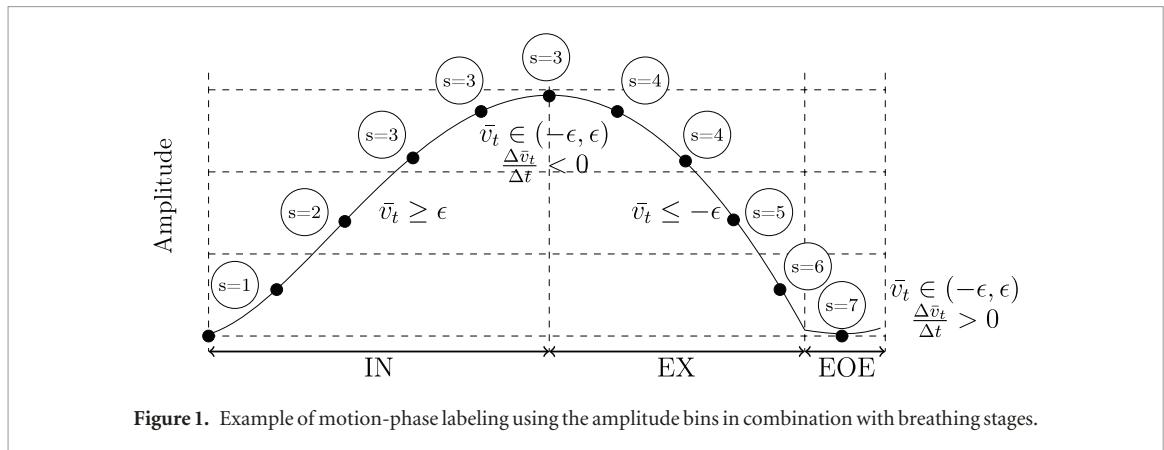
The surrogate motion signal used in this study is a time series consisting of a sequence of amplitude readings at equally spaced points in time indexed by t , where the time step depends on the acquisition rate of the MR frames (e.g. 4 frames per second). The motion signal generated for regions of interest in the thorax and abdomen demonstrates a cyclic motion due to respiration (Lee and Motai 2014). Each measurement in a motion signal is classified according to two attributes: its amplitude, which is binned into discrete ranges, and its breathing stage, which is either *inhale* (IN), *exhale* (EX), or *end-of-exhale* (EOE). We identify the breathing stage at each time step t using the moving average velocity denoted by \bar{v}_t . In particular, the IN and EX stages correspond to the increasing and decreasing parts of the motion cycle, where $\bar{v}_t \geq \epsilon$ and $\bar{v}_t \leq -\epsilon$, respectively, in which $\epsilon > 0$ is a fixed user-specified threshold and is not tuned for individual patients. To distinguish between transitioning from IN to EX or from EX to EOE, for $\bar{v}_t \in (-\epsilon, \epsilon)$ we consider the rate of change in the moving average velocity. In particular, $\frac{\Delta \bar{v}_t}{\Delta t} > 0$ suggests that the signal has transitioned into the EOE stage; otherwise, the signal has transitioned into the EX stage. The motion-phase labeling process is illustrated in figure 1, which shows measurements being assigned to seven possible motion-phase labels. We denote the set of motion phases using the index set \mathcal{S} .

2.2. Modeling motion cycle transitions

A motion cycle is a sequence of labeled motion phases associated with IN, EX, and EOE. To represent the shape and duration of a motion cycle, we employ an m -parameter model, denoted by vector $\vec{f} = [f_1, f_2, \dots, f_m]$, where \vec{f} describes the cycle properties. Let $\{\vec{f}_n : n \in \mathcal{N}\}$ be the set of observed cycles in a database of motion signals. These cycles can be grouped into a set of finite clusters, indexed by $k \in \mathcal{K}$, such that on the average, cycles within each cluster are more similar than cycles across clusters (Lee and Motai 2014). We assume that the motion signal transitions between these $|\mathcal{K}|$ different cycle types in a probabilistic manner. Let $\{\underline{C}_n : n = 0, 1, \dots\}$ be the sequence of observed cycle types in the motion signal, where $\underline{C}_n \in \mathcal{K}$ is a categorical random variable that describes the cluster of the n th cycle in the signal (the underline notation shows random variables). We model $\{\underline{C}_n : n = 0, 1, \dots\}$ as a *discrete-time Markov chain* (DTMC). The set of states is \mathcal{K} and the transition probability matrix is \mathbf{P} , where $p_{kk'} = \Pr\{\underline{C}_{n+1} = k' | \underline{C}_n = k\}$ is the probability of transitioning from cycle type k to k' . A schematic of DTMC is shown in figure 2. Note that the transition probabilities are estimated based on the frequency of transitions between observed cycles (Brémaud 1999). In this study, we assume that these transition probabilities are stationary, i.e. they do not change over time.

2.3. Modeling motion phase transitions

Within each cycle, the signal transitions between motion phases $s \in \mathcal{S}$. The sequence of motion phases visited within a cycle and the time spent at each motion phase, so-called *holding time*, may significantly vary from cycle to cycle. Therefore, we model the sequence of motion phases visited at different time steps in the cycle as a stochastic process $\{\underline{S}_t : t = 0, 1, \dots\}$, where $\underline{S}_t \in \mathcal{S}$ is a categorical random variable that describes the motion phase at time step t in the cycle. We have empirically observed that the transitions between motion phases depend on the cycle



type. Additionally, for a given cycle type, the holding times follow different probability distributions depending on the motion phases within the cycle. Therefore, we describe $\{\underline{S}_t : t = 0, 1, \dots\}$ by employing a homogeneous *discrete-time semi-Markov model* (DSMM) for each cycle type $k \in \mathcal{K}$. The DSMM for cycle type k is expressed as a sequence of pairs of random variables $\{(\underline{\Omega}_m^k, \underline{\Gamma}_m^k) : m = 1, 2, \dots\}$ in which $\underline{\Omega}_m^k$ and $\underline{\Gamma}_m^k$ are the m th motion phase visited in the cycle and its corresponding holding time, respectively. It is well known that for a DSMM, the sequence of random variables $\{\underline{\Omega}_m^k : m = 1, 2, \dots\}$ follows a DTMC, which is often referred to as the *embedded* Markov chain (Ibe 2009). The state space of the embedded Markov chain is the set of motion phases \mathcal{S} and the transition probability matrix is $\tilde{\mathbf{P}}^k$, where $\tilde{p}_{ss'}^k$ is the probability that the DSMM will next visit motion phase s' once it leaves motion phase s , that is, $\tilde{p}_{ss'}^k = \Pr\{\underline{\Omega}_{m+1}^k = s' | \underline{\Omega}_m^k = s\}$. By construction, the embedded Markov chain only considers transitions out of motion phases and hence does not include any self-transitions, that is, $\tilde{p}_{ss}^k = 0$ ($s \in \mathcal{S}$). Furthermore, $\underline{\Gamma}_m^k$ is a discrete random variable that describes the holding time of the m th transition, which depends on the current motion phase as well as the motion phase that will be visited next. Therefore, associated with every pair of motion phases s and s' with $\tilde{p}_{ss'}^k > 0$, we define a *probability mass function* (PMF) $\Phi_{ss'}^k(\tau) = \Pr\{\underline{\Gamma}_m^k = \tau | \underline{\Omega}_m^k = s, \underline{\Omega}_{m+1}^k = s'\}$ to describe the distribution of holding times.

Figure 3 shows a schematic of a DSMM with different motion phases and possible transitions. Two initial and terminal dummy states are added to represent the start and end of the motion cycle. In particular, we define the dummy terminal state to be an *absorbing* state. All other states in the DSMM are then considered to be *transient* states. The duration of the motion cycle is equivalent to the time it takes for the DSMM to start from the initial state and be absorbed by the terminal state. Moreover, the sequence of motion phases visited in a cycle is equivalent to the states visited by the DSMM starting from the initial state until absorption.

2.4. Real-time motion prediction

The goal of real-time motion prediction is to predict signal amplitude at a future time. In our framework, this requires knowing the current cycle type as well as the cycle type of subsequent cycles. Note that the type of the current cycle cannot be known for certain until it is completed. Nevertheless, we can calculate the probability that the current cycle is of type k , that is, $\Pr\{\underline{C}_n = k\}$, using the DTMC model developed in section 2.3 and the knowledge of the previous cycle type. Moreover, we can update this probability based on the observed

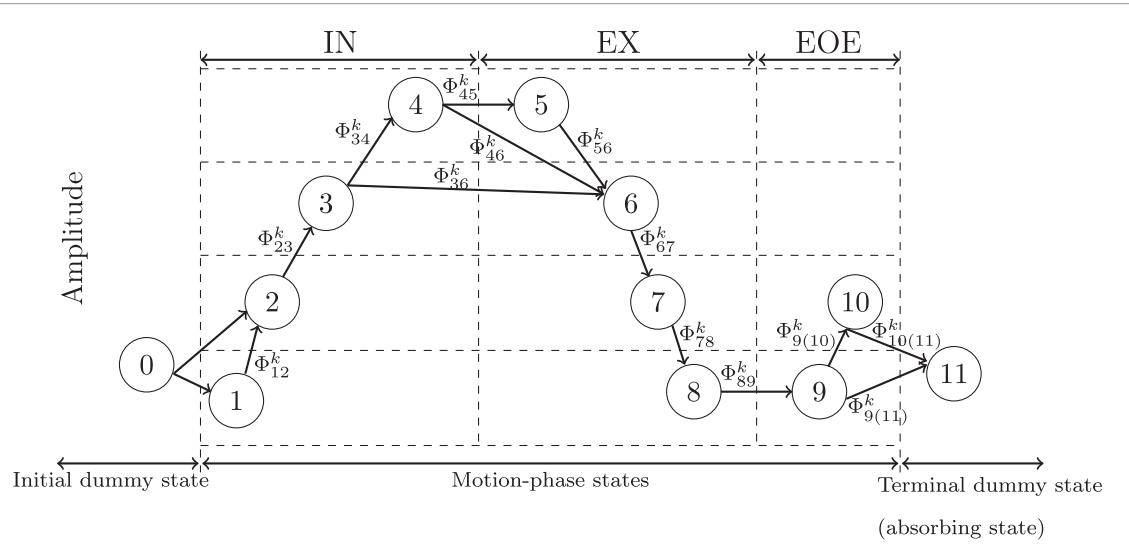


Figure 3. Example of motion-phase transition diagram of DSMM for cycle type k .

motion-phase trajectory within the current cycle. To this end, we use the Bayes' rule to update these probabilities in a dynamic manner as each MRI frame becomes available. More formally, at an intermediate time step t in the current cycle, let $\{s_1, s_2, \dots, s_t\}$ be the history of motion-phase trajectory observed since the beginning of the cycle. The probability that the current cycle is of type k , denoted by π_t^k , can be computed as

$$\begin{aligned} \pi_t^k &= \Pr \{ \underline{C}_n = k | \underline{S}_1 = s_1, \dots, \underline{S}_t = s_t \} \\ &= \frac{\Pr \{ \underline{S}_1 = s_1, \dots, \underline{S}_t = s_t | \underline{C}_n = k \} \Pr \{ \underline{C}_n = k \}}{\sum_{k' \in \mathcal{K}} \Pr \{ \underline{S}_1 = s_1, \dots, \underline{S}_t = s_t | \underline{C}_n = k' \} \Pr \{ \underline{C}_n = k' \}}. \end{aligned} \quad (1)$$

In the equation above, the *prior* probabilities, i.e. $\Pr \{ \underline{C}_n = k' \}$ ($k' \in \mathcal{K}$), are obtained using the DTMC. The probability of observing the trajectory $\{s_1, s_2, \dots, s_t\}$ for cycle type k , that is, $\Pr \{ \underline{S}_1 = s_1, \dots, \underline{S}_t = s_t | \underline{C}_n = k \}$, is obtained using the DSMM developed in section 2.3. More specifically, the observed trajectory $\{s_1, s_2, \dots, s_t\}$ may be expressed, with a slight abuse of notation, in terms of the sequence of motion phases and their holding times as $\{(\underline{\Omega}_1^k, \underline{\Gamma}_1^k) = (s_1, \tau_1), \dots, (\underline{\Omega}_{m-1}^k, \underline{\Gamma}_{m-1}^k) = (s_{m-1}, \tau_{m-1}), (\underline{\Omega}_m^k, \underline{\Gamma}_m^k) = (s_m, \geq \tau_m)\}$. Note that for the the last motion phase in the sequence s_m , only a lower bound on its holding time τ_m is known. Next, using the DSMM embedded transition probability matrix and the corresponding PMF's, the desired probability is calculated as the product of the sequence of transition and holding-time probabilities as follows:

$$\begin{aligned} &\Pr \{ (\underline{\Omega}_1^k, \underline{\Gamma}_1^k) = (s_1, \tau_1), \dots, (\underline{\Omega}_{m-1}^k, \underline{\Gamma}_{m-1}^k) = (s_{m-1}, \tau_{m-1}), (\underline{\Omega}_m^k, \underline{\Gamma}_m^k) = (s_m, \geq \tau_m) \} \\ &= \tilde{p}_{s_1, s_2}^k \Phi_{s_1, s_2}^k(\tau_1) \tilde{p}_{s_2, s_3}^k \Phi_{s_2, s_3}^k(\tau_2) \dots \tilde{p}_{s_{m-1}, s_m}^k \Phi_{s_{m-1}, s_m}^k(\tau_{m-1}) \Pr \{ \underline{\Gamma}_m^k \geq \tau_m | \underline{\Omega}_m^k = s_m \}. \end{aligned} \quad (2)$$

The last term $\Pr \{ \underline{\Gamma}_m^k \geq \tau_m | \underline{\Omega}_m^k = s_m \}$ is obtained by conditioning on the motion phase that the DSMM will be visiting next. Finally, in addition to the current cycle type, the obtained *posterior* probabilities, i.e. π_t^k ($k \in \mathcal{K}$), will be used to infer subsequent cycle types (i.e. $\underline{C}_{n+1}, \underline{C}_{n+2}, \dots$) using the DTMC model.

After inferring the cycle types, the second step of the real-time prediction involves finding the expected amplitude at a future time within each cycle type $k \in \mathcal{K}$. To predict the signal amplitude at τ time steps later in the future, starting from the current time step t with the observed motion phase s_t , we first determine the conditional expected motion phase at time step $t + \tau$, that is, $E[\underline{S}_{t+\tau} | \underline{S}_t = s_t, \underline{C}_n = k]$. This is achieved by calculating the expected remaining trajectory of the DSMM from motion phase s_t to absorption (i.e. the end of the cycle) and selecting the motion phase visited after τ time steps. More specifically, let \mathbf{Q}^k be the submatrix of $\tilde{\mathbf{P}}^k$ that contains embedded transition probabilities restricted to only transient states of the k th DSMM. The expected number of visits to each transient state in the remaining trajectory before absorption is obtained using the s_t th row of the *fundamental* matrix of the embedded Markov chain, that is, $(\mathbf{I} - \mathbf{Q}^k)^{-1}$ in which \mathbf{I} is the identity matrix (we refer the interested reader to Feldman and Valdez-Flores (2010) for a review of absorbing Markov chains). Additionally, since the transition diagram of the DSMM is acyclic, each motion phase will be visited at most once before absorption. Hence, the s_t th row of the fundamental matrix here also shows the probability of visiting each motion phase in the remaining trajectory before absorption. Furthermore, upon visiting any motion phase $s \in \mathcal{S}$, the expected holding time in that motion phase can be calculated by conditioning on the motion phase that the DSMM will visit next as follows:

$$\begin{aligned} \mathbb{E} \left[\underline{\Gamma}_m^k \mid \underline{\Omega}_m^k = s \right] &= \sum_{s' \in \mathcal{S}} \mathbb{E} \left[\underline{\Gamma}_m^k \mid \underline{\Omega}_m^k = s, \underline{\Omega}_{m+1}^k = s' \right] \Pr \left\{ \underline{\Omega}_{m+1}^k = s' \mid \underline{\Omega}_m^k = s \right\} \\ &= \sum_{s' \in \mathcal{S}} \sum_{\tau \geq 0} \tau \Phi_{ss'}^k(\tau) \tilde{p}_{ss'}^k. \end{aligned} \quad (3)$$

Using the probability of visiting each motion phase before absorption as well as its expected holding time as described above, we construct the expected remaining trajectory for each cycle type $k \in \mathcal{K}$. Next, the expected motion phase after τ time steps into the remaining trajectory is identified and converted into an amplitude value for each cycle type. By taking the weighted average of amplitude values across all cycle types using the weights π_t^k ($k \in \mathcal{K}$), we obtain the predicted amplitude at time step $t + \tau$. If the prediction horizon spans across multiple motion cycles, then the expected motion phase at time step $t + \tau$ is obtained by conditioning on visiting a sequence of consecutive cycles $\{\underline{C}_n = k_n, \dots, \underline{C}_{n+m} = k_{n+m}\}$. In that case, the conditional expected motion phase at time step $t + \tau$, that is, $\mathbb{E} [\underline{S}_{t+\tau} \mid \underline{S}_t = s_t, \underline{C}_n = k_n, \dots, \underline{C}_{n+m} = k_{n+m}]$, will be determined by concatenating the (remaining) trajectories of individual cycles in the sequence.

Finally, we note that using the DSMMs for different cycle types, we can find the probability of any given future trajectory that starts from current time step t and spans τ time steps, that is, $\Pr \{ \underline{S}_{t+1} = s_{t+1}, \dots, \underline{S}_{t+\tau} = s_{t+\tau} \mid \underline{S}_t = s_t \}$, via conditioning on the current cycle type as

$$\begin{aligned} \Pr \{ \underline{S}_{t+1} = s_{t+1}, \dots, \underline{S}_{t+\tau} = s_{t+\tau} \mid \underline{S}_t = s_t \} \\ = \sum_{k \in \mathcal{K}} \Pr \{ \underline{S}_{t+1} = s_{t+1}, \dots, \underline{S}_{t+\tau} = s_{t+\tau} \mid \underline{S}_t = s_t, \underline{C}_n = k \} \pi_t^k, \end{aligned} \quad (4)$$

where the conditional probabilities in the summation can be found similar to (2).

2.5. Anomaly detection

Due to inter- and intra-patient variability in the breathing pattern, it is possible that some motion patterns have not been observed before. Irregular breathing in particular will lead to previously unknown motion trajectories, potentially leading to a poor prediction accuracy. To identify and safeguard against these irregularities, the proposed motion predictive model is equipped with an anomaly detection mechanism to monitor changes in the breathing pattern. In particular, we adopt an anomaly detection approach previously used in other Markov modeling applications (Ye 2000, Haque *et al* 2017). The approach is based on the premise that the probability of the occurrence of an irregular trajectory calculated using the motion predictive model will be smaller than the probability of a sample trajectory that already exists in the training data set. More specifically, at each time step t we calculate the probability of observing the most recent motion-phase trajectory over the past τ time steps, that is, $\Pr \{ \underline{S}_{t-\tau} = s_{t-\tau}, \underline{S}_{t-\tau+1} = s_{t-\tau+1}, \dots, \underline{S}_t = s_t \}$ via conditioning on the current cycle type similar to equation (4) and, in turn, calculating each conditional probability similar to equation (2). In doing so, if we encounter a new motion phase or a new transition that has a probability of occurrence of zero according to the DSMM, then a very small positive probability value (e.g. 10^{-10}) is used instead to ensure that the product of the probabilities for the entire trajectory does not vanish. A probability value below a user-specified threshold indicates an anomaly. Finally, we exclude the predictions made at time steps in which anomalies are detected.

2.6. Data preparation and surrogate motion signal generation

The performance of the proposed motion predictive model is tested using cine-MRI data acquired on the MRIdian system from nine stomach cancer cases with ten fractions for each case. Cine-MR images for each cancer case are acquired on the MRIdian system in a sagittal plane at four frames per second. Images are acquired via a balanced steady-state free precession sequence (TrueFISP, Siemens). In-plane resolution and slice thickness of the images are 3.5 and 7 mm, respectively. A user-specified contour can be tracked across an acquisition via on-board software that maps the contour frame-by-frame using a *deformable image registration* (DIR) technique loosely based on a previously published method (Lu *et al* 2004).

Reconstructed images acquired by the scanner for ten fractions each from nine stomach cancer cases were collected. A motion signal for each fraction was generated by deformably registering each image to a selected reference frame (nominally the first suitable image of an acquisition). The DVF between each image and the reference frame were produced using the open-source software *Plastimatch* (Sharp *et al* 2010). Specifically, two-stage B-spline DIR was applied to native images with control point spacing of 14 mm, regularization of 0.1, and 11 maximum iterations per stage. Once DVFs were prepared, the stomach was manually contoured on the reference image for every fraction and the average *superior-inferior* (SI) component of the DVF for all pixels enclosed by the contour was calculated, as shown in figure 4. DVF-generated motion in these cases was complicated by cardiac motion stemming from the inferior aspect of the heart being abutted against the stomach. To restrict analysis to the respiratory component of observed motion, we applied a one-sided *Gaussian-weighted moving average*

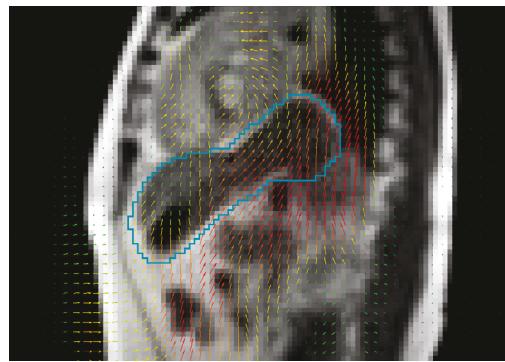


Figure 4. The average SI displacement of all pixels enclosed in the stomach contour is used to generate the motion signal.

Table 1. Range and standard deviation of the generated motion signal for 9 stomach cancer cases.

Patient	Range (mm)	Std dev (mm)
1	11.45	1.57
2	10.21	1.42
3	10.29	1.52
4	11.75	1.27
5	13.34	1.87
6	11.98	1.94
7	14.47	2.54
8	11.57	1.92
9	13.81	1.84

(GWMA) filter to suppress cardiac contributions to measured signals. The choice of the one-sided GWMA filter ensures no use of future information during prediction.

Stability of the DIR results with respect to additive Gaussian noise was estimated by adding synthetic noise to the original cine frames and generating the surrogate motion signal following the same procedure as discussed above. Accuracy of the DIR signal (calculated using the mean absolute difference between the original motion signal and the signal generated from the noisy frames) is within ± 0.29 mm when the *contrast-to-noise ratio* (CNR) is reduced by half. The prediction accuracy of the surrogate motion signal may deteriorate in the presence of significant additive image noise due to an increased level of noise in the DIR signal.

Because the average baseline and amplitude of the motion signal varies between and within each fraction, we standardize the signal amplitude at each time step by subtracting the moving average of the amplitudes over the past 30 time steps. To standardize motion signals for all patients and fractions, analysis was restricted to the first 1640 images (410 s) of each acquisition. Table 1 shows the range and the standard deviation of the motion signal for each patient case.

2.7. Model calibration

We employ the *leave-one-out cross-validation* strategy for training and testing the proposed motion predictive model, where at each step, one patient case is left out as the testing data set and the remaining patient cases are used as the training data set. The model uses eight parameters to characterize the motion cycles, as shown in figure 5, where the first five parameters describe the amplitude at the beginning of the IN stage, 2 time steps into the IN stage, beginning of the EX stage, 2 time steps into the EX stage, and beginning of the EOE stage, respectively. The last three parameters describe the duration of the IN, EX, and EOE stages, respectively. Motion cycles in the training data set are clustered into $|\mathcal{K}| = 15$ cycle types using the K-means clustering algorithm (Hartigan and Wong 1979, Kanungo *et al* 2002). To estimate the transition probability matrix \mathbf{P} of the DTMC model developed in section 2.2, for every pair of cycle types k and k' , we calculate the relative frequency of observed transitions from k to k' in the training data set. This procedure is guaranteed to produce an unbiased estimator of the transition probabilities (Brémaud 1999).

We calibrate the DSMM for each cycle type $k \in \mathcal{K}$ in section 2.3 using a total of 39 motion phases, where 14 amplitude bins are used for each of the IN and EX stages and 11 amplitude bins are used for the EOE stage. This ensures a fine amplitude resolution (around 1 mm) for the motion phases given the observed motion range in the patient cohort. To estimate the embedded transition probability matrix $\tilde{\mathbf{P}}^k$ of the k th DSMM, for every pair of motion phases s and s' , we calculate the relative frequency of observed transitions from s to s' in all cycles

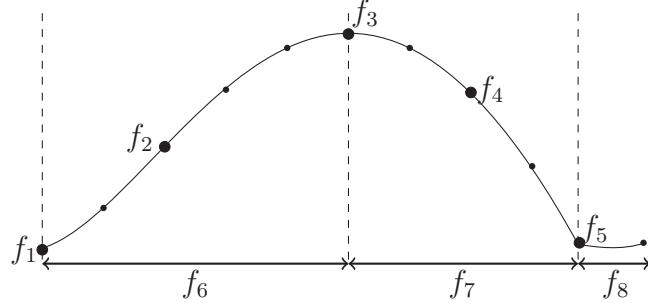


Figure 5. An eight-parameter model $\vec{f} = (f_1, \dots, f_8)$ is used to characterize and categorize cycles.

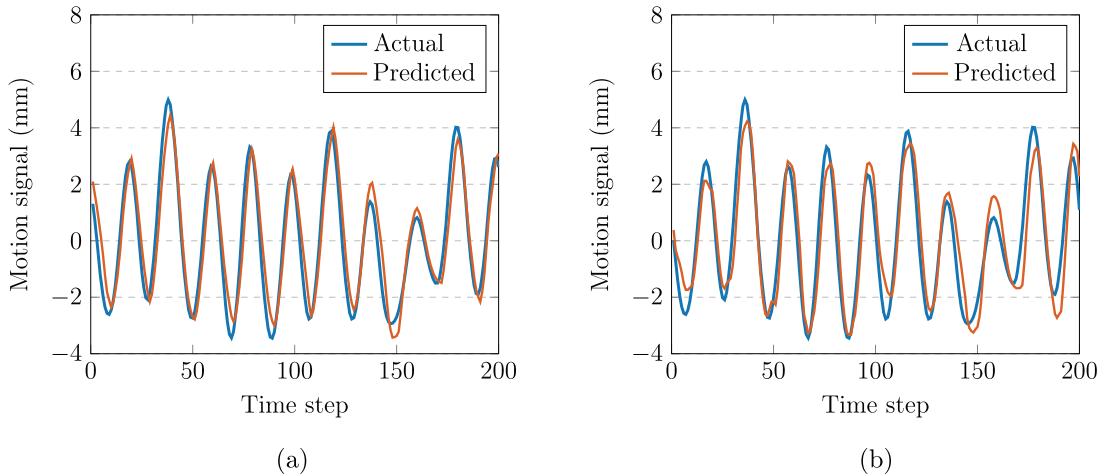


Figure 6. Comparing the actual versus predicted motion signal with prediction horizon of (a) 500 ms and (b) 1000 ms.

belonging to cluster k . Moreover, associated with each of such transitions, we calculate the holding time as the number of time steps that DSMM spent in s before transitioning to s' . The empirical distribution of the holding times is then used to obtain the PMF $\Phi_{ss'}^k$.

3. Experimental results

In the following, we report the computational results obtained from applying the motion predictive model to the patient cohort described in section 2.6.

3.1. Analysis of prediction results

Given the cine-MRI acquisition rate of 4 frames per second (a 250 ms interval), we test the performance of our motion predictive model using prediction horizons ranging from 250 to 3000 ms. Figure 6 compares the actual versus predicted motion signals using two prediction horizons of 500 and 1000 ms for a representative sample interval within the first fraction of patient case 2. As can be seen in the figure, as we prolong the prediction horizon to 1000 ms, the quality of the prediction deteriorates at the motion-phase level, but the motion cycles are still correctly predicted. The major discrepancy between the actual and predicted signals occurs at transitions between breathing stages, particularly at longer prediction horizons (1000 ms). In particular, the motion model may predict that the signal is transitioning into a new breathing stage while the actual signal is still in the previous stage, or vice versa. However, those incorrect predictions are corrected quickly owing to the continuous updating of the predicted trajectory.

Table 2 shows the *mean absolute error* (MAE), a commonly used measure of forecasting accuracy (Makridakis *et al* 1983), per fraction for different patient cases using prediction horizons of 1000 and 3000 ms. All MAE values are below 1.5 and 2.5 mm for a prediction horizon of 1000 and 3000 ms, respectively. To investigate the potential benefit of employing a motion predictive model, we also compare the prediction results obtained by the proposed method against the case of no prediction in the last row of table 2 for which the mean absolute deviation of the actual signal amplitude from the moving average of the past 30 time steps is reported. Additionally, for short prediction horizons below 500 ms, we investigate the performance of a naive prediction in which the most recent

Table 2. MAE values (mm) obtained for prediction horizons of 1000 and 3000 ms.

Fraction	P1	P2	P3	P4	P5	P6	P7	P8	P9	Fraction	P1	P2	P3	P4	P5	P6	P7	P8	P9		
1000 ms	1	0.97	1.28	0.86	1.24	1.31	1.31	1.10	0.70	1.16	1	1.52	1.92	1.52	1.74	1.86	1.88	2.02	1.25	1.68	
	2	0.80	1.46	0.81	1.22	1.14	1.14	1.23	0.62	1.05	2	1.51	2.24	1.33	1.55	1.47	1.77	2.29	1.16	1.40	
	3	0.89	1.45	0.74	1.50	1.05	1.05	1.03	0.74	0.97	3	1.77	1.96	1.25	1.93	1.53	1.52	2.04	1.17	1.36	
	4	0.92	1.18	0.89	1.19	1.26	1.26	1.09	0.94	1.27	4	1.66	1.87	1.42	1.89	1.80	1.42	1.92	1.42	2.17	
	5	1.14	1.23	1.08	1.05	1.35	1.35	0.96	0.96	1.51	5	2.14	2.00	1.74	1.60	1.77	1.61	1.87	1.50	2.22	
	6	0.85	1.38	0.80	1.14	1.09	1.09	1.03	1.08	1.45	6	1.78	2.20	1.55	1.78	1.61	1.88	1.83	1.80	2.46	
	7	0.92	1.01	0.87	1.14	1.07	1.07	1.22	0.76	1.07	3000	7	1.99	1.97	1.44	1.85	1.42	1.56	2.30	1.24	1.33
	8	1.36	1.14	1.06	1.08	1.38	1.38	1.24	0.85	1.13	ms	8	1.97	1.70	1.51	1.66	1.75	1.87	2.07	1.39	1.58
	9	0.90	1.14	1.14	1.14	1.07	1.07	1.41	0.83	0.98	9	1.52	1.79	1.60	1.61	1.53	1.47	2.26	1.32	1.41	
	10	1.51	1.46	0.96	1.16	1.04	1.04	1.20	0.87	1.05	10	2.16	2.05	1.61	1.73	1.52	1.64	2.09	1.49	1.61	
Average		1.03	1.27	0.92	1.19	1.18	1.18	1.15	0.84	1.16	Average		1.80	1.97	1.50	1.73	1.62	1.66	2.07	1.37	1.72
No prediction		3.02	2.64	2.87	2.39	3.18	3.03	3.20	2.97	2.95	No prediction		3.32	2.93	2.86	2.75	3.25	3.07	3.22	3.08	3.02

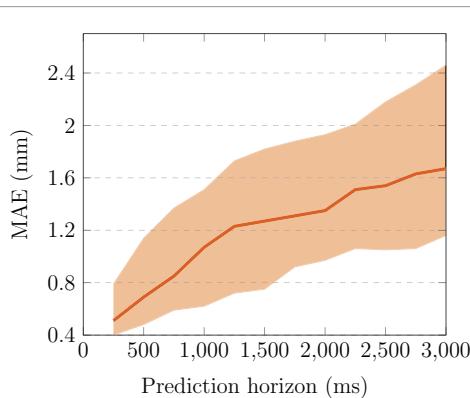


Figure 7. Median and range of prediction error (measured using MAE) in the patient cohort at different prediction horizons.

observed signal amplitude is used as a predictor of the signal amplitude for the next two time steps. When applied across all fractions of patient cases in the cohort, the naive predictor yields a median prediction accuracy of 1.01 and 1.88 mm for prediction horizons of 250 and 500 ms, respectively, compared to the prediction accuracy of 0.51 and 0.69 mm obtained by the proposed motion predictive model. Figure 7 shows the minimum, median, and maximum of average MAE values calculated for each treatment fraction of the 9 patient cases considered in the cohort using different prediction horizons.

For the patient cohort used in this study, the total computation time for model calibration and training is, on the average, 9 minutes. This includes offline computation of average trajectories for each DSMM starting from different motion phases. The computation time required for online prediction and anomaly detection at each time step is negligible (0.01 s).

3.2. Population-based versus patient-specific model calibration

To study the potential benefit of customizing the motion predictive model to each individual patient compared to employing a population-based model, we train and test the model using patient-specific data in three different calibration scenarios. The training data set in these three scenarios consists of three, six, and nine fractions, respectively, while the last fraction (i.e. fraction 10) is always used as the testing data set. Figure 8 compares the MAE obtained by the population-based and patient-specific calibration scenarios tested on the same fraction using different prediction horizons. The figure shows that the patient-specific calibration using at least six fractions (out of ten fractions) outperforms the population-based calibration. However, if the patient-specific training data set is relatively too small (e.g. three fractions), then the population-based calibration yields superior accuracy.

3.3. Anomaly detection results

Anomaly detection is performed by continuous calculation of the probability of observing the motion-phase trajectory of the most recent 3 time steps. To determine an appropriate threshold value for anomalies, we first calculated the probability of occurrence of every motion-phase trajectory of length 3 time steps observed in the training data set using the motion predictive model. We then compared the performance of 4 different threshold values corresponding to 1, 3, 5, and 10 percentiles of the distribution of the probability values. The predictions made at time steps in which anomalies were detected were then removed and the MAE was recalculated accordingly. Figure 9(a) compares the MAE excluding anomalies detected using different threshold values. The figure shows that there is a diminishing gain in the prediction accuracy as the threshold value increases beyond 5%. Figure 9(b) shows time steps with detected anomaly in a sample motion-phase trajectory using a threshold value corresponding to 5 percentile. Note that when applied to the testing data set, anomalies are detected in real time using the threshold value found from the training data set. That is, presented analysis does not correspond to retrospective elimination of data points.

We also investigate the performance of the real-time anomaly detection mechanism by calculating the false positive and false negative rates when applied to the testing data set. False negatives are irregular cycles in the testing data set that were not flagged as anomalies by the proposed method. Conversely, false positives are regular cycles that were flagged as anomalies. To that end, we first classify all cycles in the testing data set retrospectively into two subgroups of regular versus irregular. This is achieved by calculating the probability that each cycle in the testing data set belongs to each of the clusters $k \in \mathcal{K}$ using the corresponding DSMM and choosing the maximum probability value among them as a surrogate to quantify the degree to which the cycle is regular. The fifth percentile is used as a classification threshold such that all cycles that have a probability value below the threshold are labeled as irregular and the rest as regular. Next, we apply the proposed method to the testing data set in a real-

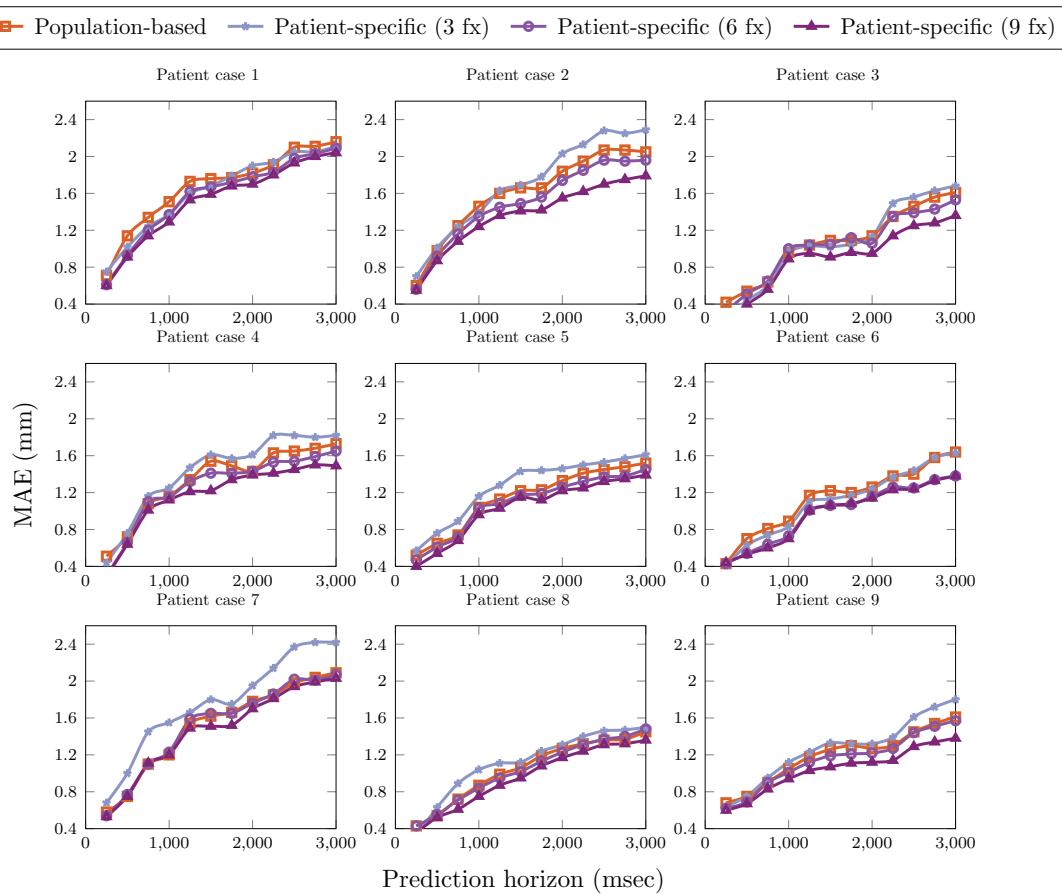


Figure 8. Comparing the performance of population-based versus patient-specific motion models for nine patient cases.

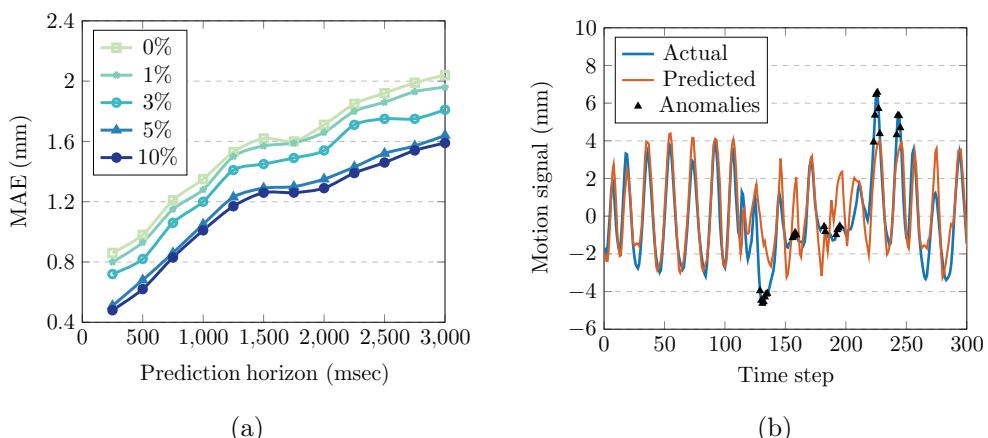


Figure 9. (a) Comparing performance of the anomaly detection method using different threshold values based on 0 to 10 percentiles; (b) Actual and predicted motion signals with detected anomalies for fraction 1 of patient case 5 using prediction horizon of 1000 ms.

time fashion to identify cycles with anomalies. We then calculate the false positive rate as the percentage of regular cycles that contain detected anomalies by the proposed method and the false negative rate as the percentage of irregular cycles that do not contain any detected anomalies.

The above approach was applied to ten fractions of Patient case 1 consisting of a total of 1000 cycles. For each of those cycles, the maximum probability of belonging to each of the 15 clusters were calculated. The 5 percentile threshold was used to classify the cycles into 50 irregular and 950 regular cycles. The proposed method detected anomalies in 46 out of 50 irregular cycles (false-negative rate = 8%). Additionally the proposed method flagged anomalies in 22 out of 950 regular cycles (false-positive rate = 2.3%). The results show that, for the patient case studied, the proposed anomaly detection method is capable of effectively detecting irregular motion cycles in real time using only the most recent segments of the cycles prior to their completion.

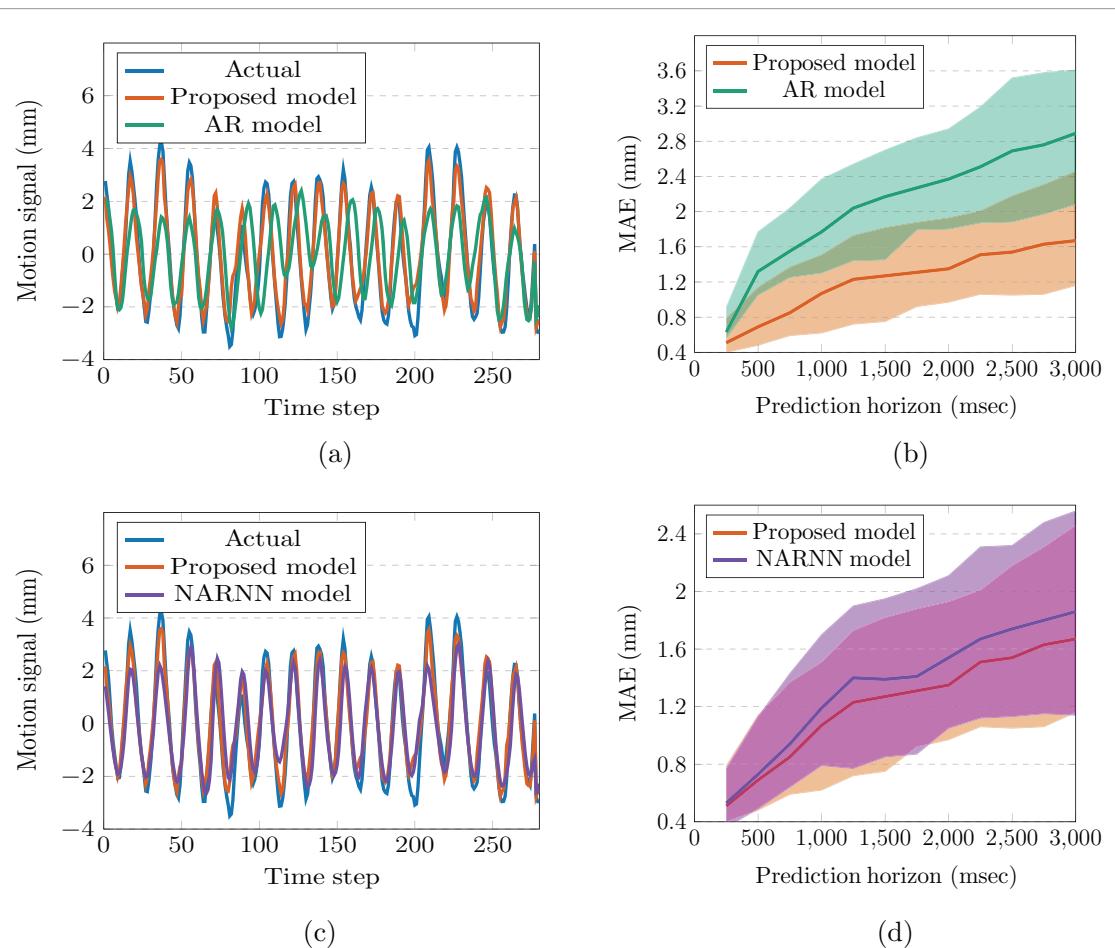


Figure 10. Comparing the performance of the proposed motion predictive model with AR and NARNN models: (a) and (c) actual versus predicted signals for a sample interval; (b) and (d) range of prediction error in the patient cohort at different prediction horizons.

3.4. Performance comparison with linear and non-linear prediction models

We also compare the performance of the proposed model against the *auto-regressive* (AR) predictive model. An AR model of order p is of the form $X_t = b + \sum_{i=1}^p a_i X_{t-i} + e_t$, where a_i ($i = 1, \dots, p$) and b are parameters of the model and e_t is white noise (Akaike 1969). For prediction horizons longer than one time step, we iteratively apply the AR model to obtain the prediction at the desired time step. The AR model was implemented using a built-in function in MATLAB (Mathworks, Natick, MA) and was calibrated and tested using the same leave-one-out cross-validation strategy. The order of the AR model was optimized to obtain the best prediction results. Figure 10(b) shows the minimum, median, and maximum of the MAE values in the patient cohort for different prediction horizons using the AR model compared to the proposed motion predictive model. Results show that the proposed model outperforms the AR model at all prediction horizons.

We also compared the performance of the proposed model against a *non-linear auto-regressive neural network* (NARNN) model. In particular, NARNN predicts the signal amplitude at a future time step using a nonlinear function of signal values over the past d time steps, that is, $X_t = f(X_{t-1}, X_{t-2}, \dots, X_{t-d})$. The nonlinear function is modeled as an artificial neural network with a single hidden layer consisting of multiple neurons. The NARNN was implemented using the *Neural Network Time Series* toolbox of MATLAB (Mathworks, Natick, MA). The Levenberg–Marquardt training algorithm was used. Sensitivity analysis was performed on the number of neurons in the hidden layer as well as the parameter d value. The NARNN model with 20 hidden neurons and $d = 40$ had superior performance among all combinations tested. Similar to the proposed method, the NARNN model was trained and tested using the leave-one-out strategy. Figures 10(c) and (d) compares the performance of the two methods. Results show that NARNN outperforms the linear auto-regressive model and its performance, though still inferior, approaches that of the proposed method. In particular, similar to the linear auto-regressive model, NARNN is not able to predict the signal peaks as accurately as the proposed method. When evaluating the prediction accuracy only at the positive and negative peaks of the signal using a prediction horizon of 1000 ms, the median MAE value in the patient cohort is 1.75 mm for the NARNN model compared to 1.29 mm for the proposed motion predictive model.

4. Discussion

In the following, we highlight the main insights gained from this study for intra-fraction motion prediction using Markov processes.

4.1. Motion signals can be accurately modeled and predicted using Markov processes

The numerical results in section 3.4 show that the proposed motion predictive model has a significantly better prediction accuracy compared to linear AR models across all prediction horizons tested in the study. It seems that the inferior performance of the AR model is due to poor prediction of the IN and EX amplitude peaks as well as difficulty in the AR model to predict the timing of these events, as shown in figure 10(a). These two issues are less prevalent in the proposed motion predictive model. Additionally, the numerical results show that the proposed method yields a better prediction accuracy, particularly at signal peaks, than a non-linear AR model developed using artificial neural networks, though to a lesser extent.

4.2. Patient-specific model calibration improves the prediction accuracy

The comparison between a population-based and patient-specific calibration in section 3.2 shows that a gain in prediction accuracy may be achieved if the proposed model is calibrated using sufficient patient-specific data. However, it is also seen that patient-specific calibration using only a few fractions may negatively impact the prediction accuracy. Sufficient patient-specific data could, in principle, be obtained through acquiring additional cine MRI prior to the start of the treatment. However, given the extent of the acquisition required, that could potentially cause work-flow inefficiency and patient inconvenience. Therefore, one possible strategy for incorporating patient-specific data would be to start with a population-based calibration at the beginning of the treatment course and switch to patient-specific calibration after acquiring sufficient cine-MRI during initial treatment fractions.

4.3. Anomaly detection helps identify irregularities in the motion signal and thus allows for proactive management of poor predictions

A central feature of the proposed motion predictive model is the ability to calculate the probability of the occurrence of any segment of the motion signal. This feature can be used to evaluate the likelihood of a recently observed motion pattern, which can assist with detecting signal anomalies. Irregularities in observed motion are correlated with poor prediction of the immediate upcoming trajectory. Therefore, in the proposed predictive model, we identify the anomalies and filter corresponding predictions. In the presented results, detecting and filtering anomalies at a level of 5% of the testing data set leads to a noticeable improvement in the prediction accuracy. In practice, this anomaly detection scheme may be implemented as an interlock that interrupts treatment when irregular motion is detected. With a 5% threshold, this approach would increase treatment time by on the order of 5%.

4.4. Study limitation and future research directions

The proposed motion predictive model is suitable for organ motion that exhibits a cyclic pattern (e.g. respiration-induced motion). To account for inter- and intra-fraction baseline shifts (e.g. drifts), the surrogate motion signal is adjusted by subtracting the moving average of the past signal amplitudes. Moreover, the method is equipped with anomaly detection to flag any spontaneous motion (e.g. due to patient's accidental moves, coughing, sneezing, etc) for which prediction accuracy may deteriorate. However, the predictive model may have limited flexibility in responding to significant intra-fraction changes in the cycle types, motion phases, and the corresponding transition patterns. The proposed anomaly detection approach helps mitigate this limitation to some extent. Ideally, the intra-fraction imaging information can be used to adjust to those changes in real time. Future research can extend this study by equipping the proposed motion predictive model with an online-learning scheme to re-calibrate the model parameters continuously based on the arriving MRI frames using a feedback loop. Prediction with our approach can in principle be applied to difference sources of motion signals. In this work, we demonstrated the approach using mean values for the SI component of the DVF within a specified ROI. An extension of these methods might attempt to extend this one-dimensional prediction to prediction of an entire image. For instance, our framework could be applied to image components derived from a dimension reduction technique such as *principal component analysis*, and prediction of these low-dimension components could then be used to reconstruct a predicted image (albeit at the expense of certain information, e.g. high frequency motion). Alternatively, our model could be applied either at the pixel-level or to mean signals extracted from pixel neighborhoods in an effort to directly predict DVFs as a brute-force approach to image prediction.

5. Conclusion

In this study, Markov and semi-Markov models were developed to predict respiration-induced intra-fraction motion using a surrogate motion signal generated from cine MRI acquired during MRgRT. To that end, the probabilistic transitions between different motion cycle types were modeled using a DTMC. Within each cycle type, a DSMM was developed to model the probabilistic transitions between different motion phases. The proposed model is used to perform real-time motion predictions, which are dynamically updated upon the arrival of new MR frames. The proposed model is equipped with an anomaly detection mechanism to find irregularities in the observed motion signal.

The proposed model was calibrated and tested using a patient cohort consisting of nine stomach cancer cases, and its performance was compared against both linear and non-linear auto-regressive models. The numerical results show that the proposed model outperforms those models across all prediction horizons. The proposed anomaly detection approach can effectively identify irregularities in the motion signal, thus providing the opportunity to proactively manage the corresponding inferior predictions. Results show that patient-specific calibration has a better prediction accuracy compared to population-based calibration only when sufficient data are available for model training. The experimental results obtained from employing Markov processes in this study are promising and show the potential of Markov modeling in intra-fraction motion prediction for MRgRT applications. The proposed motion predictive model can be further extended by incorporating an online-learning scheme to dynamically re-calibrate the model parameters, thus allowing for adjusting to intra-fraction changes in the motion pattern in real time.

Acknowledgment

This research was funded in part by the National Science Foundation through Award #1662819.

References

Acharya S *et al* 2016 Online magnetic resonance image guided adaptive radiation therapy: first clinical applications *Int. J. Radiat. Oncol. Biol. Phys.* **94** 394–403

Akaike H 1969 Fitting autoregressive models for prediction *Ann. Inst. Stat. Math.* **21** 243–7

Brémaud P 1999 Discrete-time Markov models *Markov Chains* (New York: Springer) pp 53–93

Cai W, Hurwitz M, Williams C, Dhou S, Berbeco R, Seco J, Mishra P and Lewis J 2015 3D delivered dose assessment using a 4DCT-based motion model *Med. Phys.* **42** 2897–907

Cheong K *et al* 2018 A proof-of-concept study for the real-time prediction of respiratory patterns: a simple Bayesian approach *J. Korean Phys. Soc.* **73** 368–76

Crijns S, Raaymakers B and Lagendijk J 2012 Proof of concept of MRI-guided tracked radiation delivery: tracking one-dimensional motion *Phys. Med. Biol.* **57** 7863–72

D’Souza W, Malinowski K and Zhang H 2009 Machine learning for intra-fraction tumor motion modeling with respiratory surrogates *Int. Conf. on Machine Learning and Applications* (IEEE) pp 463–7

Ernst F and Schweikard A 2009 Forecasting respiratory motion with accurate online support vector regression (SVRpred) *Int. J. Comput. Assist. Radiol. Surg.* **4** 439–47

Ernst F, Dürichen R, Schlaefer A and Schweikard A 2013 Evaluating and comparing algorithms for respiratory motion prediction *Phys. Med. Biol.* **58** 3911–29

Fallone B, Murray B, Rathee S, Stanescu T, Steciw S, Vidakovic S, Blosser E and Tymofichuk D 2009 First MR images obtained during megavoltage photon irradiation from a prototype integrated linac-MR system *Med. Phys.* **36** 2084–8

Fast M, Van de Schoot A, Van de Lindt T, Carbaat C, Van der Heide U and Sonke J 2018 Tumor trailing for liver SBRT on the MR-linac *Int. J. Radiat. Oncol. Biol. Phys.* **103** 468–76

Feldman R and Valdez-Flores C 2010 *Applied Probability and Stochastic Processes* (Berlin: Springer) (<https://doi.org/10.1007/978-3-642-05158-6>)

Green O *et al* 2018 First clinical implementation of real-time, real anatomy tracking and radiation beam control *Med. Phys.* **45** 3728–40

Haque A, DeLucia A and Baseman E 2017 Markov chain modeling for anomaly detection in high performance computing system logs *Proc. of the 4th Int. Workshop on HPC User Support Tools* (Denver, CO: ACM) pp 1–8

Hartigan J and Wong M 1979 Algorithm as 136: a k-means clustering algorithm *J. R. Stat. Soc. C* **28** 100–8

Heerkens H, Van Vulpen M, Van den Berg C, Tijssen R, Crijns S, Molenaar I, Van Santvoort H, Reerink O and Meijer G 2014 MRI-based tumor motion characterization and gating schemes for radiation therapy of pancreatic cancer *Radiother. Oncol.* **111** 252–7

Henke L *et al* 2018 Phase I trial of stereotactic MR-guided online adaptive radiation therapy (SMART) for the treatment of oligometastatic or unresectable primary malignancies of the abdomen *Radiother. Oncol.* **126** 519–26

Hong S, Jung B and Ruan D 2011 Real-time prediction of respiratory motion based on a local dynamic model in an augmented space *Phys. Med. Biol.* **56** 7137–52

Ibe O 2009 *Markov Processes for Stochastic Modeling* (Waltham, MA: Elsevier)

Kakar M, Nyström H, Aarup L, Nøttrup T and Olsen D 2005 Respiratory motion prediction by using the adaptive neuro fuzzy inference system (ANFIS) *Phys. Med. Biol.* **50** 4721–8

Kalet A, Sandison G, Wu H and Schmitz R 2010 A state-based probabilistic model for tumor respiratory motion prediction *Phys. Med. Biol.* **55** 7615–31

Kanungo T, Mount D, Netanyahu N, Piatko C, Silverman R and Wu A 2002 An efficient k-means clustering algorithm: analysis and implementation *IEEE Trans. Pattern Anal. Mach. Intel.* **24** 881–92

Keall P, Colvill E, O'Brien R, Ng J, Poulsen P, Eade T, Kneebone A and Booth J 2014 The first clinical implementation of electromagnetic transponder-guided MLC tracking *Med. Phys.* **41** 020702

Keall P, Kini V, Vedam S and Mohan R 2001 Motion adaptive x-ray therapy: a feasibility study *Phys. Med. Biol.* **46** 1–10

Keall P, Kini V, Vedam S and Mohan R 2002 Potential radiotherapy improvements with respiratory gating *Australas. Phys. Eng. Sci. Med.* **25** 1–6

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** 5303–17

Langendijk J, Raaymakers B, Raaijmakers A, Overweg J, Brown K, Kerkhof E, Van der Put R, Hårdemark B, Van Vulpen M and Van der Heide U 2008 MRI/linac integration *Radiother. Oncol.* **86** 25–9

Lamb J *et al* 2017 Online adaptive radiation therapy: implementation of a new process of care *Cureus* **9** e1618

Lee S and Motai Y 2014 *Prediction and Classification of Respiratory Motion* (Heidelberg, NY: Springer) (<https://doi.org/10.1007/978-3-642-41509-8>)

Lu W, Chen M, Olivera G, Ruchala K and Mackie T 2004 Fast free-form deformable registration via calculus of variations *Phys. Med. Biol.* **49** 3067–87

Luterstein E, Cao M, Lamb J, Raldo A, Low D, Steinberg M and Lee P 2018 Stereotactic MRI-guided adaptive radiation therapy (SMART) for locally advanced pancreatic cancer: a promising approach *Cureus* **10** e2324

Makridakis S, Wheelwright S and McGee V 1983 *Forecasting Methods and Applications* 2nd edn (New York: Wiley) (<https://doi.org/10.1002/for.3980030408>)

Mazur T, Fischer-Valuck B, Wang Y, Yang D, Mutic S and Li H 2016 SIFT-based dense pixel tracking on 0.35 T cine-MR images acquired during image-guided radiation therapy with application to gating optimization *Med. Phys.* **43** 279–93

McCall K and Jeraj R 2007 Dual-component model of respiratory motion based on the periodic autoregressive moving average (periodic ARMA) method *Phys. Med. Biol.* **52** 3455–66

McPartlin A *et al* 2016 MRI-guided prostate adaptive radiotherapy—a systematic review *Radiother. Oncol.* **119** 371–80

Menten M, Wetscherek A and Fast M 2017 MRI-guided lung SBRT: present and future developments *Phys. Med.* **44** 139–49

Metcalfe P, Liney G, Holloway L, Walker A, Barton M, Delaney G, Vinod S and Tome W 2013 The potential for an enhanced role for MRI in radiation-therapy treatment planning *Technol. Cancer Res. Treat.* **12** 429–46

Murphy M and Pokhrel D 2009 Optimization of an adaptive neural network to predict breathing *Med. Phys.* **36** 40–7

Mutic S and Dempsey J 2014 The ViewRay system: magnetic resonance-guided and controlled radiotherapy *Seminars in Radiation Oncology* vol 24 (Amsterdam: Elsevier)

Noel C, Parikh P, Spencer C, Green O, Hu Y, Mutic S and Olsen J 2015 Comparison of onboard low-field magnetic resonance imaging versus onboard computed tomography for anatomy visualization in radiotherapy *Acta Oncol.* **54** 1474–82

Park S, Lee S, Weiss E and Motai Y 2016 Intra-and inter-fractional variation prediction of lung tumors using fuzzy deep learning *IEEE J. Transl. Eng. Health Med.* **4** 1–12

Putra D, Haas O, Mills J and Burnham K 2008 A multiple model approach to respiratory motion prediction for real-time IGRT *Phys. Med. Biol.* **53** 1651–63

Ren Q, Nishioka S, Shirato H and Berbeco R 2007 Adaptive prediction of respiratory motion for motion compensation radiotherapy *Phys. Med. Biol.* **52** 6651–61

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** 5735–48

Sharp G, Jiang S, Shimizu S and Shirato H 2004 Prediction of respiratory tumour motion for real-time image-guided radiotherapy *Phys. Med. Biol.* **49** 425–40

Sharp G, Li R, Wolfgang J, Chen G, Peroni M, Spadea M, Mori S, Zhang J, Shackleford J and Kandasamy N 2010 Plastimatch—an open source software suite for radiotherapy image processing *Proc. of the XVIth Int. Conf. on the use of Computers in Radiotherapy (Amsterdam, Netherlands)*

Song H *et al* 2017 Respiratory motion prediction and prospective correction for free-breathing arterial spin-labeled perfusion MRI of the kidneys *Med. Phys.* **44** 962–73

Tacke M, Nill S, Krauss A and Oelfke U 2010 Real-time tumor tracking: automatic compensation of target motion using the siemens 160 MLC *Med. Phys.* **37** 753–61

Vedam S, Keall P, Docef A, Todor D, Kini V and Mohan R 2004 Predicting respiratory motion for four-dimensional radiotherapy *Med. Phys.* **31** 2274–83

Verma P, Wu H, Langer M, Das I and Sandison G 2011 Survey: real-time tumor motion prediction for image-guided radiation treatment *Comput. Sci. Eng.* **13** 24–35

Wu H, Sharp G, Salzberg B, Kaeli D, Shirato H and Jiang S 2004 A finite state model for respiratory motion analysis in image guided radiation therapy *Phys. Med. Biol.* **49** 5357–72

Yang Y *et al* 2016 Longitudinal diffusion MRI for treatment response assessment: Preliminary experience using an MRI-guided tri-cobalt 60 radiotherapy system *Med. Phys.* **43** 1369–73

Ye N 2000 A Markov chain model of temporal behavior for anomaly detection *Proc. of the 2000 IEEE Workshop on Information Assurance and Security* vol 166 (West Point, NY: United States Military Academy)

Yun J, Mackenzie M, Rathee S, Robinson D and Fallone B 2012 An artificial neural network (ANN)-based lung-tumor motion predictor for intrafractional MR tumor tracking *Med. Phys.* **39** 4423–33