

Quantifying Heterogeneity in Dynamic Contrast-Enhanced MRI Parameter Maps

C.J. Rose¹, S. Mills¹, J.P.B. O’Connor¹, G.A. Buonaccorsi¹, C. Roberts¹, Y. Watson¹, B. Whitcher², G. Jayson³, A. Jackson¹, and G.J.M. Parker¹

¹ Imaging Science & Biomedical Engineering,
The University of Manchester, United Kingdom

² MRI Modelling, GlaxoSmithKline, London, United Kingdom

³ Cancer Research UK Dept. of Medical Oncology, Christie Hospital, Manchester,
United Kingdom

Abstract. Simple summary statistics of Dynamic Contrast-Enhanced MRI (DCE-MRI) parameter maps (e.g. the median) neglect the spatial arrangement of parameters, which appears to carry important diagnostic and prognostic information. This paper describes novel statistics that are sensitive to both parameter values and their spatial arrangement. Binary objects are created from 3-D DCE-MRI parameter maps by “extruding” each voxel into a fourth dimension; the extrusion distance is proportional to the voxel’s value. The following statistics are then computed on these 4-D binary objects: surface area, volume, surface area to volume ratio, and box counting (fractal) dimension. An experiment using 4 low and 5 high grade gliomas showed significant differences between the two grades for box counting dimension computed for extruded v_e maps, surface area of extruded K^{trans} and v_e maps and the volume of extruded v_e maps (all $p < 0.05$). An experiment using 18 liver metastases imaged before and after treatment with a vascular endothelial growth factor (VEGF) inhibitor showed significant differences for surface area to volume ratio computed for extruded K^{trans} and v_e maps ($p = 0.0013$ and $p = 0.045$ respectively).

1 Introduction

Dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) has been used in both clinical and research evaluation of tumour biology and response to therapy [3]. Using DCE-MRI tracer-kinetic modelling [7], tumour voxels can be parameterised by K^{trans} (the volume transfer coefficient of contrast agent between the blood pool and the extravascular space), v_p (the relative volume of the blood plasma space) and v_e (the relative volume of the extravascular extracellular space). These quantities can be mapped to their spatial locations to yield parametric maps of tumour characteristics (see Figure 1).

Most studies that incorporate DCE-MRI report either baseline values of, or changes in, simple summary statistics of the kinetic parameters (e.g. median K^{trans}). These approaches neglect spatial heterogeneity within the tumours, despite this being an important feature of most solid tumours that may relate to

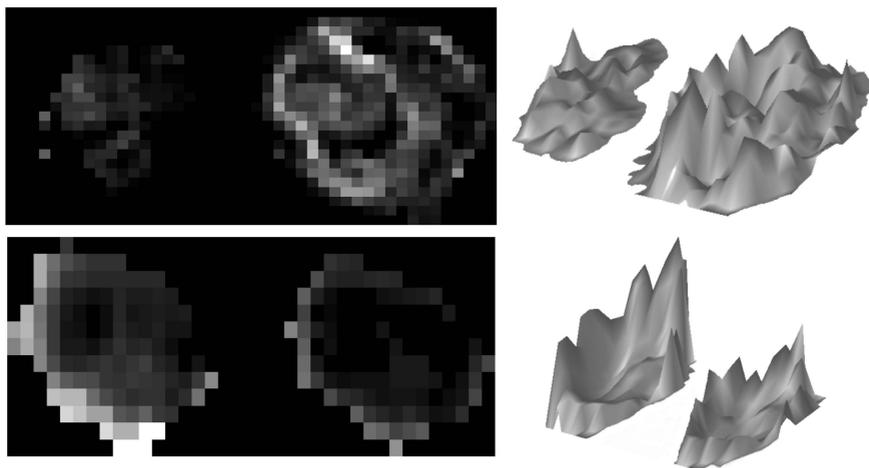


Fig. 1. K^{trans} maps. Top row: slices through a low grade (left) and a high grade (right) glioma (data range: 0-0.16 min^{-1}). Bottom row: corresponding slices through a liver metastasis before (left) and after (right) treatment with a VEGF inhibitor (data range: 0-0.45 min^{-1}). The right-most column shows the tumours rendered as 3-D surfaces with height representing the K^{trans} value at each location in the slices.

treatment choice and drug efficacy. High values of K^{trans} indicate relatively high perfusion and vessel permeability. Figure 1 shows slices through K^{trans} maps of low and high grade gliomas and a liver metastasis before and after treatment with a VEGF inhibitor. Compared to the low grade glioma, the high grade is larger, has a wider range of K^{trans} values and—importantly—has a much more complex internal structure. The liver metastasis shows clear non-uniform changes in K^{trans} values after treatment, particularly in the periphery. Statistics that capture heterogeneity will enable change in parameter values and/or their spatial arrangement to be quantified.

Texture analysis underpins much of image analysis and is clearly of interest to researchers working on MRI [2]. However, an assumption is made when utilising texture measures that the signal being analysed is, to a large degree, a repeating pattern subject to some degree of randomness. Figure 1 demonstrates that tumour DCE-MRI parameter maps do not satisfy this assumption and so traditional texture analysis methods may be unsuitable. An alternative strategy is to treat parameter maps as objects rather than textures and seek statistics which describe the form of these objects [1]. Figure 2 shows two signals that have the same mean but differ in heterogeneity. Rendering the signals as surfaces (curves in this 1-D example) shows how measures of object complexity may distinguish them. This paper develops heterogeneity statistics based on such a representation and applies them in two experiments using real data.

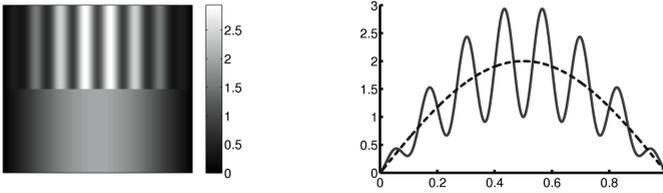


Fig. 2. Two 1-D signals that have the same mean value. The image on the left shows the grey-level values of the signals as an image; the plot on the right shows the signals as curves. The top signal on the left (plotted as a solid line on the right) is more heterogeneous than the other.

2 Quantifying Heterogeneity

Signals defined on a d -dimensional lattice can be represented as surfaces in $d + 1$ dimensions, where surface height is proportional to the signal value. DCE-MRI parameter maps are defined on 3-D lattices. Therefore we create a 4-D object by “extruding” each enhancing voxel of the original 3-D parameter map into a fourth dimension, using a 4-D binary array to represent the extruded object. The distance each voxel is extruded is proportional to its DCE-MRI model kinetic parameter (analogous to the lower-dimensional examples shown in Figure 1 and Figure 2).

The distributions of DCE-MRI tracer kinetic modelling parameters are typically heavily positively skewed. To make optimal use of the available dynamic range of the extrusion dimension, we map from model parameter to surface height using non-linear functions $f_{K^{\text{trans}}}$, f_{v_p} and f_{v_e} . For example $f_{v_p}(v_p) = \lceil r D_{v_p}(v_p) \rceil$ where: r is the maximum number of elements in the extrusion dimension; D_{v_p} is the cumulative distribution function for v_p , learned from a training set of approximately 20,000 DCE-MRI parameters; and $\lceil x \rceil$ is the ceiling function. Given a 4-D object we compute the following statistics: surface area, volume, ratio of surface area to volume and box-counting dimension [6]. The methods used to compute these statistics are explained in Section 2.1 and Section 2.2.

2.1 Computing Surface Area

While surface area is conceptually simple, a suitable algorithm for computing this quantity for a 4-D object may not be immediately obvious. The surface area of an object composed of N elements is maximised when all object elements are mutually unconnected. This maximum surface area is $s_M = N s_T$, where each element has a total surface area of s_T . In general, however, object elements are connected to one another and some potential area, s_C , is “lost” to these internal connections. The surface area of the object is therefore $s = s_M - s_C$ and we seek a way of computing $s_M - s_C$. Figure 3 shows a 2-D object which has isotropic pixels of unit size (i.e. each pixel has a total surface area of 4). Each pixel is labelled with the surface area that the pixel cannot contribute to the surface area

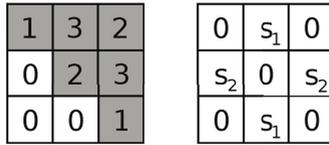


Fig. 3. A 2-D object is shown on the left (shaded region) and each object pixel is labelled with the surface area that is “lost” due to neighbouring pixels (this example uses isotropic pixels of unit size). A convolution kernel is shown on the right (see Eqn. 2).

of the object as a whole due to connections with neighbouring object pixels. The surface area of the object is therefore $s_M - s_C = (6 \times 4) - 12 = 12$ units.

Let X be an n -dimensional array where a value of 1 indicates the element belongs to the object and 0 indicates otherwise. Let F be an n -dimensional $3 \times 3 \times \dots \times 3$ convolution kernel with each connecting element set to the surface area that would be lost due to a connection in the corresponding direction—as shown on the right in Figure 3. The array K containing the surface areas involved in connections—as shown on the left in Figure 3—is computed as:

$$K = X \times (X * F) \tag{1}$$

where \times denotes element-wise multiplication and $*$ denotes convolution (after padding the edges with zeros)¹.

The total surface area of an element— s_T —is conceptually straightforward, but it is worth noting how this quantity can be computed in arbitrary dimensions. An n -cuboid has n pairs of faces of equal area. s_T can be determined by computing the areas for each pair of faces. Each area is the product of the side lengths of the n -cuboid’s sides, excluding one of these lengths (corresponding to a connection direction). Let u be a vector of the side lengths of an n -cuboid. The surface area of a face in direction i is therefore:

$$s_i = \prod_{j=1, j \neq i}^n u_j. \tag{2}$$

The total surface area of the n -cuboid is therefore:

$$s_T = 2 \sum_{i=1}^n s_i. \tag{3}$$

Given an array K computed for an object with N elements described by X , the surface area, s , of the object is:

$$s = s_M - s_C = \sum_{i=1}^N s_T - k_i \tag{4}$$

where i is simply an index into the object elements of K .

¹ A correction to Equation 1 is required if the object contains single elements which are not connected to the object but are nevertheless considered to be part of it, but further exposition is beyond the scope of this paper.

2.2 Computing Additional Complexity Statistics

In addition to surface area, we compute the volume of the 4-D object, the ratio of surface area to volume and the box-counting fractal dimension [6]. Volume is simply computed as $v = N \prod_{i=1}^n u_i$ (i.e. the number of elements multiplied by the volume of a single element). The box-counting dimension, d_0 , is a well-established statistic and is computed by investigating how the number of elements, $C(\epsilon)$, required to represent an object changes with their size, ϵ :

$$d_0 = - \lim_{\epsilon \rightarrow 0} \frac{\log C(\epsilon)}{\log \epsilon}. \quad (5)$$

The array which represents the 4-D object is composed of the natural 3-D imaging grid and that of the imposed extrusion dimension. Box counting is performed by subsampling the 4-D array and proceeds from the finest scale to the coarsest.

3 Experiment I: Grading Gliomas

Gliomas are graded based upon their histological appearance—necessitating a highly invasive procedure—using criteria set out by the World Health Organisation (WHO) [4]. Grade has been shown to relate to DCE-MRI parameter summary statistics [5] but does not—in general—correlate with tumour volume and grade cannot currently be determined using conventional routine clinical imaging. Subjectively, however, grade appears to relate to DCE-MRI parameter map heterogeneity. High grades are characterised by both hypocellularity (reduced cell density due to necrosis) and hypercellularity (increased cell density due to malignancy), leading to heterogeneity in histology and, we assume, in imaging. Gliomas therefore provide a good candidate for the kind of problem described in Section 2. We hypothesised that the methods described above would be sensitive to this heterogeneity.

3.1 Patients and Imaging Protocol

Nine adult patients with gliomas were recruited and underwent T_1 -weighted DCE-MRI on a 3T Philips Achieva MR scanner (Philips Medical Systems, Best, Netherlands)². The study was approved by the local Research Ethics Committee and all patients gave written informed consent. All MR imaging was performed before surgery (for tumour resection or debulking). All tumours were histologically confirmed to be gliomas and were graded according to criteria set out by the WHO (low grade (II) $n = 4$, high grade (III and IV) $n = 5$). While not generally true of the population, in this study the high grade tumours were approximately 1 cm in diameter larger than the low grade tumours. All data were subjected to a routine quality assurance (QA) procedure.

² The authors would like to acknowledge Sha Zhao, who developed the 3T DCE-MRI protocol.

3.2 Method

The nine tumours were manually segmented by an expert radiologist and values of K^{trans} , v_e and v_p were computed at each tumour voxel from the DCE-MR image sequences using the following tracer kinetic model [7]:

$$C_t(t) = v_p C_p(t) + K^{\text{trans}} \int_0^t C_p(t') e^{\frac{-K^{\text{trans}}(t-t')}{v_e}} dt'. \quad (6)$$

The functions C_t and C_p give the concentration of tracer in the voxel as a whole and in the voxel blood plasma respectively at time t . Each resulting parameter map was converted into its 4-D representation as described in Section 2 and surface area, volume, surface area to volume ratio and box-counting dimension were computed. Differences between the statistics for the two groups (low and high grades) were investigated using Wilcoxon rank sum tests.

3.3 Results

There were significant differences between the groups for: box dimension computed for extruded v_e parameter maps; surface area of extruded K^{trans} and v_e parameter maps; and the volume of extruded v_e parameter maps (all $p < 0.05$ without correction for multiple comparisons); see Figure 4(a). The other differences were not significant.

3.4 Discussion

In this experiment, all statistics have been used as absolute estimates of heterogeneity (cf. Section 4.4). Surface area- and volume-based statistics will be biased by tumour size, and so may not in general distinguish glioma grade. Box-counting dimension should be invariant to tumour size, although in practice object size can limit the precision with which this statistic can be estimated.

We investigated the effect that object size has on box-counting dimension. 1000 random 4-D binary arrays were generated, using a similar array size to those for the high grade (larger) gliomas. The box-counting dimension was then computed for each object. The objects were then subsampled down to approximately the same size as the low grade (smaller) gliomas, and box-counting dimension was computed for these. Pairwise differences in box-counting dimension were then computed. The differences in box-counting dimension were approximately a quarter of the size of those observed for the real glioma data, suggesting that tumour size cannot explain the differences between the two groups. Box-counting dimension on extruded v_e maps may therefore allow gliomas to be graded using a non-invasive procedure.

4 Experiment II: Quantifying Change in Heterogeneity

The following experiment demonstrates how the statistics described in the paper could be used in the drug trial setting to investigate heterogenous changes within the tumour microvasculature following administration of a VEGF inhibitor.

4.1 Patients and Imaging Protocol

Four patients with a total of 25 liver metastases ($n_1 = 10, n_2 = 1, n_3 = 7, n_4 = 7$) were recruited and underwent T_1 -weighted DCE-MRI imaging on a Philips Intera 1.5T scanner. The study was approved by the local Research Ethics Committee and all patients gave written informed consent. Two pre-treatment baseline scans were performed in the week preceding the administration of the VEGF inhibitor (though the initial baseline was missing for one patient, $n_4 = 7$). A post-treatment scan was then performed. All data were subjected to a routine QA procedure. Seven tumours were rejected due to poor data quality, leaving 18 tumours. Three of the remaining 54 scans were also rejected by QA due to substantial motion artefacts.

4.2 Method

The tumours were manually segmented by an expert radiographer. Maps of K^{trans} , v_e and v_p were computed from the DCE-MR image sequences and converted to their 4-D representations as described in Section 3.2. Surface area, volume, surface area to volume ratio and box-counting dimension were then computed for each tumour at each time point. The tumours were assumed to be independent and non-parametric ANOVAs (Kruskal Wallis) were performed to investigate group differences. Post hoc Wilcoxon sign rank tests were performed as appropriate.

4.3 Results

There were significant differences in surface area to volume ratio for extruded K^{trans} and v_e maps ($p = 0.0013$ and $p = 0.045$ respectively). Post hoc testing revealed significant differences between first baseline and post-treatment scans

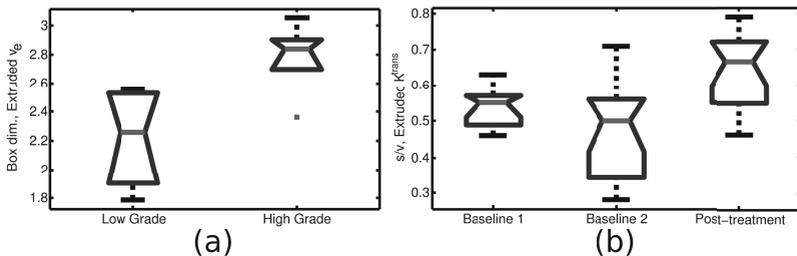


Fig. 4. (a) Boxplots showing the difference in distribution between box-counting dimension statistics computed for extruded v_e parameter maps for low- and high-grade gliomas. (b) Boxplots showing the difference in distribution between surface area to volume ratio statistics computed for extruded K^{trans} parameter maps before and after treatment with a VEGF inhibitor. Note that baseline 1 lacks data for one patient (7 data points); removing the corresponding data points from baseline 2 confirms the missing data to be the cause of the tighter distribution at baseline 1.

($p = 0.004$) and second baseline and post-treatment ($p < 0.001$) scans (without correction for multiple comparisons). As expected, there were no significant differences in heterogeneity statistics between the baseline scans. See Figure 4(b). The other differences were not significant.

4.4 Discussion

All statistics in this experiment have been used as relative estimates of heterogeneity (i.e. we are interested in a change in a statistic for a particular tumour; cf. Section 3.4). While the effects of the VEGF inhibitor can be observed using simple summary statistics, the results clearly illustrate that heterogeneity statistics are sensitive to known treatment effects. Heterogeneity statistics may therefore be useful in the evaluation of pharmaceutical therapies.

5 Conclusions

This paper has described DCE-MRI-based statistics of heterogeneity. Significant differences were found in some of these between low and high grade gliomas, suggesting that it might be possible to grade gliomas using a non-invasive imaging-based technique (e.g. as an adjunct to histology). Heterogeneity statistics were also sensitive to the effects of a VEGF inhibitor. However, further work is required to determine their ability to provide information that is unavailable from “standard” DCE-MRI summary statistics.

There was no combination of model parameter and heterogeneity statistic common to both experiments that provided significant discrimination. However, it is interesting that heterogeneity statistics based upon v_e were able to provide significant discrimination in both experiments— v_e is often considered to be of little physiological importance compared to K^{trans} and v_p . We hypothesise that the observed differences in heterogeneity are due to differences in cellular packing density. Further work is required to elucidate the generality of this observation.

References

1. Dzik-Jurasz, A., Walker-Samuel, S., Leach, M.O., Brown, G., Padhani, A., George, M., Collins, D.J.: Fractal parameters derived from analysis of DCE-MRI data correlates with response to therapy in rectal carcinoma. In: Proc. Intl. Soc. for Mag. Reson. Med., p. 2505 (2004)
2. Hajek, M., Dezortova, M., Materka, A., Lerski, R. (eds.): Texture Analysis for Magnetic Resonance Imaging. Med4publishing (2006) ISBN: 80-903660-0-7
3. Jackson, A., Buckley, D.L., Parker, G.J.M.: Dynamic Contrast-Enhanced Magnetic Resonance Imaging in Oncology. In: Medical Radiology: Diagnostic Imaging, Springer, Heidelberg (2005)

4. Kleihues, P., Louis, D.N., Scheithauer, B.W., et al.: The WHO Classification of Tumors of the Nervous System. *J. Neuropathol. Exp. Neurol.* 61(3), 215–225 (2002) Discussion 226–9
5. Patankar, T.F., Haroon, H.A., Mills, S.J., et al.: Is Volume Transfer Coefficient (K^{trans}) Related to Histological Grade in Human Gliomas? *Am. J. Neuroradiol.* 26(10), 2455–2465 (2005)
6. Peitgen, H.-O., Jürgens, H., Saupe, D.: *Chaos and Fractals*. Springer, Heidelberg (2004)
7. Tofts, P.S.: Modeling Tracer Kinetics in Dynamic Gd-DTPA MR Imaging. *J. Magn. Reson. Imaging* 7(1), 91–101 (1997)