

1 **Deep learning diffusion fingerprinting to detect brain**
2 **tumour response to chemotherapy**

3 Tom A. Roberts¹, Ben Hipwell¹, Giulia Agliardi¹, Angela d'Esposito¹, Valerie Taylor¹, Mark F.
4 Lythgoe¹, Simon Walker-Samuel¹

5 ¹Centre for Advanced Biomedical Imaging, University College London, UK.
6

7

8

9

10 **Corresponding Author:**

11 Simon Walker-Samuel

12 Centre for Advanced Biomedical Imaging

13 Paul O'Gorman Building,

14 72 Huntley Street,

15 London,

16 WC1E 6DD

17 E: simon.walkersamuel@ucl.ac.uk, T: 0044 20 7679 6329

18

19

20 **Competing Interests Statement:** The authors declare no competing interests.

21

22 **Body text word count:** 2,831

1 **ABSTRACT**

2 Artificial neural networks are being widely implemented for a range of different biomedical
3 imaging applications. Convolutional neural networks are by far the most popular type of deep
4 learning architecture, but often require very large datasets for robust training and evaluation.
5 We introduce deep learning diffusion fingerprinting (DLDF), which we have used to classify
6 diffusion-weighted magnetic resonance imaging voxels in a mouse model of glioblastoma
7 (GL261 cell line), both prior to and in response to Temozolomide (TMZ) chemotherapy. We
8 show that, even with limited training, DLDF can automatically segment brain tumours from
9 normal brain, can automatically distinguish between young and older (after 9 days of growth)
10 tumours and that DLDF can detect whether or not a tumour has been treated with
11 chemotherapy. Our results also suggest that DLDF can detect localised changes in the
12 underlying tumour microstructure, which are not evident using conventional measurements
13 of the apparent diffusion coefficient (ADC). Tissue category maps generated by DLDF
14 showed regions containing a mixture of normal brain and tumour cells, and in some cases
15 evidence of tumour invasion across the corpus callosum, which were broadly consistent with
16 histology. In conclusion, DLDF shows the potential for applying deep learning on a pixel-wise
17 level, which reduces the need for vast training datasets and could easily be applied to other
18 multi-dimensional imaging acquisitions.

19

20 **Keywords:** imaging, MRI, brain tumour, treatment response, Temozolomide, deep learning,
21 machine learning, classification

22 **Abbreviations:** ANN, artificial neural network; CT, x-ray computed tomography; PET,
23 positron emission tomography; CNN, convolutional neural network; HARDI, high angular
24 resolution diffusion weighted imaging; NODDI, neurite orientation dispersion and density
25 imaging; VERDICT, vascular, extracellular and restricted diffusion for cytometry in tumours;
26 DLDF, deep learning with diffusion fingerprinting; TMZ, Temozolomide; PFA,
27 paraformaldehyde; H&E, hematoxylin and eosin; GFAP, glial fibrillary acidic protein;

1 **1. INTRODUCTION**

2 Research into the use of deep learning with artificial neural networks (ANN) is being widely
3 undertaken. A key application is in radiology, particularly in the classification and
4 segmentation of biomedical images (1-4). ANNs have been trained and evaluated using x-
5 ray computed tomography (CT) (5-7) and positron emission tomography (PET) imaging (8-
6 10), and several magnetic resonance imaging (MRI) data types including structural T1- and
7 T2-weighted images (11), perfusion images (12), MR spectroscopy (13) and diffusion MRI
8 (14).

9 In the last five years, convolutional neural networks (CNNs) have become by far the most
10 popular choice of deep learning architecture (15). CNNs allow subtle and abstract features to
11 be characterised, but often require very large, labelled data sets (16). As an alternative
12 approach, we propose using data-rich diffusion-weighted MRI acquisitions to define a
13 fingerprint in each voxel, which can easily be classified according to tissue type and used to
14 train a deep neural network. This approach aims to produce accurate diagnostic information
15 with a relatively small number of subjects, as every voxel is treated as an independent data
16 point to train the ANN, rather than using whole images (or patches from within images) as
17 used in CNN architectures.

18 Contrast in diffusion weighted MRI is dependent on the random motion of water molecules
19 and the structure of the underlying tissue microenvironment. Various multi-direction and
20 multi-b-value diffusion MRI protocols have been proposed which are better at resolving
21 complicated tissue microstructure compared to basic diffusion tensor imaging (DTI). HARDI
22 (high angular resolution diffusion weighted imaging) and NODDI (neurite orientation
23 dispersion and density imaging) use multiple diffusion directions (17,18) to improve
24 measurement of neuronal fibre direction, whilst VERDICT (vascular, extracellular and
25 restricted diffusion for cytometry in tumours) protocols employ multiple directions and
26 multiple sets of diffusion weighted scans (19,20) to make estimates of cancer cell
27 microstructure and density. These types of diffusion MRI protocols are particularly amenable

1 to pixel-wise deep learning because large amounts of data are acquired in individual
2 subjects.

3 A key advantage of deep learning over techniques which require mathematical models to fit
4 the acquired diffusion signals, is that no assumptions are made about the underlying tissue
5 microstructure. This is especially important when imaging pathology where, often, the lesion
6 greatly alters or interferes with the tissue microstructure in the affected organ. Since most
7 mathematical models are optimised to characterise one type of tissue, regions of pathology
8 can often be poorly characterised because they simply correspond to regions where the
9 model fitting is less robust. For example, NODDI is not designed to characterise brain
10 tumour microstructure and VERDICT is not designed to characterise normal brain tissue.
11 Characterising microstructure in transitional regions containing both neuronal and cancer
12 cells is even more challenging.

13 This therefore underpins our motivation for developing deep learning with diffusion
14 fingerprinting (DLDF). In this study, we have evaluated DLDF in a set of diffusion MRI data
15 from a mouse glioma model (GL261), both prior to and following treatment with
16 Temozolomide (TMZ) chemotherapy. We investigate the ability of DLDF to distinguish brain,
17 background, control and treated tumour pixels, independent of pixel location, and without
18 applying mathematical models of tissue microstructure.

1 **2. METHODS**

2 **2.1 Ethics statement**

3 All in vivo experiments were performed in accordance with the UK Home Office Animals
4 Scientific Procedures Act, 1986 and United Kingdom Coordinating Committee on Cancer
5 Research (UKCCCR) guidelines, and with the approval of the University College London
6 Animal Ethics Committee.

7 **2.2 Mouse Glioma Model and Therapy**

8 Female, C57BL/6 nude mice (6-9 weeks old) were sourced from Harlan (UK) and were
9 housed in groups of 5 in isolated ventilated cages (IVCs) containing environmental
10 enrichment. Mice had access to food and water ad libitum.

11 Mice were injected in the right striatum with 2×10^4 GL261 mouse glioma cells, and randomly
12 assigned to control and treatment (Temozolomide (TMZ)) groups ($n = 12$ per group).

13 Following 13 days of tumour growth (day 0), to an average tumour volume of $8.1 \pm 0.8 \text{ mm}^3$,
14 three doses of TMZ (130mg/kg) or vehicle were administered by gavage, on consecutive
15 days.

16 **2.3 MRI setup**

17 MRI measurements were performed on a 9.4T horizontal bore scanner (Agilent
18 Technologies, Santa Clara, USA) with a 205/120/HD 600mT/m gradient insert. RF
19 transmission was performed with a 72mm inner diameter volume coil and a 2-channel
20 receiver coil (RAPID biomed, Ripmar, Germany). Mice were anaesthetised with 1.5-2.0%
21 isoflurane in 2 l/min oxygen and positioned prone in a cradle for imaging. The head was
22 positioned within an MR-compatible head-holder and secured with plastic ear bars to
23 minimise motion artefacts. Body temperature was maintained at physiological temperature
24 using a hot water system and monitored using a rectal probe (SA Instruments). Respiration
25 was monitored using a neonatal apnoea pad taped to the abdomen of the animal. For
26 gadolinium-based scans, an intraperitoneal infusion line was inserted into the mouse.

1 2.4 Diffusion MRI

2 Gliomas were localised using a structural spin-echo sequence. For generating the diffusion
 3 fingerprints, diffusion-weighted images were acquired in a coronal orientation using a 3-shot
 4 spin-echo echo planar imaging (EPI) sequence, with the following parameters: TR = 3s, TE
 5 = min, DM = 64², FOV = 20mm², shots = 3, slice thickness = 0.5mm, slices = 5, averages =
 6 2. In total, 46 diffusion weightings in 3 directions were acquired in addition to a 42 direction
 7 DTI acquisition ($b = 1000 \text{ s/mm}^2$). Specific gradient combinations are detailed in Table 1. TE
 8 was minimised for all scans to maximise signal-to-noise. To correct for signal changes
 9 caused by this variation in TE, an accompanying $b = 0 \text{ s/mm}^2$ (B_0) image was acquired for
 10 every combination of diffusion gradients. Total acquisition time for all DWIs was 70 minutes.
 11 Following diffusion imaging, mice were injected with gadolinium-DTPA (Magnevist). After 10
 12 minutes – to allow for the contrast agent to circulate – slice-matched T1-weighted spin echo
 13 EPI images were acquired. Tumour regions of interest (ROI) were drawn based on these
 14 images and applied to the diffusion data.

15

δ (ms)	Δ (ms)	G (G/cm)	b-value (s/mm ²)
3	10 / 20 / 30 / 40	3.6	8 / 16 / 24 / 33
3	10 / 20 / 30 / 40	7.2	30 / 63 / 97 / 130
3	10 / 20 / 30 / 40	10.8	68 / 143 / 218 / 293
3	10 / 20 / 30 / 40	14.4	120 / 254 / 387 / 521
3	10 / 20 / 30 / 40	18.0	188 / 397 / 606 / 814
3	10 / 20 / 30 / 40	21.6	270 / 571 / 872 / 1173
3	10 / 20 / 30 / 40	25.2	368 / 778 / 1187 / 1596
3	10 / 20 / 30 / 40	28.8	481 / 1016 / 1550 / 2085
3	10 / 20 / 30 / 40	32.4	609 / 1285 / 1962 / 2639
3	10 / 20 / 30 / 40	36.0	752 / 1587 / 2422 / 3257
10	30 / 40	4.0	305 / 420
10	30 / 40	8.0	1221 / 1680
10	30 / 40	12.0	2749 / 3780

16

17 **Table 1: Diffusion gradient combinations used in mouse brain imaging for creating diffusion**
 18 **fingerprints.**

1 **2.5 Histology**

2 To compare the DLDF category maps with a ground truth, histology was performed on the
3 mouse brains. After the final MRI scan, mice were terminally anaesthetised via
4 intraperitoneal injection of 100 mg/kg sodium pentobarbital (Pentoject, Animalcare, York,
5 UK) and then transcardially perfusion-fixed with 10 ml heparinised saline followed by 10 ml
6 4% paraformaldehyde (PFA). For histology, mouse brains were extracted, sliced in a coronal
7 orientation and stained with hematoxylin and eosin (H&E) and glial fibrillary acidic protein
8 (GFAP).

9 **2.6 Labelling and Diffusion Fingerprint Preparation**

10 Four tissue categories were manually identified on contrast-enhanced images: 1)
11 background (corresponding to any pixel that was not part of any other region), 2) normal-
12 appearing brain, 3) untreated tumour (regions showing contrast-enhancement in all animals
13 at day 0 and only control animals at day 9) and 4) treated tumour at day 9 (Figure 1a).
14 Diffusion fingerprints consisting of 232 ordered data points, from individual pixels (Figure 2)
15 were produced for each pixel, and were normalised to an acquisition with no diffusion
16 encoding.

17 **2.7 Deep Learning of Diffusion Fingerprints**

18 We created a deep neural network with 5 hidden layers in Keras (a model-level library,
19 Python 3.5.2) using TensorFlow as the numerical backend (21). Each hidden layer was
20 regularized with 20% dropout (to avoid overfitting), and used rectified linear unit (ReLU)
21 activation and normally-distributed initialisation. The number of nodes for each layer was
22 100, 150, 200, 512 and 1024, respectively. The output layer contained 6 nodes
23 corresponding to the categories defined above. Diffusion fingerprints from individual pixels
24 were passed to the network input layer.

1 **2.8 Training**

2 50% of the data were randomly assigned to a training set, which were further refined to
3 include equal numbers of fingerprints from each category (a total of 6,608 fingerprints).
4 Training of the network was performed on a Nvidia Titan-X GPU, using stochastic gradient-
5 descent optimisation (learning rate, 0.01; decay, 1e-6; momentum, 0.9), with mean squared
6 error loss function. Training was performed for 50,000 epochs, with 10,000 steps in each.

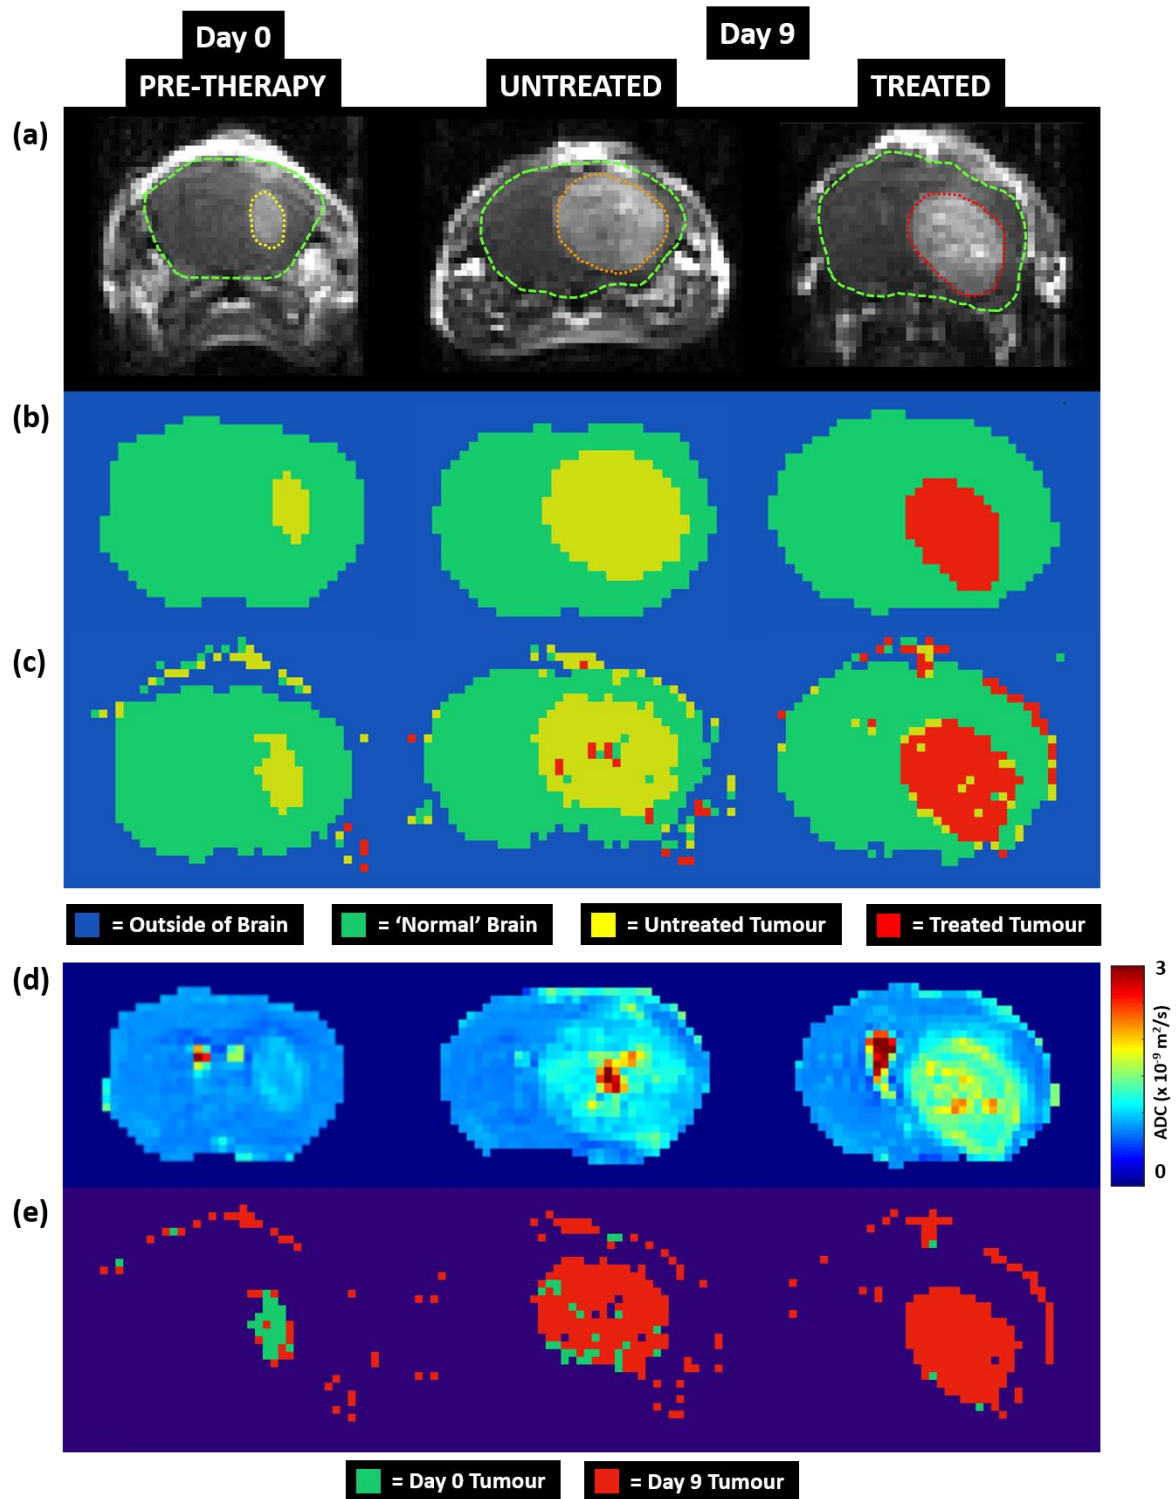
7 **2.9 Evaluation**

8 The evaluation data set contained 5 animals per group. Each diffusion fingerprint was
9 categorised by the deep learning network and reconstructed as category maps. Dropout
10 regularisation was used to estimate classification probability and variance via Monte Carlo
11 sampling (22).

12 **3. RESULTS**

13 **3.1 Brain Region Classifiers and Diffusion Fingerprint Generation**

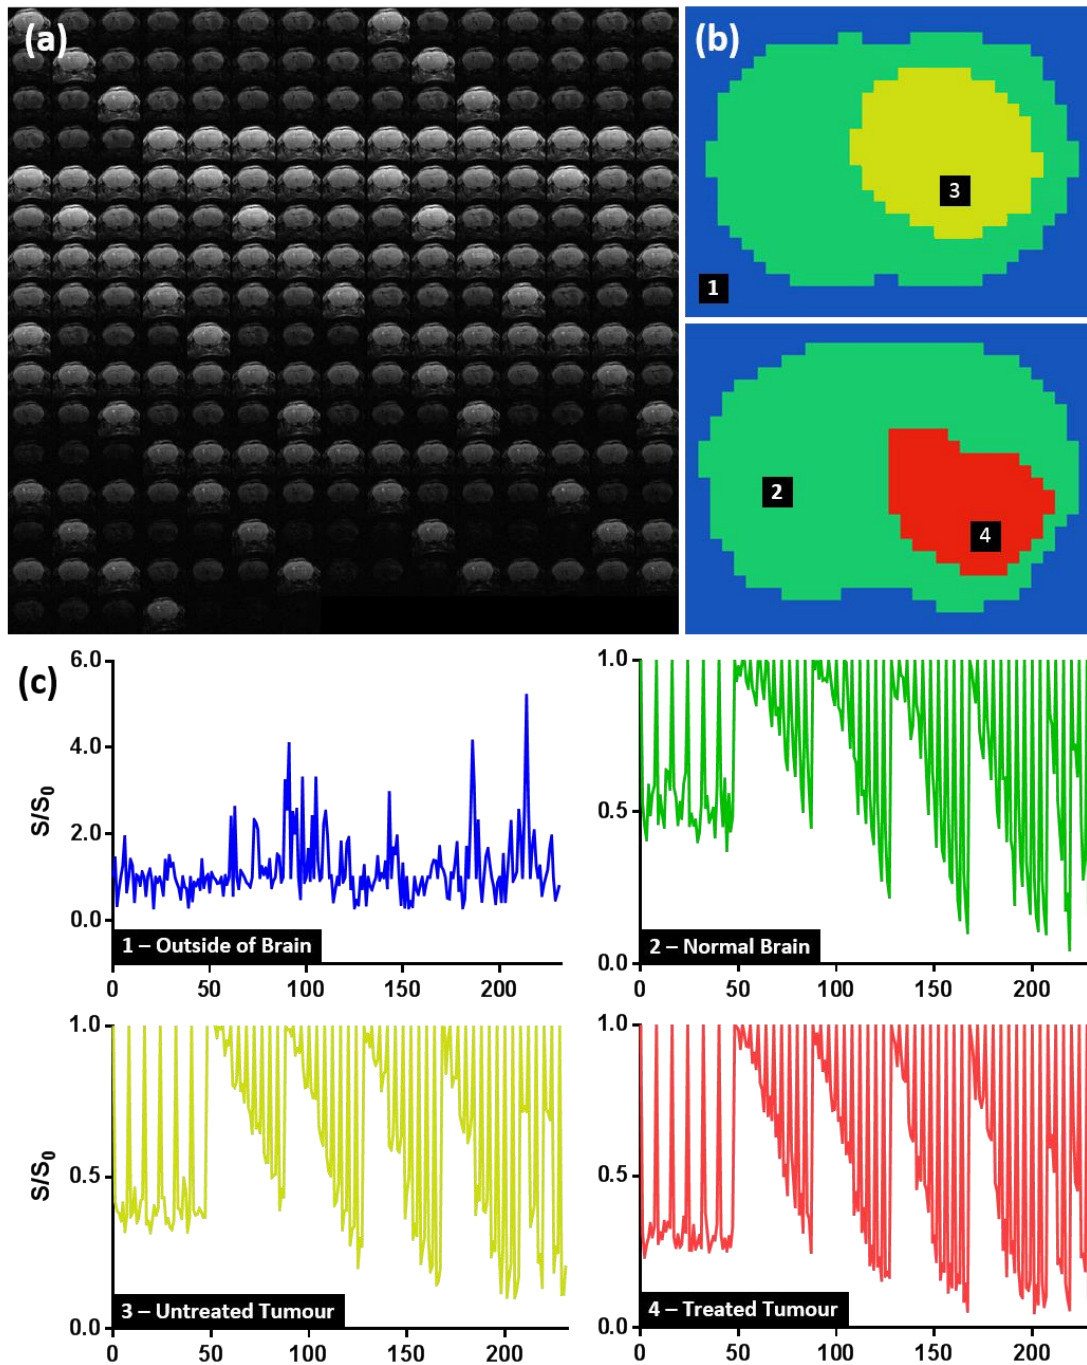
14 Gliomas were hyperintense on T1-weighted post-gadolinium images compared to normal
15 brain at both day 0 and day 9 (Figure 1a). These images were used for manual delineation
16 of brain and tumour regions (Figure 1a, dashed lines) resulting in four categories (Figure 1b):
17 outside of brain, 'normal' brain, untreated tumour (all day 0 tumours and day 9 tumours
18 which received no treatment) and treated tumour (day 9 tumours, which were treated with
19 Temozolomide).



1

2 Figure 1: Coronal images from mouse brains at day 0 and day 9. (a) T1-weighted post-
3 gadolinium slices showing manual segmentation (dashed lines). (b) Masks showing
4 classification regions used for training the deep learning neural network. (c) DLDF category
5 maps from the evaluation phase. (d) ADC maps. (e) Maps of tumour age from the evalua-
6 tion phase.

1 The category masks were applied to the slice-matched diffusion data to generate the
2 diffusion fingerprints required for training the ANN. Figure 2a shows a complete dataset of
3 diffusion MR images from one coronal slice through a mouse brain. Image intensity
4 decreases as the level of diffusion-weighting is increased. Figure 2c shows diffusion
5 fingerprints from four different regions. Diffusion fingerprint 1 (blue) comes from a region
6 outside of the animal, hence the plot is a noisy trace. Diffusion fingerprint 2 (green) comes
7 from normal brain, whilst diffusion fingerprints 3 (yellow) and 4 (red) originate from untreated
8 and treated tumour, respectively. The first 48 points in each plot represent the DTI (plus six
9 B_0) scans. The remaining points represent different combinations of diffusion gradients
10 (detailed in Table 1). Upon visual inspection, some differences can be observed in the
11 diffusion fingerprint from the normal brain compared to the fingerprints from the tumours. For
12 example, S/S_0 for the DTI scans is ~ 0.5 , compared to ~ 0.3 – 0.4 in the tumours. Furthermore,
13 S/S_0 decays more quickly in the tumours. The differences between the untreated (yellow)
14 and treated tumour (red) fingerprints are subtler, however, demonstrating the need for deep
15 learning to tease out the features within these fingerprints.



1

2 **Figure 2: (a) Complete set of diffusion MR images from a single coronal slice in the brain. (b)**

3 **Classification masks from an Untreated tumour and a Treated tumour. (c) Example diffusion**

4 **fingerprints taken from the corresponding numbered region in panel (b), showing the**

5 **normalised diffusion signal taken from a single voxel.**

1 **3.2 Model Accuracy**

2 In total, 6,608 fingerprints from 50% of the animals were used to train the ANN. Training
3 lasted 80 hours, and resulted in a model accuracy of 98.65% (where pixels were classified
4 as the same category determined by manual segmentation). Based on evaluation of the
5 ANN using the remaining 50% of the imaging data, accuracy was 98.54% (once again,
6 compared to manually segmented data).


7 **3.3 Voxel Classification**

8 A strong agreement was observed between manually segmented category maps (Figure 1b)
9 and DLDF maps generated from the evaluation phase (Figure 1c) demonstrating the ability
10 of the ANN to automatically segment tumour from normal brain. Predicted tumour volumes
11 were 2.1 ± 0.2 times larger than manually-segmented volumes. This is due in part to the
12 identification of tumour deposits that had grown beyond the skull. Ghosting artefact in
13 background regions was also miscategorised as brain or tumour tissue.

14 The ANN was able to reliably identify diffusion fingerprints corresponding to normal brain
15 (green), untreated tumour (yellow) or post-treatment tumour (red). Furthermore, the ANN
16 was able to distinguish between tumours based on age (Figure 1e): voxels within day 0
17 tumours were predominantly classified as day 0 pixels (green, 92% accuracy), whilst
18 untreated and treated tumours were mostly classified as day 9 pixels (red, 93% accuracy).
19 Regardless of age, tissue outside of tumours was predominantly classified as normal brain
20 (purple, 97% accuracy).

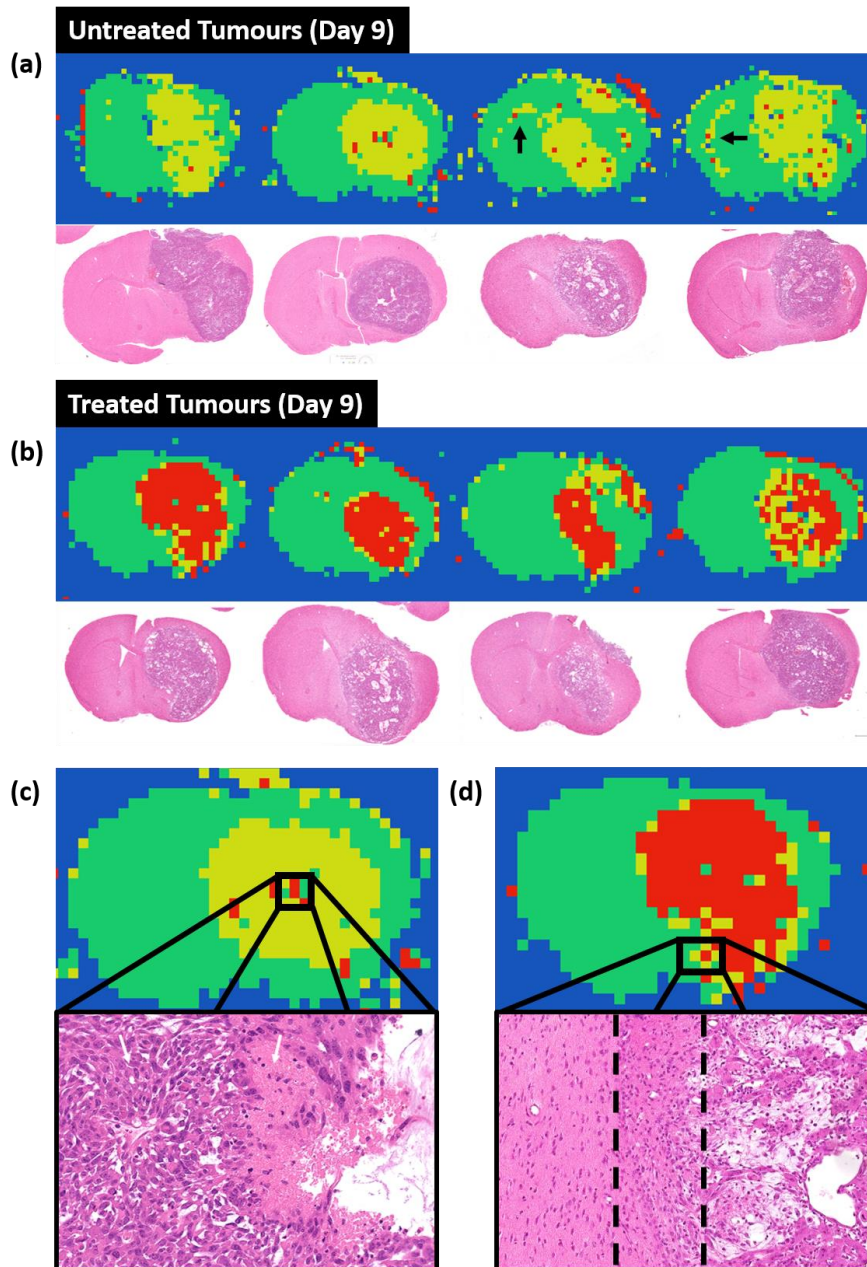
21 In comparison, ADC alone was unable to correctly stratify all pre-therapy tumours from
22 untreated tumours by simple thresholding (Figure 3b). Whilst the mean ADC value in the
23 pre-therapy tumours was significantly lower than the untreated tumours at day 9 ($p < 0.001$),
24 several tumours in both groups exhibited ADC values between 1.05×10^{-9} and 1.20×10^{-9}
25 m^2/s (blue band, Figure 3b). The mean ADC value in the treated group at day 9 was 47%

- 1 higher than baseline and all mice could be distinguished from pre-therapy and control
- 2 tumours.



3

- 4 **Figure 3: (a) Tumour volume measured from T2-weighted images. (b) Mean ADC**
- 5 **measurements from mouse brain tumours ($p < 0.001$, Kruskal-Wallis one-way ANOVA). The**
- 6 **blue region denotes a range of ADC values where pre-therapy and untreated tumours cannot**
- 7 **be stratified based on ADC alone.**

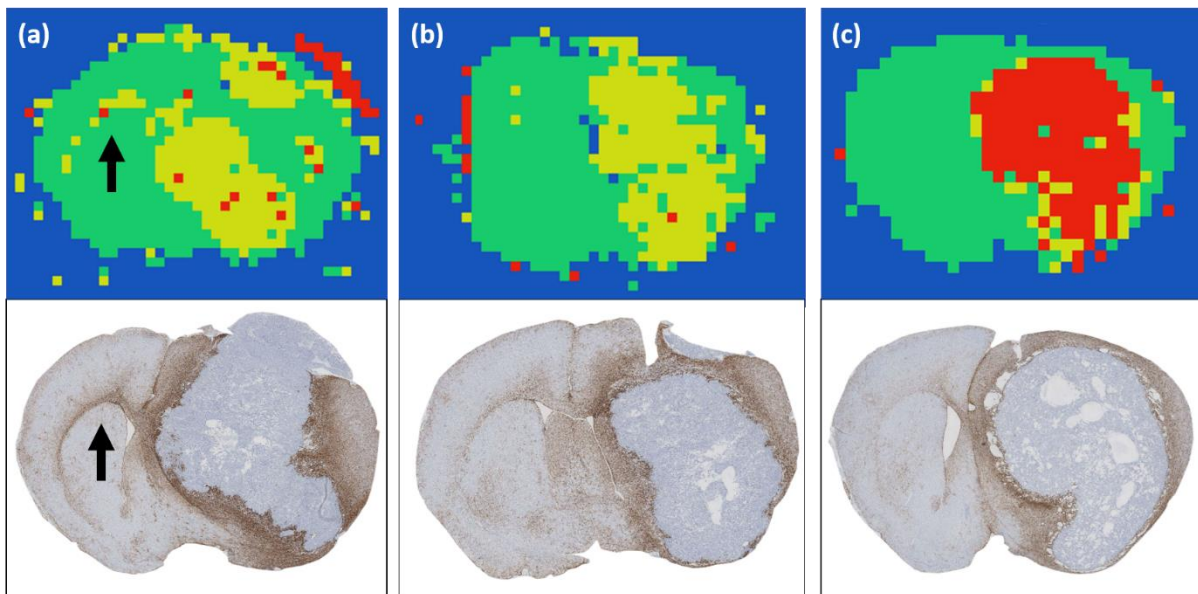


1

2 **Figure 4: DLDF classification maps and corresponding H&E histological sections at day 9 in**
3 **(a) Untreated tumours and (b) treated tumours. DLDF could correctly distinguish between all**
4 **untreated and treated tumours at Day 9. (c) and (d) Localised regions in DLDF classification**
5 **maps and the corresponding regions on histology. The untreated tumour in (c) shows voxels**
6 **classified as normal brain, which correspond approximately with neuronal tissue on histology.**
7 **The treated tumour in (d) shows a transition in voxel classification (left to right) from normal**
8 **brain (green) to untreated tumour (yellow) to treated tumour (red), which corresponds to three**
9 **distinct regions on histology: normal brain parenchyma, a region of tumour infiltration and the**
10 **bulk tumour tissue.**

1 Visual inspection of DLDF category maps from untreated tumours at day 9 (Figure 4a)
2 revealed a more spiculated appearance than treated tumours at day 9 (Figure 4b) and
3 showed a strong spatial accordance with H&E-stained histological sections. Centralised
4 regions of normal brain pixels (green) and 'treated' pixels (red) were often observed inside
5 the untreated tumours. These regions could correspond to areas on histology (Figure 4c)
6 which showed a mix of tumour cells, neuronal tissue and voids likely corresponding to cell
7 death.

8 DLDF category maps from treated animals generally showed a distinctive rim of 'untreated'
9 pixels (yellow) at the tumour periphery (Figure 4b), which could correspond to border regions
10 in H&E images where the tumour transitions into the normal brain (Figure 4d). Whether
11 these are regions that truly show less sensitivity to TMZ therapy requires further
12 investigation.



13

14 **Figure 5: DLDF category maps compared with GFAP-stained histological sections from the**
15 **same brains. (a) Evidence of tumour invasion across the corpus callosum in the DLDF map of**
16 **an untreated tumour corresponds to a region of enhanced GFAP staining (black arrows). DLDF**
17 **maps of (b) other untreated and (c) treated tumours showed no indication of corpus callosum**
18 **invasion, which was reflected in histology.**

1 In some of the untreated animals (Figure 5a, black arrows), the predicted category maps
2 showed tumour invasion into the contralateral hemisphere along the corpus callosum.
3 Histological sections from the same brains showed increased GFAP-staining along the
4 corpus callosum (Figure 5a), whereas much less GFAP-staining was visible in untreated
5 (Figure 5b) and treated (Figure 5c) mouse brains, which displayed no evidence of tumour
6 invasion on DLDF category maps.

7 **4. DISCUSSION**

8 The aim of this study was to evaluate the ability of DLDF in a mouse brain tumour model to
9 categorise normal brain, pre-treatment, untreated and Temozolomide-treated tumour pixels.
10 We have shown that DLDF can automatically segment tumours from normal brain, DLDF
11 can automatically distinguish between young and older (after 9 days of growth) tumours and
12 that DLDF can detect whether or not a tumour has been treated with chemotherapy. We
13 have also shown that, spatially, the category maps generated by DLDF were broadly
14 consistent with histology and that interesting features are detected which may correspond to
15 transitional regions within tumours where the microstructure is a mixture of normal brain and
16 tumour cells.

17 The ANN used for DLDF was trained on only 6,608 fingerprints and achieved good
18 prediction and evaluation accuracies. Broadly, the shape of the tumours delineated using
19 DLDF were consistent with manually-segmented training data. Interestingly, during manual
20 segmentation of T1-weighted images, regions of enhancement above the bulk of the tumour
21 or even external to the skull, which were likely caused by deposition of tumour cells during
22 retraction of the inoculation needle, were not segmented as tumour, and yet these regions
23 were correctly labelled as tumour on DLDF category maps.

24 Likewise, we also found evidence in DLDF category maps for tumour invasion into the
25 contralateral hemisphere, via the corpus callosum, but only in control mice. This was
26 accompanied by a much more spiculated appearing boundary, which contrasted with the

1 smoother appearance of TMZ-treated tumours. No evidence of invasion was detected in
2 treated mouse brains, most likely as a direct result of TMZ causing growth inhibition. GFAP-
3 stained histological slices from the same brains confirmed these regions of invasion, where
4 more intense staining was seen along the corpus callosum in the untreated animals. This
5 finding is of particular interest, as these invasion effects were not included in the manually
6 segmented ROIs used to train the ANN. It also presents a challenge in regard to quantifying
7 the accuracy of the ANN, as these, and other, correctly segmented and informative regions
8 actually serve to reduce the reported accuracy of the ANN.

9 Further deviations between manual segmentation and the DLDF category maps were
10 evident, and which were consistent with histology. Untreated tumours at day 9 contained
11 numerous voxels that were categorised as normal brain, and which were positioned in a
12 similar location to regions of brain tissue within the tumours. Likewise, voxels at the
13 periphery of TMZ-treated tumours were often misclassified as 'untreated'. This could be a
14 straightforward misclassification error by the ANN, or evidence of regions of tumour that
15 were less prone to changes in microstructure induced by Temozolomide, or even with a
16 potential for resistance. This could be an interesting area for further research, as
17 noninvasive biomarkers of treatment resistance would be particularly useful in the clinic for
18 monitoring response (23,24).

19 The DLDF results are also interesting as they appear to outperform an alternative, more
20 conventional approach to treatment monitoring, which is to apply simple thresholding to
21 mean tumour ADC measurements. Untreated brain tumours could be distinguished from
22 treated tumours based on their mean ADC, which was significantly higher in treated mice.
23 However, assessment of localised treatment response was not possible using ADC alone. In
24 treated mice, ADC was broadly elevated across the whole extent of tumour, however,
25 untreated mice exhibited localised regions of similarly elevated ADC, most likely
26 corresponding to regions of necrosis or (non-chemotherapy related) apoptosis. DLDF

1 provides a localised map of treatment response, possibly highlighting regions of non-
2 responding, or chemotherapy resistant, tumour tissue.

3 The results presented in this study, although preliminary, show some of the potential for
4 deep learning on an individual voxel level, and could be applied to other types of multi-
5 dimensional acquisitions. Even with limited training DLDF was able to accurately classify
6 normal brain, treated tumours and control tumours at two different time points, and in some
7 regards, was able to out-perform manual segmentation by identifying regions of tumour
8 invasion, mixed tissue types and, potentially, tumour regions that were less sensitive to
9 TMZ-treatment. The ability of DLDF to predict response to therapy and to 'age' tumour pixels
10 will be the focus of further investigation.

11 **ACKNOWLEDGEMENTS**

12 Wellcome Trust Senior Research Fellowship (grant WT100247MA), KCL and UCL
13 Comprehensive Cancer Imaging Centre CR-UK & EPSRC (In association with the DoH
14 England), the EPSRC-funded UCL Centre for Doctoral Training in Medical Imaging
15 (EP/L016478/1).

16 **REFERENCES**

- 17 1. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep
18 convolutional neural networks. 2012. p 1097-1105.
- 19 2. Roth HR, Lu L, Farag A, Shin H-C, Liu J, Turkbey E, Summers RM. Deeporgan:
20 Multi-level deep convolutional networks for automated pancreas segmentation. arXiv
21 preprint arXiv:150606448 2015.
- 22 3. Anthimopoulos M, Christodoulidis S, Ebner L, Christe A, Mougiakakou S. Lung
23 pattern classification for interstitial lung diseases using a deep convolutional neural
24 network. IEEE transactions on medical imaging 2016;35(5):1207-1216.
- 25 4. Havaei M, Davy A, Warde-Farley D, Biard A, Courville A, Bengio Y, Pal C, Jodoin P-
26 M, Larochelle H. Brain tumor segmentation with deep neural networks. Medical
27 image analysis 2017;35:18-31.
- 28 5. Matsuki Y, Nakamura K, Watanabe H, Aoki T, Nakata H, Katsuragawa S, Doi K.
29 Usefulness of an artificial neural network for differentiating benign from malignant
30 pulmonary nodules on high-resolution CT: evaluation with receiver operating
31 characteristic analysis. American Journal of Roentgenology 2002;178(3):657-663.
- 32 6. Fukushima A, Ashizawa K, Yamaguchi T, Matsuyama N, Hayashi H, Kida I, Imafuku
33 Y, Egawa A, Kimura S, Nagaoki K. Application of an artificial neural network to high-
34 resolution CT: usefulness in differential diagnosis of diffuse lung disease. American
35 Journal of Roentgenology 2004;183(2):297-305.

- 1 7. Kuruvilla J, Gunavathi K. Lung cancer classification using neural networks for CT
2 images. *Computer methods and programs in biomedicine* 2014;113(1):202-209.
- 3 8. Sharif MS, Abbod M, Amira A, Zaidi H. Artificial neural network-based system for
4 PET volume segmentation. *Journal of Biomedical Imaging* 2010;2010:4.
- 5 9. Nogueira MA, Abreu PH, Martins P, Machado P, Duarte H, Santos J. An artificial
6 neural networks approach for assessment treatment response in oncological patients
7 using PET/CT images. *BMC medical imaging* 2017;17(1):13.
- 8 10. Toney LK, Vesselle HJ. Neural networks for nodal staging of non-small cell lung
9 cancer with FDG PET and CT: importance of combining uptake values and sizes of
10 nodes and primary tumor. *Radiology* 2014;270(1):91-98.
- 11 11. Pereira S, Pinto A, Alves V, Silva CA. Brain tumor segmentation using convolutional
12 neural networks in MRI images. *IEEE transactions on medical imaging*
13 2016;35(5):1240-1251.
- 14 12. Stier N, Vincent N, Liebeskind D, Scalzo F. Deep learning of tissue fate features in
15 acute ischemic stroke. 2015. *IEEE*. p 1316-1321.
- 16 13. Ranjith G, Parvathy R, Vikas V, Chandrasekharan K, Nair S. Machine learning
17 methods for the classification of gliomas: Initial results using features extracted from
18 MR spectroscopy. *The neuroradiology journal* 2015;28(2):106-111.
- 19 14. van der Burgh HK, Schmidt R, Westeneng H-J, de Reus MA, van den Berg LH, van
20 den Heuvel MP. Deep learning predictions of survival based on MRI in amyotrophic
21 lateral sclerosis. *NeuroImage: Clinical* 2017;13:361-369.
- 22 15. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciampi F, Ghafoorian M, van der Laak JA,
23 van Ginneken B, Sánchez CI. A survey on deep learning in medical image analysis.
24 *arXiv preprint arXiv:170205747* 2017.
- 25 16. Greenspan H, van Ginneken B, Summers RM. Guest editorial deep learning in
26 medical imaging: Overview and future promise of an exciting new technique. *IEEE*
27 *Transactions on Medical Imaging* 2016;35(5):1153-1159.
- 28 17. Zhang H, Schneider T, Wheeler-Kingshott CA, Alexander DC. NODDI: practical in
29 vivo neurite orientation dispersion and density imaging of the human brain.
30 *Neuroimage* 2012;61(4):1000-1016.
- 31 18. Tuch DS, Reese TG, Wiegell MR, Makris N, Belliveau JW, Wedeen VJ. High angular
32 resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity.
33 *Magnetic resonance in medicine* 2002;48(4):577-582.
- 34 19. Panagiotaki E, Walker-Samuel S, Siow B, Johnson SP, Rajkumar V, Pedley RB,
35 Lythgoe MF, Alexander DC. Noninvasive quantification of solid tumor microstructure
36 using VERDICT MRI. *Cancer research* 2014;74(7):1902-1912.
- 37 20. Panagiotaki E, Chan RW, Dikaios N, Ahmed HU, O'Callaghan J, Freeman A,
38 Atkinson D, Punwani S, Hawkes DJ, Alexander DC. Microstructural characterization
39 of normal and malignant human prostate tissue with vascular, extracellular, and
40 restricted diffusion for cytometry in tumours magnetic resonance imaging.
41 *Investigative radiology* 2015;50(4):218-227.
- 42 21. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A,
43 Dean J, Devin M. Tensorflow: Large-scale machine learning on heterogeneous
44 distributed systems. *arXiv preprint arXiv:160304467* 2016.
- 45 22. Gal Y, Ghahramani Z. Dropout as a Bayesian approximation: Representing model
46 uncertainty in deep learning. 2016. p 1050-1059.
- 47 23. Sharma P, Hu-Lieskovan S, Wargo JA, Ribas A. Primary, adaptive, and acquired
48 resistance to cancer immunotherapy. *Cell* 2017;168(4):707-723.
- 49 24. Wargo JA, Reddy SM, Reuben A, Sharma P. Monitoring immune responses in the
50 tumor microenvironment. *Current opinion in immunology* 2016;41:23-31.

51

1 **TABLE CAPTIONS**

2 **Table 1: Diffusion gradient combinations used in mouse brain imaging for creating diffusion**
3 **fingerprints.**

4

5 **FIGURE CAPTIONS**

6 **Figure 1: Coronal images from mouse brains at day 0 and day 9. (a) T1-weighted post-**
7 **gadolinium slices showing manual segmentation (dashed lines). (b) Masks showing**
8 **classification regions used for training the deep learning neural network. (c) DLDF category**
9 **maps from the evaluation phase. (d) ADC maps. (e) Maps of tumour age from the evaluation**
10 **phase.**

11 **Figure 2: (a) Complete set of diffusion MR images from a single coronal slice in the brain. (b)**
12 **Classification masks from an Untreated tumour and a Treated tumour. (c) Example diffusion**
13 **fingerprints taken from the corresponding numbered region in panel (b), showing the**
14 **normalised diffusion signal taken from a single voxel.**

15 **Figure 3: (a) Tumour volume measured from T2-weighted images. (b) Mean ADC**
16 **measurements from mouse brain tumours ($p < 0.001$, Kruskal-Wallis one-way ANOVA). The**
17 **blue region denotes a range of ADC values where pre-therapy and untreated tumours cannot**
18 **be stratified based on ADC alone.**

19 **Figure 4: DLDF classification maps and corresponding H&E histological sections at day 9 in**
20 **(a) Untreated tumours and (b) treated tumours. DLDF could correctly distinguish between all**
21 **untreated and treated tumours at Day 9. (c) and (d) Localised regions in DLDF classification**
22 **maps and the corresponding regions on histology. The untreated tumour in (c) shows voxels**
23 **classified as normal brain, which correspond approximately with neuronal tissue on histology.**
24 **The treated tumour in (d) shows a transition in voxel classification (left to right) from normal**
25 **brain (green) to untreated tumour (yellow) to treated tumour (red), which corresponds to three**
26 **distinct regions on histology: normal brain parenchyma, a region of tumour infiltration and the**
27 **bulk tumour tissue.**

1 **Figure 5: DLDF category maps compared with GFAP-stained histological sections from the**
2 **same brains. (a) Evidence of tumour invasion across the corpus callosum in the DLDF map of**
3 **an untreated tumour corresponds to a region of enhanced GFAP staining (black arrows). DLDF**
4 **maps of (b) other untreated and (c) treated tumours showed no indication of corpus callosum**
5 **invasion, which was reflected in histology.**