Deep learning-based deformable registration of dynamic contrast enhanced MR images of the kidney

Huang, James, Guo, Junyu, Pedrosa, Ivan, Fei, Baowei

James Huang, Junyu Guo, Ivan Pedrosa, Baowei Fei, "Deep learning-based deformable registration of dynamic contrast enhanced MR images of the kidney," Proc. SPIE 12034, Medical Imaging 2022: Image-Guided Procedures, Robotic Interventions, and Modeling, 1203410 (4 April 2022); doi: 10.1117/12.2611768

Event: SPIE Medical Imaging, 2022, San Diego, California, United States

Deep Learning-based Deformable Registration of Dynamic Contrast-Enhanced MR Images of the Kidney

James Huang [1,2], Junyu Guo [3], Ivan Pedrosa [3], and Baowei Fei [1,2,3,*]

[1]University of Texas at Dallas, Dept of Bioengineering, Richardson, TX
[2]University of Texas at Dallas, Center for Imaging and Surgical Innovation, Richardson, TX
[3]University of Texas Southwestern Medical Center, Dept of Radiology, Dallas, TX

*E-mail: bfei@utdallas.edu, Website: https://fei-lab.org

ABSTRACT

Respiratory motion is a major contributor to bias in quantitative analysis of magnetic resonance imaging (MRI) acquisitions. Deformable registration of three-dimensional (3D) dynamic contrast-enhanced (DCE) MRI data improves estimation of kidney kinetic parameters. In this study, we proposed a deep learning approach with two steps: a convolutional neural network (CNN) based affine registration network, followed by a U-Net trained for deformable registration between two MR images. The proposed registration method was applied successively across consecutive dynamic phases of the 3D DCE-MRI dataset to reduce motion effects in the different kidney compartments (*i.e.,* cortex, medulla). Successful reduction in the motion effects caused by patient respiration during image acquisition allows for improved kinetic analysis of the kidney. Original and registered images were analyzed and compared using dynamic intensity curves of the kidney compartments, target registration error of anatomical markers, image subtraction, and simple visual assessment. The proposed deep learning-based approach to correct motion effects in abdominal 3D DCE-MRI data can be applied to various kidney MR imaging applications.

Keywords: deformable image registration; deep learning; convolutional neural network (CNN); dynamic contrast enhanced (DCE) MRI; kidney

1. INTRODUCTION

The kidneys are intricate organs tasked with maintaining normal human body functions by keeping a sensitive balance of body fluids which plays a large and direct role in the performance of the cardiovascular and autonomic nervous systems [1]. The renal system preserves this delicate homeostasis through filtration and excretion of metabolic waste products, as well as moderation of fluid volume, blood pressure, and pH levels within the body [2]. Renal cell carcinoma (RCC) arising from the renal epithelium is the most common type of kidney cancer, accounting for about 85% of cases [3]. Small renal mass (SRM) account for over 50% of renal masses and encompass a broad disease spectrum from benign tumors (20%) to aggressive malignancies [4]. Differentiating benign tumors from aggressive tumors is crucial in active surveillance (AS) [5]. Magnetic resonance imaging (MRI) is commonly used in clinical practice for diagnosis of various renal diseases. Dynamic contrast enhanced (DCE) MRI represents a portion of the MRI exam where images are obtained before, during and after the intravenous (I.V.) administration of a gadolinium-based contrast agent (GBCA). The temporal resolution and total acquisition time of DCE-MRI varies depending on the application. In clinical practice, only three or four post-contrast phases are commonly acquired for clinical evaluation of the renal pathology, such as in the characterization of renal masses [6-8]. In contrast, multiple rapid acquisitions (*e.g.* 30-40 phases) are typically obtained with by repeatedly acquiring a 3D T1-weighted gradient echo sequence over several minutes, creating a 4D MRI volume [9]. Following quantitative analysis of the contrast enhancement signal over time, various useful kinetic parameters can be extracted from the data, including blood flow, interstitial volume, and estimated glomerular filtration rate [2,6,7].

Respiratory motion remains one of the main sources of suboptimal image quality and data corruption in quantitative analysis of abdominal MRI acquisitions [8]. Kidney imaging is highly susceptible to respiratory motion and several strategies have been developed for reducing motion artifacts such as external triggering techniques and image-based

Medical Imaging 2022: Image-Guided Procedures, Robotic Interventions, and Modeling,
edited by Cristian A. Linte, Jeffrey H. Siewerdsen, Proc. of SPIE Vol. 12034, 1203410
© 2022 SPIE · 1605-7422 · doi: 10.1117/12.2611768

navigators [10]. Rigid and non-rigid registration-based approaches, including optical flow and demons, have also been utilized and found to be relatively successful at reducing motion artifacts in abdominal DCE-MRI and other MRI data [2,11-13]. However, co-registration of images present additional challenges beyond artifact reduction. Indeed, in addition to displacement, the presence of kidney rotation, tilting, and deformation due to respiration is common. Moreover, many non-rigid registration algorithms result in alteration of the pixel signal, which confounds subsequent quantitative analysis. The advent of deep learning can further improve upon these registration algorithms by increasing registration speed and universality while overcoming barriers such as lack of training data through data augmentation. Deep learning has been applied to such image-guided interventions relying on high registration speed and spatial accuracy – for example, in MR-US fusion for brachytherapy catheter placement and CT to x-ray registration for intraoperative surgeries [14]. Intra and inter-patient registration of brain 3D MR scans has greatly benefitted within the last few years from newly developed deep learning software, including VoxelMorph and Quicksilver [15-17]. These deep learning solutions have not been applied yet to abdominal MRI registration, where large organ displacements during image acquisition can occur [10]. Kidney DCE-MRI data can potentially benefit for image correction through deep-learning based registration [18]. The goal of this study is to evaluate the performance of a novel unsupervised deep learning-based approach for correcting kidney motion in DCE-MRI data.

2. METHODS

2.1 DCE-MRI data

The abdominal 3D DCE-MRI data used in this project were acquired from 20 patients with known small renal masses enrolled in a prospective trial. All patients signed informed consent prior to imaging. DCE-MRI datasets, each including 39 dynamic phases, were acquired using 3D T1-weighted gradient echo sequence on Philips 3T MRI scanner. The images were in 12-bit greyscale and had varying spatial dimensions of the same height-to-width ratio. Images were manually cropped to remove border, reduced to 8-bit greyscale, and resized to (224, 384) using bilinear interpolation.

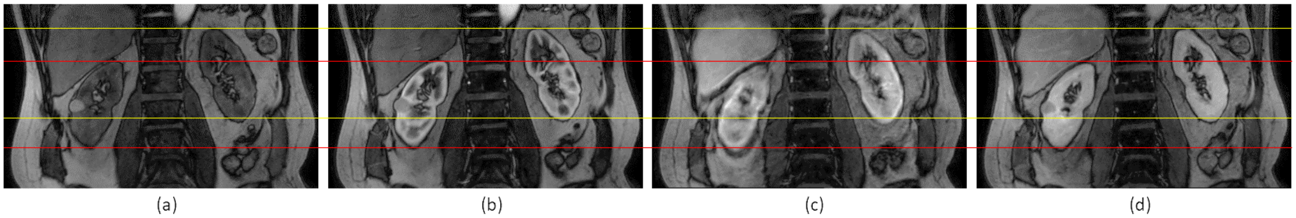

Figure 1. DCE MR images at four different dynamic phases in chronological order during the perfusion process. Yellow lines intersect upper and lower limits of left kidney in (a) and red lines intersect upper and lower limits of right kidney in (a). Images depict (a) phase 3 (pre-contrast), (b) phase 5 (during the corticomedullary phase), (c) phase 8 (early nephrographic), and (d) phase 13 (nephrographic) of the same MRI slice.

Figure 1 demonstrates the need for deformable registration in the kidneys. While most DCE-MRI frames acquired do not present any major organ shifting or motion artifacts using breath holding, the frames that do present significant amounts of deformation, such as in Figure 1c, detract important information from the scan which negatively affects the quantitative analysis. The transformations due to respiration and other factors of the right and left kidneys are independent from one another, causing global affine registration to be insufficient as a sole solution.

In this project, manual kidney delineations were utilized for both training the networks and performance evaluation. The patient images were manually segmented in AnalyzePro (AnalyzeDirect, Inc., Overland Park, KS, 66085) to delineate the kidneys. Figure 2 demonstrates the high-level workflow for deformable kidney registration and analysis.

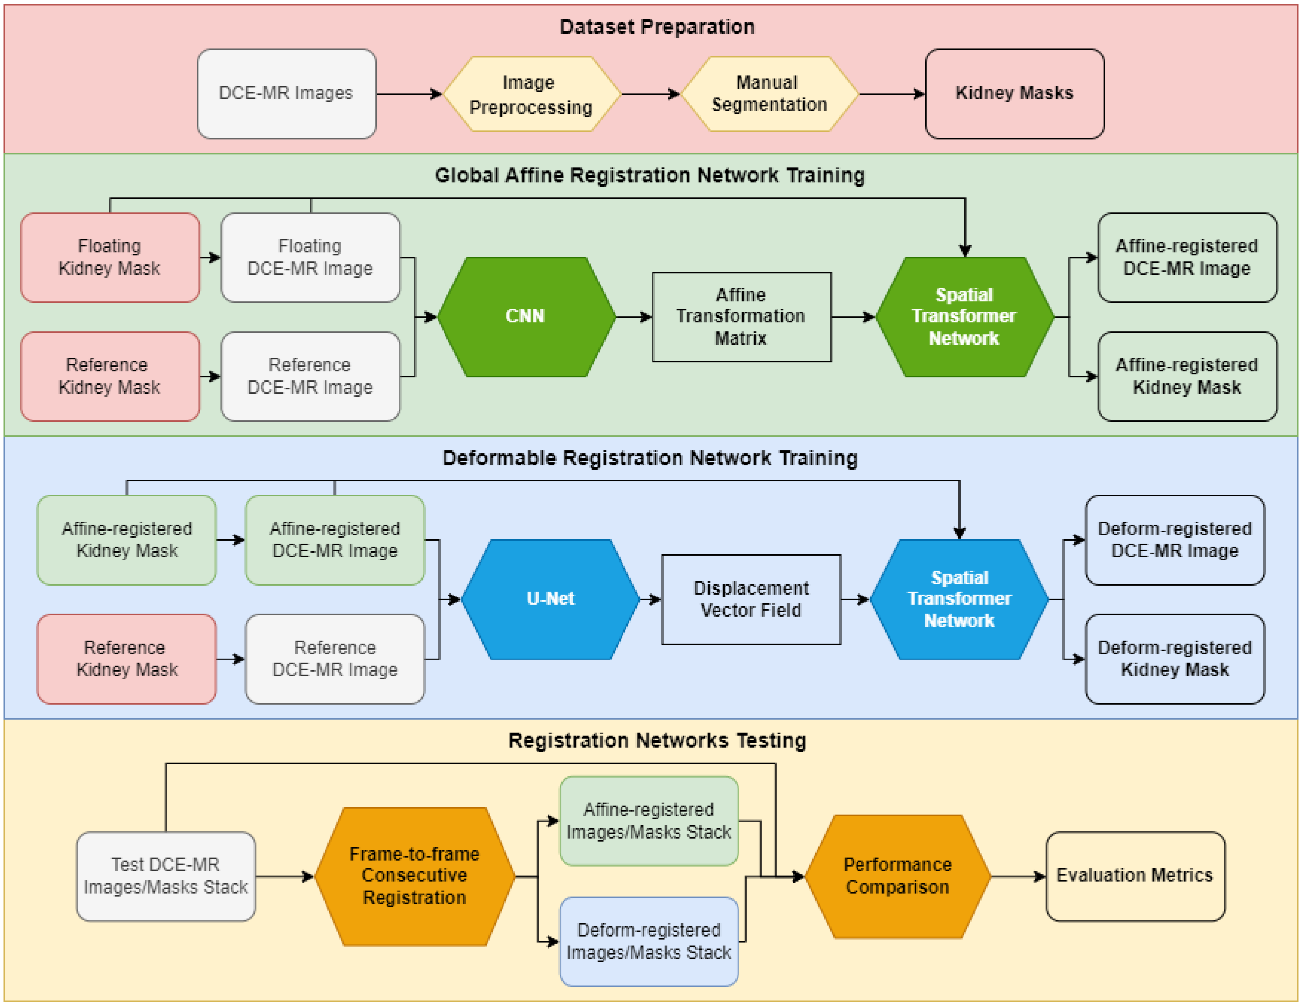

Figure 2. Graphical abstract of deep-learning based registration of kidney DCE-MR images.

2.2 Affine registration network

The affine registration network aims to perform preliminary alignment of DCE-MR images globally to improve the performance of the subsequent deformable registration step. As shown in Figure 3, the network takes a floating DCE-MR image with its corresponding segmentation as an input and a reference pair of an image and segmentation as the desired output. The predicted output is the affine-transformed floating image pair which should closely resemble the reference image pair.

Similar affine registration network architectures have been explored in the AIRNet designed by Chee et al. [19] and ConvNet proposed by de Vos et al. [20]. While our method of feature extraction from the images may differ slightly from AIRNet and ConvNet (github.com/BDdeVos/TorchIR.git), all three networks implement a convolutional neural network (CNN) encoder architecture to output a matrix of transformation parameters for affine registration [19,20]. Following concatenation of images and segmentations, each input image underwent a series of five 3x3 convolutions and max pooling with a pool size of 2 in the encoder. The resultant matrix was flattened and densely connected to reduce it to an array of six values dictated the various affine transformations of translation, rotation, scaling, and shearing. The decoder can be described as a variation of the popular spatial transformer network, proposed by Jaderberg et al. [21], which converts a regular spatial grid of the same size as the input image to a sampling grid and applies the sampler to the input images to produce a warped output image. Bilinear interpolation was used to fill in unknown values.

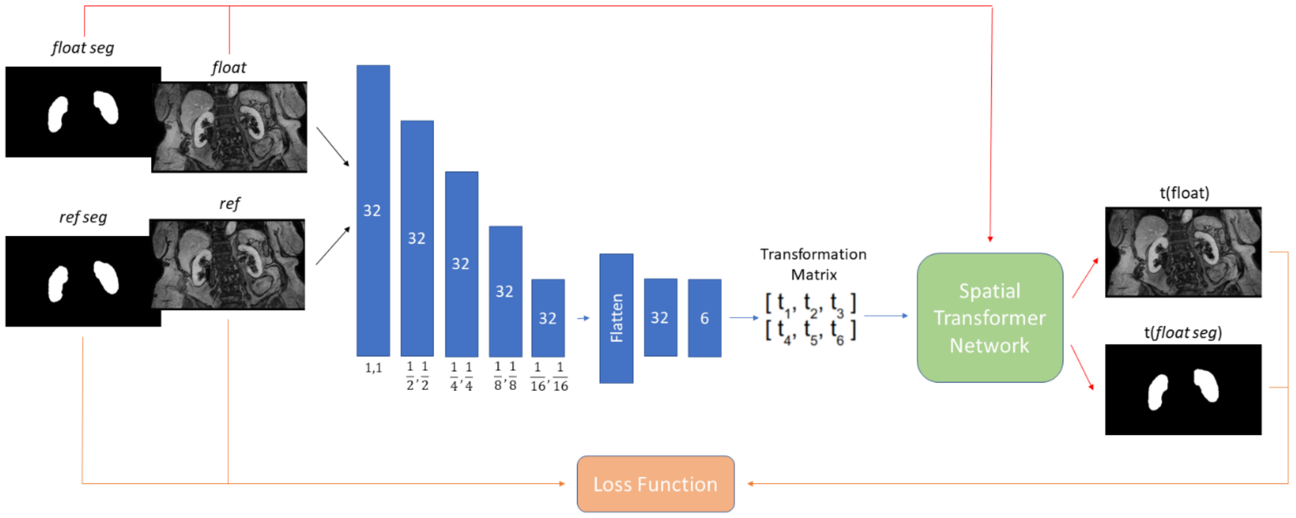

Figure 3. Network architecture for 2D affine image registration. Inputs for the network are the floating MR images with corresponding labels and the reference images and labels, and predicted outputs are the registered images and labels.

A loss function was optimized for backpropagation to the encoder of the affine registration network. Affine loss was updated as a function of the similarity between the registered image and reference image, as well as the registered label and reference label. The affine loss can be represented by:

$$L_{affine} = \alpha L_{img} + \beta L_{seg} \tag{1}$$

where L_{img} is a similarity metric between the registered image and reference image multiplied by weight α and L_{seg} is a similarity metric between the registered label and the reference label multiplied by weight β. Although the segmentations were vital for performance evaluation of the deformable registration network, it was noted that there was an insignificant difference in affine registration fidelity between two network variations programmed for inclusion and exclusion of the segmentation information as inputs into the affine registration network. Therefore, the L_{seg} term can be considered as negligible in the affine case.

2.3 Deformable registration network

The 2D deformable registration network was adapted from the open-source deep learning registration library, VoxelMorph designed by Balakrishnan et al. VoxelMorph (github.com/voxelmorph/voxelmorph.git) employs a U-Net architecture that takes two input volumes and outputs a one-to-one mapping of all voxels from one *n*-D volume to the other as a displacement vector field (DVF). This state-of-the-art unsupervised network learns and optimizes a single global function during the training phase without the need for ground truth deformations which can be rapidly and simply evaluated on any testing pair of volumes with the same number of dimensions [15,16].

Inputs to the deformable registration network are the floating image and floating label, and outputs consist of the registered image, registered label, and the DVF, denoted by ϕ in Figure 4. An additional input array of zeros of the same size as the images, named ϕ_0 is used, as detailed in the VoxelMorph papers, to provide a target DVF to encourage smooth and realistic deformations [15,16]. The network's feature extraction step was constructed as a standard U-Net, which consists of successive 3x3 convolutions in the encoder stages followed by max pooling by a factor of 2 [22]. As displayed in the network architecture in Figure 4, two dropout layers are included following the fourth and fifth convolutions to reduce overfitting. Decoder layers consist of deconvolutions and concatenations with skip connections to layers in the encoder with the corresponding spatial sizes to rebuild the spatial awareness of learned features [22]. The final output layer has n filters, where n is the number of spatial dimensions in the input data. This output can be considered as a displacement vector field, with a deformation vector at each voxel [15,16]. As with the affine registration network, a spatial transformer network was adapted to decode the DVF into a grid to resample the floating images and segmentations, using bilinear interpolation to fill in unknown values.

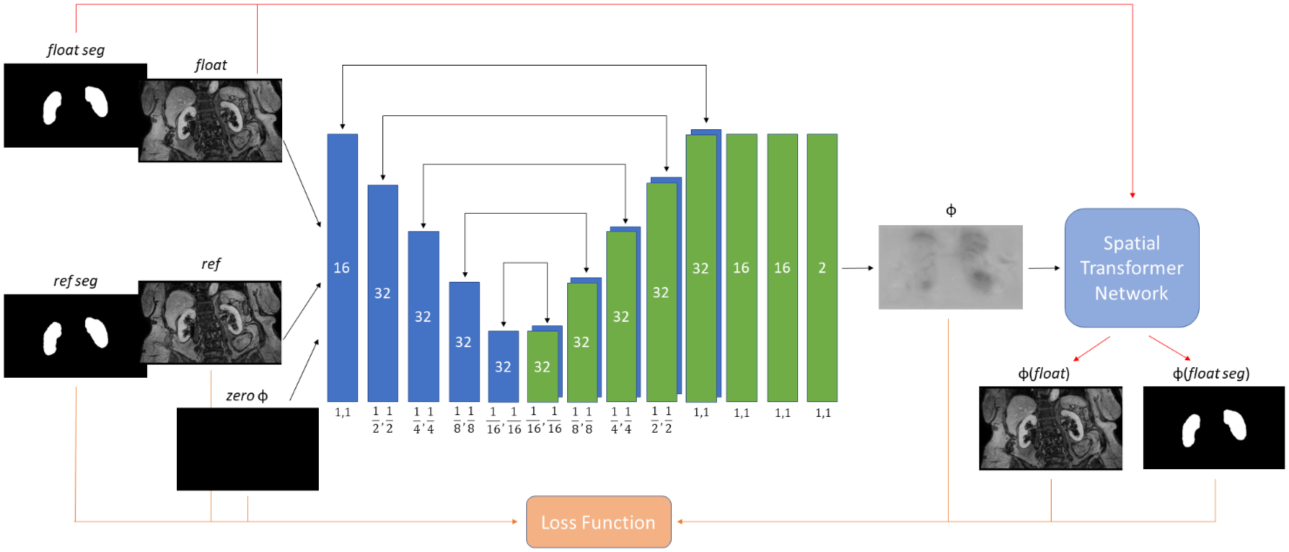

Figure 4. Network architecture for 2D deformable image registration. Inputs for the network are the floating image/segmentation and the reference image/segmentation and ϕ_0, and predicted outputs are the registered image/segmentation and the displacement vector field, ϕ.

Applying the DVF ϕ to the floating image using grid resampling and comparing the newly registered image to the reference image allows us to generate an unsupervised similarity loss. The same grid can be used for the floating segmentation to generate the registered segmentation and a more localized similarity metric for the kidneys. To achieve an anatomically significant deformation field in medical image registration applications, a smoothing term is applied in the loss function which adds the needed stability to the registration algorithm [2,15,16]. The loss function for deformable kidney registration is given by:

$$L_{deform} = \alpha L_{img} + \beta L_{seg} + \lambda L_{smooth} \tag{2}$$

where L_{img} and L_{seg} correspond to similarity metrics between the registered images/segmentations and reference images/segmentations. The third term, L_{smooth} refers to the similarity metric between the generated DVF from the U-Net and the empty array, ϕ_0. Each term is multiplied by a set coefficient to scale the weights of each of the loss components. These loss weights were discovered through a simple grid search.

2.4 Training

4D DCE-MR images from 20 patients were divided randomly into training (60%), validation (20%), and testing (20%) datasets. 4D images from 12 patients were used for training, 4 were used for validation, and the last 4 were reserved for testing. With each MR image containing between 20 and 30 slices, this equated to approximately 500 slices in total and 39 phases, or frames, for each slice for a total of 19500 2D images. For each training pair, the reference image was chosen at random from a randomly selected slice between phases 1 and 39 and the floating image was selected from the remaining phases within the same slice.

The affine registration network was trained using mean square error (MSE) as the image dissimilarity loss with a weight coefficient of α=1 and the Sørensen distance, or the complement to the dice similarity coefficient (DSC), with a weight of β=0.1 as the dissimilarity loss between the segmentations. The deformable registration network was trained using MSE as the image dissimilarity loss with a weight of α=1, the Sørensen distance as the segmentation dissimilarity loss with a weight of $\beta=0.15$, and the average squared gradient of the outputted DVF as the smoothing loss with a weight coefficient of $\lambda=0.05$.

Training for the model was done on a high-performance computer containing an 8-stack of Nvidia GeForce GTX Titan XPs. The weighted model whose results are evaluated in Section 3 was trained over 100 epochs with 25 steps per epoch and a batch size of 64. Training time was approximately two hours total at 32 seconds per epoch, including both affine and deformable registration models.

2.5 Validation and Testing

DCE-MRI data from four patients were used for testing registration. Validation was performed during training by tracking the normalized cross correlation, sum of squared difference, and mean squared error of the validation dataset. After training, testing of the model's efficacy was conducted via a randomly- selected slice from the testing dataset which contained the kidneys. Successive registration using the model was applied from phase to phase of that slice, and a video of the registered phases was generated. A side-by-side comparison of a video of the original phases next to the registration video was used as a means for preliminary evaluation of performance. Furthermore, subtraction images were built to display the difference between the registered and original images and the Jacobian of the DVF was shown as well. Simple visual inspection for the purpose of verifying good model functionality assumed three things: 1) The image registration successfully warped the original image to be closer to the reference image without losing valuable intensity data, 2) the degree of motion-related misregistration was decreased, and 3) the deformation field was smooth and seemed anatomically significant.

Additional quantitative evaluation involved four methods to assess registration efficacy. The first was generating the time intensity curve (TIC) of certain selected regions of interest (ROI) within the kidney. The TIC is the intensity signal within the same ROI across each of the phases. It is through the analysis of the TIC that important physiological properties for kidney disease diagnosis can be derived [23]. Motion artifacts and organ displacement due to respiration and other factors can cause a volatile and inaccurate TIC which makes diagnosis difficult [2]. Thus, successful registration of the kidneys should encourage a smoother, less jagged TIC while retaining the important intensity information as shown in Figure 6. The second quantitative metric was calculating the target registration error (TRE) of anatomical landmarks through manually placing several markers at the same set of locations on a reference image, affine-registered image, and deformable-registered image from the same slice and phase. To evaluate the fidelity of kidney delineations, the dice similarity coefficient and Hausdorff distance (HD) were calculated.

3. RESULTS AND DISCUSSION

3.1 Registration performance

After training, a set of predictions were generated through a series of 38 successive registrations, starting from registering phase 2 to phase 1, then registering phase 3 to the registered phase 2, and so on until phase 39 was aligned to the registered phase 38.

Two divisions of registration performance evaluation must be defined based on the chosen phases to compare. The first division can be referred to as successive comparison, in which the evaluation methods calculate the differences between consecutive phases or frames, such as the dice similarity coefficient between phase 12 and phase 11 as demonstrated in Figure 5a. The second type of registration assessment can be referred to as static comparison, in which the differences between each phase and a single selected phase are analyzed. In this project, the static image was chosen to be that of the frame of peak enhancement, or the phase in which the kidneys display the maximum signal intensity in DCE-MRI following contrast agent injection. While successive comparison provides information on organ shifting between frames, static comparison provides a useful technique for analyzing the overall stability of the registrations across all frames. It is important to note that all predictions were done using successive registration and static registration is only an evaluation method. Figure 5b provides an example of a static comparison between phase 9 and the peak phase of the same slice following deformable registration.

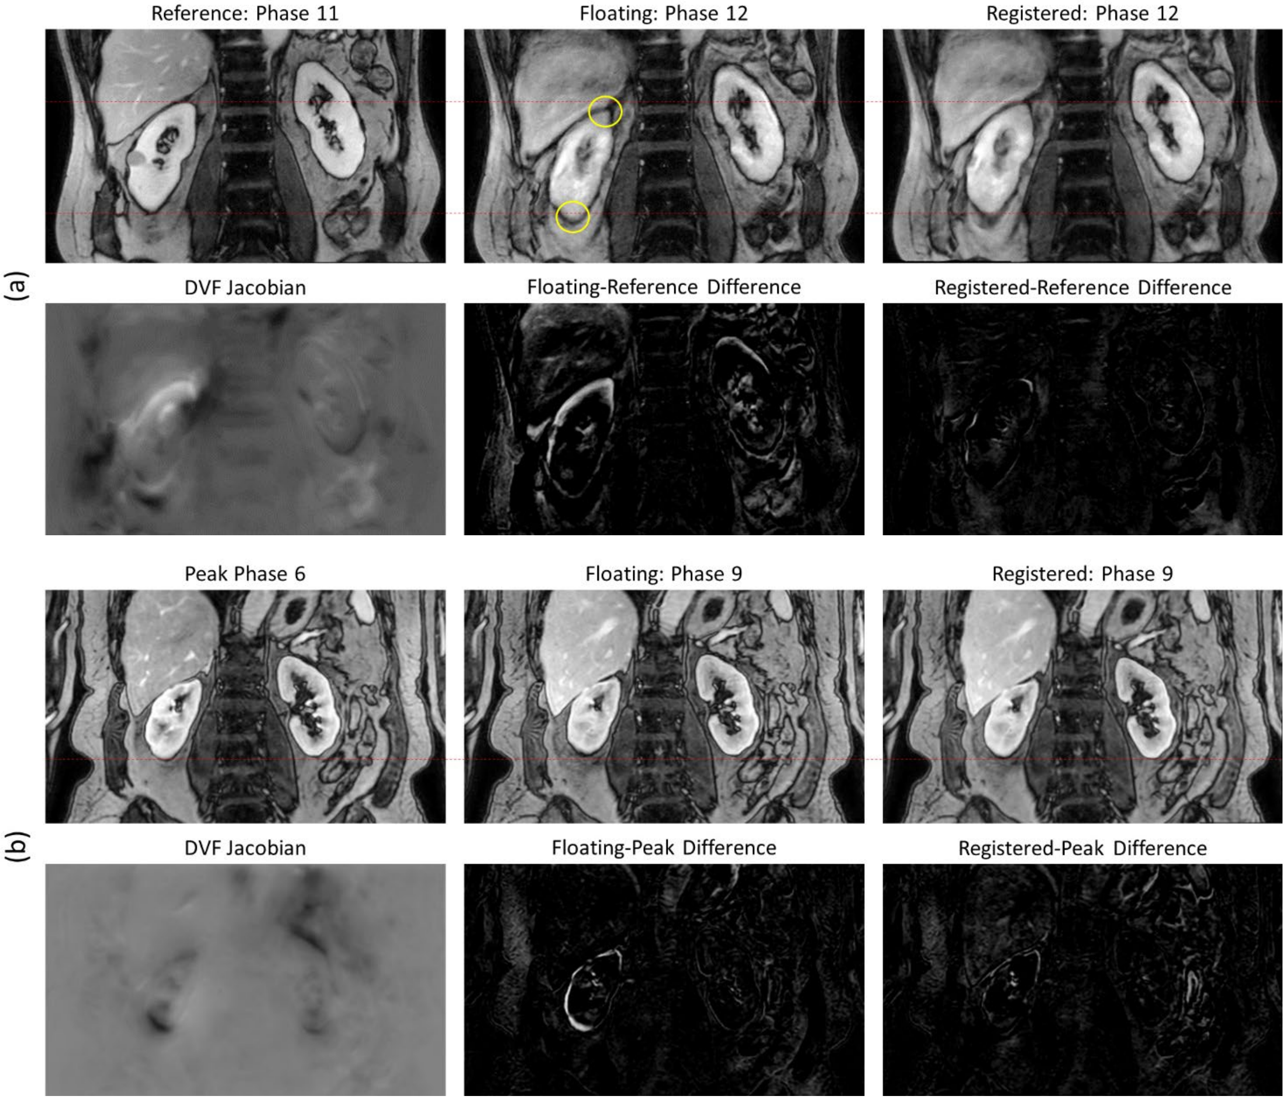

Figure 5. Two sample registrations with floating and registered phases as compared to a reference image from the same slice. The floating images represent the original image before both affine and deformable registration. The registered images represent the post-affine and post-deformable registered result. (a) Successive comparison of registration performance example between phase 12 and phase 11 of the same slice. (b) Static comparison of registration performance example between phase 9 and peak enhancement phase 6 of the same slice.

Based on simple visual inspection of the two consecutive phases in Figure 5a with their corresponding difference and registration deformation fields, it can be shown that there is a reduction in the motion artifacts caused by patient respiration in phase 12. In this case, there is a clear improvement in registration on the right kidney in both the inferior and superior regions as compared to the original while the intensity information within the cortex and medulla are still preserved after registration. In Figure 5b, where phase 9 is compared to the peak phase 6, it can be seen, especially in the given difference images that there is an improvement in the alignment of the inferior region on the right kidney as a result of registration.

Target registration errors were collected using AnalyzePro, with two selected targets on anatomical landmarks per kidney for all test slices, including original, post-affine, and post-deformable registration cases. Both successive and peak evaluation methods for Dice Similarity Coefficient, Hausdorff Distance (HSD), and TRE from all test slices were utilized to generate the metrics given in Tables 1 and 2. DSC and HSD are based on kidney segmentation masks.

Table 1. Dice Similarity Coefficient and Hausdorff Distance presented as mean and Target Registration Error presented as mean ± standard deviation for each of the original, affine, and final-warped images compared to the immediately preceding phase in the same slice.

Metric	Original	Post-Affine	Post-Deformed
DSC	0.927	0.937	0.948
HSD (mm)	2.96	2.68	2.09
TRE (mm)	3.09 ± 2.51	3.04 ± 1.76	2.15 ± 1.34

Table 2. Dice Similarity Coefficient and Hausdorff Distance presented as mean and Target Registration Error presented as mean ± standard deviation for each of the original, affine, and final-warped images compared to the static peak phase in the same slice.

Metric	Original	Post-Affine	Post-Deformed
DSC	0.928	0.933	0.949
HSD (mm)	2.97	2.61	2.40
TRE (mm)	3.18 ± 2.58	2.82 ± 2.06	1.09 ± 1.39

By comparing the metrics across each row in Tables 1 and 2, it can be shown that both the affine and deformable registration steps offered a significant improvement over the original DCE-MR images. In addition, paired two-tailed t-tests between the post-deformed and original TREs for successive and static peak cases were conducted, giving a p-value of 0.004 in the successive case and 0.001 in the static case. Time intensity curves for three example slices for the original and post-deformable registration cases are shown in Figure 6. The TICs demonstrate that the deformable registration produced a desired smoothing behavior while maintaining vital intensity information, such as time to peak enhancement.

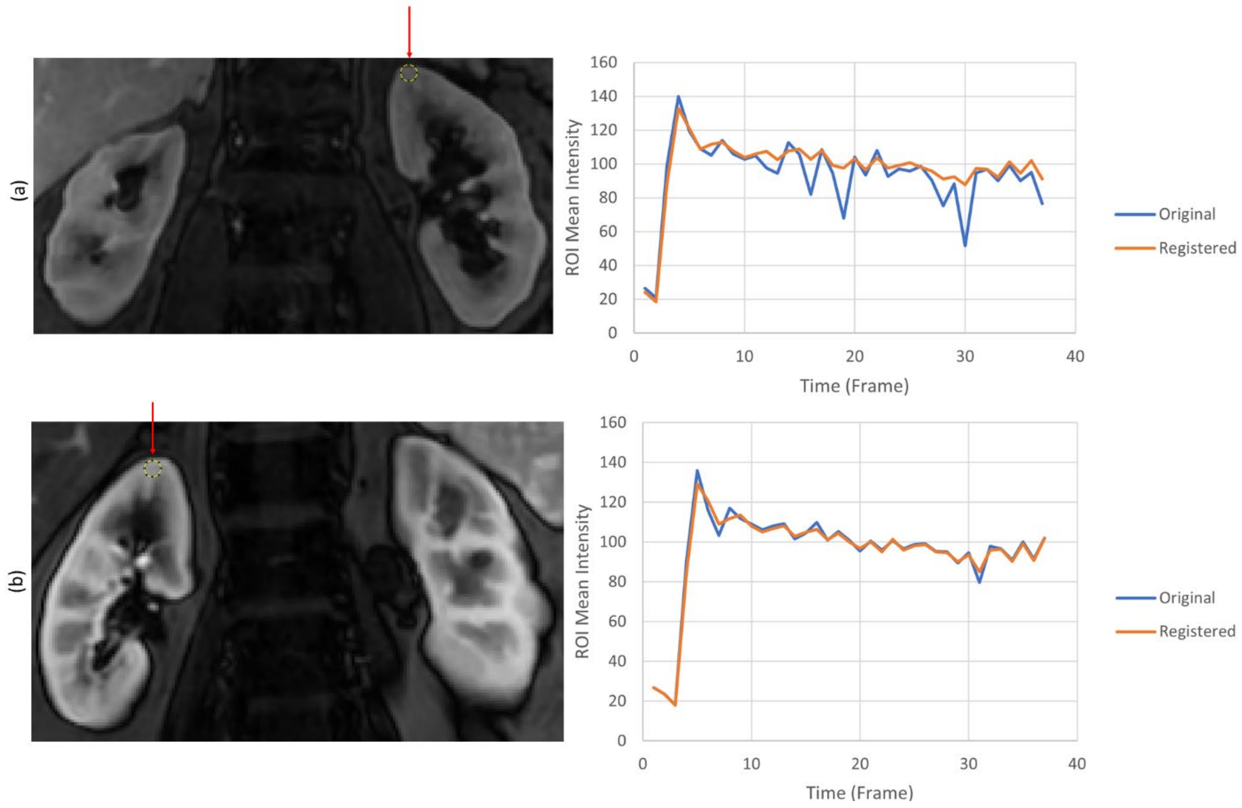

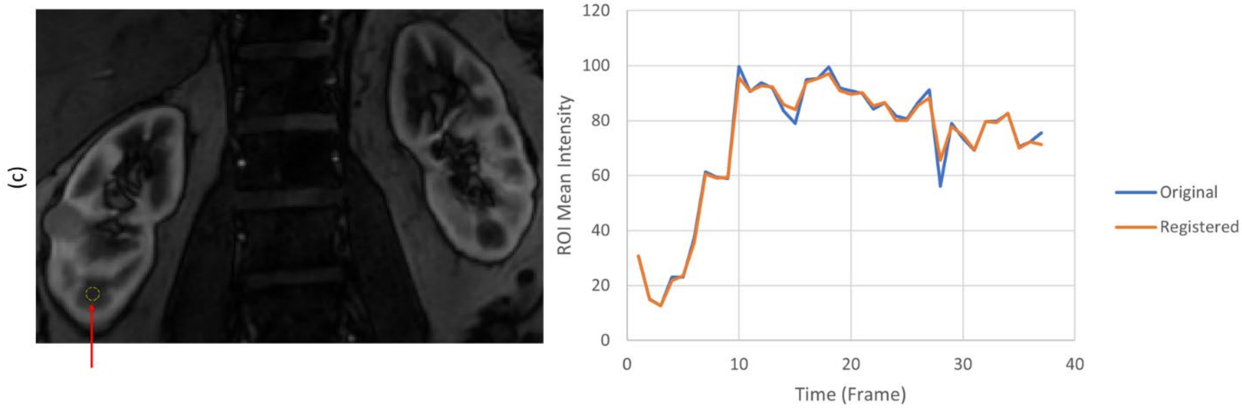

Figure 6. Evaluation of the registration performance based on the comparison between dynamic signal intensity curves within a ROI (left) before registration and after registration. X axis is the dynamic phase and Y axis is the mean signal intensity within the ROI. The ROIs chosen are given by the dotted yellow circles on the corresponding MR images. These correspond to radial regions (a) within the cortex of the left kidney, (b) within the cortex of the right kidney, and (c) within the medulla.

4. CONCLUSIONS AND DISCUSSIONS

We implemented a deep learning registration for motion correction of kidney DCE-MRI. We conducted visual qualitative as well as quantitative analysis of registration performance through extracting such metrics as dynamic signal intensity curves and target registration error at anatomical landmarks. The registration method is able to improve the image alignment non-rigidly for motion correction of the kidneys.

Deep learning registration for motion correction of DCE MRI exams of the kidneys is a new application. Such techniques have been applied and previously reported to improve registration with the liver and lungs to remove motion effects. Additionally, traditional registration methods such as a Demons algorithm have been utilized for kidney registration in DCE-MRI with variable success. We confirm that an unsupervised neural network can be utilized to improve the performance and efficiency of registration for motion correction in kidney imaging. This tool could be applied to a variety of clinical scenarios such as operative planning, radiation therapy, and longitudinal assessment of renal function on dedicated MRI exams.

ACKNOWLEDGMENTS

This research was supported in part by the U.S. National Institutes of Health (NIH) grants (R01CA156775, R01CA204254, R01HL140325, R01CA154475 and R21CA231911), by the Cancer Prevention and Research Institute of Texas (CPRIT) grant RP190588, and by the Career Enhancement Program supported by the UT Southwestern SPORE.

REFERENCES

[1] L. Robson, “The kidney - an organ of critical importance in physiology,” *The Journal of Physiology*, vol. 592, no. 18, pp. 3953–3954, 2014.
[2] X. Yang, P. Ghafourian, P. Sharma, K. Salman, D. Martin, and B. Fei, “Nonrigid registration and classification of the kidneys in 3D dynamic CONTRAST enhanced (dce) MR images,” *Medical Imaging 2012: Image Processing*, 2012.
[3] H. T. Cohen and F. J. McGovern, “Renal-cell carcinoma,” *N Engl J Med*, vol. 353, no. 23, pp. 2477–90, 2005.
[4] S. Weikert and B. Ljungberg, “Contemporary epidemiology of renal cell carcinoma: Perspectives of primary prevention,” *World Journal of Urology*, vol. 28, no. 3, pp. 247–252, 2010.

[5] C. Sebastià, D. Corominas, M. Musquera, B. Paño, T. Ajami, and C. Nicolau, “Active surveillance of small renal masses,” *Insights into Imaging*, vol. 11, no. 1, 2020.

[6] S. P. Sourbron, H. J. Michaely, M. F. Reiser, and S. O. Schoenberg, “MRI-Measurement of Perfusion AND glomerular filtration in the human kidney with a separable compartment model,” *Investigative Radiology*, vol. 43, no. 1, pp. 40–48, 2008.

[7] N. Hackstein, H. Kooijman, S. Tomaselli, and W. S. Rau, “Glomerular filtration rate measured using the Patlak plot technique AND contrast-enhanced dynamic MRI with different amounts of gadolinium-DTPA,” *Journal of Magnetic Resonance Imaging*, vol. 22, no. 3, pp. 406–414, 2005.

[8] U. I. Attenberger, S. P. Sourbron, H. J. Michaely, M. F. Reiser, and S. O. Schoenberg, “Retrospective respiratory triggering renal perfusion MRI,” *Acta Radiologica*, vol. 51, no. 10, pp. 1163–1171, 2010.

[9] S. E. Yuksel, A. El-Baz, A. A. Farag, M. E. Abo El-Ghar, T. A. Eldiasty, and M. A. Ghoneim, “Automatic detection of renal rejection after kidney transplantation,” *International Congress Series*, vol. 1281, pp. 773–778, 2005.

[10] Y. X. J. Wang, X. Wang, P. Wu, Y. Wang, W. Chen, H. Chen, and J. Lee, “Topics on quantitative liver magnetic resonance imaging,” *Quantitative Imaging Medical Surgery*, vol. 9, no. 11, pp. 1840–1890, Nov. 2019.

[11] B. Fei, J. L. Duerk, D. B. Sodee, and D. L. Wilson, “Semiautomatic nonrigid registration for the prostate and Pelvic MR volumes1,” *Academic Radiology*, vol. 12, no. 7, pp. 815–824, 2005.

[12] B. Fei, C. Kemper, and D. L. Wilson, “A comparative study of warping and rigid body registration for the prostate and pelvic MR volumes,” *Computerized Medical Imaging and Graphics*, vol. 27, no. 4, pp. 267–281, 2003.

[13] B. Fei, Z. Lee, D. T. Boll, J. L. Duerk, J. S. Lewin, and D. L. Wilson, “Image registration and fusion for interventional MRI guided thermal ablation of the prostate cancer,” *Lecture Notes in Computer Science*, pp. 364–372, 2003.

[14] Y. Fu, Y. Lei, T. Wang, W. J. Curran, T. Liu, and X. Yang, “Deep learning in medical image registration: a review,” *Physics in Medicine and Biology*, vol. 65, no. 20, p. 20TR01, Oct. 2020.

[15] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, “VoxelMorph: A learning framework FOR Deformable medical image registration,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 8, pp. 1788–1800, 2019.

[16] G. Balakrishnan, A. Zhao, M. R. Sabuncu, A. V. Dalca, and J. Guttag, “An unsupervised learning model for deformable medical image registration,” *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

[17] X. Yang, R. Kwitt, M. Styner, and M. Niethammer, “Quicksilver: Fast predictive image registration – a deep learning approach,” *NeuroImage*, vol. 158, pp. 378–396, 2017.

[18] S. Y. Huang, R. T. Seethamraju, P. Patel, P. F. Hahn, J. E. Kirsch, and A. R. Guimaraes, “Body MR imaging: Artifacts, k-space, and solutions,” *RadioGraphics*, vol. 35, no. 5, pp. 1439–1460, 2015.

[19] E. Chee and Z. Wu, “AIRNet: Self-Supervised Affine Registration for 3D Medical Images using Neural Networks,” *ArXiv*, vol. abs/1810.02583, 2018.

[20] B. D. de Vos, F. F. Berendsen, M. A. Viergever, M. Staring, and I. Išgum, “End-to-end unsupervised deformable image registration with a convolutional neural network,” *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pp. 204–212, 2017.

[21] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, “Spatial Transformer Networks,” *Advances in Neural Information Processing Systems*, 2015.

[22] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” *Lecture Notes in Computer Science*, pp. 234–241, 2015.

[23] Y. Gordon, S. Partovi, M. Müller-Eschner, E. Amarteifio, and T. Bäuerle, “Dynamic contrast-enhanced magnetic resonance imaging: fundamentals and application to the evaluation of the peripheral perfusion,” *Cardiovascular Diagnosis and Therapy*, vol. 4, no. 2, pp. 147–164, 2014.