

3D Reconstruction in Canonical Co-ordinate Space from Arbitrarily Oriented 2D Images

Benjamin Hou, Bishesh Khanal, Amir Alansary, Steven McDonagh, Alice Davidson, Mary Rutherford, Jo V. Hajnal, Daniel Rueckert, Ben Glocker and Bernhard Kainz

Abstract—Limited capture range and the requirement to provide high quality initializations for optimization-based 2D/3D image registration methods can significantly degrade the performance of 3D image reconstruction and motion compensation pipelines. Challenging clinical imaging scenarios, that contain significant subject motion such as fetal in-utero imaging, complicate the 3D image and volume reconstruction process.

In this paper we present a learning based image registration method capable of predicting 3D rigid transformations of arbitrarily oriented 2D image slices, with respect to a learned canonical atlas co-ordinate system. Only image slice intensity information is used to perform registration and canonical alignment, no spatial transform initialization is required. To find image transformations we utilize a Convolutional Neural Network (CNN) architecture to learn the regression function capable of mapping 2D image slices to the 3D canonical atlas space.

We extensively evaluate the effectiveness of our approach quantitatively on simulated Magnetic Resonance Imaging (MRI), fetal brain imagery with synthetic motion and further demonstrate qualitative results on real fetal MRI data where our method is integrated into a full reconstruction and motion compensation pipeline. Our learning based registration achieves an average spatial prediction error of 7 mm on simulated data and produces qualitatively improved reconstructions for heavily moving fetuses with gestational ages of approximately 20 weeks. Our model provides a general and computationally efficient solution to the 2D-3D registration initialization problem and is suitable for real-time scenarios.

I. INTRODUCTION

RECONSTRUCTING a 3D volume from misaligned and motion corrupted 2D images is a challenging task. The process often involves labour intensive pre-processing steps, such as manual landmark matching. Such steps exhibit both inter and intra-observer variance. Pre-processing of this nature is however important and necessary to achieve acceptable input for the following intensity-based pose optimization step of a volume reconstruction process. The optimization facilitates alignment and combination of intensity data from multiple image sources into a common co-ordinate system.

In this context, slice-to-volume registration methods [1]–[7] are effective in cases where a coarsely aligned initialization of the 3D volume is provided to initialize the reconstruction process. This initial 3D reference volume is used as a 2D/3D registration target to seed the iterative estimation of the slice orientation and intensity data combination. Reasonably good

initial coarse alignment of 2D image slices is critical to form the seed reference volume. Arbitrary subject motion can invalidate slice alignment assumptions that are based on the scanner co-ordinate system and manual intervention may be necessary. Manual correction often becomes unfeasible in practice due to the magnitude of image data involved.

Image registration is also required for applications such as atlas-based segmentation [8], tracking [9], image fusion from multiple modalities [10] and clinical analysis of images visualized in an anatomical standard co-ordinate system [11]. All of these applications suffer from poor initialization of automatic registration methods, which must be alleviated by manual pre-processing.

Two distinct registration strategies in previous work can be categorized as volume-to-slice and slice-to-volume techniques. The former are concerned with aligning a volume to a given image, *e.g.*, aligning an intra-operative C-Arm X-Ray image to a pre-operative volumetric scan. While the latter is concerned with aligning multiple misaligned 2D slices into a unique co-ordinate system of a reference volume. A recent review of slice-to-volume registration techniques is given in [12].

In practice, volume-to-slice registration is often more feasible than slice-to-volume registration. Manual alignment of one or a few 3D volumes to a single 2D slice or projection is less time consuming than manual alignment of hundreds of individual slices into a common co-ordinate system. Landmark-based techniques can help to automate this process, but this approach is heavily dependent on detection accuracy and robustness of the calculated homography between locations and the descriptive power of the used landmark encoding. 2D slices also do not provide the required 3D information to establish robust landmark matching, therefore this technique cannot be used on applications such as motion compensation in fetal imaging.

Traditional 2D to 3D slice-to-volume intensity-based registration methods [3] involve solving the inverse problem of the super-resolution from slice acquisitions, shown in Eq. 1:

$$y_i = D_i B_i S_i M_i x + n_i; \quad i = 1, 2, \dots, N \quad (1)$$

y_i denotes the i th low resolution image obtained during scan time, D_i is the down-sampling matrix, B_i is the blurring matrix, S_i is the slice selection matrix, M_i is a matrix of motion parameters, and x is the high resolution 3D volume with n_i as a noise vector. More commonly, D_i , B_i and S_i are grouped together in a single matrix W_i . Obtaining the true

B. Hou, B. Khanal, A. Alansary, S. McDonagh, D. Rueckert, B. Glocker and B. Kainz are with Biomedical Image Analysis Group, Department of Computing, Imperial College London, UK

B. Khanal, A. Davidson, M. Rutherford, and J.V. Hajnal are with the Department of Biomedical Engineering King's College London, UK

high-resolution volume x is ill-posed, as it requires inversion of the large, ill-defined matrix W_i .

The optimization methods employed in this problem domain typically do not guarantee a globally optimal registration solution from arbitrarily seeded slice alignment. The function that maps each 2D slice to its correct anatomical position in 3D space may be subject to local minima and the requirement for small initial misalignment typically improves result quality. Previous work have attempted to make this optimization robust by introducing appropriate objective functions and outlier rejections based on robust statistics [5], [13]. Despite these efforts, good reconstruction quality still depends on having good initial alignment of adjacent and intersecting slices.

The robustness of (semi-)automatic 2D-3D registration methods is characterized by their *capture range*, which is the maximum transformation misalignment from which a method can recover good spatial alignment. When the data available is limited, as is common in the 2D-3D case, this task becomes more challenging.

We provide a solution for the capture range problem, which is applicable to many medical imaging scenarios and can be used with any registration method. We demonstrate the capabilities of the method in the current work using in-utero fetal MRI data as a working example. During gestation, high-quality and high-resolution slices can be acquired through fast scanning techniques such as Single Shot Fast Spin Echo (ssFSE) [2]. Slices can be obtained in a fraction of a second, thus freezing the in-plane motion of the subject. Random motion (e.g. a fetus that is awake and active during a scan) or oscillatory motion (e.g. maternal breathing), is likely to cause adjacent slices to become incoherent and corrupt a 3D scan volume. State-of-the art slice-to-volume registration (SVR) methods [5]–[7] are able to compensate for this motion and to reconstruct a consistent volume from overlapping motion-corrupted orthogonal stacks of 2D slices. These methods tend to fail for large initial misalignment between the 2D input slices. Thus, SVR works best for neonates and older fetuses who have less space to move than for very young fetuses. However, for early diagnostics detailed anatomical reconstructions are required from an early age (younger than 25 weeks). Current volume reconstruction pipelines are often unable to successfully fully automate the 3D fetal organ reconstruction process of young and highly mobile fetuses.

Related Work: Image registration methods that can compensate for large initial alignment offsets usually require robust automatic or manual anatomical landmark detection [14]–[16] with subsequent 3D location matching. This often relies on the use of fiducial markers [16]–[19] involving special equipment and/or invasive procedures. Manual annotation of landmarks by a domain expert is the current clinical practice to initialize automatic image registration [16]. Fully automatic methods are difficult to be integrated into the clinical workflow because of their limited reliability, complex nature, long computation times, susceptibility to error and lack of generalization.

Recent work [9], related to the proposed system, uses Convolutional Neural Networks (CNNs) to estimate the automatic spatial arrangement of landmarks in projection images. The method requires a robust landmark detection algorithm,

which is domain and scanner specific and the detection quality degrades for motion-corrupted data. An early version of this work [20] has demonstrated the potential of CNNs for tackling the volume initialization problem for slice-to-volume 3D image reconstruction. The network architecture in [20] showed promising results of initializing scan slices for fetal brain in-utero volume reconstruction, and pose estimation of Digitally Reconstructed Radiograph (DRR) scan images. However, it does not provide a means of knowing incorrect predictions. Failing to account for grossly misaligned slices, that constitute outlying samples, hinder reconstruction performance and may even result in volume reconstruction failure.

We extend the previous work of [20] by rigorously evaluating several network architectures for the volume initialization problem and introduce Monte Carlo Dropout [21] for the purpose of facilitating a prediction confidence metric.

Contributions: In this paper, we introduce a learning based approach that automatically learns the slice transform model of arbitrarily sampled slices relative to a canonical co-ordinate system (*i.e.*, our approach learns the mapping function that maps each slice to a volumetric atlas). This is achieved using only the intensity information encoded in the slice without relying on image transformations in scanner co-ordinates. Our CNN predicts 3D rigid transformations, which are elements of a Special Euclidean group $SE(3)$, to re-orient the slices back into atlas space. This provides an accurate initialization for subsequent automatic registration refinement using intensity-based optimization methods. New statistical analysis and metrics [22], specific for Lie groups, are being incorporated to give a more accurate measure of how far apart a predicted slice is from the ground truth. This is combined with traditional image similarity metrics such as Cross Correlation (CC) and Structural Similarity Index (SSIM).

We evaluate the prediction performance of several CNNs architectures with real and synthetic 2D slices that are corrupted with extreme motion for quantitative comparison. Synthetic slices, with known ground truth locations, are extracted from 3D MRI fetal brain volumes of approximately 20 weeks Gestational Age (GA). We additionally evaluate the approach qualitatively on heavily motion corrupted fetal MRI where ground truth slice transformations and/or 3D volumes are not available. By providing 2D slices that are aligned using our model to initialize the subsequent reconstruction and motion compensation steps, we can qualitatively assess the improvement our methods provide for the volume reconstruction task.

We implemented Monte Carlo dropout sampling at inference to show the model's epistemic uncertainty, and provide a prediction confidence for each slice. This can be used as a metric for outlier rejection (*i.e.*, if the model is not confident about a prediction, then it can be discarded before subsequent use during reconstruction).

Our approach can be generalized to 3D-3D volumetric registration by predicting the transformation parameters of a few selected slices to be used as landmark markers. This is demonstrated in [20] by predicting thorax phantom slices where no organ specific segmentation is performed. It is also applicable to projective 2D images, which is highly valuable for X-Ray/CT registration.

II. METHOD

To fully evaluate and assess the performance of 2D/3D registration via a learning based approach, we incorporated it into a full 3D reconstruction pipeline. This features three modular components: (I) approximate organ localization, (II) canonical slice orientation estimation, and (III) intensity-based 3D volume reconstruction refinement. Organ localization (I) is concerned with localization of a Region of Interest (RoI). This can be achieved through rough manual segmentation, organ focused scan sequences or automatic methods [23]–[25].

For 3D reconstruction refinement (III), and to compensate for the remaining small misalignments between single slices caused by prediction inaccuracies, we use a modified iterative SVR method [6] incorporating super-resolution image reconstruction techniques [6], [7], [26]. To provide a sufficiently high number of samples for SVR, multiple stacks of 2D-slices need to be acquired, ideally in orthogonal orientations. We modified [6] so that instead of generating an initial reference volume from all slices oriented in scanner co-ordinates the acquired slice orientations are replaced with predicted, canonical atlas co-ordinates and continue the iterative intensity-based reconstruction process from there. Since (I) and (III) can be achieved using the state-of-the-art techniques, we focus in the remainder of this section on the canonical slice orientation estimation (II), specifically the learning and prediction of 3D rigid transforms using a variety of network architectures.

At its core, our method uses a CNN, called *SVRNet*, to regress and predict transformation parameters $\hat{T}_i$, such that $\hat{T}_i = \psi(\omega_i, \Theta)$. Θ are the learned network parameters and ω_i is a 2D image extracted from a series of slices that are acquired from a 3D volume Ω . Each Ω encloses a desired RoI, organ or particular use case scenario, such that $\omega_i \in \Omega$.

To train *SVRNet*, slices of varying orientations and their corresponding ground truth locations are required. This can be obtained from existing organ atlases or from collections of motion-free 3D volumes, e.g., pre-interventional scans or successfully (partly manually) motion compensated subjects.

Rigid Body Transformation Parameterization: The motion of a rigid body in 3D has the six Degrees of Freedom (6DoF), namely three parameters for translation (T_x , T_y , T_z) and three parameters for rotation (R_x , R_y , R_z). To model how each slice is able to move in 3D space, we divide the parameters into two categories; in-plane transformation T_x , T_y and R_z and out-of-plane transformation T_z , R_x and R_y (see Fig. 1). If each DoF is allowed ten interval delineation, this would result in 10^6 slices per fetal volume.




Fig. 1. Transformation Parameters Covering 6 Degrees of Freedom

The content of a slice (e.g. segmented brain) is always center aligned during organ localization, thus vastly degrading the in-plane motion parameters T_x and T_y . Similar to [9], we reduce the amount of slices required to create the training

and validation datasets by simplifying to a sample space constrained by the parameters: T_z , R_x , R_y and R_z . We can further reduce a portion of slices that yields little or no content at the extremities T_z of the volume.

Pre-processing and Generating Data: Pre-processing the image intensities, via standardization or Z-score normalization, is a necessary step in training neural networks. It keeps the features of an image in a common scale, with a mean of zero and a standard deviation of one, and similar features across different images consistent.

During scan-time, the intensity and bias ranges are set by the radiologists based on the visual appealing for diagnosis purposes. However, this causes each scanned volume to be biased differently and an intensity range that varies from scan to scan. To be precise, all volumes should be mean centered and have a standard deviation of 1. However, this can be particularly problematic if dealing with cases involving a specific RoI. Standardizing a RoI only requires standardization of the desired content, and not the background. This would therefore require an accurate mask for RoI segmentation, either created manually or through automatic methods. Alternatively, a simpler and quicker approximation can be made by performing only intensity normalization, provided that the intensity histogram for all volumes are of similar ranges. This method does not require a mask for RoI cases.

Every Ω , used in our experiments, has been intensity normalized and standardized to a cubic volume of length L with isotropic spacing. For training and validation, ω_i is extracted from 3D motion corrected volumes, which are of similar intensity range and bias as a result of the reconstruction process [6]. $L/4$ or $L/2$ number of sampling planes of size $L \times L$ are evenly distributed along the Z-axis. For testing, ω_i is a raw slice from a motion-corrupted 3D scanned image.

For a 3D volume containing an organ or RoI that is accurately delineated, the intensities in slices are thresholded (whether they are included in the training and validation set or not) by its variance, $\sigma^2(x)$. This is determined by the threshold variable t , where $t = K \cdot \max(\sigma^2(\omega_i))$, $\forall i \in \Omega$ and $\sigma^2(x) = 1/N \sum_i^{N-1} |x_i - \bar{x}|^2$. If $\sigma^2(\omega_i) < t$, then it is omitted. A larger K value will restrict ω_i to the central portion of the volume, geometrically. In our experiments, $K \approx 0.2$, which samples the central 80% of the volume.

For volumes with inaccurate delineation of the RoI, z is constrained by distance so that it can sample approximately the central 70% of the volume. Therefore, ω_i are extracted between $-0.7 \times L/2 \leq z \leq 0.7 \times L/2$. In such cases we choose to threshold more aggressively in order to reduce the inclusion of edge case slices, where their content can be recovered from intersecting orthogonal slices.

To create a comprehensive training set, ω_i would need to cover a vast amount of different transformation permutations in Ω . Since the transformations parameters are constrained to varying T_z , R_x , R_y and R_z , it is only required to rotate the entire stack of sampling planes in many orientations, where the offset in Z-axis accounts for varying T_z .

A straightforward method is to simply iterate the rotation through Euler angles, where $R(x, y, z)$ and $[x, y, z] \in \mathcal{U}(-\pi/2, \pi/2)$. However, it may not give the best balance of




Fig. 2. Slice plane normals w.r.t the origin via different generation methods. The vector from the origin to each point cross through each slice origin. Note that it is not possible to visualize in plane rotation in this visualization

training slice distribution as seen in Fig. 2b. In this figure, each dot represents the normal vector of the stack of sampling planes from the origin. The second method is to use polar co-ordinates, $P(\phi, \theta)$. Uniform sampling of polar co-ordinates causes an imbalance of samples, with a higher density of samples near the poles. This can be seen in Fig. 2a. A better alternative is to use Fibonacci Sphere Sampling [27], which allows each point to represent approximately the same area as seen in Fig. 2c. The sampling normals are calculated by $P(\phi_i, \cos^{-1}(z_i))$, where $\phi_i = 2\pi i / \Phi$ and $z_i = 1 - (2i + 1)/n$, $i \in 0, 1, 2, \dots, n - 1$. Φ is the golden ratio, as $\Phi^{-1} = \Phi - 1$, and is defined as $\Phi = (\sqrt{5} + 1)/2$.

The rotation of the stack of sampling planes here is defined by the rotation required to transform a 3D vector A onto a 3D vector B . Where A is the starting normal, which is a unit vector in z (i.e. $(0, 0, 1)$), and B is the target sampling normal. However, this does not account for any in-plane rotation, R_z . The entire stack is further rotated around the z -axis through a uniform distribution of angles.

For the validation set, slices are generated with the polar co-ordinates method using random normals and random in-plane rotation angles that are within the bounds of the training set. This is to simulate continuous motion, as test slices will not lie on a discrete training sampling interval, as shown in Fig. 2d.

Loss Functions and Ground Truth Labels: In the literature [28], [29], the most commonly used loss function for regression problems are Euclidean norms: $\|\hat{x} - x\|_n$. However, they may not be suitable when the regression target variables lie in a non-Euclidean space, such as a manifold. For our proposed method, the slices are being transformed rigidly in 3D space, parameterization of each slice therefore lies within the bounds of the SE(3) Lie Group. This includes both a rotation as well as a translation component. There are numerous ways of representing this transformation such as, for rotation; Euler angles, quaternions, rotation matrices.

To address the aforementioned challenge, Kendall et al. [30] proposed the loss in Eq. 2 for PoseNet, which is used to regress the pose of the camera in 6DoF. x and q are the predicted Cartesian translation and quaternion rotation parameters, whilst $\hat{x}$ and $\hat{q}$ are the respective ground truth values. The equation combines the Euclidean distance of translational loss with a weighted Euclidean distance of the rotation loss. β is a tuning parameter that is used to determine the contribution between the losses, by normalizing numbers with different

scales. Quaternions can represent a rotation by using four numbers between $+1$ and -1 , however Cartesian co-ordinates span from $-\infty$ and $+\infty$.

$$\text{Loss} = \|\hat{x} - x\|_2 + \beta \left\| \hat{q} - \frac{q}{\|q\|} \right\|_2 \quad (2)$$

Xu et al. [31] have proposed a framework based on separating loss functions that has an advantage of alleviating over-fitting. The fully connected layer of the network is split into several branches, where each branch is terminated with a separate loss function. Instead of manually tuning the contribution of discrete components in a combined loss function, the network incorporates the tuning parameter within the connection weights. As a result, the network is able to learn from multiple representations of the ground truth labels, for instance, Euler-Cartesian parameters (3 for rotation and 3 for translation) and Quaternion-Cartesian parameters (4 for rotation and 3 for translation).

We introduce a novel labelling system where the rotation and translation components of the labels are combined together. Any three points in a 3D Euclidean space form a plane, while their order defines the orientation. We therefore call them *Anchor Points*.

These 3 Anchor Points can be defined anywhere on the $L \times L$ 2D slice, as long as it is consistent throughout all slices in the dataset. For simplicity, we defined P_1 to be the bottom-left corner $(-L, -L, z)$, P_2 to be the origin $(0, 0, z)$ and P_3 the bottom right corner $(L, -L, z)$. z corresponds to the slice's offset in the Z -axis (see Fig. 3). Each point is multiplied with the rotation matrix that is used to orient the stack of sampling planes. This moves each point, and therefore the slice, to its appropriate location in world space with respect to the volume. Consequently, Anchor Point labels consists of 9 parameters ($P_1(x, y, z)$, $P_2(x, y, z)$ and $P_3(x, y, z)$). As each point is Cartesian, it can be calculated with the standard L2-norm loss function. Incorporating the multi-loss framework [31], losses for P_1 , P_2 and P_3 are calculated independently.

Modeling Network Uncertainties: In this work, we make use of different state-of-the-art network architectures, namely; CaffeNet [32], GoogLeNet [33], Inception [34], NIN [35], ResNet [36] and VGGNet [37]. Where these networks are treated as black boxes, taking ω_i as inputs whilst computing the loss of various labelling methods. We then evaluate the

performance of each architecture on regressing the proposed Anchor Point labels in Sec. IV.

A common strategy during the training of such large, state-of-the-art networks involve the use of dropout [38]. This entails muting components of the true signal, provided to individual neurons. The technique essentially provides a form of model averaging and dropout constitutes a well-understood regularization technique to reduce overfitting. As a result of dropout, neurons have the ability to produce different outputs upon successive activations. During inference, dropout is usually disabled such that network consistency is not undermined. Regression networks are therefore commonly deterministic models at inference time and do not allow for the modeling of uncertainty. Implementing fully probabilistic models that account for uncertainty in both (1) the data and (2) the model parameters (Aleatoric, Epistemic uncertainties respectively [39]) may introduce high computational cost [40].

Gal et al. [21] recently show that dropout layers in Neural Networks can be interpreted as a Bayesian approximation to probabilistic models, and can be implemented by applying a dropout before every weight layer. This is shown to be mathematically equivalent, as it approximately integrates over the models weights. The authors in [41] further show that, for the same input, performing multiple predictions during test time with dropout and taking a mean of the predictions improves result for CNN based networks. This process of performing multiple predictions from the same input by using dropout layers is called as Monte Carlo Dropout sampling which also provides model uncertainty for the given input data.

Using this technique, our experimental work in the following section focuses on taking Epistemic uncertainty into consideration in order to gauge alignment prediction confidence in real-world test cases. Slice alignment requires high precision, and we investigate the idea that a measure of prediction confidence is important to aid reconstruction quality.

Metrics on non-Euclidean Manifolds: Computing the mean prediction of the Monte Carlo Dropout samples, requires an accurate method of averaging the network output which, in our case, is a rigid transformation. However, rigid transformations do not lie in standard Euclidean spaces but in a manifold. In particular, the rigid transformations are Special Euclidean groups $SE(3)$ and is a smooth manifold where intrinsic mean and the corresponding variance can be computed [42].

Consider N rigid transformations predicted by the network: $\{x_i\}, i = 1 \dots N$. We can compute a Riemannian center of mass that minimizes geodesic distances between all data points:

$$m = \underset{y \in \mathcal{M}}{\operatorname{argmin}} \mathbb{E} [\operatorname{dist}(x, y)^2] \quad (3)$$

where $\mathcal{M}$ is a Riemannian manifold and $\operatorname{dist}(x, y)$ defines a geodesic distance between two points x and y on this manifold. A Gauss-Newton iterative algorithm on rigid transformations can be used to compute such a mean, m , from the available data points x_i 's by using a left-invariant metric: [43]

$$m_{t+1} = m_t \circ \exp_{Id} \left(\frac{1}{n} \sum \log_{Id} (m_t^{-1} \circ x_i) \right) \quad (4)$$

where, $\exp_{Id}$ and $\log_{Id}$ are the exponential and logarithmic mappings from the identity defined with left-invariant Riemannian

metric, and $\circ$ is the group composition operator. Once the mean is computed, variance is straightforward to compute with $\sigma = \mathbb{E}[(x - m)^2]$, where the logarithmic operator defines $x - m$ as $\log_m(x)$. [43] provides a detailed treatment of these notions and algorithms.

Transformation Recovery: We transform the predicted Anchor Points, Euler-Cartesian and Quaternion-Cartesian parameters to a rotation matrix and Cartesian offset. We then transform the corresponding slice to its proper location in 3D space. However, the network may introduce a prediction error, which causes the Anchor Points not to be in a perfect triangular formation. To overcome this problem, we assume:

- P_2 defines the Cartesian offset of the slice, i.e. the centre point of the slice in world space.
- The vector joining points P_1 and P_3 aligns with the bottom edge of the slice (this defines the in-plane rotation)
- The Anchor Points together define the plane, which contains the slice. It is used to calculate the rotation matrix to reorient the slice (see Fig. 3).




Fig. 3. Anchor Point Slice parameterization in 3D Space

The Cartesian offset is P_2 . For the orientation, we calculate three orthogonal vectors that defines the new co-ordinate system. Concatenating these vectors gives the rotation matrix, which transforms an identity plane to the newly predicted orientation. $v_1 = P_3 - P_1$ and $v_2 = P_2 - P_1$ are 2 intermediate vectors. v_1 cross v_2 gives the normal of the plane n_1 , this defines the new Z-axis. n_1 cross v_1 gives the new Y-axis n_2 , and v_1 directly defines the new X-axis. Finally, v_2 , n_2 and n_1 are concatenated together to get the rotation matrix.

Evaluation Metrics: The first and foremost naïve method is to monitor the training and validation loss to ensure that the network is learning. To examine a slice in detail, we present the network with a ω_i extracted from Ω . Using the parameters obtained from the network during inference, we extract a new slice from the same Ω and compare it to ω_i . This is completed via several image similarity metrics defined below.

Normalized Cross Correlation (NCC) - measures the similarity of two series (or images) as a function of the displacement of one relative to the other. This searches for features that are similar in both images by examining the pixel intensities.

$$\text{NCC}(\hat{f}, \hat{g}) = \frac{1}{mn} \sum_i \sum_j \hat{f}(i, j) \hat{g}(i, j) \quad (5)$$

where N is the number of pixels in the image and

$$\hat{f} = \frac{f - \bar{f}}{\sqrt{\sum (f - \bar{f})^2}} \quad \text{and} \quad \hat{g} = \frac{g - \bar{g}}{\sqrt{\sum (g - \bar{g})^2}} \quad (6)$$

The Peak Signal-to-Noise Ratio (PSNR) - (based on the MSE metric) is a ratio between the maximum possible power of a signal and the power of corrupting noise that affects the fidelity of its representation. This is the delta-error that is estimated by the network during regression.

$$\text{PSNR} = 10 \cdot \log_{10} \left(\frac{\text{MAX}_I^2}{\text{MSE}} \right) \quad (7)$$

where MAX_I is the maximum possible intensity (pixel value) of the image and

$$\text{MSE}(f, g) = \frac{1}{mn} \sum_i^{m-1} \sum_j^{n-1} (f(i, j) - g(i, j))^2 \quad (8)$$

The Structural Similarity (SSIM) index is an improved metric upon PSNR and MSE, and uses a combination luminance, contrast and structure to assess the image quality. Metrics such as MSE can give a wide variety of degraded quality images with drastically different perceptual quality, which is an undesirable trait.

$$\text{SSIM}(f, g) = \frac{(2\mu_f \mu_g + c_1)(2\sigma_{fg} + c_2)}{(\mu_f^2 + \mu_g^2 + c_1)(\sigma_f^2 + \sigma_g^2 + c_2)} \quad (9)$$

where μ_f and μ_g are the average pixel intensities of images f and g respectively; σ_f^2 and σ_g^2 are the variances of f and g respectively; σ_{fg} is the co-variance of f and g ; $c_1 = (k_1 L)^2$ and $c_2 = (k_2 L)^2$ are two variables that stabilizes the division with weak denominator (i.e. cases where $\mu_f^2 + \mu_g^2 \rightarrow 0$ or $\sigma_f^2 + \sigma_g^2 \rightarrow 0$); L is the maximum possible intensity of the image; with $k_1 = 0.01$ and $k_2 = 0.03$ are default constants.

ta The geodesic distance in the prediction manifold - provides an intrinsic distance measure of how far the predicted rigid transformation is from the ground truth transformation. If x is a ground truth rigid transformation and y is a predicted rigid transformation, a left-invariant geodesic distance with a metric at the tangent space of x can be computed as [43]:

$$\text{dist}(x, y) = \|\log_x(y)\|_x \quad (10)$$

where $\log$ uses the same notion as described in 'Modeling Network Uncertainties'.

III. IMPLEMENTATION

All CNNs were trained using the Caffe [32] package, on a computer equipped with an Intel i7 6700K CPU and a Nvidia Titan X Pascal GPU.

33 randomly chosen fetal volumes were used in our experiments. The volumes were intensity normalized between 0 and 255, and resampled to a dimension of $120 \times 120 \times 120$ with isotropic spacing $0.75 \times 0.75 \times 0.75$ mm. From the 33 volumes, 28 volumes were selected for the training set, with the remaining 5 for validation.

All slices are of size 120×120 (i.e., $L = 120$). 6 separate datasets were generated to cover combinations of slice generation with label representations. 3 datasets were generated via the Euler angle iteration, with the remaining 3 generated using the Fibonacci Sphere Sampling method. Of each slice generation method, there is a dataset for Euler-Cartesian labels, Quaternion-Cartesian labels and Anchor Point labels.

For the Euler generation method, angles are iterated through 18° intervals from -90° to $+90^\circ$. This gives 10 sampling intervals for each axis. Due to the fetal volumes being roughly delineated, variance thresholding is not possible as slices containing only maternal tissue may be included in the dataset. 40 slices are taken in the z-axis, between -40 and +40 with 2 mm interval. This samples the middle 66% of the volume. In total, the Euler iteration method generates 2.24M images for the training set.

For the Fibonacci Sphere Sampling method, 300 sampling normals were chosen. This gives approximately 8° separation between every normal. An additional 10 images, between 0° and 180° with 18° intervals, were generated at each normal to account for in-plane rotation. Slices are also taken between -40 and +40 in the Z-axis with 4 mm interval. This generates in total 3.36M images for the training set.

The validation slices are generated in similar fashion, except at random intervals within the bounds of the training set via random sampling to model continuous fetal motion.

We also adapted the code into Matlab and Python from [22] for mean, variance and geodesic distance computations on SE(3) groups of rigid transformations.

IV. EXPERIMENTATION AND EVALUATION

Network Architecture Performance: The six networks described in Section II were explored to examine if their architectures affect the accuracy on regression, and if so, which architecture gives the best performance to training time ratio. All networks were trained using Caffe with the Adam optimizer, max iteration of 200000, batch size of 64, learning rate of 0.0001, momentum of 0.9, and drop in learning rate of 10% every 20000 iterations. Only ResNet has a batch size of 32 as it was not able to fit onto the GPU memory.

Each network is trained with the dataset generated by the Euler angle iteration method with Anchor Point labels. Fig. 4 shows the training and testing loss during the training process, where each curve represents the loss of each Anchor Point. Table II shows the amount of parameters within each network and the training time for 200000 iterations.

We further analyze the performance of the network with image similarity metrics as well as the geodesic distance from the predicted point to the ground truth location. Since we know the ground truth location, incorrectly predicted slices made by the network are first discarded (e.g., a slice that is predicted outside the volume). We therefore define these slices to have a value that is more than three scaled Median Absolute Deviations (MAD) from the median in the geodesic metric.

$$\text{MAD} = K \cdot \text{median}(|x_i - \text{median}(x)|); \quad i = 1, \dots, N \quad (11)$$

Table I shows the average error of 1000 randomly selected validation slices that were selected from each of the 5 test patients. It can be seen that VGGNet attained the smallest geodesic distance error, as well as correlation and MSE. This is closely followed by GoogLeNet, which managed to attain the smallest PSNR, SSIM and Average Euclidean Distance error (shows the average prediction error of a single Anchor Point to the corresponding ground truth location). NIN is the least




Fig. 4. Loss graphs of Training (rows 1, 3) and validation (rows 2, 4) of the six network architectures used for Anchor Point predictions.

TABLE I
TABLE SHOWING THE MEAN ERROR AND STADNRD DEVIATION OF EACH EVALUATION METRIC

	Correlation		MSE		PSNR		SSIM		Avg E.D. Error (mm)		Avg G.D. Error	
	$\bar{\mu}$	σ	$\bar{\mu}$	σ	$\bar{\mu}$	σ	$\bar{\mu}$	σ	$\bar{\mu}$	σ	$\bar{\mu}$	σ
CaffeNet	0.8137	0.1014	1085.1	584.8715	18.4963	2.5625	0.5327	0.1705	8.3558	4.4449	20.1427	19.1897
GoogLeNet	0.8378	0.1014	935.04	576.4873	19.3564	2.9482	0.5845	0.1775	7.1498	4.8049	15.7504	15.8088
Inception v4	0.7911	0.099	1200.8	562.6443	17.9101	2.283	0.5023	0.1619	9.6816	5.6134	23.9151	21.6426
NIN	0.6734	0.1198	1935.4	701.686	15.6214	1.771	0.3646	0.1333	16.9835	7.5857	74.3866	65.9751
ResNet	0.7829	0.1094	1258.9	620.7798	17.7946	2.5017	0.4913	0.1811	8.6024	4.7808	26.4935	24.1246
VGGNet	0.8384	0.0955	928.51	546.3647	19.3142	2.8288	0.5824	0.174	7.1588	4.3919	14.9142	14.0342

TABLE II
PARAMETER COUNT AND TRAINING DURATION FOR DIFFERENT NETWORK ARCHITECTURES

Network	Number of Parameters	Training Time
CaffeNet	32639449	1 hr 40 mins
GoogLeNet	10253307	5 hrs 55 mins
Inception v4	42783001	21 hrs 20 mins
NIN	6257395	16 hrs 30 mins
ResNet	66566681	25 hrs 40 mins
VGGNet	69159385	12hrs

performing network with much greater errors. For efficiency, the GoogLeNet is the ideal choice, as it can attain almost just as good performance as VGGNet in half the time. If accuracy can be more relaxed, CaffeNet can also be used. ResNet, Inception and NIN however takes much longer to train and perform worse compared to VGGNet. Fig. 5 shows some examples of predictions made by the network.

Regression Labels: In this experiment, we use GoogLeNet for its efficiency. Here, we test the performance of the various datasets, in particular the different types of label parameterization. All 6 datasets adopts the Multi-Loss framework, and as before, the same evaluation methodology is used. 5000 slices from the validation set (1000 per fetal subject) have been randomly selected and passed through the network. Table III shows the average errors, for each dataset.

Here in every metric, Anchor Point labels were able to




Fig. 5. Top: Validation slices that are presented to the network. Bottom: Slices extracted from the respective fetal volume using parameters predicted by the network. (a) to (f) compares the ground truth slices with predicted slice in order of increasing Geodesic Distance Error.

yield a greater accuracy compared to Euler-Cartesian and Quaternion-Cartesian labels in all test cases. A two-tails independent T-test was conducted to examine the statistical significance of Anchor Point labels to Euler-Cartesian and Quaternion-Cartesian, as shown in Tab. IV. As there are 5000 samples in each dataset, the DoF is regarded as infinity. The p-values for all tests are therefore infinitesimally small.

Synthetic Reconstructions: Here, we evaluate the proposed pipeline for reconstructing a 3D MRI fetal brain with synthetically motion corrupted slices. 200 random slices are extracted from the validation dataset in order to initialize the SVR [6]. For all the 5 fetal subjects, we calculate the PSNR between the original and reconstructed 3D volumes: SVR (4 iterations)

TABLE III
AVERAGE ERRORS FOR DIFFERENT FETAL DATASETS

	Correlation	MSE	PSNR	SSIM	G.E.
Anchor Point	0.8378	935.0	19.3564	0.5845	15.7504
Eul-Cart	0.8148	1077.8	18.5254	0.537	18.749
Quat-Cart	0.8199	1046.4	18.6509	0.5448	18.1708
Euler Angle Generation					
	Correlation	MSE	PSNR	SSIM	G.E.
Anchor Point	0.8379	939.6	19.2297	0.5789	13.6885
Eul-Cart	0.8078	1110	18.3034	0.5235	20.4581
Quat-Cart	0.8147	1074.1	18.5055	0.5383	17.7344
Fibonacci Point Generation					

TABLE IV
T-TEST SCORES COMPARING EULER-CARTESIAN AND QUATERNION-CARTESIAN LABELS TO ANCHOR POINT LABELS

	Correlation	MSE	PSNR	SSIM	G.E.
Eul-Cart	10.2973	11.0646	13.8516	12.4705	8.0784
Quat-Cart	8.0766	8.8238	11.7188	10.4969	6.6121
Euler Angle Generation					
	Correlation	MSE	PSNR	SSIM	G.E.
Eul-Cart	14.471	14.0921	16.3974	15.248	18.9279
Quat-Cart	11.2659	11.1481	12.5023	11.0639	12.3679
Fibonacci Point Generation					

and Gaussian Average (no SVR), see Table V.

TABLE V
PSNR OF VOLUMES RECONSTRUCTED FROM SYNTHETIC SLICES COMPARED TO GROUND TRUTH VALIDATION VOLUMES

	#1	#2	#3	#4	#5
Gaussian Average	17.393	16.136	15.112	15.155	15.537
SVR refinement	18.933	19.184	18.560	18.028	20.531

Fetal Reconstruction: We test SVRNet on a real case where both manual and automatic reconstruction methods have failed. With no ground truth target to compare to, reconstruction quality can only be validated qualitatively. Fig. 8a,b,c show the raw scan stacks, and the degree of motion corruption.

Excessive motion can cause ambiguity for previous automatic methods. Fig. 6 shows a sequential sagittal stack of slices where the fetus has turned its head almost 90° , causing a sagittal slice to be a coronal view. This unexpected slice does not fit in the stack, and is normally rejected by robust statistics implemented in SVR. Rejecting too many slices will cause a lack of scan data, while accepting too many slices will cause a corrupt reconstruction volume as seen in Fig. 8d.

As these slices are no longer synthetic, they may be more prone to image artifacts, speckle noise, and shadows. All these factors, combined with no bias correction and smoothing from the reconstruction process, may induce a greater prediction error. Test slices are intensity scaled $[0 - 255]$ to match the intensity range known by the network, no further pre-processing are performed.

Monte Carlo dropout sampling is used in this experiment as




Fig. 6. Sequential scan slices from a sagittal image stack of a fetus with extreme motion, it can be observed that the fetus has rotated its head 90° , causing slice #14 to be a coronal view.

a measure of the network's prediction confidence. Each slice is fed through the network 100 times. The final prediction is obtained by computing the Riemannian center of mass of all 100 predicted transformations. The variance can therefore be calculated to influence the decision whether or not to include the slice in subsequent reconstruction methods, as a large variance suggests the network is not confident with the prediction. Fig. 7 shows the prediction confidence of each slice in the 4 stacks obtained from the scanner (1 coronal, 1 sagittal and 2 axial). As we have generated slices to train the network from the central portion of the fetal volume geometrically, this effect is reflected in the confidence of prediction. It can be seen that network is less confident with edge case slices compared to central slices.




Fig. 7. Slice prediction confidence of the 4 stacks obtained from scanner

The decision of whether or not to include a slice in subsequent reconstruction depends on the prediction confidence and the robustness of the chosen reconstruction algorithm for (III). Prediction confidence can be thresholded and if the reconstruction algorithm is very robust, like [6], we can make multiple predictions per slice and let the reconstruction algorithm handle further outlier rejection, which allows for a greater margin of error (see Fig. 8).

V. DISCUSSION

In this paper, we show that a learning based approach (SVRNet) are able to greatly increase the capture range for 2D-3D image registration, and can provide a robust initialization for subjects with extreme motion. The accuracy of the network predictions is influenced by efficient and novel parameterization of labels and loss functions. In particular, when dealing with parameters that do not lie on an Euclidean manifold. For instance, we found that Euler- and Quaternion-Cartesian labels attain similar performance. However, with a unique parameterization combined with the introduction of Anchor Points, we were able to further increase this accuracy.

During the training of the network, we were able to train six different architectures with fixed hyperparameter configurations achieving satisfactory accuracy to a degree. This shows that for regressing transformation parameters, hyperparameter




Fig. 8. Reconstruction attempt of a fetal brain at approx. 20 weeks GA, from heavily motion-corrupted stacks of T2-weighted ssFSE scans. It was not possible to reconstruct this scan with a significant amount of manual effort by two of the most experienced data scientists in the field using state-of-the-art methods. (a), (b) and (c) are 3 example orthogonal input stacks, with scan plane direction in coronal, axial and sagittal respectively. The view direction is initialized to Stack0 as shown in (a). Stacks (b) and (c) are not perfectly orthogonal, as they are taken at different time points and the fetus has moved. (d) Patch to Volume Registration (PVR) [7] with large patches, *i.e.* improved SVR, using the input stacks directly. (e) Gaussian average of all slices that are predicted and realigned to atlas space by SVRNet. (f) Reconstructed volume after 4 iterations of Slice to Volume Registration (SVR), initialized with slices predicted by SVRNet. The arrow points to an area where insufficient data has been acquired. The fetus moved in scan direction, thus slices are missing in this area necessary for an accurate reconstruction. (g) Training atlas representation of the slices in (e)-(f). Note that (d) is reconstructed in patient space whereas (e) and (f) are reconstructed in atlas space.

searching can be extremely time consuming and computationally expensive, whilst providing little improvement in prediction accuracy. Generating synthetic slices for training is very challenging due to the extensive search through a very large space of parameters (6DoF). We constrain the parameter space by leveraging the effects of the in-plane transformations, based on the center-aligned content as a result of organ localization.

Training a network for different organs, use cases scenarios or modalities (*e.g.*, MRI T1, T2, X-Ray exposure, etc.) can be particularly problematic without organ atlases, or existing 3D reference volumes. SVRNet requires test images to be formatted in the same way as training, this includes identical intensity ranges, spacing and translation offset removal when pre-processing 3D volumes. In our experiments, we show that the training region does not need to be delineated accurately. Our method is not restricted with respect to the used imaging modality and scenario, as seen in [20] where the network is trained on DRR images as well as whole thorax phantoms. This is valuable for 3D to 2D alignment as the whole volume can be aligned to individual 2D slices.

Another issue, which is difficult even for human experts, is to determine left-right asymmetry of a given slice without additional information. To tackle this issue, oversampling and capturing lots of slices during scan time can allow a greater margin for mis-predicted slices. Robust Statistics [6] is able to reject slices predicted in the wrong hemisphere.

The proposed method can be formulated as a classification task, where each rigid transformation can be quantized as a

class. Quantizing the permutations used in our experiments would result in 40K to 50K classes. With only 28 fetal examples per class, this will lead to a high class-imbalance introducing new difficulties in training. Reducing the number of classes, however, will decrease the prediction resolution.

VI. CONCLUSION

SVRNet is able to predict slice transformations relative to a canonical atlas co-ordinate system, using only the intensity information in the slice alone. This allows motion compensation for highly motion corrupted scans, *e.g.*, MRI scans of very young fetuses. It allows to incorporate images that have been acquired during examination and temporal proximity is not essential for good initialization of intensity-based registration methods, as it is the case in state-of-the-art methods.

We have evaluated a wide range of state-of-the-art and popular network architectures to examine their performance on prediction accuracy. We found that VGGNet, in our experiments, attained the smallest regression errors. However, GoogLeNet is more efficient to train for repeating experiments. It can achieve similar results to VGGNet with half the training time, whilst occupying 85% less memory space.

Our work leverages the computational framework to do statistics on SE(3) Lie groups, performing Bayesian Inference and Monte Carlo dropout sampling on the rigid transformation predictions of the network. This approach can also be beneficial in other applications such as [30], where CNNs are trained to produce output transformations that are not in

Euclidean space. We have shown that by calculating geodesic distances of rigid 3D transformations on a non-Euclidean manifold provides means to assess the predicted transformation parameters more accurately. This paves the way to propagate uncertainty downstream in a pipeline that uses the network output to perform other tasks in future work.

VII. ACKNOWLEDGEMENTS

We thank the volunteers, the radiographers J. Allsop and M. Fox, MITK [44], IRTK (CC public license from IXICO Ltd.), the NIHR Biomedical Research Centre. This work was supported by the Wellcome Trust IEH Award [102431]; ERC 319456; NVIDIA (GPU donations); EPSRC EP/N024494/1. Data access only in line with the informed consent of the participants, subject to approval by the project ethics board and under a formal Data Sharing Agreement. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health.

REFERENCES

- [1] F. Rousseau and et al., "Registration-Based Approach for Reconstruction of High-Resolution In Utero Fetal MR Brain Images," *Acad Radiol*, vol. 13, no. 9, pp. 1072–1081, 2006.
- [2] S. Jiang, H. Xue, A. Glover, M. Rutherford, D. Rueckert, and J. Hajnal, "MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): application to fetal, neonatal, and adult brain studies," *IEEE Trans. Med. Imag.*, vol. 26, no. 7, pp. 967–980, 2007.
- [3] A. Gholipour, J. A. Estroff, and S. K. Warfield, "Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI," *IEEE Trans. Med. Imag.*, vol. 29, no. 10, pp. 1739–1758, 2010.
- [4] K. Kim and et al., "Intersection Based Motion Correction of Multislice MRI for 3D in Utero Fetal Brain Image Formation," *Trans. Med. Imag.*, vol. 29, no. 1, pp. 146–158, 2010.
- [5] M. Kuklisova-Murgasova, G. Quaghebeur, M. A. Rutherford, J. V. Hajnal, and J. A. Schnabel, "Reconstruction of Fetal Brain MRI with Intensity Matching and Complete Outlier Removal," *Medical Image Analysis*, vol. 16, no. 8, pp. 1550–60, 2012.
- [6] B. Kainz, M. Steinberger, W. Wein, M. Murgasova, C. Malamateniou, K. Keraudren, T. Torsney-Weir, K., P. Aljabar, M. Rutherford, J. Hajnal, and D. Rueckert, "Fast Volume Reconstruction from Motion Corrupted Stacks of 2D Slices," *IEEE Trans. Med. Imag.*, vol. 34, no. 9, pp. 1901–13, 2015.
- [7] A. Alansary, B. Kainz, M. Rajchl, M. Murgasova, M. Damodaram, D. F. A. Lloyd, A. Davidson, S. G. McDonagh, M. A. Rutherford, J. V. Hajnal, and D. Rueckert, "PVR: patch-to-volume reconstruction for large area motion correction of fetal MRI," *CoRR*, vol. abs/1611.07289, 2016. [Online]. Available: <http://arxiv.org/abs/1611.07289>
- [8] P. Aljabar and et al., "Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy," *NeuroImage*, vol. 46, no. 3, pp. 726 – 738, 2009.
- [9] S. Miao and et al., "A CNN Regression Approach for Real-Time 2D/3D Registration," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1352–1363, 2016.
- [10] S. P. Constantinos, M. S. Pattichis, and E. Micheli-Tzanakou, "Medical imaging fusion applications: An overview," in *Conference Record of Thirty-Fifth Asilomar Conference on Signals, Systems and Computers (Cat.No.01CH37256)*, vol. 2, Nov 2001, pp. 1263–1267 vol.2.
- [11] (2017) Ixi dataset. [Online]. Available: <http://brain-development.org/>
- [12] E. Ferrante and N. Paragios, "Slice-to-volume medical image registration: A survey," *Medical Image Analysis*, vol. 39, pp. 101 – 123, 2017.
- [13] A. Gholipour and et al., "Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI," *IEEE Trans. Med. Imag.*, vol. 29, no. 10, pp. 1739–1758, 2010.
- [14] H. J. Johnson and G. E. Christensen, "Consistent landmark and intensity-based image registration," *IEEE Trans. Med. Imag.*, vol. 21, no. 5, pp. 450–461, May 2002.
- [15] P. Markelj, D. Tomaev, B. Likar, and F. Pernu, "A review of 3d/2d registration methods for image-guided interventions," *Medical Image Analysis*, vol. 16, no. 3, pp. 642 – 661, 2012, computer Assisted Interventions.
- [16] F. P. Oliveira and J. M. R. Tavares, "Medical image registration: a review," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 17, no. 2, pp. 73–93, 2014, pMID: 22435355.
- [17] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2-D-3-D medical image registration," *IEEE Trans. Med. Imag.*, vol. 17, no. 4, pp. 586–595, Aug 1998.
- [18] R. C. Susil, C. Ménard, A. Krieger, J. A. Coleman, K. Camphausen, P. Choyke, G. Fichtinger, L. L. Whitcomb, C. N. Coleman, and E. Atalar, "Transrectal Prostate Biopsy and Fiducial Marker Placement in a Standard 1.5T Magnetic Resonance Imaging Scanner," *The Journal of Urology*, vol. 175, no. 1, pp. 113 – 120, 2006.
- [19] B. Kainz, M. Grabner, and M. Rütther, "Fast Marker Based C-Arm Pose Estimation," in *MICCAI'08, Part II*. Springer, 2008, pp. 652–659.
- [20] B. Hou, A. Alansary, S. G. McDonagh, D. Rueckert, B. Glocker, and B. Kainz, "Predicting slice-to-volume transformation in presence of arbitrary subject motion," *CoRR*, vol. abs/1702.08891, 2017. [Online]. Available: <http://arxiv.org/abs/1702.08891>
- [21] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *Proc. ICML*, ser. Proc. Machine Learning Research, vol. 48. PMLR, 20–22 Jun 2016, pp. 1050–1059.
- [22] N. Miolane, "Defining a mean on Lie group," Master's thesis, Imperial College London, Sep. 2013. [Online]. Available: <https://hal.inria.fr/hal-00938320>
- [23] E. Konukoglu, B. Glocker, D. Zikic, and A. Criminisi, "Neighbourhood approximation using randomized forests," *Medical Image Analysis*, vol. 17, no. 7, pp. 790 – 804, 2013.
- [24] K. Keraudren, M. Murgasova, V. Kyriakopoulou, C. Malamateniou, M. Rutherford, B. Kainz, J. Hajnal, and D. Rueckert, "Automated Fetal Brain Segmentation from 2D MRI Slices for Motion Correction," *NeuroImage*, vol. 101, pp. 633–643, 2014.
- [25] K. Keraudren, B. Kainz, O. Oktay, V. Kyriakopoulou, M. Rutherford, J. Hajnal, and D. Rueckert, "Automated Localization of Fetal Organs in MRI Using Random Forests with Steerable Features," in *MICCAI'15, Part III*, 2015, pp. 620–627.
- [26] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *Signal Processing Magazine, IEEE*, vol. 20, no. 3, pp. 21–36, 2003.
- [27] Á. González, "Measurement of Areas on a Sphere Using Fibonacci and Latitude–Longitude Lattices," *Mathematical Geosciences*, vol. 42, no. 1, p. 49, 2009.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [29] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [30] A. Kendall, M. Grimes, and R. Cipolla, "Posenet: A convolutional network for real-time 6-dof camera relocation," in *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [31] C. Xu and et al., "Multi-loss Regularized Deep Neural Network," *IEEE Trans. Cir. and Sys. for Video Technol.*, vol. 26, no. 12, pp. 2273–2283, Dec. 2016.
- [32] Y. Jia and et al., "Caffe: Convolutional architecture for fast feature embedding," *arXiv:1408.5093*, 2014. [Online]. Available: <http://arxiv.org/abs/1408.5093>
- [33] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Computer Vision and Pattern Recognition (CVPR)*, 2015. [Online]. Available: <http://arxiv.org/abs/1409.4842>
- [34] C. Szegedy, S. Ioffe, and V. Vanhoucke, "Inception-v4, inception-resnet and the impact of residual connections on learning," *CoRR*, vol. abs/1602.07261, 2016. [Online]. Available: <http://dblp.uni-trier.de/db/journals/corr/corr1602.html>
- [35] M. Lin, Q. Chen, and S. Yan, "Network in network," *CoRR*, vol. abs/1312.4400, 2013. [Online]. Available: <http://arxiv.org/abs/1312.4400>
- [36] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03385>
- [37] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [38] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.

- [39] A. Kendall and Y. Gal, “What uncertainties do we need in bayesian deep learning for computer vision?” *CoRR*, vol. abs/1703.04977, 2017. [Online]. Available: <http://arxiv.org/abs/1703.04977>
- [40] G. F. Cooper, “The computational complexity of probabilistic inference using bayesian belief networks,” *Artificial Intelligence*, vol. 42, no. 2, pp. 393 – 405, 1990. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/000437029090060D>
- [41] Y. Gal and Z. Ghahramani, “Bayesian convolutional neural networks with Bernoulli approximate variational inference,” in *4th International Conference on Learning Representations (ICLR) workshop track*, 2016.
- [42] X. Pennec, “Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements,” *Journal of Mathematical Imaging and Vision*, vol. 25, no. 1, pp. 127–154, 2006.
- [43] —, “Statistical computing on manifolds for computational anatomy,” Ph.D. dissertation, Université Nice Sophia Antipolis, 2006.
- [44] I. Wolf, M. Vetter, I. Wegner, T. Bittger, M. Nolden, M. Schbinger, M. Hastenteufel, T. Kunert, and H.-P. Meinzer, “The medical imaging interaction toolkit,” *Medical Image Analysis*, vol. 9, no. 6, pp. 594 – 604, 2005, iTK. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361841505000344>