Synthesizing VERDICT maps from standard DWI data using GANs

Eleni Chiou[1,2], Vanya Valindria[1,2], Francesco Giganti[3,4], Shonit Punwani[5], Iasonas Kokkinos[2], and Eleftheria Panagiotaki[1,2]

[1] Centre for Medical Image Computing, UCL, London, UK
[2] Department of Computer Science, UCL, London, UK
[3] Department of Radiology, UCLH NHS Foundation Trust, London, UK
[4] Division of Surgery & Interventional Science, UCL, London, UK
[5] Centre for Medical Imaging, Division of Medicine, UCL, London, UK
`eleni.chiou.17@ucl.ac.uk`

Abstract. VERDICT maps have shown promising results in clinical settings discriminating normal from malignant tissue and identifying specific Gleason grades non-invasively. However, the quantitative estimation of VERDICT maps requires a specific diffusion-weighed imaging (DWI) acquisition. In this study we investigate the feasibility of synthesizing VERDICT maps from standard DWI data from multi-parametric (mp)-MRI by employing conditional generative adversarial networks (GANs). We use data from 67 patients who underwent both standard DWI-MRI and VERDICT MRI and rely on correlation analysis and mean squared error to quantitatively evaluate the quality of the synthetic VERDICT maps. Quantitative results show that the mean values of tumour areas in the synthetic and the real VERDICT maps were strongly correlated while qualitative results indicate that our method can generate realistic VERDICT maps that could supplement mp-MRI assessment for better diagnosis.

Keywords: VERDICT maps, DWI-MRI, Prostate Cancer, Generative Adversarial Networks (GANs)

1 Introduction

Multi-parametric (mp)-MRI, consisting of T2-weighted imaging, diffusion-weighted imaging (DWI) and dynamic contrast enhanced (DCE) imaging, provides non-invasive assessment of the prostate improving the detection and characterization of prostate cancer. However, despite its merits, mp-MRI has some important limitations. In particular, it is characterized by low specificity, provides equivocal findings for around 30% of the patients and correlates moderately with Gleason grade [1].

Towards addressing these limitations, advanced, model-based imaging techniques focus on extracting quantitative metrics that characterize the underlying tissue microstructure in-vivo by modeling the DWI signal [4]. In particular,

VERDICT (Vascular, Extracellular and Restricted Diffusion for Cytometry in Tumours) MRI [20,19], which has been recently in clinical trial [14] to supplement the standard mp-MRI for prostate cancer diagnosis, is a model-based, DWI technique that captures the main microstructural properties of cancerous tissue. VERDICT MRI has shown promising results discriminating normal from malignant tissue [19,8,9,25] and identifying specific Gleason grades in-vivo [15,26].

VERDICT MRI combines an optimized DWI acquisition protocol [18] and a mathematical model to estimate microstructural features such as cell size, density, and vascular volume fraction, all of which change in malignancy (Fig. 1). The general model characterizes water diffusion in three primary compartments allowing the estimation of intracellular (fIC), extracellular-extravascular (fEEs) and vascular (fVASC) volume fractions, and cell radius (R). However, the quantitative estimation of these parameters requires a specific DWI acquisition, different to the one widely used for the standard mp-MRI. Specifically, VERDICT MRI requires multiple and higher b-values with different diffusion times (90, 500, 1500, 2000, 3000 s/mm^2) in different directions to derive accurate estimates of the microstructure parameters. Higher b-values improve tumour conspicuity and characterization but require longer scan times. Thus, methods that can estimate the VERDICT maps using standard DWI data from mp-MRI acquisitions would be beneficial for improving diagnostic accuracy without increasing scan time and patient discomfort.
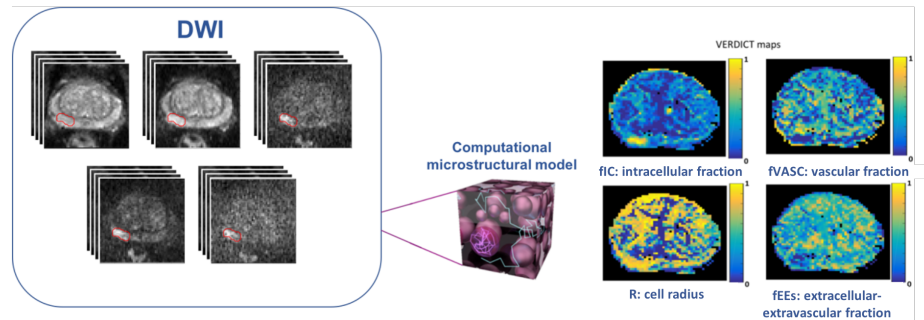

Fig. 1. VERDICT MRI framework. It combines an optimized diffusion-weighted imaging (DWI) protocol and a mathematical model to estimate microstructural features of tumours in-vivo.

Recently, several machine learning methods have been proposed to map an input image from one domain to an output image from a different domain aiming to improve the quality of data coming from routine, low-cost acquisitions or scanners or to eliminate the need for multi-modality scanning. Alexander et al. [2] proposed a general framework for image quality enhancement based on patch regression and demonstrated its effectiveness in super-resolution of brain diffusion tensor images and estimation of parametric maps from limited measurements.

They further extended this approach with probabilistic deep learning formulation and showed that modeling uncertainty allows for better generalization [23,24]. Oktay et al. [17] proposed an image super-resolution approach based on a residual convolutional neural network (CNN) to reconstruct high resolution 3D volumes from a 2D image enabling more accurate analysis of cardiac morphology. In addition, approaches relying on generative adversarial networks (GANs) [11] have been proposed for super-resolution of structural brain MRI [7,22], endomicroscopy [21] and musculoskeletal MRI [6]. Nie et al. [16] proposed a GAN-based approach to generate CT images from MRI images to eliminate multi-modality scanning. Wolterink et al. [29] used CycleGAN [30] to translate MRI to CT in the absence of paired samples. Wang et al. [27] proposed a semi-supervised approach to synthesize ADC images from T2 images to boost performance of clinical tasks in settings where there is limited supervision. Chiou et al. [10] relied on stochastic translation to translate DWI from mp-MRI to raw diffusion VERDICT MRI to improve segmentation performance.

In this work we also rely on a GAN-based approach [12] to generate VERDICT maps from standard DWI data from clinical mp-MRI acquisitions to obtain microstructural information without requiring a specialized acquisition protocol.

2   Methods

2.1   Datasets

This study has been performed with local ethics committee approval as part of the INNOVATE clinical trial [14]. The study involved a cohort of 67 men who provided informed written consent.

All participants underwent a standard mp-MRI with a 3.0-T MRI system (Achieva, Philips Healthcare, NL) as part of their standard clinical care. The DWI data was acquired with diffusion-weighted echo-planar imaging sequences. The DWI sequence was acquired with the following imaging parameters: a repetition time msec/echo time msec, 2753/80; field of view, $220{\times}220$ mm; section thickness, 5 mm; no intersection gap; acquisition matrix, $168{\times}169$mm; b values, $0, 150, 500, 1000, 2000\,\mathrm{s/mm^2}$. The total imaging time for the clinical diffusion-weighted sequences was 5 minutes 16 seconds.

VERDICT MRI data was acquired with pulsed-gradient spin-echo sequence (PGSE) using an optimised imaging protocol for VERDICT prostate characterization with 5 b-values $(90, 500, 1500, 2000, 3000\,\mathrm{s/mm^2})$ in 3 orthogonal directions [18]. Images with $\mathrm{b} = 0\,\mathrm{s/mm^2}$ were also acquired for each b-value. The DWI sequence was acquired with the following imaging parameters: repetition time msec/echo time msec, 2482–3945/50–90; voxel size, $1.25{\times}1.25{\times}5\,\mathrm{mm^3}$; slice thickness, 5 mm; slices, 14; field of view, $220{\times}220\,\mathrm{mm^2}$. The images were reconstructed to a $176{\times}176$ matrix size. The total imaging time was 12 minutes 25 seconds.

VERDICT MRI maps were generated by using the accelerated microstructure imaging via convex optimization, or AMICO framework [3]. The model has three

independent unknown parameters: fIC, R and fEES. fVASC is calculated as fVASC $= 1 - \mathrm{fIC} - \mathrm{fEES}$, and the diffusion and pseudo-diffusion coefficients are fixed to dIC $=$ dEES $= 2 \times 10^{-9} \mathrm{m}^2/\mathrm{s}$, P $= 8 \times 10^{-9} \mathrm{m}^2/\mathrm{s}$. As in [15], in this study we use the fIC, fEES, fVASC and R.

The regions of interest (ROIs) corresponding to Prostate Imaging Reporting and Data System (PI-RADS) [28] score 3, 4 and 5 were contoured on VERDICT MRI using mp-MRI for guidance by an experienced radiologist reporting more than 2000 prostate MR scans per year.

2.2   Proposed Model

Let $x \in \mathbb{R}^{H \times W \times C_{dwi}}$, where H and W are the height and the width of the DWI data and $C_{dwi} = 5$, the number of input channels corresponding the different b-values. Let also $y \in \mathcal{Y}^{H \times W \times C_{maps}}$, where $\mathcal{Y} = [0, 1]$, H and W are the height and the width of the VERDICT maps and $C_{maps} = 4$, the number of the maps. Our goal is to train a model which takes as input 2D DWI slices (5 b-values) and generates the corresponding VERDICT maps (4 maps).
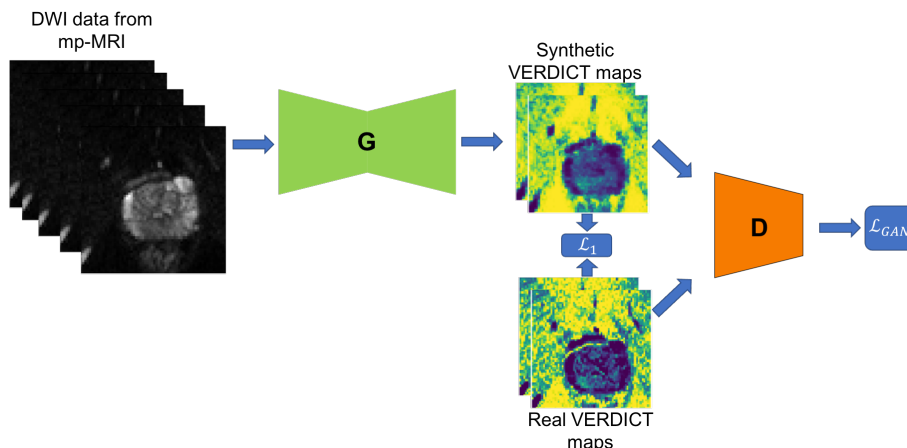

Fig. 2. Schematic representation of the proposed framework for synthesizing VERDICT maps form standard DWI data from mp-MRI acquisitions. The discriminator D is trained to discriminate between real and synthetic VERDICT maps while the generator G maps DWI data to synthetic VERDICT maps $(G : x \rightarrow y)$ that cannot be distinguished from real VERDICT maps by the discriminator D. A GAN-type loss is used to push the distribution of the synthetic maps closer to the ground truth while the L1 loss ensures that the global and local structure of the synthetic maps do not deviate significantly from the real images.

In this work we use pix2pix framework [12], which has shown great success in natural images, to map standard DWI data from mp-MRI acquisitions to

VERDICT maps. As it is illustrated in Figure 2 the framework consists of a generator network G and a discriminator network D. The discriminator D is trained to discriminate between real and synthetic VERDICT maps while the generator G maps DWI data to synthetic VERDICT maps ($G : x \rightarrow y$) that cannot be distinguished from real VERDICT maps by the discriminator D. The adversarial loss, L_{GAN}, can be expressed as

$$\mathcal{L}_{GAN}(G, D) = \mathbb{E}_{x,y}[\log(D(x, y))] + \mathbb{E}_x[\log(1 - D(G(x)))]. \tag{1}$$

A generator trained solely using the adversarial loss function can synthesize realistic-looking maps which however do not preserve the global and local structure and deviate significantly from the real images. To address this problem and generate maps that both fool the discriminator and are close to the real ground truth maps we use a pixel reconstruction loss, i.e., $L1$ distance. The $\mathcal{L}_1$ loss can be expressed as

$$\mathcal{L}_1(G) = \mathbb{E}_{x,y}[\|G(x) - y\|_1]. \tag{2}$$

The final training objective can be written as

$$\min_G \max_D \mathcal{L}_{GAN}(G, D) + \lambda \mathcal{L}_1(G), \tag{3}$$

where λ is the weight controlling the importance of the reconstruction loss.

2.3   Network architecture

The generator is an encoder-decoder convolutional network based on the U-Net architecture. The encoder consists of 6 convolutional layers followed by batch normalization layers, dropout layers and leaky rectified linear activation units (LeakyRELU). The decoder consists of 6 transposed convolutional layers followed by batch normalization layers, dropout layers and RELU. The last transposed convolutional layer is followed by tanh activation. The output of layer i of the encoder is concatenated with the output of the $n - i$ layer of the decoder, where n is the total number of layers, and it is given as input to the next layer of the decoder. The discriminator consists of 3 convolutional layers followed by batch normalization and LeakyRELU. The convolutional layers are 4x4 spatial filters applied with stride 2 and padding 1.

2.4   Implementation details

We implement the framework using Pytorch. We train both the generator and discriminator networks using mini-batch stochastic gradient descent and apply the Adam solver with a mini-batch size of 32, and momentum parameters $\beta_1 = 0.5$, $\beta_2 = 0.999$. We train the networks for 10000 epochs with an initial learning rate at 0.0001 that starts decreasing linearly to 0 after 5000 epochs. We employ dropout as a regularization strategy with dropout rate 50%. We use 60% of the patients for training, 20% for validation and 20% for testing.

Table 1. Mean squared error (MSE) calculated on the entire maps and on the prostate region only for the four maps (fIC, fEES, fVASC, R). The results are given in mean (±std) format.

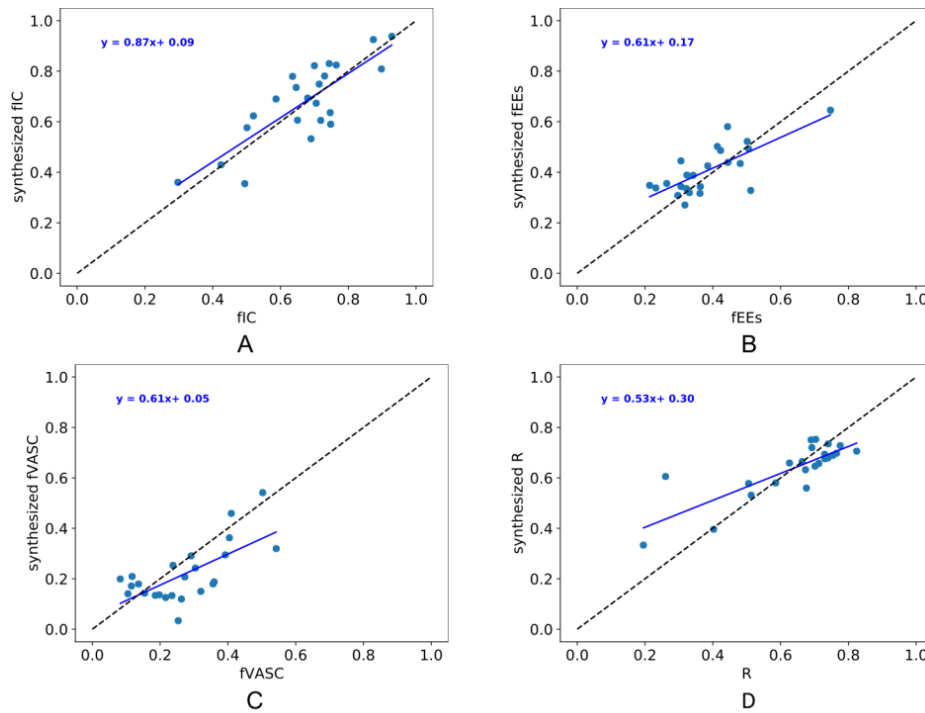|	MSE	MSE (prostate)
fIC	0.18 (0.05)	0.13 (0.03)
fEES	0.12 (0.04)	0.16 (0.04)
fVASC	0.21 (0.05)	0.18 (0.04)
R	0.23 (0.07)	0.19 (0.06)

Fig. 3. A) Mean values of ROIs calculated from the real fIC as a function of the values calculated from the synthetic fIC. B) Mean values of ROIs calculated from the real fEES as a function of the values calculated from the synthetic fEES. C) Mean values of ROIs calculated from the real fVASC as a function of the values calculated from the synthetic fVASC. D) Mean values of ROIs calculated from the real R as a function of the values calculated from the synthetic R.

3   Results

We evaluate the quality of the synthetic maps based on the mean squared error (MSE). Table 1 shows the mean/std of the MSE over 13 test subjects computed on both the entire maps and the prostate region only.
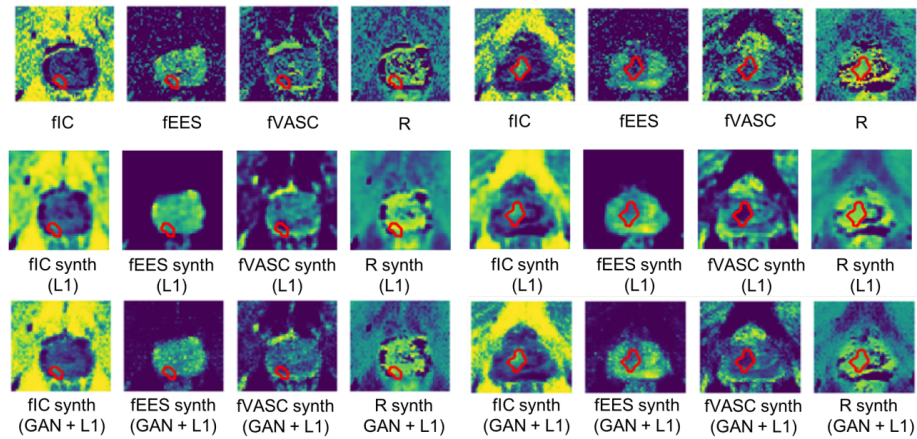
Fig. 4. fIC, fEES, fVASC, R maps and the corresponding synthetic maps for two patients with prostate lesions in the transition zone and the central zone respectively. The first row shows the ground truth maps. The second row shows the synthetic maps obtained using the L1 loss alone while the third row shows the maps obtained adding the L1 and GAN losses together.

We also calculate the mean value of each ROI on the real and the synthetic maps. Then we compute the correlations between the mean values of the ROIs by computing the Pearson's correlation coefficient and perform linear regression to quantify the relationships between the mean values in the ROIs. The relationship between the values calculated from the real and synthetic fIC maps is shown in Figure 3 (A). The values show a linear relationship following the regression line $fIC_{syn} = 0.87fIC_{real} + 0.09$ and the Pearson correlation coefficient is 0.81 (P < 0.05), indicating that there is a strong correlation between the values. Figure 3 (B) shows the relationship between the mean values obtained by the real and the synthetic fEEs maps. The linear relationship is the regression line $fEES_{syn} = 0.61fEES_{real} + 0.17$ and the Pearson correlation coefficient is 0.74 (P < 0.05). The real and the synthetic fVASC values have a liner relationship given by the regression line $fVASC_{syn} = fVASC_{real}0.61 + 0.05$ and the Pearson correlation coefficient is 0.67 (P < 0.05). The real and the synthetic R values exhibit a linear relationship $R_{syn} = 0.53R_{real} + 0.30$ and the Pearson correlation coefficient is 0.82 (P < 0.05).

Figure 4 demonstrates an example of real and synthetic VERDICT parametric maps for two patients with prostate lesions in the transition zone and the central zone respectively. L1 loss alone leads to reasonable but blurry maps; adding the GAN loss gives much sharper results.

4   Discussion

In this study, we investigate a GAN-based approach for image synthesis for prostate cancer characterization. Our results demonstrate that the proposed approach is viable for generating realistic VERDICT maps from standard DWI data from mp-MRI acquisitions.

Table 1 gives the average MSE calculated between the synthetic and real VERDICT maps. As we can see in the table there is a small difference between the MSE computed using the whole image and using only the prostate region. This indicates that our approach is stable among all regions and all maps, especially the most important one (fIC) that has low error.

In addition the ROI measurements of the synthetic and real VERDICT maps are highly correlated. In particular, there is high correlation for fIC maps (0.81 ($P < 0.05$)), which have been shown to be the most important in differentiating specific Gleason grades.

Figure 4 shows that the synthetic maps have realistic appearance and preserve important quantitative information. In alignment with the real maps, the synthetic fIC, fEEs, fVASC and R maps clearly depict the lesions which are characterized by high signal in fIC and R maps and low signal in fEES and fVASC maps.

Our objective is to investigate the feasibility of generating VERDICT maps from the clinically-available DWI data from mp-MRI acquisitions to bring the advantages of VERDICT maps in clinical practice. We demonstrate that GAN-based methods have the potential to generate realistic VERDICT maps that preserve important clinical information without requiring specified DWI acquisition protocols, but only using the widely available DWI data from mp-MRI acquisitions. Obtaining VERDICT maps using standard DWI data would reduce acquisition time and patient discomfort. This would also allow us to use already acquired mp-MRI data to get microstructural information.

Despite the good quality of the synthetic maps there are still some limitations. Specifically, the synthetic images are smoother compared to real ones which means that for small ROIs the quantitative values could be wiped out. Methodological improvements that enforce semantic consistency before and after translation [5,13] could resolve this issue and allow the synthesis of high quality VERDICT maps.

5   Conclusion

In this work we present an approach for synthesizing realistic VERDICT maps from standard DWI data from prostate mp-MRI acquisitions. Our results indicate that the synthetic maps have realistic appearance and preserve important quantitative information. This could allow the exploitation of VERDICT maps for improved prostate cancer diagnosis without increasing acquisition time and patient discomfort.

References

1. Ahmed, H.U., Bosaily, A.E.S., Brown, L.C., Gabe, R., Kaplan, R., Parmar, M.K., Collaco-Moraes, Y., Ward, K., Hindley, R.G., Freeman, A., et al.: Diagnostic accuracy of multi-parametric MRI and TRUS biopsy in prostate cancer (PROMIS): a paired validating confirmatory study. The Lancet (2017)

2. Alexander, D.C., Zikic, D., Ghosh, A., Tanno, R., Wottschel, V., Zhang, J., Kaden, E., Dyrby, T.B., Sotiropoulos, S.N., Zhang, H., et al.: Image quality transfer and applications in diffusion MRI. NeuroImage (2017)

3. Bonet-Carne, E., Johnston, E., Daducci, A., Jacobs, J.G., Freeman, A., Atkinson, D., Hawkes, D.J., Punwani, S., Alexander, D.C., Panagiotaki, E.: VERDICT-AMICO: Ultrafast fitting algorithm for non-invasive prostate microstructure characterization. NMR in Biomedicine (2019)

4. Bourne, R., Panagiotaki, E.: Limitations and prospects for diffusion-weighted MRI of the prostate. Diagnostics (2016)

5. Cai, J., Zhang, Z., Cui, L., Zheng, Y., Yang, L.: Towards cross-modal organ translation and segmentation: A cycle and shape consistent generative adversarial network. MedIA (2019)

6. Chaudhari, A.S., Fang, Z., Kogan, F., Wood, J., Stevens, K.J., Gibbons, E.K., Lee, J.H., Gold, G.E., Hargreaves, B.A.: Super-resolution musculoskeletal MRI using deep learning. MRM (2018)

7. Chen, Y., Shi, F., Christodoulou, A.G., Xie, Y., Zhou, Z., Li, D.: Efficient and accurate MRI super-resolution using a generative adversarial network and 3D multi-level densely connected network. In: MICCAI (2018)

8. Chiou, E., Giganti, F., Bonet-Carne, E., Punwani, S., Kokkinos, I., Panagiotaki, E.: Prostate cancer classification on VERDICT DW-MRI using convolutional neural networks. In: MLMI (2018)

9. Chiou, E., Giganti, F., Punwani, S., Kokkinos, I., Panagiotaki, E.: Automatic classification of benign and malignant prostate lesions: A comparison using VERDICT DW-MRI and ADC maps. In: ISMRM (2019)

10. Chiou, E., Giganti, F., Punwani, S., Kokkinos, I., Panagiotaki, E.: Harnessing uncertainty in domain adaptation for MRI prostate lesion segmentation. In: MICCAI (2020)

11. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014)

12. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)

13. Jiang, J., Hu, Y.C., Tyagi, N., Zhang, P., Rimner, A., Mageras, G.S., Deasy, J.O., Veeraraghavan, H.: Tumor-aware, adversarial domain adaptation from CT to MRI for lung cancer segmentation. In: MICCAI (2018)

14. Johnston, E., Pye, H., Bonet-Carne, E., Panagiotaki, E., Patel, D., Galazi, M., Heavey, S., Carmona, L., Freeman, A., Trevisan, G., et al.: INNOVATE: a prospective cohort study combining serum and urinary biomarkers with novel diffusion-weighted magnetic resonance imaging for the prediction and characterization of prostate cancer. BMC Cancer (2016)

15. Johnston, E.W., Bonet-Carne, E., Ferizi, U., Yvernault, B., Pye, H., Patel, D., Clemente, J., Piga, W., Heavey, S., Sidhu, H.S., et al.: VERDICT MRI for prostate cancer: intracellular volume fraction versus apparent diffusion coefficient. Radiology (2019)

16. Nie, D., Trullo, R., Lian, J., Wang, L., Petitjean, C., Ruan, S., Wang, Q., Shen, D.: Medical image synthesis with deep convolutional adversarial networks. IEEE Transactions on Biomedical Engineering (2018)

17. Oktay, O., Bai, W., Lee, M., Guerrero, R., Kamnitsas, K., Caballero, J., de Marvao, A., Cook, S., O'Regan, D., Rueckert, D.: Multi-input cardiac image super-resolution using convolutional neural networks. In: MICCAI (2016)

18. Panagiotaki, E., Ianus, A., Johnston, E., Chan, R., Atkinson, D., Alexander, D., et al.: Optimised verdict MRI protocol for prostate cancer characterisation. In: ISMRM (2015)

19. Panagiotaki, E., Chan, R.W., Dikaios, N., Ahmed, H.U., O'Callaghan, J., Freeman, A., Atkinson, D., Punwani, S., Hawkes, D.J., Alexander, D.C.: Microstructural characterization of normal and malignant human prostate tissue with vascular, extracellular, and restricted diffusion for cytometry in tumours magnetic resonance imaging. Investigative Radiology (2015)

20. Panagiotaki, E., Walker-Samuel, S., Siow, B., Johnson, S.P., Rajkumar, V., Pedley, R.B., Lythgoe, M.F., Alexander, D.C.: Noninvasive quantification of solid tumor microstructure using VERDICT MRI. Cancer Research (2014)

21. Ravì, D., Szczotka, A.B., Pereira, S.P., Vercauteren, T.: Adversarial training with cycle consistency for unsupervised super-resolution in endomicroscopy. MedIA (2019)

22. Sánchez, I., Vilaplana Besler, V.: Brain MRI super-resolution using generative adversarial networks. In: MIDL (2018)

23. Tanno, R., Worrall, D.E., Ghosh, A., Kaden, E., Sotiropoulos, S.N., Criminisi, A., Alexander, D.C.: Bayesian image quality transfer with CNNs: exploring uncertainty in dMRI super-resolution. In: MICCAI (2017)

24. Tanno, R., Worrall, D.E., Kaden, E., Ghosh, A., Grussu, F., Bizzi, A., Sotiropoulos, S.N., Criminisi, A., Alexander, D.C.: Uncertainty modelling in deep learning for safer neuroimage enhancement: Demonstration in diffusion MRI. NeuroImage (2021)

25. Valindria, V., Palombo, M., Chiou, E., Singh, S., Punwani, S., Panagiotaki, E.: Synthetic Q-space learning with deep regression networks for prostate cancer characterisation with verdict. In: ISBI (2021)

26. Valindria, V., Singh, S., Palombo, M., Chiou, E., Mertzanidou, T., Kanber, B., Punwani, S., Panagiotaki, E.: Non-invasive gleason score classification with VERDICT-MRI. In: ISMRM (2021)

27. Wang, Z., Lin, Y., Cheng, K.T.T., Yang, X.: Semi-supervised mp-MRI data synthesis with StitchLayer and auxiliary distance maximization. MedIA (2020)

28. Weinreb, J.C., Barentsz, J.O., Choyke, P.L., Cornud, F., Haider, M.A., Macura, K.J., Margolis, D., Schnall, M.D., Shtern, F., Tempany, C.M., Thoeny, H.C., Verma, S.: PI-RADS Prostate Imaging – Reporting and Data System: 2015, Version 2. European Urology (2016)

29. Wolterink, J.M., Dinkla, A.M., Savenije, M.H., Seevinck, P.R., van den Berg, C.A., Išgum, I.: Deep MR to CT synthesis using unpaired data. In: SASHIMI (2017)

30. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)