IPEM conferences

Artificial Intelligence in MRI

online

PROGRAMME

	Introduction
09.30 – 10.00	**AI in Radiology: The Story Behind the Data** Ben Glocker, Imperial College London
10.00 – 10.20	**AutoConfidence: per-patient reference-free validation for clinical confidence in deep-learning for radiotherapy** Michael G. Nix, Leeds Teaching Hospitals NHS Trust
10.20 – 10.40	**Fat Suppression in T2-weighted MRI images using Cycle-Consistent Generative Adversarial Networks** Shouxin Wang, University of Manchester
10.40 – 10.55	**Break**
10.55 – 11.25	**AI in CMR image reconstruction** Jennifer Steeden, University College London
11.25 – 11.55	**AI in CMR post-processing** Avan Suinesiaputra, University of Leeds
11.55 - 12.15	**OSL presentation**
12.15 – 13.30	**Break**
13.30 – 14.00	**MONAI - deep learning in healthcare imaging** Benjamin Murray, King's College London
14.00 – 14.30	**Radiomics and MRI** Simon Doran, Institute of Cancer Research
14.30 – 14.50	**QNICE - Clinical deployment of novel quantitative image analysis techniques using Python and containerisation technology.** James Moggridge, National Hospital for Neurology and Neurosurgery
14.50 – 15.05	**Break**
15.05 – 15.25	**Cartilage-MRI radiomics model for prediction of knee replacement surgery in patients with osteoarthritis: data from the osteoarthritis initiative.** E Peake, University of Nottingham
15.25 – 15.45	**Training and evaluating deformable image registration without manual ground truth** A Tattersall, University of Edinburgh
15.45 – 16.15	**Debate** Debate on the role, training required and department support for Clinical Scientists working with AI

AI in Radiology: The Story Behind the Data

Ben Glocker, Imperial College London

In this talk, we will discuss how the language of causality can help to shed new light on two major challenges in image-based predictive modelling: scarcity of high-quality annotated data, and mismatch between the development data and the target environment. We will argue that the underlying causal relationship between images and annotations should be established and considered when making decisions on the machine learning strategies. For illustration, we will look at some worked clinical examples which aim to demonstrate both the benefits and challenges of causal reasoning in medical imaging.

Biography
Ben Glocker is Reader in Machine Learning for Imaging at the Department of Computing at Imperial College London where he co-leads the Biomedical Image Analysis Group. He also leads the HeartFlow-Imperial Research Team and is scientific advisor for Kheiron Medical Technologies and a Visiting Researcher at Microsoft Research Cambridge. He holds a PhD from TU Munich and was a postdoc at Microsoft and a Research Fellow at the University of Cambridge. His research is at the intersection of medical imaging and artificial intelligence aiming to build computational tools for improving image-based detection and diagnosis of disease.

AutoConfidence: per-patient reference-free validation for clinical confidence in deep-learning for radiotherapy

Michael G. Nix, Leeds Teaching Hospitals NHS Trust

Michael G. Nix, PhD,[*,†] David Bird, MSc,[*] Ane Appelt, PhD, Louise Murray, MD, Hazel McCallum, PhD, Bashar Al-Qaisieh, PhD,[*] and Ali Gooya, PhD[*,†]

*michael.nix@nhs.net

[*] Leeds Cancer Centre, Leeds Teaching Hospitals NHS Trust, Leeds, United Kingdom;

[†]Radiotherapy Research Group, Leeds Institute of Medical Research and Biomedical Imaging, University of Leeds, Leeds, United Kingdom;

**Northern Centre for Cancer Care, Newcastle upon Tyne Hospitals NHS Foundation Trust, UK; Centre for Cancer, Newcastle University, UK

Purpose and Objectives:

AI methods are often seen as 'black-boxes' with unknown mechanism and failure-modes. As such, clinical confidence in AI predictions is hard to achieve, especially in high risk workflows, necessitating laborious and error-prone human checking, which negates workflow benefits. AutoConfidence provides a framework for co-training independent validation networks with predictive networks in a range of AI applications affecting high-stakes clinical decision making. These validation networks provide reference-free validation of AI predictions based only on input and prediction, independent of the predictive networks they are trained against. Here, we demonstrate AutoConfidence in the context of deep-learning image-translation (synthetic CT generation) for radiotherapy.

Materials and Methods:

AutoConfidence architecture: Generative Adversarial Networks (GANs) conventionally use a 'Discriminator' as a critic to separate AI-generated examples from real ones, improving the generator performance adversarially. pix2pix is a recent conditional GAN which enables accurate image translation from one domain (modality) to another, via a patch-wise discriminator which critiques parts of the image. AutoConfidence builds on and inverts the cGAN concept, to leverage the potential of the spatially resolved discriminator which can identify 'unconvincing' regions of the generator prediction. Discriminator performance is improved by a) full resolution shallow U-net architecture, b) training against known generator errors, and c) adaptive focal-loss. Herein, generator and discriminator were adversarially co-trained, such that the discriminator becomes an 'expert critic', able to identify regions which are 'unlike real images' on a pixel-resolved level. Discriminator-forgetting was alleviated by continuous retraining on a random sample of predictions from early epochs. Following training, the AutoConfidence discriminator operates on predictions from any source, enabling independent, ground-truth free, patient-specific quality assessment for AI in radiotherapy. Image-synthesis: 32 T2w-SPACE RT position pelvic MRs (2240 slices), from anorectal cancer patients, were deformably registered to planning CT and used to adversarially train a cGAN (fig. 1). The discriminator network was trained to predict Hounsfield unit error in the generator output, relative to reference CT. Discriminator performance was measured in a separate validation cohort, via mean absolute error (MAE) and mean error (ME).

Results:

Reference free prediction of residual errors in high-quality sCT (MAE=35.1 HU), via AutoConfidence was achieved with validation-cohort MAE of 5 (3.8-8.3) % and ME of -1.4 (-4.5 to 0.0) %. Error prediction on low-quality sCT (MAE=45 HU) from early training yielded AutoConfidence MAE and ME of 5.6 (3.7-11.2) % and 1.8 (0.0-4.8) %. Fig 2. shows examples of 'good' and 'poor' sCT slices with errors vs. ground truth and AutoConfidence maps, identifying underestimated bone density (c and d) as well as missing or erroneous bone. AutoConfidence can assess sCT from independent sources, producing reference-free confidence maps based *only* on input MRI and predicted sCT.

Conclusions:

1. AutoConfidence can identify low-confidence examples or regions of AI predictions, independent of ground truth and prediction method, enabling automated patient-specific verification of 'black box' AI predictions for improved clinical confidence. Areas requiring intervention can be highlighted for review, increasing clinical confidence in high-stakes clinical decision making and facilitating efficient automated workflows.

Fat Suppression in T2-weighted MRI images using Cycle-Consistent Generative Adversarial Networks

[1]Wang S, [1]Aznar M, [1]Cobben D, [2]McDaid L, [1]van Herk M, [1]Hanson H, [1]Rowland BC
[1]Faculty of Biology, Medicine & Health, University of Manchester, UK.
[2]The Christie NHS Foundation Trust, Manchester, UK

Background. Fat suppression (FS) is an important technique used for fatty tissue detection and tumour evaluation in magnetic resonance imaging (MRI) applications. However, patients often complain about increased scanning time, scanner noise and motion restrictions because of additional FS images, especially for patients with claustrophobia[3]. Furthermore, because of the heterogeneity of the magnetic field in the MRI scanner, it can be difficult to acquire homogeneous FS images. In this study, in order to reduce the discomfort of patients and get clearer images, we trained Cycle-Consistent Generative Adversarial Networks (CycleGAN) to generate the FS images by retrospectively removing fat signal from standard T2-weighted images [1, 5].

Methods. We collected T2-weighted TSE images from 21 lung cancer patients (7 males and 14 females) totalling 959 slices with fat, and 956 FS slices. Different CycleGAN networks were trained varying three parameters: the batch size, the number of epochs and the number of discriminator layers. We compared the histograms of the generated images with real FS images to determine the network with the best performance. Four experienced observers evaluated 20 blinded pairs of real and generated images (scoring from 1=strongly prefer generated images to 5=strongly prefer real images).

Results. The network with the best performance had batch size of 4, 400 epochs and the standard discriminator (70×70 PatchGAN). The mean scores of 20 generated images given by each observer ranged from 2.80 to 3.60. The overall mean score was 3.20, indicating only a very slight preference for real images.

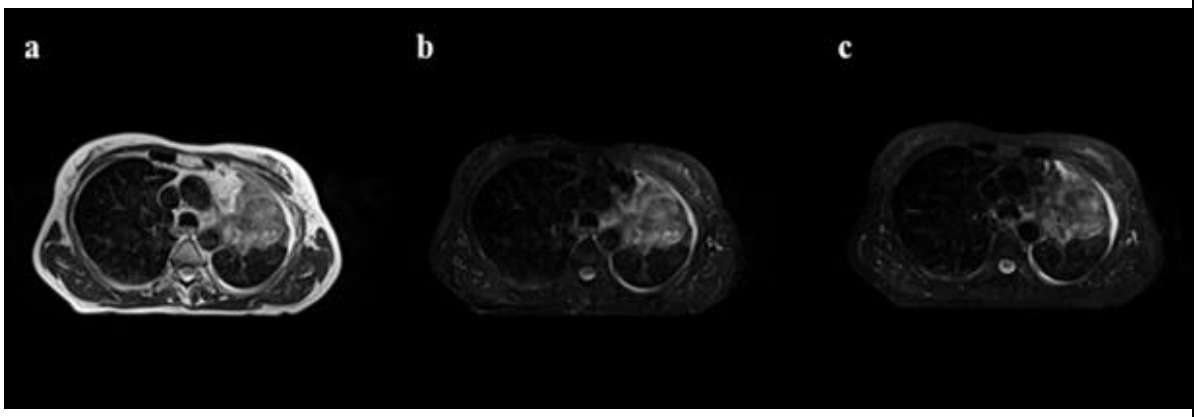

An original T2-weighted MRI image **(a)**; the generated fat-suppressed image **(b)**; the real fat-suppressed image **(c)**

Discussion. Most AI image translation techniques require paired images, which are hard to acquire accurately because of patient motion. The power of CycleGAN lies in its ability to use unpaired images, avoiding the blurring which results from imperfect image pairs. CycleGAN has already been widely employed in MR-to-CT translation[2, 4], here we demonstrate it can also be applied in MR-to-MR translation problems as well.

Conclusion. CycleGAN has the potential to generate fat-suppressed T2-weighted images from T2-weighted images for lung MRI.

Key references.

2. Isola, P., Zhu, J.-Y., Zhou, T. & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017. pp. 1125-1134.
3. Lei, Y., Harms, J., Wang, T., Liu, Y., Shu, H. K., Jani, A. B., Curran, W. J., Mao, H., Liu, T. & Yang, X. (2019). 'MRI-only based synthetic CT generation using dense cycle consistent generative adversarial networks', Medical physics, 46(8), pp. 3565-3581.
4. Oztek, M. A., Brunnquell, C. L., Hoff, M. N., Boulter, D. J., Mossa-Basha, M., Beauchamp, L. H., Haynor, D. L. & Nguyen, X. V. (2020). 'Practical considerations for radiologists in implementing a patient-friendly MRI experience', Topics in Magnetic Resonance Imaging, 29(4), pp. 181-186.
5. Yi, X., Walia, E. & Babyn, P. (2019). 'Generative adversarial network in medical imaging: A review', Medical image analysis, 58, p. 101552.
6. Zhu, J.-Y., Park, T., Isola, P. & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision, 2017. pp. 2223-2232.

AI in CMR image reconstruction

Jennifer Steeden, University College London

AI in CMR post-processing

Avan Suinesiaputra, University of Leeds

OSL presentation

“Real-world” Radiomics and MRI

Simon Doran, Institute of Cancer Research

Texture analysis has a long history in MRI. However, it remained a somewhat niche topic until a series of key papers in the mid 2010s led to a paradigm-shift in our methods of data analysis. The realisation that statistical modeling methods originally developed in the -omics fields could be applied to large numbers of automatically-generated imaging features led to the field of *radiomics*. A number of encouraging initial results have been followed by a period of technical consolidation, most recently culminating in work by the Image Biomarker Standardisation Initiative.

MRI poses numerous challenges to such efforts in standardisation, challenges that the IPEM community is well placed to address. In contrast to the case of x-ray CT (where the Hounsfield number of a tissue is a standardised quantity that is relatively easy to measure with a known degree of accuracy), quantitative images in MRI are already difficult to standardise, even before the additional complexities of radiomics feature generation are introduced. Furthermore, the range of acquisition options available is vast. Differences in spatial resolution, slice thickness, slice profiles, relaxation parameters, sequence timings, k-space trajectories, multi-coil combinations, reconstruction methodologies and, now, AI-based acquisition acceleration strategies all have the potential to change the measured feature values.

Amid this array of potential pitfalls, how do we make progress in MRI radiomics? This talk will explore one approach, which we have termed "real-world" radiomics and I will present the results of an MRI radiomics study in breast cancer that exemplifies some of these issues.

The speaker

Dr Simon Doran is a senior staff scientist in the Division of Radiotherapy and Imaging at The Institute of Cancer Research, London, where he leads the effort in advanced imaging informatics. Dr Doran gained his first degree in physics and theoretical physics at the University of Cambridge and obtained a PhD in quantitative magnetic resonance imaging at the laboratory of Professor Laurie Hall. After a hugely enjoyable period of postdoctoral research in ultra-rapid MR imaging with the team of Prof Michel Décorps (INSERM U438) in Grenoble, he lectured for 11 years in the Department of Physics at the University of Surrey before moving to the ICR in 2006.

Simon manages the UK-wide Repository Unit of CRUK's National Cancer Imaging Translational Accelerator (NCITA) and, locally, he works on the development of the ICR’s Research PACS platform for archiving and visualising image data. Both of these efforts are centred around the XNAT image archive platform, which is rapidly gaining traction as a leading open-source solution for academic imaging data management. The philosophy behind these developments is described in Doran *et al.* Radiographics 32, 2135–2150 (2012). Ongoing work in the group is customising the XNAT to meet the needs of the cancer imaging community, with a particular focus on the areas of online visualisation of image data, radiolog-radiotherapy integration, clinical trials, multimodality imaging and preclinical imaging.

Simon's other interests include 3D optical computed tomography (CT) microscopy and high-resolution MRI in the imaging of cancer tissues. His research explores use of 3D radiation dosimetry using optical CT in combination with radiochromic dosimeters, particularly PRESAGE®. This work concentrates on high-resolution microimaging of doses delivered in synchrotron microbeam radiotherapy. Dr Doran is a member of the EU’s SYRA3 COST Action TD1205 – a multidisciplinary network to develop synchrotron radiotherapy and radiosurgery techniques to treat brain tumours and other diseases of the central nervous system.

MONAI - deep learning in healthcare imaging

Benjamin Murray, King's College London

MONAI is a free available community-supported framework for deep learning in healthcare imaging, created by the open-source foundation Project MONAI. Its goal is to provide the tools for training and deploying novel AI solutions to biomedical imaging problems, and to facilitate the adoption of deep learning in clinical settings. MONAI helps provide solutions to the complex data handling requirements that are particular to biomedical imaging, and it provides a growing library of predefined networks and reference implementations of published research. MONAI can be used to train and deploy neural networks and scale from commodity hardware to dedicated clusters. This talk will give both an overview of MONAI and some MR-focused clinical applications of deep learning and how MONAI helps (or will help in upcoming releases) enable these applications.

Ben Murray is a Senior Research Associate at Biomedical Engineering & Imaging Sciences, King's College London. Ben has returned to academia after spending many years in industry, working as a software engineer in games, medical imaging, actuarial modelling, and bioinformatics, amongst other fields. Ben has been involved in the design of MONAI from its inception, with a focus on how MONAI interacts with the rest of the biomedical imaging deep learning ecosystem. He also sits on a number of the MONAI working groups. Ben's research is focused on making deep learning more robust in a clinical setting.

QNICE - Clinical deployment of novel quantitative image analysis techniques using Python and containerisation technology.

[1]Moggridge J, [1,2]Vos SB, [2]Prados F, [1,2]Goodkin O, [2]Pemberton H, [3]Matthews D, [3]Schmidt P, [1,2]Barkhof F, [1]Yousry T, [1]Thornton J

[1]Neuroradiology, National Hospital for Neurology and Neurosurgery, UK.

[2]Centre for Medical Image Computing, University College London, UK

[3]Research IT Services, University College London, UK

Background. The [1]Quantitative Neuroradiology Initiative (QNI) was established as a model framework to facilitate the transition of image quantification software from the research domain into hospital neuroradiology workflow. Quantitative tools are developed under this framework and delivered as containerised packages using Docker. A proof of concept software, the QNI Central Engine (QNICE) was developed to bring these containerised packages into clinical deployment.

Methods. The QNICE was written using Python using the Flask framework to provide a RESTful API and user interface with Celery and RabbitMQ handling asynchronous task management. The engine consists of an interface and a management component. The interface handles incoming DICOM series, series conversion into nifti format and the generation of an instruction file containing the additional the metadata required by the quantitative tool e.g. patient date of birth. The generated files are served to the management component which: runs the containerised analysis module; monitors the progress of the analysis; and notifies the interface once analysis is complete and output data is available (either as a pdf report or annotated image series). The interface then feeds the report or series back into radiological workflow by sending it to PACS and making it available on a web page served by the interface.

The QNICE was used to perform pre-clinical deployment evaluation on an advanced machine-learning based tool for measuring regional brain volume in suspected dementia patients. Fifteen patient MRI studies were sent to the engine for analysis with performance and workflow monitored by the development team.

Results. Out of the 15 studies sent to the QNICE, 12 were processed successfully with 3 failing at the analysis stage. Performance in those completing analysis was as expected but the time taken for an individual analysis of approximately 18 hours raised concerns regarding timely delivery of reports for radiological reporting.

Conclusion. The QNICE successfully ran the analysis tool but testing revealed two potential issues: analysis failed in 3 cases with no feedback to the end user; and the time taken for completion of analysis cast doubt as to whether the tool could be used without compromising reporting times. For the former of these issues, preliminary investigations suggest that image quality was responsible - this has been addressed via inclusion of an image quality control step in the QNICE workflow with a report being generated regardless of the outcome to provide feedback to the clinicians. Regarding the latter, speed of processing is an inherent limitation of the analysis tool, and an important consideration when introducing any novel analysis method that could disrupt radiological workflow. The introduction of parallelisation in the QNICE management component has allowed multiple studies to be processed simultaneously allowing production of the analysis report prior to reporting without incurring significant delays.

Key References.

[1]The Quantitative Neuroradiology Initiative Framework: Application to Dementia.

Goodkin O, Pemberton H, Vos SB, Prados F, Sudre CH, Moggridge J, Cardoso MJ, Ourselin S, Bisdas S, White M, Yousry T, Thornton J, Barkhof F. Br J Radiology. Sep 2019

Cartilage-MRI radiomics model for prediction of knee replacement surgery in patients with osteoarthritis: data from the osteoarthritis initiative.

[1,2]Peake E, [1,2]Dorothee A
[1]NIHR Nottingham Biomedical Research Centre, QMC, University of Nottingham, UK
[2]Division of Clinical Neuroscience, University of Nottingham, UK.

Background. Timely identification of individuals likely to need joint replacement surgery may support personalized patient care, enabling tailored management based on individual risk assessment. In patients with early OA at high risk of needing joint replacement, additional emphasis on proactive preventative treatments may prolong the good health of the knee delaying the need for surgery (1).

Manual segmentation of cartilage tissue from MRI has shown sufficient sensitivity to detect changes, including early progression (2), however it is time consuming and susceptible to inter-rater variability (3). Cartilage segmentation using deep learning has shown excellent performance (4,5), with the potential for rapid automated segmentation of large datasets. In this study, we propose a radiomics driven model using standardized features extracted from automatically segmented knee cartilage to predict joint replacement subjects with symptomatic knee osteoarthritis.

Methods. In this retrospective study, data from the Osteoarthritis Initiative for subjects with baseline osteoarthritis between 2006-2015 were reviewed. Cartilage segmentations were automatically generated using a U-net ensemble for 1433 subjects with three-dimensional double echo steady state MRI. Linear discriminant analysis was used to predict if a joint with osteoarthritis would require replacement surgery within 5-years based on radiomic features and clinical information. Cox regression was used to evaluate the prediction model.

Results. Our radiomics signature for cartilage had a multivariate cox proportional hazard of 5.07 (CI 4.71 – 5.44, p<0.001) using imaging only and 7.34 (CI 6.90 – 7.79, p<0.001) using imaging and clinical information. Imaging features are most predictive for early OA based on radiographic stratification of disease, with univariate cox of 15.09 (13.52 – 16.67) for mild OA and 1.96 (1.47 -2.44) for severe OA, respectively.

Discussion. This white-box model provides individualized predictions based on cartilage components using imaging features from the Image Biomarker Standardisation Initiative. This study leverages the power of deep learning for fully automated segmentation and demonstrates how radiomics can predic clinically meaningful ‘hard’ outcomes for patients.

Conclusion. Radiomic markers from cartilage are useful for predicting incidence of knee replacement surgery within 5-years of magnetic resonance imaging. The greatest benefit is seen in early OA where cartilage damage is less extensive and the prediction of TKR has the most potential to moderate disease progression.

Key references.

1. Bowden JL, Hunter DJ. Therapy: Are you managing osteoarthritis appropriately? Vol. 13, Nature Reviews Rheumatology. Nature Publishing Group; 2017. p. 703–4.
2. Cicuttini F, Eckstein F,Peterfy C, Raynauld JP, Waterton JC. Magnetic resonance imaging (MRI) of articular cartilage in knee osteoarthritis (OA): morphological assessment. Osteoarthr Cartil. 2006;14(SUPPL. 1):46–75.
3. Gatti AA. NEURALSEG: state-of-the-art cartilage segmentation using deep learning – analyses of data from the osteoarthritis initiative. Osteoarthr Cartil. 2018 Apr; 26: S47–8.
4. Arthofer C, Auer DP, Chevasson R, Peake EJ, Pszczolkowski S. Ensemble learning for robust knee cartilage segmentation: data from the osteoarthritis initiative. bioRxiv. 2020 Sep 1
5. Klein S, Nieminen MT, Panfilov E, Saarakkala S, Tiulpin A. Improving Robustness of Deep Learning Based Knee MRI Segmentation: Mixup and Adversarial Domain Adaptation. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW). IEEE; 2019. p. 450–9.

Training and evaluating deformable image registration without manual ground truth

[1] Tattersall AG, [1]Semple SIK, [1]Kershaw LE, [2]Goatman KA,

[1] Edinburgh Imaging, University of Edinburgh, Edinburgh, UK.

[2] Canon Medical Research Europe, Edinburgh, UK.

Background: Artificial Intelligence algorithms based on convolutional neural networks (CNNs) have advanced state-of-the-art performance in many areas of image analysis. However, such algorithms typically require large numbers of annotated images to adequately train and evaluate them. We use a CNN to improve registration of dynamic contrast enhanced (DCE) magnetic resonance (MR) image sequences of the abdomen, which are affected by large intensity differences and motion between frames. This type of registration is difficult due to intensity change being mistaken as motion and can create unrealistic images [5]. An unsupervised registration algorithm, VoxelMorph [2], does not require ground truth annotation. However, the problem of how to tune hyperparameters and evaluate the algorithm remains; we address this problem in this small pilot study.

Method: In addition to VoxelMorph, two traditional registration packages are included for comparison: Elastix [3,4] and SyN [1]. Abdominal DCE-MR sequences from 45 subjects were included, each having 150 frames. For all experiments, the first image in the DCE acquisition was treated as the reference image that all other frames were registered to. We measured absolute registration error using 42 corresponding manually placed landmarks on the spine, uterus and bladder. The sum of squared differences (SSD) and the normalized mutual information (NMI) between registered images were also recorded. Additionally, known deformations were applied to the first image in the series to create a synthetic dataset, which was registered with the three algorithms and evaluated by calculating the mean squared error of the predicted and known displacement fields (DFMSE).

Results: Table 1 lists the registration accuracy measurements. Despite large qualitative differences, quantitative differences in both SSD and NMI were minimal. The manually placed landmarks did not reflect the visual improvement in alignment, but rather indicated worsening of the alignment. In contrast the results from the synthetic deformations follow the qualitative assessment more closely.

Discussion: These experiments demonstrate the challenge of evaluating non-rigid registration. Manual landmarks should be the gold standard metric, but they are extremely time-consuming and in the case of abdominal images without sharp, distinct landmarks, insufficiently accurate to be useful. Indirect measures of accuracy such as SSD and NMI can easily give a false impression of accuracy, particularly when they are the metric optimised by the registration algorithm. VoxelMorph for example uses the SSD in its loss function, however there are many poorly regularised ways of reducing the SSD without necessarily improving alignment (as shown in these experiments). The most promising technique was the synthetic deformation, which is known and can be controlled to be as physically plausible as required.

Type of Registration	Landmark RMSD (mm)	SSD (x106)	NMI	DFMSE
None	1.343	162+/-0.17	0.345+/-1.3e-6	-
VoxelMorph	3.067	44+/-0.029	0.616+/-9.1e-7	3.132
SyN	2.022	56+/-0.022	0.614+/-8.7e-7	2.357
Elastix	2.096	62+/-0.027	0.614+/-8.9e-7	2.691

Table 1 – Registration results for RMSD between Landmarks, SSD (± sd), NMI (± sd), and DFMSE

In future work, we will improve the augmentations to create realistic motion and include a change in intensity in the synthetic images to simulate a DCE-MR series.

References

[1] B. B. Avants et al. Med. Image Anal. 12, 26-41 (2008). [2] G. Balakrishnan et al. IEEE Trans. Med. Imaging. 38, 1788-1800 (2019). [3] J.M. Guyader et al. J. Magn. Reson. Imaging. 42, 315-330 (2015). [4] S. Klein et al. IEEE Trans. Med. Imaging. 29, 196-205 (2010). [5] V. Hamy et al., Med. Image Anal. 18, 301–313 (2014).