

Towards Real-Time Motion Compensation in Radio-Transparent Robotic Radiation Therapy

Olalekan Ogunmolu and Rodney D. Wiersma

Abstract Stereotactic radiosurgery based on frame-based positioning is a well-established technique for treating benign and malignant lesions. In recent years, there has been a trend in frameless systems which use image-guided robot positioning for managing patient motion compensation in radiation therapy. A majority of these works utilize rigid-links in Stewart-Gough kinematic mechanisms for patient motion correction. A fundamental weakness of a majority of these approaches is that the rigid mechanical components of these robots interfere with the therapeutic radiation beam. In addition, they share their complete workspace with the patient's body (risking safety), and their constant curvature components are hardly suitable for manipulating the soft tissues of the human body. In this report, we highlight some recent advancements that aim to stem these issues in our line of work; these are systems that utilize soft mechanisms for patient motion compensation in robotic radiation therapy.

1 Introduction

Across the world, cancer remains an existential burden on rich and poor nations alike. In 2019 alone, an estimated 1,762,450 new cancer cases will be diagnosed in the United States, whereupon 606,880 will lead to fatality [1]. At a cost projection of \$147.3 billion it consumes about 4.2% of overall health care spending. Similarly, the International Agency for Research on Cancers estimates that the highest rates of cancer incidence over the coming decades will fall on low- and middle-income countries [2].

Olalekan Ogunmolu and Rodney Wiersma.
3620 Hamilton Walk, John Morgan Building, Perelman School of Medicine, The University of Pennsylvania, Philadelphia, PA 19104.
e-mail: ([olalekan.ogunmolu](mailto:olalekan.ogunmolu@pennmedicine.upenn.edu), [rodney.wiersma](mailto:rodney.wiersma@pennmedicine.upenn.edu))@pennmedicine.upenn.edu



Fig. 1: The Cyberknife and 6-DOF robotic couch system. ©Accuray Inc.




Fig. 2: A multi-leaf collimator (MLC) used in IMRT and 3DCRT. ©Varian Medical Systems.

Means of treating cancers may include one or a combination of drugs, radiation therapy, immunotherapy, stem cell transplant, targeted therapy, precision therapy, chemotherapy, or surgery. Radiation Therapy (RT), sometimes in conjunction with surgery and chemotherapy, can be an invaluable single cancer treatment modality: it is very cost-effective (accounting for only 5% of the total cost of cancer care [3]), and it has an advanced mode of radiation production and delivery. By shaping the geometry of high-energy radiation it allows radiation escalation to tumor targets while simultaneously sparing organs-at-risk (OARs). The importance of RT is underscored by the fact that half of all cancer patients undergo RT treatment during the course of their illness; in fact, an estimated 40% of all curative cancer treatment modality are

performed with RT [4]. Thus, because of its advanced radiation delivery method, RT is often the most suitable treatment modality for H&N cancers.

1.1 Robots in Radiation Therapy

Traditionally in conventional high-energy RT, a non-static beam is used as a surgical instrument so that the radiation source can move along primitive geometrical paths, irradiating the patient on a treatment couch in the process (see Figure 1). Using rectangular fields, blocks and wedges to specify flatness and symmetry, uniform radiation intensity can be produced [5]. In order to adaptively reposition the radiation beam during treatment, modern approaches utilize a six degree-of-freedom (DOF) robot arm to overcome the cross-sectional radiation delivery limitations of conventional systems.

1.1.1 Robotic Beam Repositioning

In RT, high-energy photons are generated in a linear accelerator (LINAC) machine; the LINAC is mounted on the end effector of an open kinematic chain robot arm. By this arrangement, the radiation beam can be delivered to the tumor target e.g. [6] within a patient lying in a supine position on a 6-DOF translational and rotatory robotic couch [7, 8].

To conform the radiation to the tumor, minimize sustained damage to normal tissues, and ensure sparing of OARs during irradiation, in an *inverse treatment planning process*, multi-leaf collimator or MLCs (shown in Figure 2) are sequenced in order to conform the geometric field of the ionizing radiation into a non-uniform field [9]. Essentially, the MLCs conform the spatial localization of a high dose volume to a target volume.

1.1.2 Patient Motion Correction During RT

An open problem in radiosurgery is that of keeping the patient's position consistent with those in the pre-calculated *treatment planning parameters*. Studies have shown that serious changes do occur in delivered dose when a patient is slightly displaced or when there is a misalignment from the registered patient's pose angle. Ling et. al [10] found that minute changes in couch angles affected target delivery results significantly more than accelerator angular changes. These uncertainties in couch translational and rotatory magnitudes may reduce the minimum target dose or increase the maximum cord dose [10]. In stereo-tactic radiosurgery (SRS), the misalignment between planned and delivered dose can cause eczema or brain lesions [11]. If there is a geometric miss, highly conformal potent dose increases the risk of underdose to tumors or undesirable high dose to critical organs and nearby tissues.



Fig. 3: Masks and Frames for Head Immobilization

Therefore, in order to avoid dose miss, guarantee precision of dose delivery, assure repeatable positioning during interfractional treatments, as well assure the efficacy of dose escalation to a target or minimize OARs' exposure to toxicity, a patient's position on the treatment couch should not fluctuate.

1.2 Limitations of Existing Immobilization Technologies

Currently in clinics, rigid frames and masks (see [Figure 3](#)) are used to keep the patient immobilized on the machine so as to mitigate patient motion errors. Mask-based immobilization uses thermoplastic masks (see left inset of [Figure 3](#)). Before the patient wears the mask, the solid thermoplast is allowed to stretch by heating it in water to a preset temperature. When the molecular bonds in the thermoplast weaken, a patient may wear it and secure their head to the couch. It reduces immobilization accuracy owing to flex and shrinkage over multiple use. For deep tumors nearby critical structures such as the brain stem, and for novel treatment modalities such as single isocenter multiple-target SRS, masks are not suitable given the high sensitivity of rotational head motion.

Frame-based immobilization involves a metal ring screwed to the skull of the patient, which is then bolted to the treatment table (see right inset of [Figure 3](#)). The invasive nature and discomfort of the frame causes poor patient compliance to trajectory and reduces clinical efficacy. Even so, for certain patients frame placement is impossible given their unique cranial anatomy or prior surgical bone flaps; the frame limits the use of multiple RT delivery as patients cannot be subjected to daily attachment and removal of the frame.

Setup errors between fractionated treatments (interfractional) or patient motion errors during a treatment session (intrafractional) often need to be corrected *in real-time* during treatment. Currently in clinics, the treatment is stopped, and the machine is recalibrated when the error is too large for this process to go on. The discomfort

caused by head and neck masks and frames in prolonged treatment (i) can increase the voluntary and involuntary motion of patients; (ii) are time-consuming to calibrate on a treatment machine since doses are usually delivered in fractions over many weeks or months; (iii) lack real-time position correction of the patient's head motion; and (iv) have been known to cause patient discomfort after treatment.

The CyberKnife system (see Fig 2), in spite of its advanced mode of beam repositioning and radiation delivery, requires a frame or an immobilization mask, and is incapable of real-time closed-loop feedback head motion corrections when the treatment beam is on. The CyberKnife Synchrony, while capable of precise, non-surgical tumor and lesions treatment in SRS and stereotactic body radiotherapy (SBRT), only executes *a-priori* trajectories; furthermore, it is only FDA-approved for lung tumors' treatment; correction requirements in systems such as this require far less accuracy, typically $< 5mm$ [12] than brain targets.

Additionally, real-time closed-loop head motion compensation for the CyberKnife system is inhibited by its high load-to-weight ratio which indirectly affects its repeatability: as an open-loop kinematic chain, it exhibits poor accuracy since the weight of the segments that follow each link in the robot and the load of the mechanical structure contributes to the large flexure of torques; its links inherently magnify errors from shoulder out to the end-effector, consequently hampering its use for sophisticated control strategies that may minimize or eliminate load-dependent errors; additionally, its setup distance from the patient is a recipe for delayed execution of control laws in ensuring that radiation beam reaches its target without significantly affecting dose delivery. Given its stiffness (it weighs 160kg), it exhibits a high load-to-weight ratio, with a complicated actuation system so that its passive bending stiffness overwhelms the degree of deformation for rapid patient repositioning.

Frameless and maskless (F&M) positioning systems are an emerging non-invasive immobilization technology in radiosurgery; they work without utilizing rigid masks and frames – reducing side effects and optimizing patient comfort with little trade-off in efficiency or effectiveness. The goal is to correct patient motion, ideally with a closed-loop feedback controller implemented in real-time on a high-precision robotic system – improving the satisfaction of patients and clinicians, and maximizing dose delivered to a tumor whilst minimizing healthy tissues' exposure to radiation. Parallel robot configurations have found good use along this research thrust. This is despite their higher number of actuated joints. In a way, this is an advantage because they distribute the weight of the load around the links of the robot, improving manipulation accuracy as a result; they also exhibit a desirable lightness property (albeit at the expense of a reduced workspace), and minimize the flexure torques that are otherwise common with open-loop kinematic chains. Thus, parallel kinematic configurations, in theory, enable greater precision with minimal control complications owing to the non-cumulative actuator errors [13].

Recent F&M research directions include the steel-cast assembled 4-DOF robot of [14] which corrected translational motion and a pitch rotational head motion; the HexaPOD parallel manipulator of [15] which utilized a system identification and model predictive control approach to correct a tumor position

on an Hexapod, or the in-house fabricated 6-DOF Stewart-Gough platform of [16]. However, these systems share common drawbacks e.g.

- given their constant-curvature end-effectors/platforms, they are incapable of providing sophisticated manipulation e.g. for the inadvertent respiratory motions that often induce deviation from a target;
- being made out of rigid bodies, the attenuation of ionizing radiation dose has to be factored into treatment plans when these systems get commissioned; and
- exhibiting planar platforms/tool frames, there is an inherent delay in head motion compensation along the patient's frontal axis that may inhibit clinical accuracy.

To improve the treatment planning process, these drawbacks need to be addressed. This would require the interdisciplinary effort of engineers, roboticists, physicists, and surgeons alike. In what follows, we present some of our work in soft robotic patient motion correction systems in RT [17–24].

2 Materials

In this section, we describe the mechanical components necessary for the soft actuation system described in [21] and [24].

- Mannequin head and torso simulator (custom-made, $155W \times 240L \times 200D$ mm)¹, 1;
- Craftsman air canister (maximum pressure: 320psi), 1;
- Off-the-shelf suction pump (max. pressure: 12 psi), 3;
- National Instruments (NI) myRIO microcontroller, 1;
- PWM voltage regulator (custom-built for DIOs on the myRIO), 6;
- 3D camera (Ensenso N35 model, mounted $\cong 45^\circ$ to the vertical above the head), 1;
- Base inflatable air bladder or IAB, 1
 - custom-made from elastomeric polymers;
 - elastomer covered with soft styrofoam to absorb reactive pressure from head manipulation;
 - size in reference configuration: $180mm \times 280mm$;
 - maximum width in current configuration: $\sim 75mm$;
 - 2 crack-resistant polyethylene tubing: 1/8" internal diameter (ID) and 1/4" outer diameter (OD).
- Side IABs, $\times 2$
 - custom-made from elastomeric polymers;
 - elastomers encased in breathable foam pads for comfort;
 - size in reference configuration: $180mm \times 140mm$;

¹ [Link to Solidworks Model](#)




Fig. 4: Hardware Description

- maximum width in current configuration: $\sim 75mm$;
- 2 crack-resistant polyethylene tubing: 1/8" ID and 1/4" OD for each IAB
- Proportional solenoid valves
 - Dakota Instruments EM valves (Model PSV0105, Orangeburg, NY, USA) for proportional torques to the soft actuators, 6.

3 Methods

The model head lies in a supine position above a planar table as shown in [Figure 4](#). We employed a 3D camera from Ensenso GmbH (model N35) to reconstruct the surface image and measure head pose. The N35 camera captures multiple image pairs during exposure; each image pair is made up of different patterns, controlled by piezo-actuators. A stereo-matching algorithm gathers the information from all image pairs after capture to produce a high-resolution point cloud (PCL) of the scene [\[25\]](#). We mounted the 3D sensor such that its lens faced the head at approximately 45° from the vertical during experiments. All vision processing, systems modeling and control laws were computed on a CORSAIR PC. We exchange the neuro-control and sensor signals via the publish-subscribe IPC of the ROS middleware installed on the PC. Adaptive control laws were sent via udp packets to the RIO microcontroller. The system setup is shown in [Figure 4](#).

The reference frame of the head is described as follows: the pitch/x-axes is along the frontal axis of the head, from the left ear out to the right ear; the yaw/z-axes is equivalent to the Sagittal axis, running from the back of the head through the forehead; while the roll/y-axes runs along the Craniocaudal axis, from the neck through the top of the head. The side IABs are positioned at both temples/cheeks of the patient respectively for controlling the roll- and x-axes motions, while the IAB underneath the head, henceforth referred to as the base IAB, controlled the pitch- and z-axes motions. What we present here is a simultaneous control of the head motion about the z , pitch and roll axes.

3.1 Face Segmentation

In order to pose a good control problem, we must define a local frame on the patient's face. This frame will then be used to determine the patient's real-time pose during control. Through spatial decomposition of the scene, we delineate the face from the table and other peripheral actuators around the head. In a stagewise process, we segment candidate objects that lack the characterization of the object we want to extract from the scene (the face of the mannequin). Our vision segmentation philosophy is inspired by spatial decomposition methods such as [26] that determine subdivisions and boundaries in a scene to allow retrieval of a region of interest given a proximity measure. Since the table's location cannot exceed a given height during experiments (the couch does not move as the IABs provide complete closed-loop control) and the camera's position is fixed relative to the head, separating objects that represent planar 2D geometric shapes from the scene would simplify the face segmentation algorithm.

Thus, we remove objects that fit a planar primitive geometry, and then cluster the remaining objects afterwards: we search for a simplified 2D planar object, fitted to the scene such that points $\mathbf{p}_i \in \mathcal{P}$ that support a 2D plane is found within a tolerance $0 \leq |d| \leq |d_{max}|$, where $|d_{max}|$ is a user-defined threshold based on the measured couch height [26]. We proceed as follows:

- The point cloud of the scene was acquired from the computed disparity map of the two raw camera images on the sensor;
- To minimize sensor noise whilst preserving 3D representation, the acquired point cloud was downsampled using a SAmple Consensus (SAC)-based robust moving least squares algorithm (RMLS) [26, §6];
- We then searched for the edges of 2D planar regions in the scene with Maximum Likelihood SAmple Consensus (MLESAC) [27], and we bound the resulting plane indices by computing their 2D convex hull;
- A model fitting stage extruded the computed hull (of objects lying above the 2D planar region) into a prism model based on a defined L_1 Manhattan distance; this gives the points whose height threshold is about the region of the face in the scene [28];

- Clustering the remaining points based on a heuristically-determined L_2 distance between points remaining within the polygonal plane, we find the largest cluster as the face.

We now describe these segmentation stages. Normalizing the coordinates of the original point cloud, $\mathcal{P}$, we ensure that the distance between points $p \in \mathcal{P}$ is upper-bounded by 1 based on the diagonal of $\mathcal{P}$'s bounding box. It computes a weighting factor, α , given by

$$\alpha = \mu_x + k \cdot \sigma_x, \quad (1)$$

where μ and σ respectively denote the mean and standard deviation of the mean distance distribution between points and k is a user-chosen variable. Through SAC, an estimate $\hat{\mathcal{P}}$ of the original point cloud is computed and represented as a set of equidistant grid points in the neighborhood of $\mathcal{P}$. The points of $\mathcal{P}$ are then projected to a local plane of reference through their k nearest neighbors to assure proximity to the surface of $\mathcal{P}$. Points $\hat{p}_i \in \hat{\mathcal{P}}$ are fitted to the surface that approximates $\mathcal{P}$ with a bivariate polynomial height function in a local Darboux frame (with orthonormal axes $\mathbf{U}$, $\mathbf{V}$, $\mathbf{N}$; $\mathbf{V}$ is chosen to be parallel to the local reference frame's normal). The polynomial weights are computed for the k nearest neighbors of $\mathbf{p}$ as

$$w_i = \exp\left(-\frac{\|\hat{\mathbf{p}} - \mathbf{p}_i\|^2}{\alpha}\right). \quad (2)$$

Surfaces in the proximity of $\hat{p}$'s neighbors are approximated using

$$n_{(u,v)} = \sum_{i=1}^N c_i \cdot f_{(u,v)}^i, \quad (3)$$

where u , v , and n are coordinates along the Darboux frame axes, $f_{(u,v)}^i$ are height function members of bivariate polynomials. We refer the reader to [26] for a more detailed treatment of the resampling algorithm. The result of the resampling algorithm is shown in the top-right image of Figure 5.

The inertial frame has the z -axis pointing up from the head along the dorsoventral axis (towards the camera), y -axis pointing right along the frontal axis; and the x -axis as the cross product of the y - and z - axes. This is chosen as $O(0, 0, 0.712m)$ from the camera's origin. To simplify the complexity of the planar structure in the scene, the table is modeled as a 2D planar geometric primitive so that finding points that fit a defined model hypothesis involves estimating a single distance to the frontal plane of the table surface rather than multiple points if the model was represented with points. Searching for horizontal planes that are perpendicular to the z -axis of the head is carried out using the maximum likelihood SAC [27] algorithm implemented in the PCL Library [29] to generate model hypotheses. The plane segmentation algorithm is defined in 1. The plane segmentation process is run once. Once the plane model is found, its indices and those of objects lying above it are separated and stored in separate data structures. Every subsequent iteration consists of (i) computing the 2D




Fig. 5: [Top-left]: Dense point cloud of the experimental setup scene. [Top right]: Downsampled cluttered cloud of the left scene. [Bottom-left]: Using RANSAC, we searched for 2D plane candidates in the scene and compute the convex hull of found planar regions. We then extrude point indices within the hull into a prismatic polygonal model to give the face region. [Bottom-right]: An additional step clusters the resultant cloud based on a Euclidean distance. The largest cluster is taken to be the face.

Algorithm 1 Plane Segmentation Algorithm

```

function PlaneSeg( $\mathcal{P}, a, b, c, d$ )
  for  $i = 1$  to  $N$  do
    sample non-collinear points  $\{p_i, p_j, p_k\}$  from  $\mathcal{P}$ 
    calculate the model coefficients  $a\mathbf{x} + b\mathbf{y} + c\mathbf{z} = d$ 
    find distances from all  $\mathbf{p} \in \mathcal{P}$  to the plane  $(a, b, c, d)$ 
    store points  $\mathbf{p}^* \in \mathcal{P}$  that satisfy the model hypothesis
       $0 \leq |d| \leq |d_{max}|$ .
    end for
  return maximum of the stored points  $\mathbf{p}^*$ .
end function

```

convex hull of point indices of objects above the table using the *Qhull* library², (ii)

² The Qhull library: <http://www.qhull.org/>

Algorithm 2 Euclidean Clustering Algorithm**Require:** a kd-tree data structure for the PCL $\mathcal{P}$ **Require:** a cluster of linked list $\mathcal{C}$, and a queue, $\mathcal{Q}$ **function** ClusterEuclid **for** points $\mathbf{p}_i \in \mathcal{P}$ **do** enqueue $\mathbf{p}_i$ to $\mathcal{Q}$ search for neighbors $\mathcal{P}_i^k$ of $\mathbf{p}_i$ in a sphere of radius $r < x$ add neighbors $\mathbf{p}_i^k \in \mathcal{P}_i^k$ that are not yet processed to $\mathcal{Q}$ add $\mathcal{Q}$ to $\mathcal{C}$ and reset $\mathcal{Q}$ to an empty list. **end for****return** $\mathcal{C}$.**end function**

with a pre-defined prismatic model candidate, we stored the extruded points to the approximate facial height above the table; and (iii) separating the face from every other point in the resulting cloud through the Euclidean clustering (EC) method of [30]. A distinct point cluster is defined if the points in cluster $C_i = \{\mathbf{p}_i \in \mathcal{P}\}$ and cluster $C_j = \{\mathbf{p}_j \in \mathcal{P}\}$ satisfy the L_2 -distance threshold

$$x \leq \min \|\mathbf{p}_i - \mathbf{p}_j\|_2 \quad (4)$$

whereupon the largest cluster returned by Algorithm 2 represents the face.

Finding the face in the scene after carrying out EC algorithm is a question of finding the largest index in the list $\mathcal{C}$ in linear time for n clusters. The face segmentation results are presented in Figure 5. We then compute the Cartesian position of the face with respect to the camera origin by taking the center of mass of the segmented facial region (bottom-right image of Figure 5). This is obtained by calculating the mean-value of all the points in the resulting cloud (≈ 600 points on average).

3.2 Head Pose Estimation

With the facial point cloud segmented, we define three points on the head. Our goal is to compute the optimal translation and rotation of the head from a model point set $\mathbf{X} = \{\vec{x}_i\}$ to a measured point set $\mathbf{P} = \{\vec{p}_i\}$, where $N_x = N_p = 3$, and the point $\vec{x}_i \in \mathbf{X}$ has the same index as $\vec{p}_i \in \mathbf{P}$. All point coordinates are with respect to a Cartesian frame in the stereo camera. We consider the world frame to have an origin at the centroid of the clustered point cloud of the head when all IABs are at ambient pressure, with the axes oriented as described in figure to the right of this page.

Following Besl and McKay's work in [31], we compute the cross-covariance matrix of $\mathbf{P}$ and $\mathbf{X}$ as Σ_{px} , extract the cyclic components of this skew symmetric matrix as Δ , and use it to form the symmetric 4×4 matrix $\mathbf{Q}(\Sigma_{px})$ as follows,

$$\mathbf{Q}(\Sigma_{px}) = \begin{bmatrix} \text{tr}(\Sigma_{px}) & \Delta^T \\ \Delta & \Sigma_{px} + \Sigma_{px}^T - \text{tr}(\Sigma_{px})\mathbf{I}_3 \end{bmatrix}. \quad (5)$$

The unit eigenvector, q_R , that corresponds to the maximum eigenvalue of $\mathbf{Q}(\Sigma_{px})$ is selected as the optimal rotation quaternion; we find the optimal translation vector as

$$\vec{q}_T = \vec{\mu}_x - \mathbf{R}(\vec{q}_R)\vec{\mu}_p \quad (6)$$

where μ_x and μ_p are the mean of point sets $\mathbf{X}$ and $\mathbf{P}$ respectively. Obtaining the roll, pitch and yaw angles from q_R is trivial and the pose of the face is described by tuples $[q_T, q_R] = \{x, y, z, \theta, \phi, \psi\}$ with respect to the world frame. Given the 3-DOF setup, we choose to control three states of the head: z, θ, ϕ (*i.e.* z , roll, and pitch).

3.3 Control Design

We propose an adaptive control strategy in a Bayesian setting, which given an initial prior distribution of controls and 3-DOF head pose, minimizes a cost criterion as the expected value of control laws that will yield a future desired head pose. We consider the pwm voltages that power the valves as input, $\mathbf{u}$, the head pose as the output, $\mathbf{x}$ and an unknown disturbance $\mathbf{w}(k)$. We first describe the nonlinear function approximator model $\hat{f}(\mathbf{u}(k-d), \mathbf{x}(k), \mathbf{w}(k))$, which is constructed from memory-based input and output experimental data that satisfy

$$Z^N = \{u(k), u(k-1), \dots, u(k-n_u), x(k), x(k-1), \dots, yx(k-n_x)\} \quad (7)$$

that satisfy the Lipschitz continuity. (7) implies an input $\mathbf{u}(\cdot)$ at time $k-d$, produces an output $\mathbf{x}(k)$ at d time instants later. The next section describes how we formulate the class of minimum error variance controllers that predict the effect of actions $\mathbf{u}(\cdot)$ on states $\mathbf{x}(\cdot)$ using a self-tuning regulator.

3.4 Adaptive Neuro-Control Formulation

We fixed a persistently exciting input signal $u_{ex} \in L_2 \cap L_\infty$ to excite the nonlinear modes of the system. We then parameterized the system with a neural network with sufficient number of neurons. The neural network (NN) provided information on the changing parameters of the system during control trials. The adjustment mechanism is computed from inverse Lyapunov analysis, where we choose adaptive laws that guarantee a nonpositive-definite Lyapunov function candidate when evaluated along the trajectories of the error dynamics.

Our contribution is the approximation of the nonlinear system by a long short-term memory (LSTM) [32], equipped with an adequate number of neurons in its hidden layers. We parameterized the last layer of the network with a fully connected




Fig. 6: Function Approximator Model

layer that outputs control torques to the valves. The neural network can be seen as a memory-based model that remembers effective controls for the adaptation mechanism in the presence of uncertainties and external disturbance.

The neural network is shown in Figure 6. Depending on the region of attraction of the system the network is approximating, it parameterizes the nonlinear dynamical system $f(\cdot)$ and maps the parameterized model to appropriate valve torques. There exists additional feedforward + feedback terms in the global controller (introduced shortly) that guarantee system stability and robustness to uncertainties. Therefore, the global controller keeps the states of the system bounded under closed-loop dynamics, ensures convergence to desired trajectories from states that are initialized outside the domain of attraction, and guarantees robust reference tracking in the presence of non-parametric uncertainties.

For the multi-input, multi-output (MIMO) adjustable system,

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{A}(\mathbf{u} - f(\mathbf{x}, \mathbf{u})) + \mathbf{w}(k) \quad (8)$$

where $\mathbf{x} \in R^n$, $\mathbf{u} \in R^m$ are known input and output vectors, and $\mathbf{A} \in R^{n \times n}$, $\mathbf{A} \in R^{m \times m}$ are unknown matrices, $\mathbf{B} \in R^{n \times m}$, $\text{sgn}(\mathbf{A})$ are known matrices, and $\mathbf{w}(k) \in R^n$ is a bounded time-varying unknown disturbance, upper-bounded by a fixed positive scalar $\mathbf{w}_{max}$. with a Hurwitz matrix $\mathbf{A}_m \in R^{n \times n}$ and $\mathbf{B}_m \in R^{n \times m}$ commanded by a reference signal $\mathbf{r} \in R^m$. For this system, we note that $n = 3$ and $m = 6$. We choose a model-reference adaptive controller (MRAC) capable of operating in the presence of parametric (ε_f), and non-parametric ($\mathbf{w}(k)$) uncertainties so as to assure the boundedness of all signals within the closed-loop system. Our controller is of the following form

$$\mathbf{u} = \hat{\mathbf{K}}_x^T \mathbf{x} + \hat{\mathbf{K}}_r^T \mathbf{r} + \hat{f}(\mathbf{x}, \mathbf{u}), \quad (9)$$

where $\hat{\mathbf{K}}_x$ and $\hat{\mathbf{K}}_r$ are adaptive gains to be designed shortly. The $\hat{\mathbf{K}}_x^T \mathbf{x}$ term keeps the states of the approximation set $\mathbf{x} \in \mathbf{B}_R$ stable, while the $\hat{\mathbf{K}}_r^T \mathbf{r}$ term causes the states to follow a given reference trajectory. The function approximator $\hat{f}(\cdot)$ ensures states that start outside the approximation set $\mathbf{x} \in \mathbf{B}_R$ converge to $\mathbf{B}_R$ in finite time (it converges non-parametric errors ε_f that puts certain states out of the approximation set into $\mathbf{B}_R$). We can generally write the NN model as

$$\hat{f}(x) = \hat{\Theta}^T \Phi(\mathbf{x}) + \varepsilon_f,$$

where $\hat{\Theta}^T$ denotes the vectorized weights of the neural network and $\Phi(\mathbf{x})$ denotes the vector of Lipschitz inputs and outputs, and ε_f is the approximation error. The closed-loop dynamics therefore become

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{A} \left(\hat{\mathbf{K}}_x^T \mathbf{x} + \hat{\mathbf{K}}_r^T \mathbf{r} + \hat{f}(\cdot) - f(\cdot) \right). \quad (10)$$

Theorem: Given correct choice of adaptive gains $\hat{\mathbf{K}}_x$ and $\hat{\mathbf{K}}_r$, the error vector $\mathbf{e}(k) = \dot{\mathbf{x}}(k) - \dot{\mathbf{x}}_m(k)$ will be uniformly ultimately bounded, and the state $\mathbf{x}$ will converge to a neighborhood of $\mathbf{r}$.

The main matter of the proof to the above theorem is in [21].

3.5 Network Design

We require accurate mapping of temporally lagged patterns in inputs to output states, a dynamic nonlinear model of valve encoder values to sensor measurements that accurately maps $f(\cdot)$ in (8). We choose a LSTM [32] due to its capacity for long-term context memorization and inherent multiplicative units that avoid oscillating weights or vanishing gradients when error signals are backpropagated in time [32, 33]. LSTMs truncate gradients in the network where it is harmless by enforcing constant error flows through their *constant error carousels*. As a result, LSTMs are robustly more powerful for adaptive sequence-to-sequence modeling or mapping data that temporally evolve in time. Their biological model makes them more suitable for adaptive robotics such as soft robots than previously used artificial NNs such as feedforward networks [34], radial basis-functions [35, 36] or vanilla RNNs [37].

The NN model takes a memory-based concatenated vector of current inputs and past outputs, propagates them through three hidden layers, with each layer made up of $\{9, 6, 6\}$ neurons each, applies 30% dropout and then maps the last layer to a fully connected layer that generates valve torques. The architecture of the neuro-controller is shown in Figure 6. The last layer is designed to generate appropriate valve torques based on an internal model of the plant. A self-tuning adaptive control law (with a feedforward regulation and state feedback component)

adapts to the internal parameters of the plant to ensure stability of the system and bounded tracking of given trajectory. The overall network has neuron connection weights and thresholds of approximately 1,400. This makes search for a suitable controller feasible.

The LSTM model estimates a model $f(\mathbf{x})$, that minimizes the mean-squared error between predicted output $\hat{\mathbf{x}}(k)$ and actual output $\mathbf{x}(k)$ according to

$$f(\mathbf{x}(k)) = \arg \min_w \mathbf{V}_N(w, \Phi(\mathbf{x})) \quad (11)$$

$$\text{where } \mathbf{V}_N(\mathbf{w}, \Phi(\mathbf{x})) = \sum_{t=1}^{\mathcal{K}} \sum_{i=1}^n \frac{1}{2} (\hat{\mathbf{x}}_i(t) - \mathbf{x}_i(t))^2, \quad (12)$$

and $\Phi(\mathbf{x})$ is the regression vector on a bounded interval $[1, N]$. (11) is minimized using stochastic gradient descent so that at each iteration, we update the parameters (weights) of the network $\mathbf{w}_i$ based on the *ordered derivatives* of $\mathbf{V}_N(\mathbf{w}, \Phi(\mathbf{x}))$ (Werbos [38]) *i.e.*

$$\mathbf{w}_{k+1} \leftarrow \eta \mathbf{w}_k - \alpha \sum_{i=1}^n \nabla_{\mathbf{w}} \mathbf{V}(\mathbf{x}_i, \hat{\mathbf{x}}_i(\theta_k)). \quad (13)$$

η (set to 1) hastens the optimization in a direction of low but steepest descent in training error, and α is a sufficiently small learning rate (set to 5×10^{-3}), and $\nabla_w \mathbf{V}(\theta, \Phi(\mathbf{x}))$ is the derivative of $\mathbf{V}$ with respect to w averaged over the k -th batch (we used a batch size of 50). We initialized the weights of Figure 6 from a one-dimensional normal distribution with zero-mean and unit variance.

4 Results from Experiments

The 3-DOF pose of the head is made up of the state tuple $\{z(k), \theta(k), \phi(k)\}$. We sample from the parameters of the trained network and we set $\hat{f}(\cdot)$ in (9) to the fully connected layer of samples from the network. We publish the control law from the neural network and subscribe in a separate node. The gains $\hat{\mathbf{K}}_x$ and $\hat{\mathbf{K}}_r$ in (??), were found by solving the ODEs iteratively using a single step of the integral of the solutions to $\hat{\mathbf{K}}_x(t)$, $\hat{\mathbf{K}}_r(t)$. Our solution is an implementation of the Runge-Kutta Dormand-Prince 5 ODE-solver available in the Boost C++ Libraries³. We found a step-size of 0.01 to be realistic. $\mathbf{x}_m$ is computed based on the solution to the forced response of the linear system,

$$\mathbf{x}_m(t) = e^{\mathbf{A}_m t} \mathbf{x}_m(0) + \int_0^t e^{\mathbf{A}_m (t-\tau)} \mathbf{B}_m \mathbf{r}(\tau) d\tau.$$

³ <https://goo.gl/l7JyYe>

We set $\mathbf{x}_m(0) = \mathbf{x}(0)$ at $t = 0$ and for a settling time requirement of $T_s = 5\text{secs}$ at which the response remains within 2% of final value, we find that

$$\mathbf{A}_m = \begin{bmatrix} -\frac{1334}{1705} & 0 & 0 \\ 0 & -\frac{1334}{1705} & 0 \\ 0 & 0 & -\frac{1334}{1705} \end{bmatrix}. \quad (14)$$

For a nonnegative $\mathbf{Q}$ and a positive definite $\mathbf{P}$, the pair $(\mathbf{Q}, \mathbf{A}_m)$ will be observable (LaSalle's theorem) so that the dynamical system is globally asymptotically stable. We choose a positive definite $\mathbf{Q} = \text{diag}(100, 100, 100)$ as the dissipation energy and set $\mathbf{A} = I_{3 \times 3}$ so that solving the general form of the lyapunov equation, we have

$$\mathbf{P} = \begin{bmatrix} -\frac{170500}{2668} & 0 & 0 \\ 0 & -\frac{170500}{2668} & 0 \\ 0 & 0 & -\frac{170500}{2668} \end{bmatrix} \quad (15)$$

The six solenoid valves operate in pairs so that two valves create a difference in air mass within each IAB at any given time. Therefore, we set

$$\mathbf{B} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \quad (16)$$

The non-zero terms in (16) denote the maximum duty-cycle that can be applied to the Dakota valves based on the software configuration of the NI RIO PWM generator.

4.1 Analysis

The three DoFs of the head are coupled and there is a limited reachable space with the IABs. It is therefore paramount that desired trajectories be ascertained as physically realizable before rolling out control trials. We therefore placed the head to physically realizable positions in open-loop control settings before testing the close-loop control system on such feasible goal poses.

Figure 7 show the performance of the controller when commanded to move the head from $[z, \theta, \phi]^T = [2.5\text{mm}, .25^\circ, 35^\circ]^T$ to $[14\text{mm}, 1.6^\circ, 45^\circ]^T$. We observe strong steady-state convergence along 2-DoFs, namely z and pitch axes with a 20 second rise time. The roll motion is however characterized by offshoots that may be caused by the coupled DOF. We perform a second experiment, seen in Figure 8, where we evaluate the performance of the controller on the roll angle of the head. We observe that the controller behaves well controlling the roll motion in isolation.

5 Conclusion




Fig. 7: Head motion correction along z, pitch and roll axes.




Fig. 8: Head motion correction along roll axis.

References

1. American Cancer Society: Cancer facts and figures (2019). URL <https://www.cancer.org/research/cancer-facts-statistics/all-cancer-facts-figures/cancer-facts-figures-2019.html>. Accessed October 23, 2019 ¹
2. IARC brochure: A unique agency (2019). URL https://www.iarc.fr/wp-content/uploads/2019/07/IARC-brochure-EN-June_2019.pdf. Accessed October 23, 2019 ¹
3. Ringborg, U., Bergqvist, D., Brorsson, B., Cavallin-Ståhl, E., Ceberg, J., Einhorn, N., Frödin, J.e., Järhult, J., Lamnevik, G., Lindholm, C., et al.: The Swedish Council on Technology Assessment in Health Care (SBU) systematic overview of radiotherapy for cancer including a prospective survey of radiotherapy practice in Sweden 2001–Summary and Conclusions. *Acta Oncologica* **42**(5-6), 357–365 (2003) ²

4. Baskar, R., Lee, K.A., Yeo, R., Yeoh, K.W.: Cancer and radiation therapy: current advances and future directions. *International journal of medical sciences* **9**(3), 193 (2012) [3](#)
5. Webb, S.: *Intensity-Modulated Radiation Therapy*. Institute of Physics Publishing Ltd, Bristol and Philadelphia (2001) [3](#)
6. Schweikard, A., Tombropoulos, R., Adler, J.R.: Robotic Radiosurgery with Beams of Adaptable Shapes. In: N. Ayache (ed.) *Computer Vision, Virtual Reality and Robotics in Medicine*, pp. 138–149. Springer Berlin Heidelberg, Berlin, Heidelberg (1995) [3](#)
7. Lee, S., Chang, K.H., Shim, J.B., Cao, Y., Lee, C.K., Cho, S.J., Yang, D.S., Park, Y.J., Yoon, W.S., Kim, C.Y.: Evaluation Of Mechanical Accuracy For Couch-based Tracking System (CBTS). *Journal of applied clinical medical physics* **13**(6), 157–169 (2012) [3](#)
8. Gevaert, T., Verellen, D., Tournel, K., Linthout, N., Bral, S., Engels, B., Collen, C., Depuydt, T., Duchateau, M., Reynders, T., et al.: Setup Accuracy Of The Novalis Exactrac 6dof System For Frameless Radiosurgery. *International Journal of Radiation Oncology* Biology* Physics* **82**(5), 1627–1635 (2012) [3](#)
9. Webb, S.: Conformal intensity-modulated radiotherapy (imrt) delivered by robotic linac-conformality versus efficiency of dose delivery. *Physics in Medicine and Biology* **45**(7), 1715 (2000) [3](#)
10. Xing, L., Lin, Z.X., Donaldson, S.S., Le, Q.T., Tate, D., Goffinet, D.R., Wolden, S., Ma, L., Boyer, A.L.: Dosimetric Effects Of Patient Displacement And Collimator And Gantry Angle Misalignment On Intensity Modulated Radiation Therapy. *Radiotherapy and Oncology* **56**(1), 97–108 (2000) [3](#)
11. Takakura, T., Mizowaki, T., Nakata, M., Yano, S., Fujimoto, T., Miyabe, Y., Nakamura, M., Hiraoka, M.: The geometric accuracy of frameless stereotactic radiosurgery using a 6d robotic couch system. *Physics in Medicine & Biology* **55**(1), 1 (2009) [3](#)
12. Keall, P.J., Mageras, G.S., Balter, J.M., Emery, R.S., Forster, K.M., Jiang, S.B., Kapatoes, J.M., Low, D.A., Murphy, M.J., Murray, B.R., et al.: The Management of Respiratory Motion in Radiation Oncology Report of AAPM Task Group 76 A. *Medical physics* **33**(10), 3874–3900 (2006) [5](#)
13. Hunt, K.H.: Structural Kinematics of In- Parallel-Actuated Robot-Arms **105**(December 1983), 705–712 (2019) [5](#)
14. Liu, X., Belcher, A.H., Grelewicz, Z., Wiersma, R.D.: Robotic stage for head motion correction in stereotactic radiosurgery. In: 2015 American Control Conference (ACC), pp. 5776–5781. IEEE (2015) [5](#)
15. Herrmann, C., Ma, L., Schilling, K.: Model Predictive Control For Tumor Motion Compensation In Robot Assisted Radiotherapy. *IFAC Proceedings Volumes* **44**(1), 5968–5973 (2011) [5](#)
16. Belcher, A.: Patient Motion Management with 6-DOF Robotics for Frameless and Maskless Stereotactic Radiosurgery. Ph.D. thesis, The University of Chicago (2017) [6](#)
17. Ogunmolu, O., Gans, N., Jiang, S., Gu, X.: An Image Guided Soft Robotic Patient Positioning System for Maskless Head And Neck Cancer Radiotherapy: A Proof of Concept Study. *Medical Physics: The International Journal of Medical Physics Research and Practice* **42**, 3266–3266 (2015) [6](#)
18. Ogunmolu, O.P., Gu, X., Jiang, S., Gans, N.R.: A real-time, soft robotic patient positioning system for maskless head-and-neck cancer radiotherapy: An initial investigation. In: Automation Science and Engineering (CASE), 2015 IEEE International Conference on, Gothenburg, Sweden, pp. 1539–1545. IEEE (2015) [6](#)
19. Ogunmolu, O., Folkerts, M., Nguyen, D., Gans, N., Jiang, S.: Deep BOO! Automating Beam Orientation Optimization in Intensity Modulated Radiation Therapy. In: Algorithmic Foundations of Robotics, XIII Workshop, Merida, Mexico. Published in Springer's Proceedings in Advanced Robotics (SPAR) book (2018) [6](#)
20. Ogunmolu, O.P., Gu, X., Jiang, S., Gans, N.R.: Vision-based control of a soft robot for maskless head and neck cancer radiotherapy. In: Automation Science and Engineering (CASE), 2016 IEEE International Conference on, Fort Worth, Texas, pp. 180–187. IEEE (2016) [6](#)
21. Ogunmolu, O., Kulkarni, A., Tadesse, Y., Gu, X., Jiang, S., Gans, N.: Soft-neuroadapt: A 3-dof neuro-adaptive patient pose correction system for frameless and maskless cancer radiotherapy.

- In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vancouver, BC, CA, pp. 3661–3668. IEEE (2017) 6, 14
22. Almubarak, Y., Aniket, J., Ogunmolu, O., Gu, X., Jiang, S., Gans, N., Tadesse, Y.: Design and development of soft robots for head and neck cancer radiotherapy. In: SPIE: Smart Structures + Nondestructive Evaluation. SPIE (2018) 6
 23. Ogunmolu, O., Liu, X., Wiersma, R.: Mechanism and Constitutive Model of a Continuum Robot for Head and Neck Cancer Radiotherapy (2019) 6
 24. Ogunmolu, O.P.: A Multi-DOF Soft Robot Mechanism for Patient Motion Correction and Beam Orientation Selection in Cancer Radiation Therapy. Ph.D. thesis, The University of Texas at Dallas; UT Southwestern Medical Center (2019) 6
 25. Ensenso: Flexview. <http://www.ensenso.com/products/flexview/>. Accessed on January 21, 2016. 7
 26. Rusu, R.B.: Semantic 3D object Maps for Everyday Manipulation in Human Living Environments. PhD thesis (2009) 8, 9
 27. Torr Philip HS and Zisserman, A.: MLESAC: A New Robust Estimator with Application to Estimating Image Geometry. *Computer Vision and Image Understanding* **78**(1), 138–156 (2000) 8, 9
 28. Rusu, R.B., Marton, Z.C., Blodow, N., Dolha, M.E., Beetz, M.: Functional Object Mapping of Kitchen Environments. 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS pp. 3525–3532 (2008). DOI 10.1109/IROS.2008.4650972 8
 29. Rusu, R.B., Cousins, S.: 3D is here: Point Cloud Library (PCL). In: IEEE International Conference on Robotics and Automation (ICRA). Shanghai, China (2011) 9
 30. Rusu, R.B., Holzbach, A., Beetz, M., Bradski, G.: Detecting and Segmenting Objects for Mobile Manipulation. In: Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on, pp. 47–54. IEEE (2009) 11
 31. Besl, Paul J.; McKay, N.D.: A Method for Registration of 3D Shapes. (1992) 11
 32. Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. *Neural computation* **9**(8), 1735–80 (1997) 12, 14
 33. Bengio, Y., et al.: Learning Long-term Dependencies with gradient Descent is Difficult. *IEEE Transactions on Neural Networks* (1994). Doi: 10.1109/72.279181 14
 34. Dinh, H., Bhasin, S., Kamalapurkar, R., Dixon, W.E.: Dynamic Neural Network-based Output Feedback Tracking Control for Uncertain Nonlinear Systems. *Journal of Dynamic Systems, Measurement, and Control* (2017) 14
 35. Patino, H., Liu, D.: Neural network-based model reference adaptive control system. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **30**(1), 198–204 (2000) 14
 36. Lavretsky, E., Wise, K.: Robust Adaptive Control with Aerospace Applications. Springer (2005) 14
 37. Wang, J.S., Chen, Y.P.: A Fully Automated Recurrent Neural Network for Unknown Dynamic System Identification and Control. *IEEE Transactions on Circuits and Systems* **53** (2006) 14
 38. Werbos, P.J.: Backpropagation Through Time: What It Does and How to Do It. *Proceedings of the IEEE* **78**(10), 1550–1560 (1990). DOI 10.1109/5.58337 15