Corpus callosum MR image classification

A. Elsayed [a,*], F. Coenen [a], C. Jiang [a], M. García-Fiñana [b], V. Sluming [c]

[a] Department of Computer Science, University of Liverpool, Ashton Building, Ashton Street, Liverpool L69 3BX, United Kingdom
[b] Centre for Medical Statistics and Health Evaluation, University of Liverpool, Shelley's Cottage, Brownlow Street, Liverpool L69 3GS, United Kingdom
[c] School of Health Sciences, Thompson Yates Building, The Quadrangle, Brownlow Hill, Liverpool L69 3GB, United Kingdom

ARTICLE INFO

ABSTRACT

An approach to classifying magnetic resonance (MR) image data is described. The specific application is the classification of MRI scan data according to the nature of the corpus callosum, however the approach has more general applicability. A variation of the "spectral segmentation with multi-scale graph decomposition" mechanism is introduced. The result of the segmentation is stored in a quad-tree data structure to which a weighted variation (also developed by the authors) of the gSpan algorithm is applied to identify frequent sub-trees. As a result the images are expressed as a set frequent sub-trees. There may be a great many of these and thus a decision tree based feature reduction technique is applied before classification takes place. The results show that the proposed approach performs both efficiently and effectively, obtaining a classification accuracy of over 95% in the case of the given application.

1. Introduction

The application of data-mining techniques to image data involves a number of challenges relating to the representation of images into an appropriate format that permits the application of data-mining techniques. This pre-processing typically involves some form of image segmentation to identify image features/objects; followed by the recasting of the image set into some appropriate format.

This paper is focused on a particular application domain, the classification of magnetic resonance (MR) image data, more specifically the classification of MR image data according to the nature of the *corpus callosum*. The corpus callosum is a highly visible structure in MR images whose function is to connect the left hemisphere of the brain to the right hemisphere, and be responsible for communication between these two hemispheres. The specific application used to illustrate the described classification process is the categorisation of MR images into one of two classes: (i) musicians and (ii) non-musicians. However, the process has more general applicability.

The classification process commences with the segmentation of the input images. With respect to the specific application reported here this is to identify the corpus callosum. For this purpose a variation of a standard image segmentation technique, *spectral segmentation*, is introduced. A registration process is then applied, so that each identified corpus callosum is founded upon the same origin. The pixel representations, for each corpus callosum, is then tessellated and stored in a quad-tree data structure, one tree per image. The advantage offered by the quad-tree representation, unlike some other representations, is that information about the spatial relationships between individual pixels is maintained. These trees are then processed to identify frequent sub-trees that occur across the input image set. For this purpose a weighted graph-mining algorithm was developed to take into consideration that greater significance should be assigned to quad-tree nodes closer to the root than nodes further away. The identified sub-trees then formed the fundamental elements of the feature space. Each image was then represented in terms of this feature space using individual feature vectors.

From experimentation it was discovered that the number of features (frequent sub-trees) could be substantial and that many features seemed to play no part in the resulting classification. To reduce the number of features to a more manageable number a feature selection mechanism was applied using a decision-tree algorithm; features that did not appear in the decision tree were removed from the input set. A second application of the decision-tree algorithm then yielded the final classifier (although any other appropriate classifier generator could have been used). The results, in terms of the specific application under investigation, were found to be very good.

* Corresponding author.
E-mail addresses: a.el-sayed@liv.ac.uk (A. Elsayed), coenen@liv.ac.uk (F. Coenen), c.jiang@liv.ac.uk (C. Jiang), m.garciafinana@liv.ac.uk (M. García-Fiñana), vanessa.sluming@liv.ac.uk (V. Sluming).

This paper makes a number of contributions:

- A new approach to image pre-processing (for image mining) founded on weighted frequent sub-graph mining, is described; an approach that has much broader potential application.
- An interesting image mining application which offers clear medical benefits with respect to medical diagnosis (e.g. eurological disorders) is described. To the best knowledge of the authors there are no corpus callosum studies that take the "shape" of the callosum into consideration (most studies reported in the literature concentrate of the size of the callosum).
- A new variation, founded on an intensity threshold, of the established spectral segmentation with multi-scale graph decomposition approach.

The rest of this paper is organised as follows. The application domain is described in Section 2 and some relevant previous work in Section 3. The proposed classification process is described in Section 4. The various steps in the process: segmentation, registration and tessellation, weighted graph mining, feature selection and image classification are described in Sections 5–9, respectively. An evaluation of the approach is reported in Section 10, followed by some conclusions in Section 11.

2. The application domain

The focus of the work described here is brain MR images, and in particular a specific structure in these images called the corpus callosum. An example image is given in Fig. 1. The corpus callosum is located in the centre of the image, the fornix is a related structure which often "blurs" into the corpus callosum and presents a particular challenge in the context of segmentation. The corpus callosum is of interest to medical researchers for a number of reasons. The size and shape of the corpus callosum have been shown to be correlated to sex, age, neurodegenerative diseases (such as epilepsy) and various lateralized behaviour in people. It is conjectured that the size and shape of the corpus callosum reflects certain human characteristics (such as a mathematical or musical ability). It is a very distinctive feature in MRI brain scans.

Several medical studies indicate that the size and shape of the corpus callosum, in humans, are correlated to sex [1,6,18], age

Fig. 1. Corpus callosum in a midsagittal brain MR image.

[18,20], brain growth and degeneration [9,14], handedness [5] and various types of brain dysfunction [7,11]. In order to find such correlations in living brains, magnetic resonance (MR) is regarded as the best method to obtain cross-sectional area (and shape) information of the corpus callosum. Since manual tracing of the corpus callosum in MRI scans is time consuming, operator dependent, and does not directly give quantitative measures of shape; there is a need for automated and robust methods for localization, delineation and shape description of the corpus callosum. This is then the motivation for the work described here.

3. Previous work

Image classification systems tend to rely on a pre-processing step, specific to an application, to extract a (reduced) set of "interesting" features from the image data. This set is then used as the input to a classification algorithms. There is a significant body of literature covering the domain of image mining and image classification with application in many areas.

There is also a substantial body of work directed at the mining of MR images. For example Chen and Herskovits [2] present a Bayesian-network for joint classification founded on voxelwise MR image analysis. This work considers associations between cerebral morphology of all brain voxels and age or sex. The approach includes an embedded feature selection phase in classifier learning. Ruan et al. [17] present a fully automatic three-dimensional classification of brain tissues for MR images into three types of main tissues: Gray Matter, White Matter, and Cerebrospinal fluid and mixtures of these brain tissues. Chun and Greenshields [3] describe a three-dimensional image classification technique, founded on the Markov and Gibbs Random Field models. A Bayesian context decision rule is adopted and an MRF–GRF stochastic model is used for the classification of multi-echo MR images. This work, and similar work, is directed at the classification of MR images in their entirety; the work described here concentrates on regions of interest within MR images.

To the best knowledge of the authors there has been little work on the application of data-mining techniques to study the corpus callosum. Two studies of note are that of Herskovits and Gerring [10] and that of Machado et al. [15]. Herskovits and Gerring describe a Bayesian-network approach to lesion-deficit analysis (LDA) that identifies associations between elements of the brain including the corpus callosum. Machado et al. apply a visual data-mining method MRI scan data to reveal differences in the callosal morphology between male and female samples.

4. Methodology

In this section an overview of the proposed MR image classification process is presented, further detail is given in the following sections. A block diagram outlining the process is given in Fig. 2 (the directed arcs indicate data flow). The process commences with image segmentation to extract the objects (regions) of interest, in the case of the application under consideration here this is the corpus callosum midsagittal slice of a 3D brain MR image. The next step is to represent the pixel defined corpus callosum segments as quad-trees (one quad-tree per image), using a predefined level of decomposition. For this representation to be effective a registration process must first be applied to the corpus callosum segments. The third step is to apply a weighted graph-mining technique to identify frequently occurring sub-trees within the quad-tree data set. The identified sub-graphs are then the attributes/features used to define each corpus callosum in terms of a *feature vector*. As there may be a great many of these, and it is conjectured that some may be redundant or superfluous, the fourth step is to apply a feature
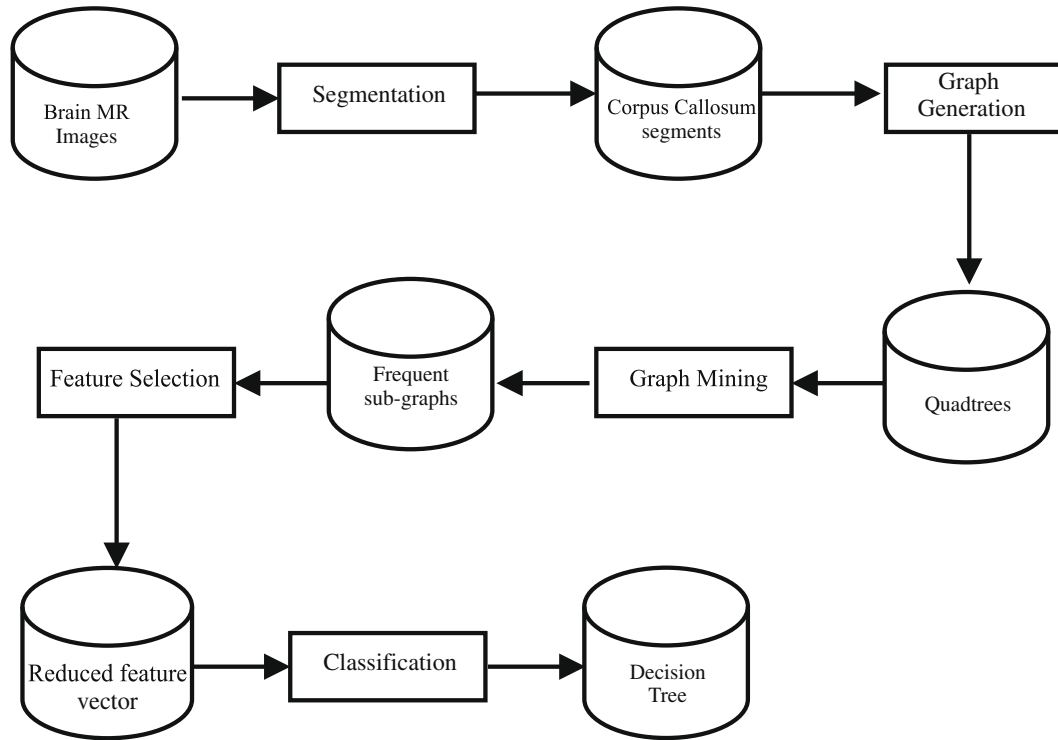
Fig. 2. Proposed methodology.

selection technique (to reduce the number of attributes by removing *irrelevant* features from the features vector). In the final step, the reduced data set is submitted to the classifier generator (a decision-tree algorithm in Fig. 2).

5. Image segmentation

The objective of image segmentation is to partition images into meaningful regions. For the work described here a variation of the *multi-scale spectral image segmentation* algorithm by Cour et al. was used [4]. Cour and Shi's algorithm works on multiple scales of the image in parallel, without iteration, to capture both coarse and fine level details, using the normalized cuts partitioning framework for each image segmentation [19]. The algorithm has been shown to work well in many studies; however it was found that, in the context of the corpus callosum, the algorithm did not work as well as expected. This was because of various reasons: firstly medical MR images have a lot of noises, secondly in MR images the intensity-level distributions between different soft tissues are not widely distributed and moreover the complexity of tissue boundaries cause many pixels to contain mixtures of tissues. For this reason, we need to enhance the contrast of the MR images. A variation of Cour and Shi's algorithm was therefore developed that applied a threshold interval to extract objects with the same intensity values (such as the corpus callosum) during the application of the segmentation. This was found to give a much improved result.

Our proposed variation of Cour and Shi's algorithm is founded on the observation that the corpus callosum, which is located at the centre of the brain, comprises *white matter tissue* (i.e. the pixel represented corpus callosum has high intensity values). Although one can visually recognize the outline of the corpus callosum (Fig. 1), portions of its boundary are indistinct, which can make it difficult to apply segmentation algorithms based on edge information alone. This is particularly the case at the top portion of the corpus callosum, and between the corpus callosum and the Fornix (see Fig. 1). A further problem is that, quite often, intensity varia-

tions within the corpus callosum can be comparable or exceed the difference with the surrounding tissues.

Fig. 3 shows a pixel intensity value histogram of the corpus callosum derived from 30 selected MR images (256 gray levels were used) where the corpus callosum was very well defined and easy to detect using Cour and Shi's algorithm. From the figure it can be seen that:

- The corpus callosum tends to have relatively high intensity values, and
- The distribution of intensity values seems to follow the normal distribution.

The latter is demonstrated in Fig. 4 which shows that the corpus callosum pixel values follow the normal distribution with mean $\overline{X} = 160$ and standard deviation $S = 20$. Fig. 5a shows that with a threshold interval of $\overline{X} \pm S$, the corpus callosum can barely be rec-


Fig. 3. Histogram of the pixel grayscale values of the corpus callosum.
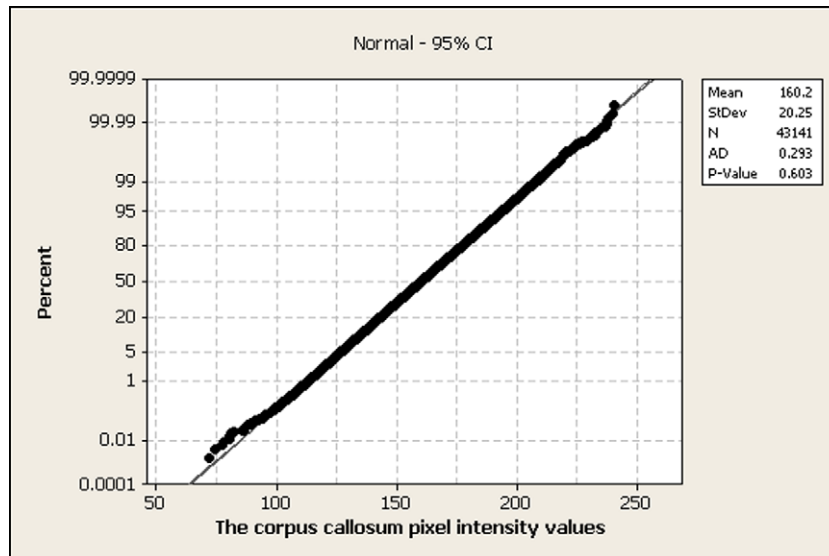
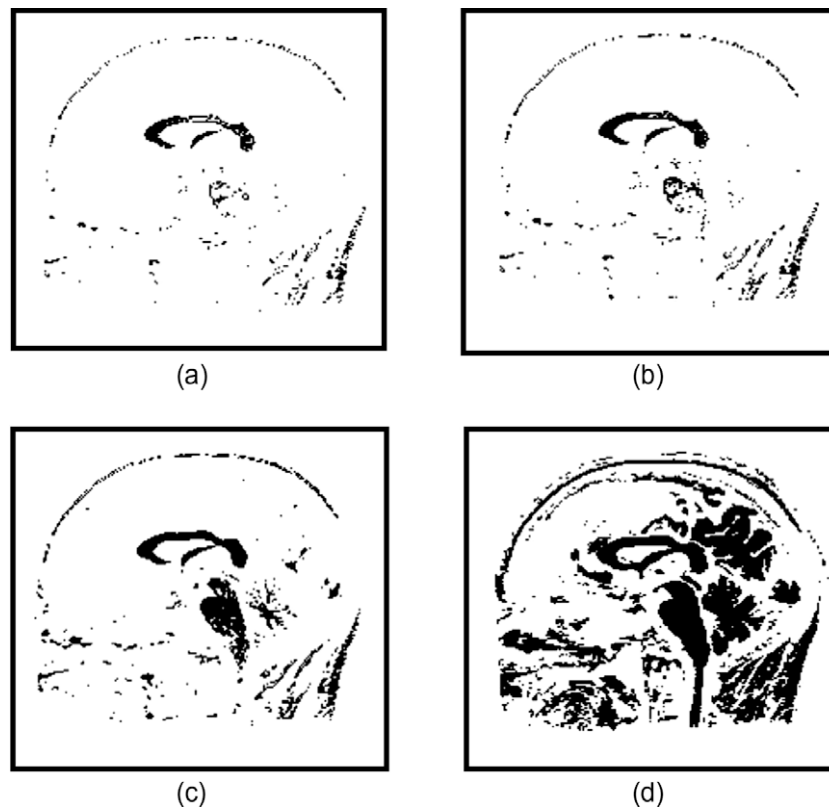Fig. 4. Probability plot of the corpus callosum pixel values.

Fig. 5. Thresholding with various threshold intervals.

ognized. With a threshold interval $\overline{X} \pm 2S$ a relatively distinct callosum shape is evident with a few other non-adjacent structures visible as shown in Fig. 5b. With the threshold interval set at $\overline{X} \pm 3S$, the corpus callosum is clearly defined, although additional non-adjacent structures are also visible (Fig. 5c). In Fig. 5d, the corpus callosum starts to "blur" into the surrounding tissues using a threshold interval wider than $\overline{X} \pm 3S$. The significance here is that although the threshold values may differ depending on individual images, the high intensity property of the corpus callosum can be exploited to yield a segmentation algorithm that is both effective and efficient across the input image set. Therefore the interval $\overline{X} \pm 3S$ was chosen, so as to exclude intensity values outside the interval. This strategy was incorporated into Cour and Shi's algorithm and used to successfully extract the corpus callosum (and other incidental objects with the same intensity values).

6. Registration and tessellation

After segmentation a registration process was applied to the pixel represented images which were then tessellated and stored

in a quad-tree representation. Prior to the registration process some *data cleaning* was also undertaken to remove the "incidental objects" discovered during segmentation. The heuristic used was that the object representing the corpus callosum can be identified using statistical measures (it is the largest object and is locate in roughly the centre of the brain). Having identified the corpus callosum object any other high intensity objects were not considered to be part of the corpus callosum and were removed.

On completion of data cleaning the registration process was undertaken by fitting each identified corpus callosum into a minimum bounding rectangle (MBR). The tessellation process then comprised recursively decomposing the given MBR space into quadrants. Each quadrant was represented by a node in the quad-tree (with colour black or white), with the root of the quad-tree represented the entire MBR space. The tessellation process continued with each of the quadrants "sofar" being again partitioned into sub-quadrants, and so on. The process was terminated when a predefined level of granularity was reached, or a particular sub-quadrant was sufficiently homogeneous (95% black or white).

The advantage of the quad-tree representation was that it maintained information about the relative location and size of groups of pixels (i.e. the shape of the corpus callosum). Nodes nearer the root of the tree represented a larger group of pixels than nodes further away from the root. The use of the quad-tree data structure was considered to be of particular relevance in the context of the corpus callosum as medical opinion suggested that the shape and size of the corpus callosum is of interest [1,6,18,20].

The next stage was to process the quad-tree represented images to identify significant features that occur across the image set (i.e. the set of quad-trees). This was achieved using a bespoke frequent sub-graph mining technique described in the following section.

7. Weighted graph mining

There are various forms of graph/tree mining. In the context of this paper the form of graph mining of interest is transaction graph mining, where the images are represented as a collection of small graphs or trees. Graph mining is broadly concerned with the discovery of interesting patterns in graph or tree data, the interesting patterns are typically frequent sub-graphs. and thus the phrase "frequent sub-graph mining" is often used. The challenge of frequent sub-graph mining is that it necessitates the repeated generation and comparison of candidate sub-graphs, the so called graph isomorphism checking problem. The basic approach is to "grow" candidate sub-graph that occur frequently in the graph set, however this is a computationally expensive process.

Given the quad-tree representation advocated in this paper, nodes nearer the root node are considered to be more significant than others. A weighted frequent sub-graph mining algorithm was therefore developed. The weightings were calculated according to the proximity of individual nodes to the root node in each tree. This weighting concept was built into a variation of the well known gSpan algorithm [21]. The algorithm operates in a depth first search manner, level by level, following a "generate, calculate support, prune" loop. Candidate sub-graphs are pruned if their *support* (frequency with which they occur across the graph set) is below a user defined "support threshold". Note that a lower threshold will identify a greater number of frequent sub-graphs. Space restrictions preclude further detailed discussion of this algorithm here, however, interested readers are referred to Jiang and Coenen [12].

Application of the Weighted gSpan algorithm allowed for the identification of frequent sub-graphs (trees), within the quad-tree represented corpus callosum segments, representing common substructures (features) within the data set. Experimentation indi-

cated that, to capture the necessary level of detail, a low support threshold was required. However this produced a large number of frequent sub-graphs many of which were redundant. A feature selection operation (discussed in the following section) was thus applied to the identified frequent sub-graphs.

8. Feature selection

Feature selection is a well understood process used in data mining for removing irrelevant features from the feature space so as to enhance computational efficiency. Feature selection has attracted a great deal of attention within the data-mining community, especially in the context of classification and prediction where the aim is to identify features that are "strong discriminators". Classic feature selection methods select individual features whose distribution has a strong correlation with individual class labels. Reported methods [22] include: frequency thresholding, information gain, mutual information, Pearson Correlation, and the χ^2 statistic. An acknowledged shortcoming of these methods is that redundant features may be selected due to ignorance of the dependency between features. In order to overcome this "wrapper methods" have been proposed that adopt a classifier model to rate the feature subsets and determine highly discriminative features [13]. However, such wrapper methods can be computational expensive. A straightforward wrapper method is to apply a decision-tree generator to the features [8]. In this context decision-tree algorithms offer the advantage that they inherently estimate the suitability of features for separation of objects representing different classes. Features that are included as "choice points" in the decision tree are thus selected, while all remaining features are discarded.

For the work described here, the well established C4.5 algorithm [16] was used. The objective was to select a sub-set of features (frequent sub-graphs) with strong discriminative power. The results presented in Section 10 below indicate that this strategy was particularly effective.

9. Classifier generation

The final stage in the process was to build the desired classifier using the identified set of discriminating features. Any appropriate classifier generator can be used for this purpose, however in the reported experiments (Section 10) Quinlan's C4.5 algorithm was used. The reasons for selecting a decision-tree algorithm were that they are: (i) a very popular and effective data-mining technique, (ii) non-parametric and (iii) computationally fast [16].

10. Experimentation and evaluation

A number of experiments were undertaken to analyse the performance of the proposed method in terms of classification accuracy. The experiments described here used an MR image set

Table 1
Classification accuracy for corpus callosum segments represented as 4-level quadtrees (maximum of 256 tree nodes).

Levels	Support	F-before	C-before	F-after	C-after
4	20	24549	46.23	15	70.75
4	30	4264	46.23	16	69.81
4	40	1193	50.94	18	68.87
4	50	639	50.94	19	71.7
4	60	262	54.72	16	68.87
4	70	151	50.94	21	61.32
4	80	86	51.89	16	52.83
4	90	54	52.83	17	50.94

Table 2
Classification accuracy for corpus callosum segments represented as 5-level quad-trees (maximum of 1024 tree nodes).

Levels	Support	F-before	C-before	F-after	C-after
5	20	16094	56.6	12	90.57
5	30	4630	51.89	12	83.96
5	40	2100	48.11	10	80.19
5	50	1155	53.77	13	85.85
5	60	637	65.09	14	80.19
5	70	405	55.66	14	81.13
5	80	252	55.66	18	80.19
5	90	130	54.72	17	70.75

Table 3
Classification accuracy for corpus callosum segments represented as 6-level quad-trees (maximum of 4096 tree nodes).

Levels	Support	F-before	C-before	F-after	C-after
6	20	35223	60.38	11	85.85
6	30	9461	66.98	10	95.28
6	40	4059	67.92	12	84.91
6	50	2260	50.00	11	83.96
6	60	1171	60.38	11	90.57
6	70	741	53.77	13	83.96
6	80	433	54.72	13	77.36
6	90	232	43.40	14	75.47

Table 4
Classification accuracy for corpus callosum segments represented as 7-level quad-trees (maximum of 16384 tree nodes).

Levels	Support	F-before	C-before	F-after	C-after
7	20	448683	48.00	13	83.80
7	30	34440	50.00	12	85.85
7	40	11998	45.28	11	89.62
7	50	6402	50.94	10	86.79
7	60	3317	60.38	13	87.74
7	70	2032	53.77	13	75.47
7	80	1117	50.94	13	76.42
7	90	476	52.83	12	78.30

comprising 106 brain images divided equally into two categories (53 images per category): musician and non-musician. It is acknowledged that, in data-mining terms, a data set of 106 records

is small. One of the challenges of the work described here is the limited amount of raw data available for experimentation. This is because of a number of difficulties that are presented when collecting MR image data, namely: (i) they are extremely expensive to produce, and (ii) the time consuming nature of MRI scanning. Consequently the MR image data sets available for research are relatively small, compared to the usual data sets used for the evaluation of data-mining techniques.

As noted above each pixel represented corpus callosum segment is translated, using the described process, into a quad-tree representation. The maximum number of leaf nodes in any quad-tree is given by 4^N where N is the quad-tree level. The quad-tree levels applied in the experiments were 4, 5, 6 and 7 (equating to a maximum number of nodes of 256, 1024, 4096 and 16384, respectively). Note that the level of detail increases with the number of quad-tree levels.

Table. 1 shows the classification results obtained using Ten Cross Validation (TCV) with a quad-tree level of 4. The *F-before* and *F-after* columns indicate the number of features before and after the application of feature selection. The *C-before* and *C-after* columns give the classification accuracy before and after feature selection. The support threshold is the minimum frequency with which a sub-graph must occur across the data set for the sub-graph to be considered "frequent".

Tables 2–4 shows the classification obtained using TCV with quad-tree levels of 5, 6 and 7, respectively. The column headings should be interpreted in the same way as for Table. 1.

Inspection of Tables 1–4 demonstrate that the overall classification accuracy improves after the application of the feature selection strategy. The best classification accuracy of 95.28% (19 correct classifications per 20 images) was obtained using a quad-tree level of 6 coupled with a 30% support threshold. Regardless of the quad-tree level, the trend of the classification accuracy improved as the threshold support decreased. This is because more frequent sub-graphs are identified as can be seen from the *F-before* columns. It is likely that as the support threshold increases, significant sub-graphs are not discovered by the graph mining algorithm. From the tables it can also be observed that accuracy increases as the quad-tree levels are increased, up to level 6 s, and then begins to fall of. It is conjectured that this is because "over fitting" starts to take place as the quad-tree representation starts to get too detailed.
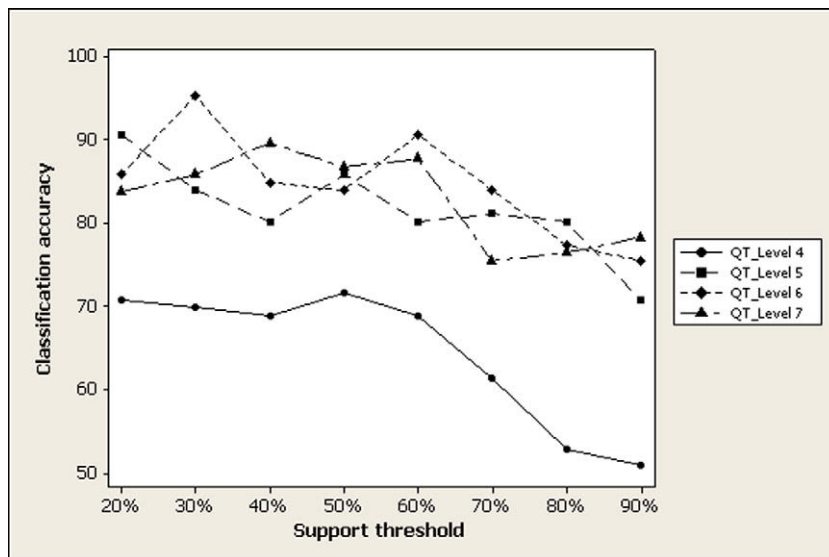

Fig. 6. Classification accuracy for brain MRI images represented in different quad-tree levels.

The results presented in Tables 1–4 are summarised by the graph presented in Fig. 6, which plots classification accuracy (Y axis) against support threshold (X axis) for the sequence of quad-tree levels featured in the experiments.

Image segmentation and graph mining are both computationally expensive processes. The time complexity for the image segmentation was about 2 min per image. For the given data set the graph mining algorithm took some 2 min to process and identify several thousand frequent sub-graphs. The graph mining algorithm took significantly longer to identify much larger numbers of frequent sub-graphs. The worst case was the 448, 683 frequent sub-graphs found when the quad-tree level was set to 7 and the support threshold to 20%. Typically the entire classification process took several minutes to process the 106 image test set.

11. Conclusion

In this paper an approach to MR image classification based on graph mining has been described. The work was directed at a particular MR image classification application, the classification of MR images according to the nature of the corpus callosum featured within these images. However, the approach has more general applicability. Of particular note is the use of feature vectors constructed from frequent sub-graphs identified using a weighted variation of gSpan. The results obtained are extremely encouraging and clearly demonstrate the utility of the approach. Future work will include the application of the process to other brain MR image applications, for example to the automated identification of disorders such as Epilepsy. The research team are also interested in alternative methods of pre-processing MR image data, and mechanism for post-processing of results. In the context of the latter the generation of explanations to support classification of new data is considered to be significant (i.e. the retracing of a classification result to the features in the input image that caused the classification).

References

[1] L. Allen, M. Richey, Y. Chain, R. Gorski, Sex differences in the corpus callosum of the living human being, Journal of Neuroscience 11 (1991) 933–942.
[2] R. Chen, E.H. Herskovits, A Bayesian network classifier with inverse tree structure for voxelwise magnetic resonance image analysis, in: Proceedings of the ACM SIGKDD 2005, 2005, pp 4–12.
[3] J. Chun, R. Greenshields, Classification algorithm for multi-echo magnetic resonance image using Gibbs distributions, in: Proceedings of the third International Conference on Image Analysis Applications and Computer Graphics, LNCS, Springer, 1995, pp. 419–426.
[4] T. Cour, F. Benezit, J. Shi, Spectral segmentation with multiscale graph decomposition, in: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cvpr'05), vol. 2, 2005, pp. 1124–1131.
[5] P. Cowell, A. Kertesz, V. Denenberg, Multiple dimensions of handedness and the human corpus callosum, Neurology 43 (1993) 2353–2357.
[6] C. Davatzikos, M. Vaillant, S. Resnick, J. Prince, S. Letovsky, R. Bryan, A computerized approach for morphological analysis of the corpus callosum, Journal of Computer Assisted Tomography 20 (1996) 88–97.
[7] R. Duara, A. Kushch, K. Gross-Glenn, W. Barker, B. Jallad, S. Pascal, D. Loewenstein, J. Sheldon, M. Rabin, B. Levin, H. Lubs, Neuroanatomic differences between dyslexic and normal readers on magnetic resonance imaging scans, Archives of Neurology 48 (1991) 410–416.
[8] K. Grabczewski, N. Jankowski, Feature selection with decision tree criterion, in: Proceedings of the Fifth International Conference on Hybrid Intelligent Systems (HIS'05), 2005, pp. 212–217.
[9] H. Hampel, S. Teipel, G. Alexander, B. Horwitz, D. Teichberg, M. Schapiro, S. Rapoport, Corpus callosum atrophy is a possible indicator of region and cell type-specific neuronal degeneration in Alzheimer disease, Archives of Neurology 55 (1998) 193–198.
[10] E.H. Herskovits, J.P. Gerring, Application of a data-mining method based on Bayesian networks to lesion-deficit analysis, in: Proceedings of the Neuroimage, 2003, pp. 1664–1673.
[11] G. Hynd, J. Hall, E. Novey, D. Eliopulos, K. Black, J. Gonzalez, J. Edmonds, C. Riccio, M. Cohen, Dyslexia and corpus callosum morphology, Archives of Neurology 52 (1995) 32–38.
[12] C. Jiang, F. Coenen, Graph-based image classification by weighting scheme, in: Proceedings of the AI'2008, Springer, 2008, pp. 63–76.
[13] R. Kohavi, G. John, Wrappers for feature subset selection, Artificial Intelligence 97 (1-2) (1997) 273–324.
[14] I. Lyoo, A. Satlin, C.K. Lee, P. Renshaw, Regional atrophy of the corpus callosum in subjects with Alzheimer's disease and multi-infarct dementia, Psychiatry Research 74 (1997) 63–72.
[15] A. Machado, J. Gee, M. Campos, Visual data mining for modeling prior distributions in morphometry, IEEE Signal Processing Magazine 21 (3) (2004) 20–27.
[16] R. Quinlan, C4.5: A Program for Machine Learning, Morgan Kaufmann, 1993.
[17] S. Ruan, C. Jaggi, J. Xue, J. Fadili, D. Bloyet, Brain tissue classification of magnetic resonance images using partial volume modeling, IEEE Transactions on Medical Imaging 19 (12) (2000) 1179–1187.
[18] D. Salat, A. Ward, J. Kaye, J. Janowsky, Sex differences in the corpus callosum with aging, Journal of Neurobiology of Aging 18 (1997) 191–197.
[19] J. Shi, J. Malik, Normalized cuts and image segmentation, IEEE Pattern Analysis and Machine Intelligence 22 (8) (2000) 888–905.
[20] S. Weis, M. Kimbacher, E. Wenger, A. Neuhold, Morphometric analysis of the corpus callosum using MRI: correlation of measurements with aging in healthy individuals, American Journal of Neuroradiology 14 (1993) 637–645.
[21] X. Yan J. Han, Gspan: graph-based substructure pattern mining, in: ICDM'02: Second IEEE Conference on Data Mining, 2002, pp 721–724.
[22] Y. Yang, J. Pedersen, A comparative study on feature selection in text categorization, in: D.H. Fisher (Ed.), Proceedings of ICML-97, 14th International Conference on Machine Learning, Nashville, US, Morgan Kaufmann Publishers, San Francisco, US, 1997, pp. 412–420.