Pattern Recognition Letters 112 (2018) 290–296

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Multi-part segmentation for porcine offal inspection with auto-context and adaptive atlases

Stephen McKenna[a,*], Telmo Amaral[b], Thomas Plötz[b,d], Ilias Kyriazakis[c]

[a] CVIP, School of Science and Engineering, University of Dundee, Dundee DD1 4HN, UK
[b] Open Lab, School of Computer Science, Newcastle University, Newcastle upon Tyne NE1 7RU, UK
[c] Agriculture, School of Natural and Environmental Sciences, Newcastle University, Newcastle upon Tyne NE1 7RU, UK
[d] School of Interactive Computing, College of Computing, Georgia Institute of Technology, Atlanta, GA, USA

ARTICLE INFO

Article history:
Received 18 December 2017
Available online 30 July 2018

MSC:
41A05
41A10
65D05
65D17

Keywords:
Multi-class segmentation
Auto-context
Atlas-based segmentation
Automated inspection

ABSTRACT

Extensions to auto-context segmentation are proposed and applied to segmentation of multiple organs in porcine offal as a component of an envisaged system for post-mortem inspection at abbatoir. In common with multi-part segmentation of many biological objects, challenges include variations in configuration, orientation, shape, and appearance, as well as inter-part occlusion and missing parts. Auto-context uses context information about inferred class labels and can be effective in such settings. Whereas auto-context uses a fixed prior atlas, we describe an adaptive atlas method better suited to represent the multimodal distribution of segmentation maps. We also design integral context features to enhance context representation. These methods are evaluated on a dataset captured at abbatoir and compared to a method based on conditional random fields. Results demonstrate the appropriateness of auto-context and the beneficial effects of the proposed extensions for this application.

1. Introduction

Segmentation of non-rigid biological objects into their constituent parts presents various challenges. Here we address a segmentation task in which parts are organs in body images captured at abbatoir. This constitutes one stage in an envisaged on-site system for screening of pathologies; these are characteristically organ-specific. The spatial arrangement of organs in an image is only weakly constrained and their shape is variable. Furthermore their appearance changes due to factors including cause of pathology, surface contaminants, and specular reflections. There can be limited control over orientation, severe occlusions between parts, and parts may be missing altogether. In this paper we describe adaptations to the auto-context (AC) segmentation algorithm to address such a task. We apply these to segment heart, lungs, diaphragm and liver in porcine offal. The groups of inter-connected organs are called *plucks*, examples of which are shown in Figs. 2 and 3.

Auto-context [3] is an iterative technique that combines *contextual* classification information with local image features. AC is relatively flexible and easy to implement, and has been applied to various biomedical imaging problems [3,4]. The context features used by AC to inform class label inference at a pixel location are posterior class probabilities produced by the previous iteration. These probabilty values are typically sampled at a fixed set of locations relative to the pixel in question. Additionally we design *integral context* features obtained by summing probability values over sets of locations. In the application considered here we argue that sums over rows and sums over the entire foreground are appropriate.

One attractive feature of AC is that a prior *atlas* can be used as a source of contextual data for the initial iteration. Such an atlas can be obtained by averaging rigidly registered manual segmentation maps. However, a single averaged map does not provide a good representation of the multi-modal map distribution that arises as a result of the variations mentioned above, such as occlusions and missing parts. We describe weighted atlas auto-context (WAAC), a method that adapts an atlas representation to be relevant to the current image. This improved atlas is used at the next iteration as an additional source of information together with the label probability maps.

In this paper we combine integrated context and WAAC into one system, extending work reported in conference papers on integral context [1] and WAAC [2]. We report a direct comparison of all of these methods applied to segmentation of multiple organs in pig offal, and we also compare with a conditional random field

* Corresponding author.
E-mail addresses: s.j.z.mckenna@dundee.ac.uk, stephen@computing.dundee.ac.uk (S. McKenna).

https://doi.org/10.1016/j.patrec.2018.07.031

(CRF) method. We evaluate performance in terms of Dice coefficient distributions, pixel-wise classification and quadratic scores.

2. Background

Post-mortem inspection is an important means of ensuring the safety and quality of meat products, enabling the detection of public health hazards and pathologies, and providing useful feedback to farmers. There are moves towards visual-only inspection of pig carcasses and offal without palpation, in order to minimise risk of cross contamination [5,6]. This along with the potential to detect a greater number of pathologies with improved reproducibility than currently possible with manual inspection [7] motivates development of automated visual inspection. Reliable segmentation of organs would constitute an important step towards this goal. In this context even modest improvements in organ segmentation could be significant as regions assigned to the wrong organ may ultimately lead to missed or falsely detected pathologies.

Applications to meat production deal mostly with estimation of proportions of muscle, fat and bone either *in vivo* and post-mortem, sometimes involving segmentation of organs without distinguishing them individually [8,9]. Tao et al. [10] segmented poultry spleen from surrounding viscera as an aid to detection of splenomegaly. Jørgensen et al. [11] segmented gallbladders in chicken livers from images acquired at two visible wavelengths. Stommel et al. [12] envisaged a system for robotic sorting of ovine offal that would involve recognition of multiple organs.

Most literature on segmentation of multiple organs deals with human abdominal organs in CT or MR imaging through techniques including level set optimisation [13], statistical shape models [14], and atlas-based methods [15,16].

Segmentation methods that incorporate spatial context information include those combining inference algorithms based on belief propagation (BP) [17] with models like conditional random fields (CRFs) [18]. Disadvantages common to many such techniques that aim to capture context information include their reliance on fixed spatial configurations with confined neighbourhood relations and complex training procedures.

There is extensive literature dealing with the construction of unbiased atlases for multi-modal data, especially in brain magnetic resonance (MR) image analysis, as in the work of Blezek and Miller [19] and Zikic et al. [20]. Some related work makes use of AC. Kim et al. [21], for example, employed an approach similar to that of Zikic et al. [20], training multiple models, each based on an individual annotated image, so that the probability map of a new image was obtained by averaging maps predicted by individual models. Zhang et al. [22] proposed a hierarchy of AC models whose bottom level is similar to the set of models used by Zikic et al. [20] and Kim et al. [21]. Given a new image, only the best models in the hierarchy are selected to contribute to the final probability map. Model training via these techniques can be computationally expensive.

3. Methods

3.1. Auto-context (AC)

We perform segmentation using methods built around the auto-context (AC) algorithm of Tu and Bai [3]. AC learns to map an input image to a multi-class segmentation map consisting of posterior probabilities over class labels. It iteratively refines the segmentation map by using the label probabilities in a given iteration as a source of contextual data for the following iteration. Label probabilities at a set of locations relative to the location to be classified are concatenated with local image features to form a combined feature vector for training the next classifier.

Let S be a set of m training images X_j together with their label maps Y_j, i.e. $S = \{(Y_j, X_j),\ j = 1..m\}$. At each iteration t we want to train a classifier that outputs the probability distribution $p_{ji}^{(t)}$ over labels $y_{ji} \in \{1..K\}$ for pixel i in image X_j, given the image patch $X_j(N_i)$ from which local features are computed, and label probability map $P_j^{(t-1)}(i)$ (see Eq. (1)).

$$p_{ji}^{(t)} = p(y_{ji}|X_j(N_i),\ P_j^{(t-1)}(i)) \tag{1}$$

In $X_j(N_i)$, N_i denotes all pixels in the image patch, and $P_j^{(t-1)}(i)$ is map $P_j^{(t-1)}$ output for image X_j at the previous iteration $t-1$, but now centred on pixel i.

AC produces a sequence of classifiers, one per iteration. Before the first iteration, all probability maps $P_j^{(0)}$ can be initialised using a prior atlas $Q^{(0)}$, obtained by averaging m training label maps:

$$Q^{(0)} = \frac{1}{m}\sum_j Y_j. \tag{2}$$

At each iteration, given pixel i in image X_j, the actual feature vector input to the classifier is composed of local image features extracted from patch $X_j(N_i)$ concatenated with context features extracted from the re-centered label probability map $P_j^{(t-1)}(i)$. Context features are the probabilities extracted from selected locations on map $P_j^{(t-1)}(i)$, including the central location that corresponds to the current image pixel i. Selected locations are typically defined by a sparse star-shaped “stencil”.

In our implementation of AC, context probabilities for a location are extracted at 90 surrounding stencil points as well as at the location itself. At the first iteration, context consists of the 5 class label probabilities provided by the prior atlas at each of the 91 associated context points; at subsequent iterations, it consists of the label probabilities output by the classifier at the previous iteration, at the same context points. This gives 91 × 5=455 context features per image point. We use multi-layer perceptron classifiers (MLPs); these can be trained to directly estimate posterior probability distributions over the class labels.

3.2. Integral context (IC)

Context data can be enhanced by including integral features, i.e. sums of class label probabilities. We augment the context features described above with two types of integral context features suitable for our application.

The relative positions of organs along the vertical direction vary little from image to image, given that each pluck hangs from a hook and the part of the pluck that is attached to the hook is very consistent across plucks. Thus, given a point on an image, class probabilities averaged over the *row* to which the point belongs provide the classifier on the next iteration with useful information as to which organs are likely to occur at that particular height. For example, a row containing heart is likely to contain also lungs, but very unlikely to contain liver.

In contrast, relative positions of organs along the horizontal direction vary considerably from image to image, given lack of control over the orientation of the pluck around the vertical axis. The heart, in particular, is sometimes fully occluded. Nevertheless, organs are fairly consistent in volume from pig to pig. Thus, class probabilities averaged over the *whole* image reflect the proportions of the pluck covered by each visible organ, and provide the next classifier with useful information on which organs are likely to be visible and how visible they are. For example, a small proportion of visible diaphragm is consistent with a hidden heart and a large proportion of lung.