



HHS Public Access

Author manuscript

IEEE Trans Med Imaging. Author manuscript; available in PMC 2021 November 01.

Published in final edited form as:

IEEE Trans Med Imaging. 2020 November ; 39(11): 3523–3534. doi:10.1109/TMI.2020.2998600.

Deep Predictive Motion Tracking in Magnetic Resonance Imaging: Application to Fetal Imaging

Ayush Singh,

Department of Radiology, Boston Children's Hospital, and Harvard Medical School, Boston, Massachusetts, USA

Seyed Sadegh Mohseni Salehi,

Hyperfine Research Inc.

Ali Gholipour [Senior Member, IEEE]

Department of Radiology, Boston Children's Hospital, and Harvard Medical School, Boston, Massachusetts, USA

Abstract

Fetal magnetic resonance imaging (MRI) is challenged by uncontrollable, large, and irregular fetal movements. It is, therefore, performed through visual monitoring of fetal motion and repeated acquisitions to ensure diagnostic-quality images are acquired. Nevertheless, visual monitoring of fetal motion based on displayed slices, and navigation at the level of stacks-of-slices is inefficient. The current process is highly operator-dependent, increases scanner usage and cost, and significantly increases the length of fetal MRI scans which makes them hard to tolerate for pregnant women. To help build automatic MRI motion tracking and navigation systems to overcome the limitations of the current process and improve fetal imaging, we have developed a new real-time image-based motion tracking method based on deep learning that learns to predict fetal motion directly from acquired images. Our method is based on a recurrent neural network, composed of spatial and temporal encoder-decoders, that infers motion parameters from anatomical features extracted from sequences of acquired slices. We compared our trained network on held-out test sets (including data with different characteristics, e.g. different fetuses scanned at different ages, and motion trajectories recorded from volunteer subjects) with networks designed for estimation as well as methods adopted to make predictions. The results show that our method outperformed alternative techniques, and achieved real-time performance with average errors of 3.5 and 8 degrees for the estimation and prediction tasks, respectively. Our real-time deep predictive motion tracking technique can be used to assess fetal movements, to guide slice acquisitions, and to build navigation systems for fetal MRI.

Index Terms—

Convolutional neural network; Recurrent neural network; Long short term memory; fetal MRI; Motion tracking; Pose estimation; Prediction; Image registration; MRI

Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

ayush.singh@childrens.harvard.edu.

I. Introduction

A. Motivation

MAGNETIC Resonance Imaging (MRI) is a relatively slow imaging technique hence it is extremely susceptible to subject motion. To deal with this limitation, when MRI scans are performed, subjects are instructed to stay completely still. To scan newborns and young children, this requires strategies such as feed-and-wrap, padding, or training, whichever is applicable, to restrain or reduce motion [1]–[3]. There has been extensive research and development in motion-robust sequences and motion correction techniques in MRI (e.g. [4]–[10]), however none of these techniques can be universally applied to all MRI sequences and all patient populations. For example none of the above-referenced techniques can be used for motion tracking in fetal MRI, as discussed next.

Among all rapidly-emerging MRI applications, fetal MRI is, arguably, one of the most challenging, due to uncontrollable, large, and irregular fetal movements [11]. In particular, in midgestation fetuses have enough space to stretch and rotate in large angles. Fetal motion is complex and cannot be monitored or tracked by external sensors or camera systems or accounted for by cardiac and/or respiratory gating. Fetal MRI motion correction techniques have thus relied upon *retrospective* image registration solely based on image information [12]–[20].

Slice-to-volume registration, which has been widely used in retrospective fetal MRI reconstruction, is inherently an ill-posed problem [21]. It has a limited capture range as it relies on iterative optimization of intensity-based similarity metrics that are only surrogate measures of alignment between a reference volume and slices. Moreover, a motion-free reference volume may or may not be readily available. To increase capture range, one may use grid search on rotation parameters along with multi-scale registration [22]; but this approach is computationally expensive as it is based on iterative numerical optimization at test time. For reference volumes, one may use age-matched atlases, e.g. [23], and perform atlas-based registration, e.g. [22], [24], however these methods are also computationally expensive for *real-time* application.

To improve capture range and the speed of subject-to-atlas image registration, in a recent work [25], deep regression convolutional neural networks (CNNs) were trained to estimate 3D pose of the fetal brain based on image slices and volumes. Partly inspired by [25], in this paper we present a novel *deep predictive motion tracking* framework based on long short term memory (LSTM) [26] recurrent neural networks (RNNs). While the technique in [25] addressed *static 3D pose estimation* only (based on regression CNNs), our work here addresses *dynamic, real-time, 3D motion tracking* in MRI, for the first time, using RNNs, exploiting LSTM modules and innovative learning strategies, that are explained in this paper. In static pose estimation we infer 3D pose of the anatomy based on one slice, whereas in dynamic motion tracking, we infer relative pose changes of the subject based on a time series of slices. Our proposed method, therefore, learns to predict motion trajectory based on MRI slice time series. While motivated by an unmet need in the application domain, our

technique was inspired by the most recent advances in computer vision, which are reviewed next, where we also review the related work in fetal MRI and MRI motion tracking.

B. Related Work

Pose estimation using 2D (digital) images and videos has been extensively researched in computer vision, where algorithms aim to find 3D pose of objects with respect to camera. Work in this area can be studied in two main groups: methods that predict key points leveraging object models to find object orientation, e.g. [27]; and methods that predict object pose directly from images to discrete pose space-bins, e.g. [28], [29] and [30]. While the majority of pose estimation techniques have been designed as classification methods, the problem has been recently modeled and solved by regression deep neural networks [31]. Deep CNNs have shown great performance in pose estimation in recent years, e.g. [31]–[34].

Three-dimensional pose estimation from 3D or stack-of-2D medical images has also been recently addressed using CNNs. For a review of the related pose estimation and registration methods we refer to [25]. For fetal MRI, in particular, deep regression CNNs were designed for slice-to-volume registration on non-Euclidean manifolds [35], and used to estimate transformation parameters for fetal head position to reconstruct fetal brain MRI volumes from slices [36]. Real-time fetal head pose estimation was achieved in [25] by multi-stage loss minimization using mean squared error and geodesic loss, and used for image-to-template and inter-subject rigid registration.

The above-referenced techniques treat image slices independently. Therefore, while they are powerful in that they learn to predict head position based on single slices (or volumes), they ignore the rich information content of stack of sequentially acquired slices and the dynamics of head motion. Consequently, these methods may be limited in their predictive performance as they ignore (or do not model) the dynamics of motion (e.g., the motion velocity). Moreover, the average 3D pose estimation error of these methods is often high for slices in the boundaries of the anatomy where image features are sparse [37]. While pose estimation methods can be combined with iterative slice-to-volume registration for head motion tracking, e.g. [19]; a natural, promising extension of this line of work is dynamic image time series modeling, which has been the subject of our work presented in this paper. In our experiments, we compared our predictive motion tracking technique with zero-velocity and auto-regressive prediction models built upon static 3D pose estimation methods.

Traditional time series prediction models such as ARIMA (auto-regressive integrated moving average; seasoned, and non-seasoned) expect data to be locally stationary. These are regression models that make strong assumptions about data to predict future values based on past observations. These models shall be paired with other techniques to effectively process and use image time series information; but this integration may not be straightforward. RNNs [38], on the other hand, can handle non-stationary and nonlinear data. They offer end-to-end framework to take images as input and make predictions, and are flexible in terms of the corresponding objectives.

Variants of RNNs such as networks based on LSTM [26] have the capacity to learn the amount of information to remember and forget from past sequences. This makes them less susceptible to unaccounted cases that cannot be easily handled by graph designer of dynamic Bayesian networks (DBNs) [39]. Compared to traditional models where error propagation leads to error accumulation in long-term prediction, advanced LSTM-based methods, such as sequence-to-sequence (Seq2Seq) learning [40], can reliably predict variable time steps with long prediction horizons.

Deep predictive motion tracking using RNNs based on video sequences has also been widely studied in robotics and computer vision, e.g. [40]–[42]. A review of these studies is beyond the scope of this paper, but we briefly review some representative methods and studies. The first group of techniques based on siamese networks detect and use regions close to object locations to track objects, e.g. [43], [44]. Large datasets can be used to train these networks for feature extraction and region proposals for simultaneous one-shot detection (classification) and online tracking (regression) [45]. Early performance gains in accuracy were obtained by passing features from an object detector to LSTMs [46]. In the LSTM category, the Real-time Recurrent Regression (Re³) network [47] combined non-differentiable cropping and warping with feature extraction using a residual network (ResNet), and passed them to LSTM for object tracking.

C. Contributions

In this paper we present, for the first time, a dynamic motion tracking framework for MRI based on deep learning. Compared to recent developments in static 3D pose estimation from MRI slices and volumes based on CNNs [25], [36], in this work we exploit RNNs for predictive dynamic motion tracking. Compared to motion tracking in computer vision, robotics, digital image and video processing, where 3D pose or projected motion of objects is modeled and estimated based on 2D+time images (videos) with respect to cameras, in this work we deal with 3D rigid motion of anatomy (in the scanner/world coordinate system) from stacks of sequentially acquired slices (3D+time image time series). Consequently, while the majority of human pose tracking or video object tracking methods are formulated and solved as classification problems in a parameter space, we solve a regression problem where 3D rigid motion parameters are estimated based on features directly extracted from MRI time series.

Our contributions are threefold: 1) We developed a learning-based, image-based, real-time dynamic motion tracking in MRI based on deep RNNs: Our model encodes motion using LSTM after extracting spatial features from sequences of input images using CNNs, estimates objectives for given images and creates a context vector that is used by LSTM decoders to regress against angle-axis representation and translation offset to predict 3D rigid body motion. The network constitutes multiple representation heads to avoid overfitting to either rotation or translation parameters. 2) We devised multi-step prediction by feeding output of previous decoder as input to current decoder combined with the context vector. 3) We trained and tested networks on sequences with masked slices that are slices lost due to intermittent fast intra-slice motion.

We developed and tested our method for fetal head motion tracking in fetal MRI, which is a very challenging problem due to the wide range of fetal head positions and motion; but the technique can be used in broader applications. The fetal brain MRI data intrinsically shows a wide feature range due to inter-subject variability and different age of fetuses at the time of MRI scans as well as rapid changes that occur to the fetal brain during gestation. To train and test models we used images of different fetuses scanned at different gestational ages. We simulated motion and also used motion trajectories from sensor recordings of head motion of volunteer subjects to test the generalization capacity of our trained network. We set up a probing task to examine temporal and spatial dependency of our trained model. Our network infers motion parameters from features extracted from 2D slice time series, therefore it does not require coverage of the entire brain in 3D and hence does not require data that are on a regular grid. Our experiments showed that our trained model not only estimated motion trajectories but also was able to make long term predictions based on sequences of fetal brain MRI slices with both simulated and real motion in the test set. The paper is organized as follows: the details of our network and methods are discussed next. Then, the experiments and experimental results are described in Section III; which are followed by a discussion in Section IV and conclusion in Section V.

II. Methods

A. Problem formulation

Our goal is to take in a sequence of slices X_1, X_2, \dots, X_n ($X_n: N \times N$) sampled sequentially (in time) from 3D fetal anatomy (usually acquired in an interleaved manner) in an MRI scan to estimate and predict 3D pose (rotation and slice position) Y_1, Y_2, \dots, Y_{n+m} of the fetal brain for current n timesteps as well as future m timesteps (timestep unit defined in Section III-A). Our technique does not put any restriction on the values of n and m . Although n is limited by the number of input slices, m can be variable *i.e.* either less, equal or greater than n . The slices are from a stack of sliced anatomy where the anatomy moves in 3D in between slice acquisitions. For the purpose of this study we assume that the fetal brain is extracted in each slice using a real-time fetal brain MRI segmentation method [48]. For the development and evaluation of predictive motion tracking, we also assume that center-aligned slices are extracted from 3D fetal brain images reconstructed and segmented using the existing techniques [15], [49].

Figure 1 shows how the data is pre-processed and prepared for fetal head motion tracking. The region-of-interest (RoI), which is the fetal brain in this study, is first extracted using a real-time brain extraction method [48] and the slices are cropped, masked, and center-aligned to form a 3D stack. For slices that are corrupted by intra-slice motion (causing full or partial signal loss), the brain extraction method does not generate brain masks that are coherent between those slices and their spatially neighboring slices. The motion-corrupted slices can, therefore, be detected by statistical or learning based methods (e.g., outlier detection [16], [37] or support vector machines [19]). Hence, fetal motion appears as inter-slice motion with occasional black (masked) slices due to intra-slice motion. The problem is formulated as finding 3D rigid transformations, T , relative to the starting slice X_1 , of the fetal head at the times corresponding to slice X_j acquisitions.

A 3D rigid-body transformation T has 6 degrees-of-freedom represented by a vector t comprising of three translation (t_x, t_y, t_z) and three rotation ($\theta(\theta_x, \theta_y, \theta_z)$) parameters. For 3D rotation representation we follow [25] which uses Euler's theorem and the Rodrigues rotation formula to represent the 3×3 rotation matrix by the angle-axis representation where the rotation axis is its unit vector and the angle in radians defines its magnitude. Since we center align the images in the pre-processing step, the translation parameters are assumed to be known a priori, which allows us to constrain our parameter space to the slice position z and the rotations θ represented by the angle-axis formalism. The methods in [25] can be used to estimate the initial pose and the a priori translation parameters.

B. Deep regression RNN for predictive motion tracking

As shown in Figure 2, our deep RNN model for predictive slice-level motion tracking in MRI is built of two main parts: an encoder and a decoder. The encoder network, which is composed of deep CNN blocks followed by unidirectional LSTM and P blocks, takes a sequence of slices X_1, \dots, X_n as input, and estimates a sequence of n transformations as well as an encoder state, which is fed into the decoder network. Conditioned on the encoder state, the decoder network, which also constitutes LSTM and P blocks, predicts transformations for future time steps m . A P block involves three representation heads, each consisting of a dense block and an activation function for regression at the output layer. The activation functions are $\pi \tanh$ for the rotation parameters θ and rectified linear unit (ReLU) for slice position shown here by z . In the sections that follow we discuss each of the network components and the details of training.

C. Encoder: Spatial

For spatial encoding, convolutions [50] are applied to each slice X_n of a sequence where n is the index of the slice in the sequence. Figure 3 shows the architecture of the spatial encoder network. Through weight sharing the same CNN is trained and applied to all slices. This means there is no dedicated network for each timestep. Instead during training, kernel weights of the same CNN are updated to account for variations in all timesteps. This allows the spatial encoder CNN to learn anatomical variations between different ages, and pass the encoded information into the temporal encoder. We used parametric rectified linear unit (PReLU) as activation function as it has shown better performance than ReLU [51]. PReLU avoids the dying ReLU problem, in which a neuron (with ReLU activation) becomes inactive when it gets negative input making the gradient of an inactive neuron zero, hence unable to pass any information via backpropagation.

D. Encoder: Temporal

Just as CNN learns spatial variations, RNN learns variations between elements in a sequence. Since vanilla RNNs face the vanishing gradient problem [52], which makes it difficult to propagate gradients back in time, we used LSTM [26], which also learns what to remember and what to forget. This is important to learn the anatomy and how it is sampled by slices over time using the gating mechanism. Based on encoded image features from the CNN, the LSTM learns to estimate the state of the anatomy, *i.e.* the 3D pose of the anatomy and its sampling. LSTM has three primary components: W, U, b , where W is the recurrent

connection between previous and current hidden layers, U connects inputs to current hidden layer and b is bias:

$$LSTM_{encoder}: X_n, h_{n-1}, c_{n-1} \rightarrow h_n, c_n \quad (1)$$

$$i_n = \sigma(W_i X_n + U_i h_{n-1} + b_i) \quad (2)$$

$$f_n = \sigma(W_f X_n + U_f h_{n-1} + b_f) \quad (3)$$

$$o_n = \sigma(W_o X_n + U_o h_{n-1} + b_o) \quad (4)$$

$$\hat{c}_n = \tanh(W_{\hat{c}} X_n + U_{\hat{c}} h_{n-1} + b_{\hat{c}}) \quad (5)$$

$$c_n = f_n \odot c_{n-1} + i_n \odot \hat{c}_n \quad (6)$$

$$h_n = o_n \odot \tanh(c_n) \quad (7)$$

For each time step n , the memory cell $c_n \in \mathbb{R}^n$ is called as it controls exposure of the previous memory c_{n-1} with current input X_n . This is done by combining c_{n-1} multiplied by the forget gate f_n , with the computed hidden state h_n multiplied by the input gate i_n . These are called gates because they squash values between 0 and 1 using the sigmoid activation function σ . The element-wise multiplication \odot controls how much of information is let through: The input gate controls how much of the current input goes through; the forget gate controls the throughput of the previous state; and the output gate controls the amount of exposure of the internal states to the next timesteps (or the downstream layers). All gates have dimensions equal to that of the hidden layer h_n , which is computed by multiplying the hyperbolic tangent \tanh of memory c_n with the output $o_n \cdot \hat{c}_n$ is the candidate hidden state that connects the current input X_n to the previous hidden state. One can ignore old memory completely (all zeros f_n) or ignore states (all zeros in i_n), but we chose to store nuances of changes in data over time thus the values were chosen to be between 0 and 1.

Flattened feature maps pass from the spatial encoder to the unidirectional LSTM network. Output of each time step of the encoder and decoder LSTM go through dense fully-connected layers to get estimated and predicted parameters. The last nonlinear function with weights $W_{\theta_{xyz}}$ on top of the dense layer is $\pi \times \tanh$ which limits the output of each element from $-\pi$ to $+\pi$ and simulates the constraints of each element of the rotation vector (θ_x , θ_y) and θ_z independently:

$$\theta_n^{xyz} = \pi \tanh(W_{\theta_{xyz}} o_n + b_{\theta_{xyz}}) \quad (8)$$

The slice index (z) estimator head with weights W_z contains a scalar, as the network tries to estimate the continuous slice index along with its orientation. For inference, the continuous index is rounded (i.e. $\lfloor z \rfloor$) to infer a discrete slice number.

$$z_n = \max(0, W_z o_n + b_z) \quad (ReLU) \quad (9)$$

E. Decoder: Modeling variable and long term predictions

The conventional approach to predict sequential data is to use n steps of the sequence from the past to predict the immediate future time step $n + 1$ and repeat recursively to make future predictions up until the desired prediction horizon. This model, however, shows limited multi-step prediction performance in applications such as image-based motion tracking as it faces issues raised by compounding errors especially when initial predictions may exhibit relatively large amounts of error. To mitigate this issue and make variable-length, long-term predictions we follow the idea of sequence to sequence learning [40]. In this approach, an LSTM encodes the input sequence of images into a fixed dimension vector, and another LSTM decodes the target sequence from this vector. The advantage of this technique is that we no longer need to rely on encoder estimates to predict variable-length time steps of the future as encoder and decoder are two separate LSTM networks. Figure 2 shows our LSTM network unrolled.

Each decoder is trained to predict parameters of the following step. Therefore, input to the first decoder is the estimation vector \hat{Y}_n of the last slice X_n from the encoder and the rest of the decoder takes output of the previous decoding step \hat{Y}_{m-1} so that over time the model learns to correct its own mistakes.

$$LSTM_{decoder}: \hat{Y}_{n+m-1}, h_{n+m-1}, c_{n+m-1} \rightarrow h_{n+m}, c_{n+m} \quad (10)$$

The goal of decoding is to model the conditional probability of $P(Y_1, \dots, Y_{n+m} | X_1, \dots, X_n)$. The decoder uses h_n, c_n from encoder as its initial state to compute $P(Y_{n+m})$. However the decoder does not directly model $P(Y | X)$, its power comes from modeling probability of current output with respect to all previous timesteps $P(Y_{n+m} | Y_{<n+m}, X_n)$ where $Y_{<n+m}$ represents output from 1 to $n + m - 1$. The posterior probability of the output state given inputs, with model parameters γ , is as follows

$$P_\gamma(Y | X) = \prod_{n=1}^{n+m} P_\gamma(Y_n | Y_{<n}, X) \quad (11)$$

F. Loss functions

The coupling between in-plane and out-of-plane rotation with the slice select direction and slice location z hinders optimization and learning [35]. To alleviate this issue, we divided the rotation θ regression heads from Equation (8) and added a hidden layer one each for θ_{xy} and θ_z as follows:

$$\theta_n^{xy} = \pi \tanh(W_{\theta_{xy}} \tanh(W_{\theta_{xyz}} o_n + b_{\theta_{xyz}}) + b_{\theta_{xy}}) \quad (12)$$

$$\theta_n^z = \pi \tanh(W_{\theta_z} \tanh(W_{\theta_{xyz}} o_n + b_{\theta_{xyz}}) + b_{\theta_z}) \quad (13)$$

which changes our loss calculation from

$$L_{Total} = L_{\theta_{xyz}} + L_z \quad (14)$$

$$L_{Total} = L_{\theta_{xy}} + L_{\theta_z} + L_z \quad (15)$$

For training, we minimized mean squared error (MSE) for both estimation and prediction $L_{Total} = L_{estimation} + L_{prediction}$ where $L = \|Y - \hat{Y}\|_2$. We used *tanh* as activation of this hidden layer as its derivative provides a stronger gradient for regression tasks compared to *ReLU* or *sigmoid* functions. In summary, we split our rotation loss in two separate layers; and regressed our rotation and slice location parameters using the backpropagation algorithm.

III. Experiments

To train, test, and evaluate our method we conducted experiments with real fetal MRI data with simulated motion and motion tracking data of volunteers who moved inside the scanner while motion parameters were recorded using an external motion tracking sensor. All fetal MRI and volunteer experiments were performed under protocols approved by the institutional review board committee, and written informed consent was obtained from all pregnant women volunteers and other volunteers. We divided our main experiments into estimation for 10 timesteps and prediction for 10 timesteps; and evaluated our trained model for generalization, robustness, and latency; and compared our results against pose estimation networks in particular those based on SVRNet [35], PoseNet [25], and our baseline models for estimation and prediction. Further, we tested our estimated motion parameters with a retrospective slice-to-volume reconstruction method [53]. In this section, we describe the fetal MRI data and its pre-processing first; and then the details of our experiments that involved generating the training data and the results of estimation and prediction for both simulated and real motion trajectories.

A. Fetal MRI dataset

The fetal MRI dataset consisted of repeated multi-planar T2-weighted single shot fast spin echo scans as well as reconstructed T2-weighted fetal brain MRI scans of 82 fetuses scanned at a gestational age (GA) between 21 and 37 weeks (mean=30.1, stdev=4.6) on 3-Tesla Siemens Skyra scanners with 18-channel body matrix and spine coils. The in-plane spatial resolution of the original scans was 1 mm, the slice thickness was 2–3 mm, and the temporal resolution for slice acquisition was equal to the repetition time (TR), which was 1.5s. This defined the time unit for slice-level motion tracking, so the timestep in motion tracking was

1.5s. Brain masks were automatically extracted on slices of the original scans using the real-time algorithm in [48]. The scans were automatically cropped around the fetal head ROI (based on the masks) and were then processed using slice-by-slice motion correction to reconstruct a super-resolved 3D volume [15], [17] at an isotropic resolution of 0.8 mm. Final 3D brain masks were then generated on the reconstructed images using Auto-Net [49] and manually corrected in ITK-SNAP [54] as needed.

Brain-extracted reconstructed volumes were then registered to a spatiotemporal fetal brain MRI atlas described in [23]. We normalized the intensity of the reconstructed images to zero-mean, unit-variance. The set of 82 scans were split into 30, 6, 40 and 6 for training, validation, test, and reconstruction, respectively; where the GA range spanned over 29 to 35 weeks for the training set, and from 26 to 37 weeks for the test set. We intentionally chose a narrower age range for the training set than the test set to examine the generalization capacity of the trained models on extended age ranges. To generalize well, the trained models had to account for both intrinsic inter-subject anatomical variations (due to different fetuses in the training and test sets) and anatomical variations due to different maturation levels of fetuses scanned at different GA ranges. The training, validation, test, and reconstruction set splits never had scans of the same subject. The GA of the reconstruction set subjects were 28, 30, 32, 32, 35, and 37; and between 6–10 (mean=7) multi-plane stacks of slices were used to reconstruct a volume for each of those subjects.

B. Generating the Training Data

To achieve our goal of predicting motion and slice position from sequences of slices, we aimed to train networks to learn the patterns of slice sampling and fetal head motion in reference to the fetal brain anatomy while it develops during gestation. To generate the training, validation, and test data for this purpose, from the pre-processed fetal MRI data, we generated sequences of fetal MRI slices with motion. This involved two sampling components: spatial sampling of slices and temporal sampling of spatial slices to model fetal motion. For slice excitation and spatial sampling, we sampled sequentially along permuted Z axes with 5 mm slice gap to account for fetal MRI acquisitions that are interleaved.

For temporal sampling to generate dynamic transformations corresponding to fetal motion, we exploited curve fitting with smoothing cubic Splines for each of the rotation angles. In this scheme, smoothing cubic splines generated different motion trajectories by interpolating curves between randomly-generated control points. The number of control points varied to control speed of motion. This was analogous to how fast or slow the fetus moved between scans. Further, to account for fast motion that disrupts slice encoding, we randomly masked a timestep in all slices. This resembled intra-slice motion as the brain masking technique in [48] generated all-zero masks for motion-corrupted slices. Figure 4 shows five 10-timestep sequences generated from the reference (GT) image sequence with random patterns and different speeds of motion.

We sampled 32 sequences for each subject in the training set 300 times (epochs). This led to $30 \text{ subjects} \times 32 \text{ sequences (1 batch size of 5 speed categories)} \times 300 \text{ times} = 288,000$ sequences for training, where speed of motion was controlled by the number of smoothing spline control points sampled from a normal distribution ($\mu = 6.4$, $\sigma = 1.36$, bounds=[4, 8]).

The initial rotation matrices were bound to $[-\pi/3, \pi/3]$ range; and the rotation parameters $\theta_x, \theta_y, \theta_z$ were sampled from a zero-mean normal distribution in the $[-\pi/6, \pi/6]$ range. This led to maximum rotation bounds of $[-\pi/2, \pi/2]$. For the validation set we followed the same sampling strategy, which led to $6 \times 32 \times 300 = 57,600$ sequences for validation.

C. Test Datasets

To test and compare algorithms, we sampled 32 sequences per speed of [4, 8] where we followed the spatial and temporal sampling strategies described in the previous section. This resulted in a total of 40 test subjects \times 32 samples = 1280 sequences of 20 timesteps (10 estimation + 10 prediction) each for test. Even though our main goal was to evaluate one-step ahead prediction, having 10 prediction timesteps allowed us to test efficacy of the model on long-term predictions. While our training data was limited to sequences generated from fetal MRI scans using the described procedure, to evaluate the generalization capacity of the trained models for new (unseen) patterns of motion, in addition to the test set described above, we used motion data recorded using head motion tracking sensors [55] from 10 volunteers. Rigid 3D transformation parameters were recorded in the scanner as volunteers moved their head with different patterns and speeds during scans. We applied these motion trajectory parameters to each of the 40 fetal test subjects, which led to a total of 400 new sequences with realistic motion patterns that differed from the motion patterns of the training data. The scans of the 6 test subjects in the reconstruction set were directly used in the reconstruction experiments. The details of the implementation and experiments are discussed next.

D. Implementation and Experimental Details

We used the mean square error (MSE) loss and the RMS-prop optimizer with initial learning rate of 0.001 ending in 0.00001 over the course of 300 epochs, decreasing the learning rate when the loss plateaued for 50 consecutive epochs. Due to the temporal nature of MRI slice acquisitions and the fact that the boundary slices did not include sufficient anatomical features, we limited estimation and prediction of motion trajectories to slices $s_i; i \in [0.4S, 0.9S]$, where S was the total number of slices in each reconstructed brain volume.

We conducted experiments and evaluated our model in both estimation and prediction tasks. For estimation, we compared our model (with 4.7M parameters) with two state-of-the-art fetal MRI pose estimation methods, *i.e.* an 18-layer residual network (ResNet) with two regression heads, one for angles θ and the other for slice location z , based on PoseNet [25] (with 11M parameters), and a VGG16-style network based on SVRNet [36] (with 14.7M parameters). Since SVRNet chose VGG16 among several other models, namely GoogLeNet, CaffeNet, Inception v4, and ResNet, we only compared against VGG16, as according to [36] it generated the lowest MSE.

For prediction, we conducted experiments for one-step and multi-step ahead predictions. To implicitly model motion states (*i.e.* to estimate motion velocity and acceleration) we needed a window size of at least three timesteps. In our experiments we used a window size of 10 for estimation and prediction each. For multi-step prediction, we limited our evaluation to 10 timesteps in the future although this was a choice and not a theoretical limit on the

prediction horizon. We compared our predictor against three baselines: 1) a naive predictor that used estimation at current time as one-step ahead prediction (referred to as zero velocity predictor); 2) an auto-regressive model that recursively used its own predictions in a sliding window of size 10 to predict multi-step motion trajectories; and 3) a predictive model that we adopted based on the network proposed in [56]. In this model (with 44M parameters), the data was passed directly into an LSTM without spatial feature encoding, thus we refer to it as directLSTM.

For the volume reconstruction experiments from multiple scans, we rearranged slices of the original fetal MRI scans (with inter-slice motion) based on slice timing, estimated 3D pose, and passed estimated parameters from our motion tracking algorithm along with volume-to-volume transformation to the canonical atlas space [25] to NiftyMIC [20], [53]. We then compared the reconstructions to reconstructions directly performed by NiftyMIC in the atlas space. We compared reconstructed images using Structural Similarity Index (SSIM) which has a range of -1 to $+1$ where 1 means a perfect match, and Normalized Root Mean Square Error (NRMSE), which ranges between 0 and 1 where 0 means perfect match (0 error).

E. Results

Figure 5 shows 10 estimated and 10 predicted timesteps for a train case and a test case compared to the ground truth slices in the top rows. The predicted rotation was accurate after multiple timesteps. Table I shows average loss of estimation and prediction tasks (defined in Section II-F) on the test data with synthetic motion, along with the standard errors computed between groups of fetuses in the test set based on the prediction timestep (time), age at scan, and speed of motion, for the ablation studies as well as the comparisons to baseline and alternative methods. We compared our "full model" trained with sequences with masked slices (resembling slices corrupted by intra-slice motion) and split loss explained in Section II-F against our "baseline" which was trained without masked slices in the training set sequences and without split heads, and "masked bl" which was trained with masked slices but without the split loss functions. The best results in each comparison, shown in bold, show that our full model outperformed the baselines and all other models in both estimation and prediction tasks. The low standard errors of our model show its consistent and robust performance with respect to the different characteristics of the test data. Figure 6 shows the rotational MSE of multi-step prediction per timestep (estimation for time 10 and predictions for times 11 to 19) on test data, where the images corresponding to time points 1 to 10 were the inputs to the model.

In the next sets of experiments, we evaluated our model for 1) its generalization performance for our test data that included subjects scanned at gestational ages not included in the training set; 2) its performance for different speeds of motion; 3) its robustness in the presence of intra-slice motion (i.e. lost slices in the input sequence due to fast motion that disrupted signal during slice encoding); and 4) its generalization and robustness to motion patterns that were different from the motion patterns in the training data (i.e. motion patterns recorded from volunteer subject experiments). Figure 7 shows boxplots of the MSE of the estimation and prediction tasks for 10 timesteps grouped by gestational age and datasets.

The consistency in error statistics across test and train datasets and GA, indicate that the trained model was robust and generalized well to the test data.

Table II shows the MSE of pose estimation, one-step and multi-step prediction for test data grouped by the location of a lost slice (due to intra-slice motion) in the input slice sequence. This table compares the performance of two models: our model trained without any missed (masked) slices in the training sequences (referred here as the baseline); and our full model trained with randomly missed (masked) slices in the training set. These results show that 1) in the baseline model, both estimation and prediction errors were higher when the lost slice was closer to the end of the input sequence; *i.e.* missing slice 10 in the sequence led to much higher errors (shown in red) compared to missing slices in earlier locations; 2) Our full model performed better than the baseline with much more consistent and robust performance; and 3) Our full model's performance degrades if the first timestep is masked because being first timestep it does not have information from past and by masking it. These show that when our model was trained with randomly masked slices in the training sequences, it learned to rely less on the last slices in the sequence to gain robustness to intra-slice motion.

We evaluated the generalization capacity of our model trained on data with synthetic motion, on motion trajectories recorded from volunteer subjects (that were never used in training). Figure 8 shows the mean squared pose prediction error for different timesteps for the test data with the recorded motion trajectories, obtained from our full model and other predictors. The results show that our model generated very low multi-step prediction errors, whereas all other methods showed high errors that increased with prediction horizon.

Our final experiment focused on end-to-end volume reconstruction from multiple stack-of-slices with motion parameters estimated by our model and reconstructed with NiftyMIC [53]. The results of the reconstructions with our estimated motion parameters for 6 subjects in the test set have been shown in Figure S1, and compared favorably with reconstructions using conventional slice-to-volume registration in terms of NRMSE and SSIM. In particular, we achieved average NRMSE of 0.151 with standard deviation of 0.023; and SSIM of 0.912 with standard deviation of 0.031 for our reconstructions. Supplementary Figure S2 shows multi-plane views of a sample case, and Figure S3 shows a case where reconstruction with conventional slice-to-volume registration failed, whereas the reconstruction was improved when we plugged in our estimated motion parameters in the first iteration of reconstruction.

The latency for prediction on our hardware (an NVIDIA GeForce 1080 Ti) was ~ 1.42 ms per data point where each sequence comprised of 10 slices and outputs were 10 estimations and 10 predictions. Considering the slice acquisition time of ~ 1.5 seconds for T2-weighted MRI and ~ 80 ms for echo-planar imaging, this is real-time.

IV. Discussion

To the best of our knowledge, this is the first report of the development of a real-time predictive motion tracking technique for fetal MRI. Up until now motion correction in fetal MRI has been done retrospectively through non-causal iterations of slice-to-volume

registration and reference volume reconstruction. Slice-to-volume registration is intrinsically an ill-posed problem [21]. To overcome this issue, retrospective fetal MRI motion estimation methods that relied on slice-to-volume registration, evolved from hierarchical [12], [13] and slice intersection-based [14] methods to progressive [15]–[17], [20], patch-based [18], and more recently, dynamic motion estimation techniques [19]. There have also been nonrigid and deformable extensions of slice-to-volume registration [57].

General-purpose, image-based, MRI motion tracking techniques sought regularization through modeling motion dynamics [7], or used robust state space models to estimate relative position of sequentially-acquired slices [37], [58]. While the underlying phenomena are nonlinear, these techniques made simplifying assumptions to linearize the problem and used image registration along with state space estimation by Kalman filtering (or its robust extensions) for motion tracking. Bayesian filtering based Kalman filters fail to model nonlinear relationships as well as non-Gaussian noise, and their extended versions also fail when dynamics are highly nonlinear. These techniques are thus difficult to scale up to real life scenarios, in particular in challenging applications such as fetal MRI.

More capable Gaussian mixture models [59], process models [60], or dynamic Bayesian networks (DBN) [39] can accommodate complex dynamics but need strong priors by experts which makes them prone to the same practical issues that exist in conventional methods especially when long term prediction is desired. As a result of using image registration, these techniques are computationally intensive and cannot be easily applied in real-time. More importantly, none of the current techniques explicitly uses image information and image recognition to model motion dynamics for 3D pose estimation. Registration-based methods are slow and offer a limited capture range, which makes them prone to failure when motion is continuous and large. In other words, even when integrated with state space estimation methods for dynamic motion tracking, registration-based techniques may not easily recover if they lose subject's position. This is especially problematic in motion estimation in fetal MRI as fetuses in the second and early third trimesters move frequently and rotate in large angles. Finally, almost all of the current methods rely on certain initialization assumptions such as the existence of a motion-free reference scan for registration, which is restrictive and unrealistic when considered for use in real-time applications, such as motion tracking for real-time navigation.

In this paper we showed predictive potential of recurrent neural networks for modeling end-to-end motion in MRI. To this end, we developed a combination of spatial encoders based on convolutional neural networks and temporal encoder and decoder networks based on CNN-LSTM to learn the spatiotemporal features of anatomy and slice sampling from imaging data to predict motion trajectories. Loss functions on multiple regression heads led to a robust model that generalized well beyond the training set to fetuses scanned at different ages and with motion patterns that were recorded from volunteers, which were characteristically different from the synthetic motion patterns that were used in training.

To resemble fetal head motion, our volunteer subjects moved their head at different speeds and in different directions to the largest possible extents while we recorded their motion. Comparing the results shown in Figure 8 (for the fetal test data with recorded motion) with

the results in Figure 7 (for the fetal test data with synthetic motion) indicates that the average prediction error on recorded data was lower than the average prediction error on synthetic data, despite the fact that the synthetic motion was generated by the same procedure that generated motion patterns in the training data. We attribute this to the fact that the recorded motion had constraints imposed by the mechanical linkage between head and neck that made it easier to predict compared to the synthetic motion.

Our approach is a learning-based technique, so its performance depends on what it learns from the training data. Our training data involved large rotations in the -90° to 90° range over 15 seconds (10 timesteps). Our training data generation methodology differed from those in earlier 3D pose estimation works, e.g., [25] [35], which randomly rotated individual slices without taking surrounding slices into account. We generated sequences of interleaved slices covering the 3D anatomy while the anatomy moved on a motion trajectory synthesized by B-Spline curve fitting. This is more realistic than moving slices independently. Yet our model may benefit from training with more realistic simulations of motion, for example using bio-mechanical models of fetal motion [61] or from ground truth motion recorded from adult volunteers. In this study we used recorded motion only for testing. Dynamic predictive motion tracking, as we proposed here, may also be useful to assess normal versus abnormal patterns of fetal movements [62] from cinematographic MRI (or 4D ultrasound), which, in-turn, may be used to assess fetal motor behavior [63], [64].

Obtaining ground truth fetal motion is difficult, especially for large ranges of motion. Motion estimates obtained from successful slice-to-volume reconstructions are typically only available (and reliable) for small ranges of motion. Slice to volume reconstruction techniques rely on 1) redundant slice acquisitions, 2) outlier detection and rejection, and 3) robust reconstructions. Therefore, they effectively filter or remove the effect of mis-registered and motion-corrupted slices [15], [16]. The transformations obtained for the remainder of the slices that are effectively used in reconstruction are typically small; and yet may not be sufficiently reliable to be used as ground truth. Therefore, to test our approach on original fetal MRI scans with motion, we used our estimated motion parameters along with a powerful slice-to-volume reconstruction method [53] to reconstruct volumes from multiple stack-of-slices. Reconstruction with our estimated motion parameters compared favorably against reconstruction with retrospective slice-to-volume registration (Supplementary Figures S1–S3).

Our model generalized well to data from subjects at ages outside of the age range of the training data and with realistic motion patterns that were never used in training. Initially we found that the model had difficulty estimating motion for large and fast movements. To resolve this we used curriculum training which trained the network on difficult samples more often than easier ones that alleviated the issue. Our initially trained models also had difficulty generalizing to unseen gestational ages with large and fast movement in the validation set. To resolve that, we added batch normalization and regularized by reducing the number of parameters in the model which resulted in better performance. Yet, since our method is a learning-based approach, its performance is expected to degrade if there is significant domain shift between the training and test data. For example, the performance of our model may significantly drop if a different modality or sequence is used as test, or if a

significantly different set of parameters are used in fetal MRI scans. To adapt the model to new domains, domain adaptation techniques or pre-processing may be used, e.g. [25]. Also, our trained model may not generalize well for tracking motion of severely abnormal anatomies. Possible remedies for this problem are to include abnormal cases in training and to use curriculum learning with appropriate data augmentation. These are excellent directions for future work.

Our model architecture is small compared to most state-of-the-art RNNs. This helped us achieve real-time performance. Curriculum training helped the network focus on more difficult samples, i.e. sequences with large and fast motion. We kept our model a causal predictive model for its intended application which is real-time motion tracking and navigation. For other applications, such as retrospective processing of image time series, using signal from the future, e.g. by bidirectional LSTM, is expected to increase performance but would break the causal nature of the model. To train our model we used the MSE loss due to its well-posed convex nature for optimization in our high-dimensional search space. For static pose estimation [25], a second stage optimization (refinement) with the geodesic loss, which is a natural Riemannian metric on the compact Lie group $SO(3)$ of orientations, improved the results. We observed a similar trend here but at a relatively lower degree. By fine tuning our model (trained with MSE) for 10 additional epochs at a learning rate of 0.0001 with geodesic loss we observed average error reduction of 0.4° in estimation, which was statistically significant; but did not see a statistically significant reduction in prediction error.

By observing a sequence of slices, our trained model predicts the relative 3D pose (motion) of the anatomy with respect to an initial pose. For real-time slice navigation, therefore, we require an estimate of the initial pose; which can be obtained by the pose estimation techniques proposed in [36] and [25]. Although those techniques can accurately estimate the 3D pose of the fetal brain in a canonical (atlas) space based on a volume or a slice (or stack of slices) close to the center of the anatomy, their estimation error is relatively high in the border slices where image features are sparse, and their predictive performance is limited for fast and large motion. Experimental results in motion tracking showed that our technique outperformed time series prediction models built upon those static pose estimation methods. Therefore, to build an effective and efficient real-time fetal MRI navigation system, a combination of initial pose estimation by techniques such as those proposed in [36] and [25] and our predictive motion tracking technique is needed. Echo-planar imaging [65] may be an appropriate choice to acquire fast volumes (as 3D localizer or navigator) to estimate initial pose at the beginning or in intervals between sequences.

V. Conclusion

We developed and presented a technique that is capable of estimating and predicting the 3D pose trajectories of the fetal brain in real-time despite large fetal movements. This technique, when augmented with other real-time components and implemented on MRI scanner platforms, may be used to track fetal head motion as slices are acquired, make recommendations for scan orientations as a decision support system or a human-in-the-loop navigation system, or to build real-time automatic fetal MRI systems, which, in-turn, can

lead to much more efficient, effective, and tolerable fetal MRI scan sessions. Real-time predictive motion tracking can also play a critical role in real-time assessment of the quality of highly motion-sensitive scans such as fetal functional MRI that are very difficult to perform, and to automatically adapt the duration of such scans to ensure data of sufficient quality is acquired for post-acquisition processing. Finally, image-based dynamic motion tracking can also be used to assess fetal movements and motor behavior *in-utero* from cine MRI and 4D ultrasound.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This study was supported in part by the Department of Radiology at Boston Children's Hospital, by the National Institutes of Health (NIH) grants R01 EB018988 and R01 NS106030, and by a Technological Innovations in Neuroscience Award from the McKnight Foundation. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH or the McKnight Foundation.

References

- [1]. Malamateniou C, Malik S, Counsell S, Allsop J, McGuinness A, and Hayat T, "Motion-compensation techniques in neonatal and fetal MR imaging," *American Journal of Neuroradiology*, vol. 34, no. 6, pp. 1124–1136, 2013. [PubMed: 22576885]
- [2]. Afacan O, Erem B, Roby DP, Roth N, Roth A, Prabhu SP, and Warfield SK, "Evaluation of motion and its effect on brain magnetic resonance image quality in children," *Pediatric radiology*, vol. 46, no. 12, pp. 1728–1735, 2016. [PubMed: 27488508]
- [3]. Boston Children's Hospital, "Partnering with families to minimize exposure to anesthesia," 2019, <https://thriving.childrenshospital.org/minimize-anesthesia-exposure/>, Last accessed on 2019-06-10.
- [4]. Pipe JG, "Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 963–969, 1999.
- [5]. Thesen S, Heid O, Mueller E, and Schad LR, "Prospective acquisition correction for head motion with image-based tracking for real-time fMRI," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 44, no. 3, pp. 457–465, 2000.
- [6]. Maclaren J, Herbst M, Speck O, and Zaitsev M, "Prospective motion correction in brain imaging: a review," *Magnetic resonance in medicine*, vol. 69, no. 3, pp. 621–636, 2013. [PubMed: 22570274]
- [7]. White N, Roddey C, Shankaranarayanan A, Han E, Rettmann D, Santos J et al., "PROMO: real-time prospective motion correction in MRI using image-based tracking," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 63, no. 1, pp. 91–105, 2010.
- [8]. Afacan O, Wallace TE, and Warfield SK, "Retrospective correction of head motion using measurements from an electromagnetic tracker," *Magnetic resonance in medicine*, 2019.
- [9]. Frost R, Wighton P, Karahano lu FI, Robertson RL, Grant PE, Fischl B et al., "Markerless high-frequency prospective motion correction for neuroanatomical MRI," *Magnetic resonance in medicine*, vol. 82, no. 1, pp. 126–144, 2019. [PubMed: 30821010]
- [10]. Wallace TE, Afacan O, Waszak M, Kober T, and Warfield SK, "Head motion measurement and correction using FID navigators," *Magnetic resonance in medicine*, vol. 81, no. 1, pp. 258–274, 2019. [PubMed: 30058216]

- [11]. Gholipour A, Estroff JA, Barnewolt CE, Robertson RL, Grant PE, Gagoski B et al., “Fetal MRI: a technical update with educational aspirations,” *Concepts in Magnetic Resonance Part A*, vol. 43, no. 6, pp. 237–266, 2014.
- [12]. Rousseau F, Glenn OA, Jordanova B, Rodriguez-Carranza C, Vigneron DB, Barkovich JA, and Studholme C, “Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images,” *Academic radiology*, vol. 13, no. 9, pp. 1072–1081, 2006. [PubMed: 16935719]
- [13]. Jiang S, Xue H, Glover A, Rutherford M, Rueckert D, and Hajnal JV, “MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): application to fetal, neonatal, and adult brain studies,” *IEEE transactions on medical imaging*, vol. 26, no. 7, pp. 967–980, 2007. [PubMed: 17649910]
- [14]. Kim K, Habas PA, Rousseau F, Glenn OA, Barkovich AJ, and Studholme C, “Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation,” *IEEE transactions on medical imaging*, vol. 29, no. 1, pp. 146–158, 2010. [PubMed: 19744911]
- [15]. Gholipour A, Estroff JA, and Warfield SK, “Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI,” *IEEE transactions on medical imaging*, vol. 29, no. 10, pp. 1739–1758, 2010. [PubMed: 20529730]
- [16]. Kuklisova-Murgasova M, Quaghebeur G, Rutherford MA, Hajnal JV, and Schnabel JA, “Reconstruction of fetal brain MRI with intensity matching and complete outlier removal,” *Medical image analysis*, vol. 16, no. 8, pp. 1550–1564, 2012. [PubMed: 22939612]
- [17]. Kainz B, Steinberger M, Wein W, Kuklisova-Murgasova M, Malamateniou C, Keradren K et al., “Fast volume reconstruction from motion corrupted stacks of 2D slices,” *IEEE transactions on medical imaging*, vol. 34, no. 9, pp. 1901–1913, 2015. [PubMed: 25807565]
- [18]. Alansary A, Rajchl M, McDonagh SG, Murgasova M, Damodaram M, Lloyd DF et al., “PVR: patch-to-volume reconstruction for large area motion correction of fetal MRI,” *IEEE transactions on medical imaging*, vol. 36, no. 10, pp. 2031–2044, 2017. [PubMed: 28880160]
- [19]. Marami B, Salehi SSM, Afacan O, Scherrer B, Rollins CK, Yang E et al., “Temporal slice registration and robust diffusion-tensor reconstruction for improved fetal brain structural connectivity analysis,” *NeuroImage*, vol. 156, pp. 475–488, 2017. [PubMed: 28433624]
- [20]. Ebner M, Wang G, Li W, Aertsen M, Patel PA, Aughwane R et al., “An automated localization, segmentation and reconstruction framework for fetal brain MRI,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 313–320.
- [21]. Ferrante E and Paragios N, “Slice-to-volume medical image registration: A survey,” *Medical image analysis*, vol. 39, pp. 101–123, 2017. [PubMed: 28482198]
- [22]. Taimouri V, Gholipour A, Velasco-Annis C, Estroff JA, and Warfield SK, “A template-to-slice block matching approach for automatic localization of brain in fetal MRI,” in *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2015, pp. 144–147.
- [23]. Gholipour A, Rollins CK, Velasco-Annis C, Ouaalam A, Akhondi-Asl A, Afacan O et al., “A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth,” *Scientific reports*, vol. 7, no. 1, p. 476, 2017. [PubMed: 28352082]
- [24]. Tourbier S, Velasco-Annis C, Taimouri V, Haggmann P, Meuli R, Warfield SK et al., “Automated template-based brain localization and extraction for fetal brain MRI reconstruction,” *NeuroImage*, vol. 155, pp. 460–472, 2017. [PubMed: 28408290]
- [25]. Salehi SSM, Khan S, Erdogmus D, and Gholipour A, “Real-time deep pose estimation with geodesic loss for image-to-template rigid registration,” *IEEE transactions on medical imaging*, vol. 38, no. 2, pp. 470–481, 2018. [PubMed: 30138909]
- [26]. Hochreiter S and Schmidhuber J, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997. [PubMed: 9377276]
- [27]. Wu J, Xue T, Lim JJ, Tian Y, Tenenbaum JB, Torralba A, and Freeman WT, “Single image 3D interpreter network,” in *European Conference on Computer Vision*. Springer, 2016, pp. 365–382.
- [28]. Pavlakos G, Zhou X, Chan A, Derpanis KG, and Daniilidis K, “6-dof object pose from semantic keypoints,” in *Robotics and Automation (ICRA), IEEE International Conference*. IEEE, 2017, pp. 2011–2018.

- [29]. Tulsiani S and Malik J, "Viewpoints and keypoints," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 1510–1519.
- [30]. Su H, Qi CR, Li Y, and Guibas LJ, "Render for CNN: Viewpoint estimation in images using cnns trained with rendered 3d model views," in Proceedings of the IEEE International Conference on Computer Vision, 2015, pp. 2686–2694.
- [31]. Mahendran S, Ali H, and Vidal R, "3D pose regression using convolutional neural networks," in IEEE International Conference on Computer Vision, vol. 1, no. 2, 2017, p. 4.
- [32]. Newell A, Yang K, and Deng J, "Stacked hourglass networks for human pose estimation," in European Conference on Computer Vision. Springer, 2016, pp. 483–499.
- [33]. Alp Guler R, Neverova N, and Kokkinos I, "Densepose: Dense human pose estimation in the wild," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 7297–7306.
- [34]. Andriluka M, Iqbal U, Insafutdinov E, Pishchulin L, Milan A, Gall J, and Schiele B, "PoseTrack: A benchmark for human pose estimation and tracking," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 5167–5176.
- [35]. Hou B, Alansary A, McDonagh S, Davidson A, Rutherford M, Hajnal JV et al., "Predicting slice-to-volume transformation in presence of arbitrary subject motion," in International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, 2017, pp. 296–304.
- [36]. Hou B, Khanal B, Alansary A, McDonagh S, Davidson A, Rutherford M et al., "3-D reconstruction in canonical co-ordinate space from arbitrarily oriented 2-D images," IEEE transactions on medical imaging, vol. 37, no. 8, pp. 1737–1750, 2018. [PubMed: 29994453]
- [37]. Marami B, Scherrer B, Afacan O, Erem B, Warfield SK, and Gholipour A, "Motion-robust diffusion-weighted brain MRI reconstruction through slice-level registration-based motion tracking," IEEE Transactions on Medical Imaging, vol. 35, no. 10, pp. 2258–2269, 2016. [PubMed: 27834639]
- [38]. Rumelhart DE, Hinton GE, and Williams RJ, "Learning representations by back-propagating errors," Cognitive modeling, vol. 5, no. 3, 1988.
- [39]. Gindele T, Brechtel S, and Dillmann R, "Learning driver behavior models from traffic observations for decision making and planning," IEEE Intelligent Transportation Systems Magazine, vol. 7, pp. 69–79, 2015.
- [40]. Sutskever I, Vinyals O, and Le QV, "Sequence to sequence learning with neural networks," in Advances in neural information processing systems, 2014, pp. 3104–3112.
- [41]. Ondruska P and Posner I, "Deep tracking: Seeing beyond seeing using recurrent neural networks," in AAAI Conference on Artificial Intelligence, 2016.
- [42]. Krebs S, Duraisamy B, and Flohr F, "A survey on leveraging deep neural networks for object tracking," 2017 IEEE 20th International Conference on Intelligent Transportation Systems, pp. 411–418, 2017.
- [43]. Held D, Thrun S, and Savarese S, "Learning to track at 100 fps with deep regression networks," in European Conference on Computer Vision. Springer, 2016, pp. 749–765.
- [44]. Valmadre J, Bertinetto L, Henriques J, Vedaldi A, and Torr PH, "End-to-end representation learning for correlation filter based tracking," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 2805–2813.
- [45]. Li B, Yan J, Wu W, Zhu Z, and Hu X, "High performance visual tracking with siamese region proposal network," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 8971–8980.
- [46]. Ning G, Zhang Z, Huang C, Ren X, Wang H, Cai C, and He Z, "Spatially supervised recurrent convolutional neural networks for visual object tracking," in 2017 IEEE International Symposium on Circuits and Systems (ISCAS). IEEE, 2017, pp. 1–4.
- [47]. Gordon D, Farhadi A, and Fox D, "Re3: Real-time recurrent regression networks for visual tracking of generic objects," IEEE Robotics and Automation Letters, vol. 3, no. 2, pp. 788–795, 2018.
- [48]. Salehi SSM, Hashemi SR, Velasco-Annis C, Ouaalam A, Estroff JA, Erdogmus D et al., "Real-time automatic fetal brain extraction in fetal MRI by deep learning," in 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). IEEE, 2018, pp. 720–724.

- [49]. Salehi SSM, Erdogmus D, and Gholipour A, "Auto-context convolutional neural network (auto-net) for brain extraction in magnetic resonance imaging," *IEEE transactions on medical imaging*, vol. 36, no. 11, pp. 2319–2330, 2017. [PubMed: 28678704]
- [50]. LeCun Y, Haffner P, Bottou L, and Bengio Y, "Object recognition with gradient-based learning," in *Shape, contour and grouping in computer vision*. Springer, 1999, pp. 319–345.
- [51]. He K, Zhang X, Ren S, and Sun J, "Identity mappings in deep residual networks," in *European Conference on Computer Vision*. Springer, 2016, pp. 630–645.
- [52]. Hochreiter S, Bengio Y, Frasconi P, Schmidhuber J et al., "Gradient flow in recurrent nets: the difficulty of learning long-term dependencies," 2001.
- [53]. Ebner M, Wang G, Li W, Aertsen M, Patel PA, Aughwane R et al., "An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain mri," *NeuroImage*, vol. 206, p. 116324, 2020.
- [54]. Yushkevich PA, Piven J, Hazlett HC, Smith RG, Ho S, Gee JC, and Gerig G, "User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability," *Neuroimage*, vol. 31, no. 3, pp. 1116–1128, 2006. [PubMed: 16545965]
- [55]. Gholipour A, Polak M, Van Der Kouwe A, Nevo E, and Warfield SK, "Motion-robust MRI through real-time motion tracking and retrospective super-resolution volume reconstruction," in *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2011, pp. 5722–5725.
- [56]. Martinez J, Black MJ, and Romero J, "On human motion prediction using recurrent neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2891–2900.
- [57]. Uus A, Zhang T, Jackson LH, Roberts TA, Rutherford MA, Hajnal JV, and Deprez M, "Deformable slice-to-volume registration for motion correction of fetal body and placenta mri," *IEEE Transactions on Medical Imaging*, 2020.
- [58]. Marami B, Scherrer B, Khan S, Afacan O, Prabhu S, Sahin M et al., "Motion-robust diffusion compartment imaging using simultaneous multi-slice acquisition," *Magnetic resonance in medicine*, vol. 81, no. 5, p. 3314, 2019. [PubMed: 30443929]
- [59]. Ammoun S and Nashashibi F, "Real time trajectory prediction for collision risk estimation between vehicles," *2009 IEEE 5th International Conference on Intelligent Computer Communication and Processing*, pp. 417–422, 2009.
- [60]. Wiest J, Höffken M, Kressel U, and Dietmayer KCJ, "Probabilistic trajectory prediction with gaussian mixture models," *2012 IEEE Intelligent Vehicles Symposium*, pp. 141–146, 2012.
- [61]. Verbruggen SW, Loo JH, Hayat TT, Hajnal JV, Rutherford MA, Phillips AT, and Nowlan NC, "Modeling the biomechanics of fetal movements," *Biomechanics and modeling in mechanobiology*, vol. 15, no. 4, pp. 995–1004, 2016. [PubMed: 26534772]
- [62]. Piontelli A et al., *Development of normal fetal movements*. Springer, 2014.
- [63]. Hayat TT, Martinez-Biarge M, Kyriakopoulou V, Hajnal JV, and Rutherford MA, "Neurodevelopmental correlates of fetal motor behavior assessed using cine mr imaging," *American Journal of Neuroradiology*, vol. 39, no. 8, pp. 1519–1522, 2018. [PubMed: 29880478]
- [64]. Hayat TT and Rutherford MA, "Neuroimaging perspectives on fetal motor behavior," *Neuroscience & Biobehavioral Reviews*, vol. 92, pp. 390–401, 2018. [PubMed: 29886176]
- [65]. Afacan O, Estroff JA, Yang E, Barnewolt CE, Connolly SA, Parad RB et al., "Fetal echoplanar imaging: Promises and challenges," *Topics in Magnetic Resonance Imaging*, vol. 28, no. 5, pp. 245–254, 2019. [PubMed: 31592991]




Figure 1: The Region-of-Interest (RoI), here the fetal brain, is extracted using a real-time segmentation technique, e.g. [48], cropped, center aligned, and intensity normalized to form a volume of stacked slices for deep predictive fetal head motion tracking.

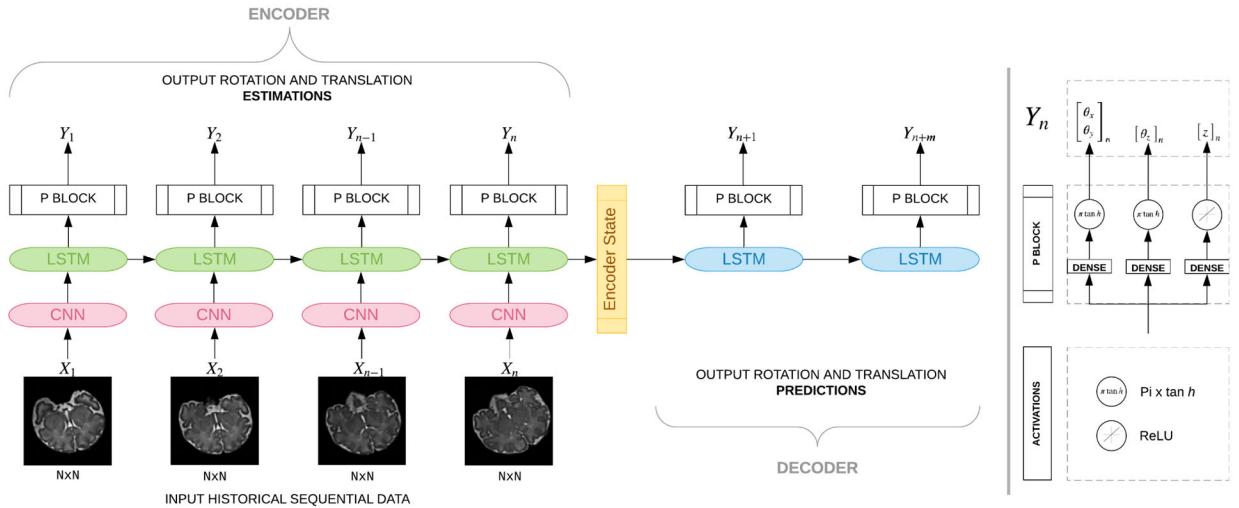


Figure 2: Our many-to-many Seq2Seq model that takes as input sequence of slices and estimates angles as well as predictions. Multiple LSTM units are shown since we unroll our network. All units of the same type and color share weights, hence they get the same gradient update during training. This model comprises of an encoder and a decoder component. The encoder, which contains spatial encoder (CNN) blocks followed by a temporal encoder that contains LSTM units and P blocks, encodes and learns sequence-of-image features to estimate position parameters. The encoder state is fed into the decoder network which comprises of LSTM units followed by P blocks. Each P block has three heads with $\pi \times \tan h$ activation for the rotation parameters and ReLU activation for the slice position.




Figure 3:

The architecture of the spatial encoder CNN blocks of our deep predictive motion tracking model shown in Figure 2. Each encoder performs 3×3 convolutions followed by batch normalization, PReLU [51] and *MaxPooling* that down-samples the image in half, extracts local dependencies and reduces computation in downstream layers. This enables fine-grained feature preservation. The number of filters are doubled in each layer until it reaches 512. Finally, features from the CNN are flattened and transferred as spatial encoding of time step n in the sequence to the LSTM layer of the encoder. Compared to the deep spatial encoder network used in [25] to infer 3D pose from a single slice, our CNN is lightweight, which boosts its real-time performance while it effectively encodes features of multiple sequentially-acquired slices and pass them to the LSTM modules to build an encoder state (Figure 2).




Figure 4:

A demo of five sequences of 10 timesteps each generated with different speeds of motion (corresponding to the number of spline control points from 4 to 8) from the 3D reconstructed fetal brain MRI scan of GA 35 weeks (shown at the top row). Randomly masked slices indicate slices corrupted by intra-slice motion.

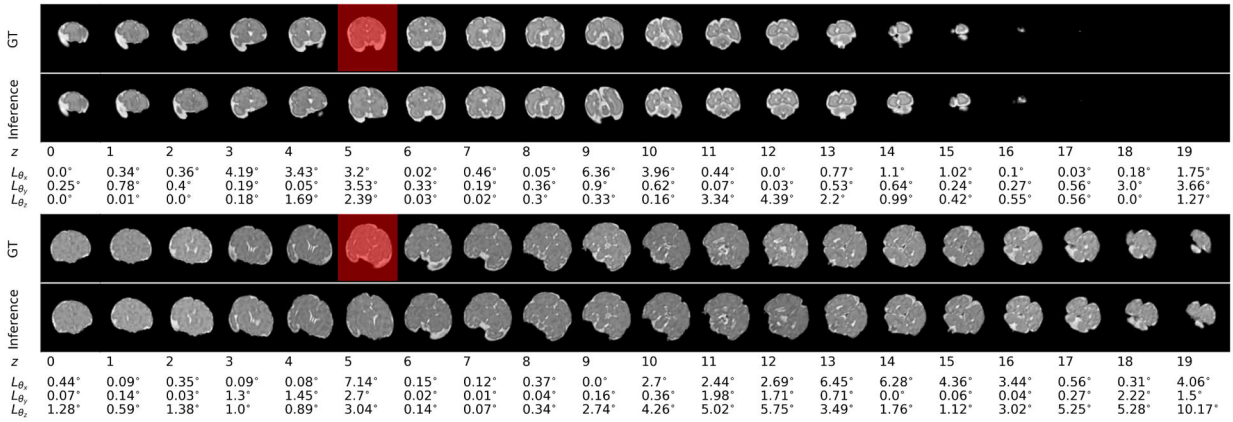


Figure 5: Inference (i.e. estimation for the first 10 timesteps and prediction for the rest of the 10 timesteps) in the bottom rows has been compared to the ground truth sequence in the top rows for scans of two fetuses: the first figure is a scan of a 28-week, and the second figure is a scan of a 36-week GA fetus from the test set. Errors based on the MSE loss (Section II-F) have been shown underneath each timestep. In these figures the slices shown with red masks were masked in the input sequence. It can be seen that the estimated slices (in the bottom rows) corresponding to the masked slices, showed relatively larger error, but the masked slices did not have a major effect on predictions. Slight increase in prediction error with prediction time horizon was seen in the test sequence, but the predictions were overall accurate.




Figure 6:

Boxplots showing the statistics of the average rotational MSE loss on test data computed for prediction per timestep. Our model outperformed all other prediction models implemented and tested here (i.e., zero velocity, auto-regressive, and directLSTM).

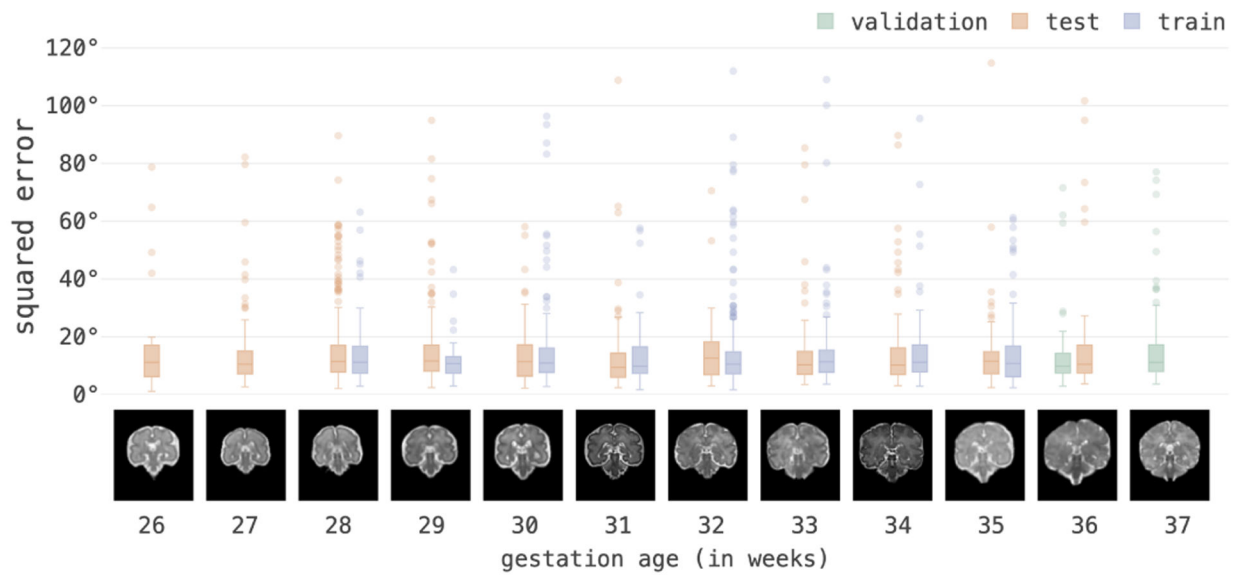


Figure 7: Average MSE of 3D pose in degrees of one-step ahead prediction tasks for 10 timesteps grouped by GA. Consistent errors show that our model generalized well to variations in anatomy and GA outside of the domain and range that it was trained on.




Figure 8:

3D pose MSE in degrees of multi-step prediction for the test data with motion trajectories recorded from volunteers, shows the generalization capacity of our model on real motion patterns. In all baseline models, the prediction error increased with the prediction steps due to compounding errors. In contrast, by passing context from encoder and prediction from previous decoding, our model maintained low error in multi-step prediction.

Table I:

Mean squared error (μ error) for estimation and prediction of 3D pose in degrees along with the overall standard error of mean (σ_μ) and the standard error of different timesteps, ages, and speed of motion for the test data. The top part of the table compares estimation models and the bottom part compares prediction models. In these comparisons we also tested our model trained without any masked slices in the sequences, referred to as the “baseline”, our second baseline trained with masked slice sequences but without the split heads and the loss function explained in Section II-F (referred to as “masked bl.”) and our “full model” trained with both masked slices and the split loss function. Significant reduction in both estimation and prediction errors was achieved by our full trained model compared to baselines and all other compared models. Low standard errors show that our model performed consistently, and was robust to variations in data, timesteps, GA, and the speed of motion.

Model	μ error	σ_μ	σ_μ time	σ_μ age	σ_μ speed
VGG16	129.33	11.74	3.72	3.48	9.51
Resnet18	82.60	5.76	3.55	1.31	3.34
Our baseline	20.19	2.57	1.21	2.23	2.06
Our masked bl.	9.10	2.31	1.11	1.92	2.45
Our full model	3.55	0.22	0.17	0.05	0.23
directLSTM	103.20	3.09	0.97	13.52	5.80
Zero velocity	74.14	1.09	0.86	1.77	1.32
Auto regressive	96.77	1.66	0.69	1.83	2.17
Our baseline	33.51	2.35	1.17	1.23	1.11
Our masked bl.	11.28	1.28	1.17	0.23	0.51
Our full model	8.07	0.72	0.42	0.39	0.59

Table II:

Results of a probing task on our full model trained with masked data against our model trained on unmasked data (baseline): 3D pose MSE in degrees of estimation (Est), one step prediction (OSP) and multi-step prediction (MSP) on test data, which have been shown based on the timestep in which a slice was masked in the test sequence (first column). Results of both models on unmasked test data (first row) were similar, however the prediction performance of the baseline model indicates that to make predictions this model put a heavy weight on slices that appeared towards the end of the sequence. On the other hand, our full model trained with randomly-masked sequences, performed more consistently and robustly with respect to the position of the masked slice in the input test sequence.

Timestep Masked	Baseline model error			Masked model error		
	Est	OSP	MSP	Est	OSP	MSP
No Mask	1.37	2.97	7.41	1.03	2.93	7.69
1	5.83	4.42	10.48	4.86	3.70	10.05
2	4.83	3.03	7.58	2.97	2.87	7.62
3	4.36	2.98	7.50	2.17	2.86	7.61
4	3.06	3.03	7.98	1.87	2.71	7.98
5	3.87	3.05	8.13	2.01	2.83	7.41
6	3.29	4.06	8.39	2.43	2.91	7.63
7	3.25	4.17	8.65	2.59	2.93	7.69
8	3.91	6.37	9.21	2.61	3.02	7.74
9	4.06	6.78	10.74	2.68	3.59	8.15
10	4.19	17.37	15.89	3.88	6.96	9.54