

Tumour segmentation in breast tissue microarray images using spin-context

Shazia Akbar¹
shaziaakbar@computing.dundee.ac.uk

Telmo Amaral²
telmoamaral@sapo.pt

Stephen J. McKenna¹
stephen@computing.dundee.ac.uk

Alastair Thompson³
a.m.thompson@dundee.ac.uk

Lee Jordan⁴
lee.jordan@nhs.net

¹ School of Computing
University of Dundee

² Institute of Systems and Robotics
University of Coimbra

³ Dundee Cancer Centre
Ninewells Hospital, Dundee

⁴ Pathology
Ninewells Hospital, Dundee

Abstract

A method for automatic segmentation of tumour regions in breast histopathology images is described. It uses auto-context to label pixels based on local image features and contextual label probabilities. We propose spin-context to compute context features that are invariant under image rotation. Quantitative evaluation is reported using spots stained for estrogen receptor. The use of context resulted in improved segmentation.

1 Introduction

Tissue microarrays (TMAs) are used extensively to analyse various types of cancer for molecular and protein markers. Annotation software for breast tissue histopathology images often requires a pathologist to partially annotate some tissue components in order for the software to then analyse a whole mount slide. When applied to TMA spots, typically 0.6mm in diameter, regions are often mislabelled due to lack of context. In this paper a method is reported that probabilistically classifies pixels as non-tumour or tumour (invasive or in-situ carcinoma). We propose a distribution-based auto-context descriptor called spin-context and report results on estrogen-receptor stained TMA images ('spots') by comparison to manual segmentation performed by a pathologist.

2 Related Work

Related work on tumour segmentation includes that of Chomphuwiset *et al.* [1], who used Hough transform-based techniques to detect cell nuclei in liver histopathology images. Chomphuwiset *et al.* also integrate a Bayesian network to combine random forest classification results, obtained from texture features, with context information from nearby nuclei and

regions. Wang *et al.* [9] proposed a method for segmentation of tumour, stroma and inflammatory cells in TMA images using tissue architecture extraction and a tumour texture learning model. Tissue architecture extraction consisted of a stain separation method and an unsupervised multistage entropy-based segmentation method. Tumour texture learning consisted of a Markov random field image segmentation system.

Auto-context has been used for medical image segmentation. Morra *et al.* [5] used Adaboost with auto-context to segment hippocampus in 3D structural MRI. Tu *et al.* [8] used auto-context to segment multiple structures in brain MRI. Tao *et al.* [7] used Gaussian mixtures with simplified auto-context to segment ground glass nodules in 3D lung CT data. Montillo *et al.* [4] segmented structures such as aorta, pelvis, and lungs in 3D CT data, proposing an extension of decision forest classifiers that incorporates semantic context in a manner similar to auto-context. Jurrus *et al.* [2] described an auto-context method to detect membranes in electron micrographs. All of the above were not distribution-based descriptors and, appropriately for those applications, were not invariant under image rotation. To the best of our knowledge, auto-context has not been applied to segmentation of 2D histopathology images.

3 Method

We address locating carcinoma (invasive or in situ) in TMA spots. This is formulated as classifying each location on a grid as being tumour or non-tumour. The image patch around each location is characterised using local features extracted at full resolution, specifically differential invariants up to 2nd order [6] and intensity spin image features [3]. We propose a method called *spin-context* which incorporates context in a rotationally invariant fashion, as the rotation of the tissue in histopathology images is arbitrary. Before describing spin-context we briefly describe auto-context classification [8] and spin image features [3].

Intensity-domain spin image features were proposed for texture representation [3]. A spin feature encodes the distribution of brightness values within a circular support region centred at a location, \mathbf{x}_0 , using a histogram representation that is invariant under image rotation. The contribution of a pixel \mathbf{x} depends on its intensity value, $I(\mathbf{x})$, and its distance from \mathbf{x}_0 , $\|\mathbf{x} - \mathbf{x}_0\|$, as shown in Equation (1). α and β are parameters that determine bin size in the two-dimensional ‘soft’ histogram, H , where each bin is indexed by the radial distance interval, d , and intensity interval, i .

$$H_{(d,i)} = \sum_{\mathbf{x}} \exp\left(-\frac{(\|\mathbf{x} - \mathbf{x}_0\| - d)^2}{2\alpha^2} - \frac{|I(\mathbf{x}) - i|^2}{2\beta^2}\right) \quad (1)$$

The auto-context method, proposed by Tu and Bai in 2009 [8], is an iterative pixel labelling technique, in which some of the label probabilities output at a given iteration are used as contextual data that are concatenated with local image features to form the input vector for the following iteration. This technique is formally described in Algorithm 1, where, for pixel n , t^n is the ground truth value (0 for background, 1 for tumour), \mathbf{z}^n is the local feature vector, \mathbf{c}^n is the context feature vector, and y^n is the predicted probability of being tumour. Classifier models are denoted by ϕ , and \mathcal{N} is the index set for all pixels (training and testing). In a cross-validation experiment (e.g. ten-fold), at the start of each fold p we (1) initialise all pixel probabilities with a uniform prior; and (2) identify the set of training pixel indices. Then, at each auto-context iteration j , we (a) compute a context feature vector from the probability values predicted at the previous iteration, for all pixels; (b) train a classifier

Algorithm 1 Auto-context.

1. initialise $y_{p,0}^n$, for $n \in \mathcal{N}$
2. identify $\mathcal{N}_p^{train} \subset \mathcal{N}$
3. for each iteration j
 - (a) compute $\mathbf{c}_{p,j}^n$ from $y_{p,j-1}^n$, for $n \in \mathcal{N}$
 - (b) train $\phi_{p,j}$ to predict t^n from \mathbf{z}^n and $\mathbf{c}_{p,j}^n$, for $n \in \mathcal{N}_p^{train}$
 - (c) use $\phi_{p,j}$ to predict $y_{p,j}^n$ from \mathbf{z}^n and $\mathbf{c}_{p,j}^n$, for $n \in \mathcal{N}$

to predict truth values from local and context feature vectors, for training pixels; and (c) use the trained classifier to predict new probability values from local and context feature vectors, for all pixels.

3.1 Spin-context

Tu and Bai used a star shaped ‘stencil’ to select context location points around the pixel being classified. The resulting context features from this stencil were not invariant under image rotation. By using an alternative *spin-context*, we can compute rotationally invariant context features for a given grid location from label probability values within a circular support region. Spin-context is extracted analogously to intensity spin features, computing a two-dimensional ‘soft’ histogram reflecting the distribution of probabilities within the support region, with rows representing probability intervals and columns representing radial distance intervals. Figure 1 illustrates spin-context for a given support region.

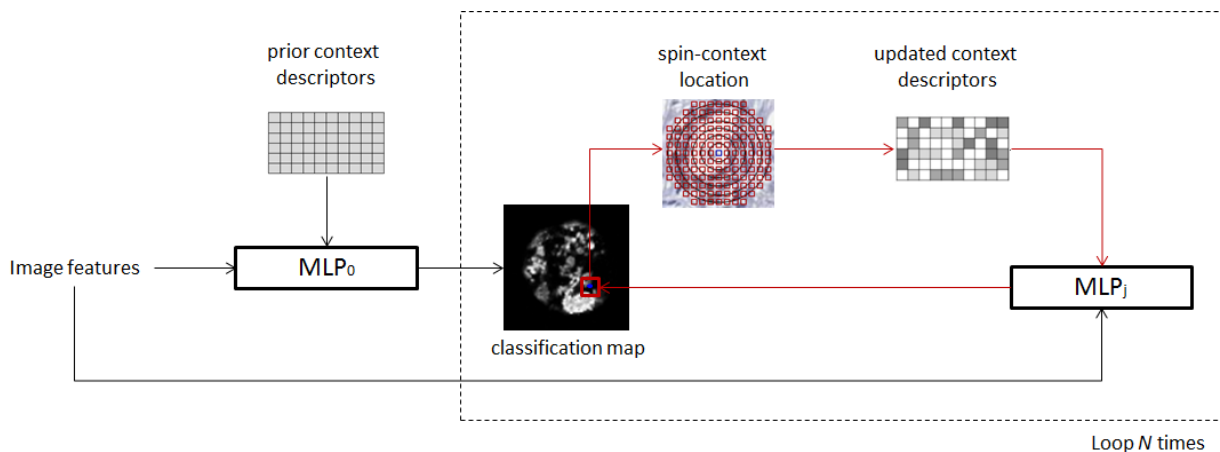


Figure 1: Spin-context constructs context descriptions for a point to be classified (the blue dot) by applying a circular support region centred on that location. The resulting classification map produced by the MLP classifier updates context descriptors iteratively.

The spin-context descriptor allows clutter outside the tissue spot’s boundary to be disregarded while considering only data within the spot region. Figure 2 illustrates the advantage of using spin-context to produce a more accurate representation of context information around the boundaries of the spot. The use of boundary information prior to context extraction allows the contributions of out-of-boundary points towards the two-dimensional spin

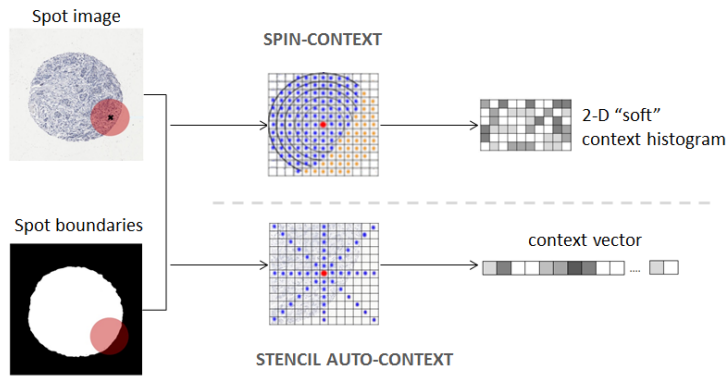


Figure 2: A binary mask is used to ignore the contributions of pixels outside the spot’s boundary towards the spin histogram. Stencil-context, however, corresponds to label probability values at all locations lying on a star-shaped stencil, regardless of spot boundaries.

histogram to be ignored. In doing so, not only is context information accurate for the current iteration but subsequent iterations also reflect on accurate information extracted from the spot region. The stencil-context descriptor, not being distribution-based, does not allow this level of flexibility to be maintained, resulting in background interference.

4 Experiments

TMA spots were subjected to nuclear staining for estrogen receptor (ER). Spot images were 4000×4000 pixels. Data consisted of 64 images, 32 of which contained tumour regions annotated by a highly experienced pathologist and 32 were confirmed to contain only healthy tissue. Example pathologist annotations are shown in Figure 4.




Figure 3: Precision-recall curves for tumour localisation. (a) Effect of six spin-context iterations on MLP classification. (b) Comparison of stencil-context and spin-context for iteration 4 on MLP classification.

Tumour labelling was evaluated using ten-fold cross-validation on the 64 spots. Multi-layer perceptron (MLP) classifiers had five hidden units, a regularisation constant of 0.1 and used scaled conjugate gradient optimisation. Local and context features were computed at points on a 76×76 grid (a grid step of 50 pixels). Differential invariant features were computed at three scales using a Gaussian pyramid and filters with a standard deviation of 8 pixels. Intensity spin local features were computed at two scales (again using a Gaussian

pyramid) with a circular support region with a radius of 50 pixels. Spin-context used a circular support region with a radius of six grid points. We also tried auto-context (non-rotationally invariant context) using a stencil in which neighbouring grid points within a radius of six grid spacings in each of the eight cardinal and inter-cardinal compass directions were used as context. Labellings obtained were compared to ground-truth segmentations provided by the pathologist to compute precision-recall curves.

5 Results

The precision-recall curve in Figure 3(a) displays the results obtained for six spin-context iterations. All six iterations of spin-context improved results, however this improvement was noticeable in the initial four iterations, whereas the subsequent two iterations improved results marginally. Figure 3(b) compares spin-context with stencil based auto-context. At lower recall values spin-context was superior. In both cases MLP classifiers were used and six context iterations were executed.




Figure 4: Tumour location probabilities obtained by spin-context. Shown for each TMA spot are the pathologist’s annotation, the labelling obtained using local image features (iteration 1) and labellings obtained after incorporating label context (iteration 2 and iteration 3). (a) shows invasive cancer labelled largely in agreement with the pathologist. (b) shows healthy tissue. (c) shows one of the worst results obtained.

Figure 4 shows three spots, two containing tumour and one not containing tumour, along with their expert annotations and the outputs of the spin-context method. In Figure 4(a), posterior probabilities within tumour regions increased at each iteration, so that after the final iteration they were above 0.9 for most tumour pixels. In Figure 4(b), non-zero probabilities occur within regions of normal tissue at the first iteration; however, their values decreased after further iterations, so that a binarisation of the labelling would result in an almost entirely empty (i.e. correct) output. Figure 4(c) shows a case of unsuccessful labelling. Initially, the entire spot region was lightly detected and gradually removed from the left after subsequent iterations. This may have been due to the higher density of tissue which is commonly found in tumour regions in training data. This suggests the need for more training examples of this type.

6 Conclusion

A method for tumour segmentation was presented incorporating rotationally invariant context features. It was validated against manual annotations provided by a pathologist. Figure 4 shows how spin-context can be useful to pathology research in locating tumour regions.

Acknowledgements

This work was supported by the Chief Scientist Office (grant no. CZB/4/761); and the Engineering and Physical Sciences Research Council (DTA grant).

References

- [1] P. Chomphuwiset, D. Magee, R. Boyle, and D. Treanor. Context-based classification of cell nuclei and tissue regions in liver histopathology. In *Proc. Medical Image Understanding and Analysis*, 2011.
- [2] E. Jurrus, A.R.C. Paiva, S. Watanabe, J.R. Anderson, B.W. Jones, R.T. Whitaker, E.M. Jorgensen, R.E. Marc, and T. Tasdizen. Detection of neuron membranes in electron microscopy images using a serial neural network architecture. *Medical Image Analysis*, 14(6):770–783, 2010.
- [3] S. Lazebnik, C. Schmid, and J. Ponce. A sparse texture representation using local affine regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1265–1278, 2005.
- [4] A. Montillo, J. Shotton, J. Winn, J. Iglesias, D. Metaxas, and A. Criminisi. Entangled decision forests and their application for semantic segmentation of CT images. In *Information Processing in Medical Imaging*, pages 184–196. Springer, 2011.
- [5] J.H. Morra, Z. Tu, L.G. Apostolova, A.E. Green, C. Avedissian, S.K. Madsen, N. Parikshak, X. Hua, A.W. Toga, C.R. Jack Jr, et al. Validation of a fully automated 3D hippocampal segmentation method using subjects with alzheimer’s disease mild cognitive impairment, and elderly controls. *Neuroimage*, 43(1):59–68, 2008.
- [6] C. Schmid and R. Mohr. Matching by local invariants. Technical Report RR-2644, INRIA, 1995.
- [7] Y. Tao, L. Lu, M. Dewan, A. Chen, J. Corso, J. Xuan, M. Salganicoff, and A. Krishnan. Multi-level ground glass nodule detection and segmentation in CT lung images. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009*, pages 715–723. Springer, 2009.
- [8] Z. Tu and X. Bai. Auto-context and its application to high-level vision tasks and 3D brain image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(10):1744–1757, 2010.
- [9] C. Wang, D. Fennel, I. Paul, K. Savage, and P. Hamilton. Robust automated tumour segmentation on histological and immunohistochemical tissue images. *PLoS ONE*, 6(2):e15818, 2011.