Automatic Classification of Human Epithelial Type 2 Cell Indirect Immunofluorescence Images using Cell Pyramid Matching

Arnold Wiliem, Conrad Sanderson, Yongkang Wong, Peter Hobson, Rodney F. Minchin, Brian C. Lovell

The University of Queensland, QLD 4072, Australia
NICTA, GPO Box 2434, Brisbane, QLD 4001, Australia
Queensland University of Technology, QLD 4000, Australia
National University of Singapore, Singapore
Sullivan Nicolaides Pathology, Australia

- Published as:

A. Wiliem, C. Sanderson, Y. Wong, P. Hobson, R.F. Minchin, B.C. Lovell.

Automatic Classification of Human Epithelial Type 2 Cell Indirect Immunofluorescence Images using Cell Pyramid Matching.
Pattern Recognition, Vol. 47, No. 7, pp. 2315–2324, 2014.

<http://dx.doi.org/10.1016/j.patcog.2013.10.014>

Abstract

This paper describes a novel system for automatic classification of images obtained from Anti-Nuclear Antibody (ANA) pathology tests on Human Epithelial type 2 (HEp-2) cells using the Indirect Immunofluorescence (IIF) protocol. The IIF protocol on HEp-2 cells has been the hallmark method to identify the presence of ANAs, due to its high sensitivity and the large range of antigens that can be detected. However, it suffers from numerous shortcomings, such as being subjective as well as time and labour intensive. Computer Aided Diagnostic (CAD) systems have been developed to address these problems, which automatically classify a HEp-2 cell image into one of its known patterns (eg. speckled, homogeneous). Most of the existing CAD systems use handpicked features to represent a HEp-2 cell image, which may only work in limited scenarios. We propose a novel automatic cell image classification method termed Cell Pyramid Matching (CPM), which is comprised of regional histograms of visual words coupled with the Multiple Kernel Learning framework. We present a study of several variations of generating histograms and show the efficacy of the system on two publicly available datasets: the ICPR HEp-2 cell classification contest dataset and the SNPHEp-2 dataset.

Keywords: indirect immunofluorescence tests; bag of visual words; HEp-2 cell classification; local features

Figure 1: Examples of strong positive ANA specimens. See Fig. 2 for images of individual cells.

Figure 2: Sample images from ICPRContest dataset [10] and SNPHEp-2 dataset.

1. Introduction

The Anti-Nuclear Antibody (ANA) test is commonly used by clinicians to identify the existence of Connective Tissue Diseases such as Systemic Lupus Erythematosus, Sjögren’s syndrome, and Rheumatoid Arthritis [22]. The hallmark protocol for doing this is through Indirect Immunofluorescence (IIF) on Human Epithelial type 2 (HEp-2) cells [22, 40]. This is due to its high sensitivity and the large range expression of antigens. Examples of specimen images are shown in Figure 1. Despite the advantages, the IIF approach is labour intensive and time consuming [4, 24]. Each ANA specimen must be examined under a fluorescence microscope by at least two scientists. This also renders the test result subjective, and thus has low reproducibility and large variabilities across personnel and laboratories [15, 34].

In recent years, there has been increasing interest in employing image analysis techniques for various routine clinical pathology tests [14, 15, 17]. Results produced by these techniques can be used to support the scientists’ manual/subjective analysis, leading to test results being more reliable and consistent across laboratories [15]. Thus, in order to address the shortcomings of the manual test procedure, one could use Computer Aided Diagnostic (CAD) systems which automatically determine the pattern in the given HEp-2 cell images of a specimen [7, 8, 15, 16, 23, 34, 35, 41].

Table 1 presents notable CAD systems proposed in the literature over the last five years. Most of these systems use carefully handpicked features which may only work in a particular laboratory environment and/or microscope configuration. To address this, several approaches employ a large number of features

Table 1: Existing CAD systems for HEP-2 cell classification.

Approach	Descriptors	Classifier
Perner et al. [23]	Textural	Decision Tree
Hiemann et al. [15]	Structural; textural	LogisticModel Tree
Elbischger et al. [8]	Image statistics; cell shape; textural	Nearest Neighbour (NN)
Hsieh et al. [16]	Image statistics; textural	Learning Vector Quantisation (LVQ)
Soda et al. [34]	Specific set of features (e.g. textural) for each class	Multi Expert System
Cordelli et al. [7]	Image statistics; textural; morphological	AdaBoost
Strandmark et al. [35]	Morphological; image statistics; textural	Random Forest
Ali et al. [2]	Biological-Inspired Descriptor	Boosted k-NN Classifier
Theodorakopoulos et al. [36]	Morphological and texture features	Kernel SVM (KSVM)
Thibault et al. [37]	Morphological and texture features	Linear Regression, Random Forest
Ghosh et al. [13]	Histograms of Oriented Gradients, image statistics and textural	SVM
Li et al. [19]	Textural and image statistics	SVM
Di Cataldo et al. [5]	GLCM and DCT features	SVM
Snell et al. [33]	Texture and shape	Multistage classifier
Ersoy et al. [9]	Local shape measures, gradient and textural	ShareBoost
Wiliem et al. [41]	Bag of visual words with dual-region structure	Nearest Convex Hull Classifier (NCH)

and apply an automated feature selection process [15]. Another approach uses Multi Expert Systems to allow the use of a specifically tailored feature set and classifier for each HEP-2 cell pattern class [34]. Nevertheless, the generalisation ability of these systems is still not guaranteed since these systems were only evaluated on a dataset with a specific setup.

One of the most popular approaches for automatic image classification, here called the bag-of-visual-words (BoW) approach, is to represent an image in terms of a set of visual words, selected from a dictionary that has been trained beforehand [18, 30, 39, 42]. In order to model an image, the BoW approach divides the image into small image patches, followed by patch-level feature extraction. An encoding process is then employed to compute a histogram of occurrences of visual words based on these patches. BoW descriptors often have higher discrimination power compared to the other image descriptors [18, 39, 41, 42]. However, the BoW descriptor has many design options. For example, one needs to determine which patch-level features and encoding technique is most suitable for the task at hand. Our previous study presents an extensive evaluation of popular BoW descriptors in the literature applied to the domain of cell classification [41].

A single histogram of visual words of an image only describes the visual word statistics and does not retain spatial information (ie. where a visual word appears in the image). Previous studies suggest that location and scale information can provide meaningful discriminative information [18, 44]. For example, the locations of visual words describing a wheel could be used to infer the type of vehicle (ie. whether it is a motorcycle, car, or truck). Spatial Pyramid Matching (SPM) was proposed to exploit this information [18]. Specifically, each image is processed as a pyramid of levels, with each level containing non-overlapping regions. The levels differ from each other through an increasing number of regions. Each region is divided into small image patches, and an average histogram of visual words is computed for each region. The histograms from all regions are then fed into a Support Vector Machine (SVM) classifier [32] that uses a specialised kernel.

Our previous work [41] proposed a Dual-Region (DR) structure within the BoW framework, specifically designed for cell images. Each cell image is divided into two regions: (1) an inner area enclosing inside the cell; and (2) an outer area containing only the cell edge. The use of two regions forces the inner and

outer cell content to be modelled and compared separately, leading to higher recognition accuracies than using only one average region (ie. single histogram) for each cell image. An advantage of this approach is that it has lower dimensionality than SPM (ie. approximately 90% less), leading to considerably lower storage requirements. However, a mixing coefficient which indicates relative region importance needs to be empirically determined.

The work presented in this paper extends our previous study by proposing a novel approach termed Cell Pyramid Matching (CPM), which incorporates the positive aspects of the SPM and DR approaches, while omitting their negative aspects. Furthermore, we show that combining the CPM approach with a learning framework known as Multiple Kernel Learning [25] (where several variants of CPM are employed concurrently) leads to state-of-the-art performance on the SNPHEp-2 dataset [41], and is comparable to the state-of-the-art on the ICPRContest dataset [10].

We continue this paper as follows. We first delineate the HEp-2 cell classification task in Section 2. In Section 3 we discuss various forms of BoW descriptors and the proposed CPM approach. Section 4 is devoted to experiments and discussions, followed by the main findings in Section 5.

2. HEp-2 Cell Classification Task

Each positive HEp-2 cell image is represented as a three-tuple (I, M, δ) which consists of: (i) the Fluorescein Isothiocyanate (FITC) image channel I ; (ii) a binary cell mask image M which can be manually defined, or extracted from the (DAPI) image channel [15]; and (iii) the fluorescence intensity $\delta \in \{\text{strong}, \text{weak}\}$ which specifies whether the cell is a strong positive or weak positive. Strong positive images normally have more defined details, while weak positive images are duller.

Let Y be a probe image $Y = (I, M, \delta)$, and ℓ be its class label. Given a gallery set $\mathcal{G} = \{(I, M, \delta)_1^\mathcal{G}, (I, M, \delta)_2^\mathcal{G}, \dots, (I, M, \delta)_m^\mathcal{G}\}$, the task of a classifier $\varphi : Y \times \mathcal{G} \mapsto \hat{\ell}$ is to produce $\hat{\ell}$, where ideally $\hat{\ell} = \ell$.

We consider six HEp-2 cell patterns [40] listed below; example images are shown in Fig. 2.

- (1) *homogeneous*: a uniform diffuse fluorescence covering the entire nucleoplasm sometimes accentuated in the nuclear periphery
- (2) *coarse speckled*: densely distributed, variously sized speckles, generally associated with larger speckles, throughout nucleoplasm of interphase cells; nucleoli are negative
- (3) *fine speckled*: fine speckled staining in a uniform distribution, sometimes very dense so that an almost homogeneous pattern is attained; nucleoli may be positive or negative
- (4) *nucleolar*: brightly clustered larger granules corresponding to decoration of the fibrillar centers of the nucleoli as well as the coiled bodies
- (5) *centromere*: rather uniform discrete speckles located throughout the entire nucleus
- (6) *cytoplasmic*: a very fine dense granular to homogeneous staining or cloudy pattern covering part or the whole cytoplasm

Figure 3: Conceptual diagram of the general approach for obtaining histograms of visual words from cell images. Both the FITC image and its corresponding mask image are divided into small overlapping patches. Patch-level features are extracted from FITC patches. Local histogram from each FITC patch-level features is obtained by an encoder employing a learned dictionary of visual words. Finally, multiple regional descriptors are then computed by pooling the local histograms of FITC patches belonging to each region.

Figure 4: Conceptual diagrams for various spatial structures to obtain multiple region descriptors.

3. Bag of Words Classification Systems

A conceptual illustration of the general approach for obtaining histograms of visual words from HEp-2 cell images is shown in Fig. 3. Each cell image is first resized into a canonical size and then divided into small overlapping patches. The patches are in turn represented by patch-level features. The *local* histogram from each patch is then extracted by using the pre-trained visual word dictionary. The local histograms located inside a region are pooled to compute the overall histogram for the region. Finally, the cell image is represented by a set of regional histograms; examples of regional structures are shown in Fig. 4.

In the following sub-sections, we first describe low-level patch-level features, followed by presenting various methods for local histogram extraction. The regional structures (ie. SPM, DR and the proposed CPM) are discussed afterwards. Finally, we overview a framework known as Multiple Kernel Learning (MKL), which combines information captured from several descriptors.

3.1. Patch-level Feature Extraction

Given a HEp-2 cell image $(\mathbf{I}, \mathbf{M}, \delta)$, both the FITC image $\mathbf{I}$ and mask image $\mathbf{M}$ are divided into small overlapping patches $\mathcal{P}_I = \{\mathbf{p}_{I,1}, \mathbf{p}_{I,2}, \dots, \mathbf{p}_{I,n}\}$ and $\mathcal{P}_M = \{\mathbf{p}_{M,1}, \mathbf{p}_{M,2}, \dots, \mathbf{p}_{M,n}\}$. The division is accomplished in the same manner of both images, resulting in each patch in the FITC image having a corresponding patch in the mask image. Let f be a patch-level feature extraction function $f: \mathbf{p}_I \mapsto \mathbf{x}$, where $\mathbf{x} \in \mathbb{R}^d$. $\mathcal{P}_I$ now can be represented as $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$.

For evaluation purposes, we selected two popular patch-level feature extraction techniques, based on the Scale Invariant Feature Transform (SIFT) and the Discrete Cosine Transform (DCT). The SIFT descriptor is invariant to uniform scaling, orientation and partially invariant to affine distortion and illumination changes [21]. These attributes are advantageous in this classification task as cell images are unaligned and have high within class variabilities. DCT based features proved to be effective for face recognition in video surveillance [30, 42]. By using only the low frequency DCT coefficients (essentially a low-pass filter), each patch representation is relatively robust to small alterations [30]. We follow the extraction procedures for SIFT and DCT as per [20] and [30], respectively.

The dictionary of visual words, denoted as $\mathcal{D}$, is trained from patches extracted in sliding window manner from training cell images. Each histogram encoding method has specific dictionary training procedure.

3.2. Generation of Local Histograms

For each patch-level feature that belongs to region r , $\mathbf{x}_j \in \mathcal{X}_r$, a local histogram $\mathbf{h}_j$ is obtained. In this work we consider three prominent histogram encoding methods: (1) vector quantisation; (2) soft assignment; (3) sparse coding. The methods are elucidated below.

3.2.1. Vector Quantisation (VQ)

Given a set $\mathcal{D}$, the dictionary of visual words, the i -th dimension of local histogram $\mathbf{h}_j$ for patch $\mathbf{x}_j$ is computed via:

$$\mathbf{h}_{ji} = \begin{cases} 1 & \text{if } i = \arg \min_{k \in 1, \dots, |\mathcal{D}|} \text{dist}(\mathbf{x}_j, \mathbf{d}_k) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $\text{dist}(\mathbf{x}_j, \mathbf{d}_k)$ is a distance function between $\mathbf{x}_j$ and $\mathbf{d}_k$, while $\mathbf{d}_k$ is the k -th entry in the dictionary $\mathcal{D}$ and $|\mathcal{D}|$ is the number of elements in $\mathcal{D}$. The dictionary is obtained via the k -means algorithm [3] on training patches, with the resulting cluster centers representing the entries in the dictionary.

The VQ approach is considered as a hard assignment approach since each image patch is only assigned to one of the visual words. Such hard assignment can be sensitive to noise [39].

3.2.2. Soft Assignment (SA)

In comparison to the VQ approach above, a more robust approach is to apply a probabilistic method [30]. Here the visual dictionary $\mathcal{D}$ is a convex mixture of Gaussians. The i -th dimension of the local histogram for $\mathbf{x}_j$ is calculated by:

$$\mathbf{h}_{ji} = \frac{w_i p_i(\mathbf{x}_j)}{\sum_{k=1}^{|\mathcal{D}|} w_k p_k(\mathbf{x}_j)} \quad (2)$$

where $p_i(\mathbf{x})$ is the likelihood of $\mathbf{x}$ according to the i -th component of the visual dictionary $\mathcal{D}$:

$$p_i(\mathbf{x}) = \frac{\exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^T \mathbf{C}_i^{-1}(\mathbf{x} - \boldsymbol{\mu}_i)\right]}{(2\pi)^{\frac{d}{2}} |\mathbf{C}_i|^{\frac{1}{2}}} \quad (3)$$

with w_i , μ_i and C_i representing the weight, mean and diagonal covariance matrix of Gaussian i , respectively. The scalar d represents the dimensionality of $\mathbf{x}$. The dictionary $\mathcal{D}$ is obtained using the Expectation Maximisation algorithm [3] on training patches.

3.2.3. Sparse Coding (SC)

It has been observed that each local histogram produced via Eqn. (2) is sparse in nature (ie. most elements are close to zero) [42]. In other words, the SA approach described in Section 3.2.2 is an *indirect* sparse coding approach. Hence, it is possible to adapt *direct* sparse coding algorithms in order to represent each patch as a combination of dictionary atoms [6, 44], which theoretically can lead to better recognition results [42].

A vector of weights $\boldsymbol{\theta} = [\theta_1, \theta_2, \dots, \theta_{|\mathcal{D}|}]^T$ is computed for each $\mathbf{x}_j$ by solving a minimisation problem that selects a sparse set of dictionary atoms. As the theoretical optimality of the ℓ_1 -norm minimisation solution is guaranteed [38], in this work we used:

$$\min \frac{1}{2} \|\mathbf{D}\boldsymbol{\theta} - \mathbf{x}_j\|_2^2 + \lambda \sum_k \|\theta_k\|_1 \quad (4)$$

where $\|\cdot\|_p$ denotes the ℓ_p -norm and $\mathbf{D} \in \mathbb{R}^{d \times |\mathcal{D}|}$ is a matrix of dictionary atoms. The dictionary $\mathbf{D}$ is trained by using the K-SVD algorithm [1], which is known to be suitable for obtaining reasonable dictionaries in similar cases, ie., using a large number of small image patches [28].

As $\boldsymbol{\theta}$ can have negative values due to the objective function in Eqn. (4), we construct each local histogram using the absolute value of each element in $\boldsymbol{\theta}$ [42]:

$$\mathbf{h}_j = [|\theta_1|, |\theta_2|, \dots, |\theta_{|\mathcal{D}|}|] \quad (5)$$

Compared to both Eqns. (1) and (2), obtaining the histogram using sparse coding is considerably more computationally intensive, due to the need to solve a minimisation problem for each patch.

3.3. Histogram Pooling

Let $\mathcal{X}_r$ be the set of patch-level features belonging to region r . The overall histogram representation for region r is then obtained via averaging local histograms [30, 42]:

$$\mathbf{H}^{[r]} = \frac{1}{|\mathcal{X}_r|} \sum_{j=1}^{|\mathcal{X}_r|} \mathbf{h}_j \quad (6)$$

where $|\mathcal{X}_r|$ is the number of elements in set $\mathcal{X}_r$. In the following subsections, we describe several possible spatial layouts for the regions and the corresponding similarity measures.

3.4. Spatial Structures for Multiple Region Descriptors

In this section we describe two existing spatial structures for using multiple regional descriptors (ie. SPM and DR), followed by the proposed CPM approach. The conceptual diagram for each approach is shown in Fig. 4.

3.4.1. Spatial Pyramid Matching (SPM)

The regions are organised similar to an image pyramid with several levels [18]. At each level l , the image is divided into $(2^l) \times (2^l)$ non-overlapping regions. For instance, at level 0 (ie. the top level), the image is divided into 1×1 region; at level 1, the image is divided into 2×2 regions. In this work, we follow Lazebnik et al. [18] by using a three-level pyramid (ie. levels 0, 1 and 2): 1×1 , 2×2 and 4×4 . In total, there are $1 + 4 + 16 = 21$ regions. The *pyramid match kernel* is used to measure the similarities between two images [18]:

$$\mathcal{K}(\mathbf{H}_1, \mathbf{H}_2) = \frac{1}{2^L} G(\mathbf{H}_1^{[0,r]}, \mathbf{H}_2^{[0,r]}) + \sum_{l=1}^L \frac{1}{2^{L-l+1}} G(\mathbf{H}_1^{[l,r]}, \mathbf{H}_2^{[l,r]}) \quad (7)$$

where $\mathbf{H}_k^{[l,r]}$ is the r -th regional histogram of levels l of the k -th image, while L is the maximum number of levels (ie. $L = 2$). $G(\cdot, \cdot)$, is a histogram intersection kernel, defined as [18]:

$$G(\mathbf{H}_1^{[l,r]}, \mathbf{H}_2^{[l,r]}) = \sum_j \min(\mathbf{H}_{1,j}^{[l,r]}, \mathbf{H}_{2,j}^{[l,r]}) \quad (8)$$

where $\mathbf{H}_{k,j}^{[l,r]}$ is j -th dimension of a regional histogram for level l and region r of image k .

3.4.2. Dual Region (DR)

Each cell is divided into an inner region, which covers the cell content, and an outer region, which contains information related to cell edges and shape [41]. To this end, each patch is first classified as either belonging to the inner or outer region by inspecting its corresponding mask patch. More specifically, let $\mathcal{X} = \mathcal{X}^{[o]} \cup \mathcal{X}^{[i]}$, with $\mathcal{X}^{[o]}$ representing the set of outer patches, and $\mathcal{X}^{[i]}$ the set of inner patches. The classification of patch $\mathbf{p}$ into a region is done via:

$$\mathbf{p}_l \in \begin{cases} \mathcal{X}^{[o]} & \text{if } \text{fg}(\mathbf{p}_M) \in [\tau_1, \tau_2) \\ \mathcal{X}^{[i]} & \text{if } \text{fg}(\mathbf{p}_M) \in [\tau_2, 1] \end{cases} \quad (9)$$

where $\mathbf{p}_M$ is the corresponding mask patch; $\text{fg}(\mathbf{p}_M) \in [0, 1]$ computes the normalised occupation count of foreground pixels from mask patch $\mathbf{p}_M$; τ_1 is the minimum foreground pixel occupation of a patch belonging to the outer region; τ_2 is the minimum pixel occupation of a patch belonging to the inner region. Note that the size of the inner and outer regions is indirectly determined via Eqn. (9). Based on preliminary experiments, we have found that $\tau_1 = 0.3$ and $\tau_2 = 0.8$ provide good results. Unlike SPM, there are only two regional histograms required to represent a cell image. As such, the DR descriptor is $(21 - 2)/21 \approx 90\%$ smaller than SPM.

The similarity between two images is defined via:

$$\mathcal{K}(\mathbf{H}_1, \mathbf{H}_2) = \exp[-\text{dist}(\mathbf{H}_1, \mathbf{H}_2)] \quad (10)$$

Adapting [30], $\text{dist}(\mathbf{H}_1, \mathbf{H}_2)$ is defined by:

$$\text{dist}(\mathbf{H}_1, \mathbf{H}_2) = \alpha^{[i]} \|\mathbf{H}_1^{[i]} - \mathbf{H}_2^{[i]}\|_1 + \alpha^{[o]} \|\mathbf{H}_1^{[o]} - \mathbf{H}_2^{[o]}\|_1 \quad (11)$$

where $\mathbf{H}_k^{[i]}$ and $\mathbf{H}_k^{[o]}$ are the inner and outer region histograms of image k , respectively; $\alpha^{[i]}$ and $\alpha^{[o]}$ are positive mixing parameters which define the importance of information contained for each region, under the constraint of $\alpha^{[i]} + \alpha^{[o]} = 1$. A possible drawback of the DR approach is that determining good settings for the τ_1 , τ_2 and $\alpha^{[i]}$ parameters is currently a time consuming procedure, where a heuristic or grid-based search is used. Furthermore, not all valid settings in such a search might be evaluated, which can lead to sub-optimal discrimination performance.

3.4.3. Cell Pyramid Matching (CPM)

The proposed CPM approach combines the advantages of both SPM and DR structures. It adapts the idea of using a pyramid structure from SPM as well as the inner and outer regions from DR. Unlike SPM, CPM has only two levels: level 0 which comprises the whole cell region, and level 1 which comprises of inner and outer regions. The advantages of this combination are two fold: (1) the CPM descriptor only requires 3 histograms to represent a cell image, and is hence $(21 - 3)/21 \approx 85\%$ smaller than SPM; (2) as the CPM follows the SPM construct, it employs the pyramid match kernel, which eliminates the mixing parameters in DR.

3.5. Multiple Kernel Learning

Fusing information provided by various image descriptors and spatial structures (each with a dedicated kernel, as shown above) may improve discrimination ability, if the given descriptors are at least partially capturing differing information. To that end we have elected to use the Multiple Kernel Learning (MKL) framework, which aims to learn the optimum mixing of various kernels [25]. Let $\{\mathbf{x}_i, \mathbf{y}_i\} \in \mathcal{G}$ be the training set, where $\mathbf{x}_i$ is a feature vector and $\mathbf{y}_i \in \{-1, +1\}$ is the corresponding groundtruth label¹.

The MKL classifier is an extended form of the SVM classifier, defined as:

$$\varphi(\mathbf{q}) = \sum_{k=1}^n \beta_k \mathcal{K}(\mathbf{q}, \mathbf{x}_k) + \mathbf{b} \quad (12)$$

where $\mathbf{q}$ is a query point, $\mathbf{x}_k \in \mathcal{G}$ is the k -th training point, β_k is the “importance” weight of the k -th training point, $\mathbf{b}$ is the bias term, and $\mathcal{K}(\cdot, \cdot)$ is a combination kernel defined as:

$$\mathcal{K}(\mathbf{a}, \mathbf{c}) = \sum_{m=1}^M w_m \mathcal{K}_m(\mathbf{a}, \mathbf{c}) \quad (13)$$

where $\mathcal{K}_m(\cdot, \cdot)$ is the m -th kernel, with w_m its corresponding mixing weight, under the constraints of $w_m \geq 0$ and $\sum w_m = 1$. Without losing generality, $\mathcal{K}_m(\cdot, \cdot)$ can be the kernel defined in Eqn. (7) or (10).

In the MKL learning scheme, the importance weights and kernel mixing weights are learned together. In this work we employ the SimpleMKL method for learning [25], which employs a convex and smooth objective function.

4. Experiments and Results

In this section we first compare the performance of six variants of the BoW descriptor, where each of the two low-level feature extraction techniques (SIFT and DCT) is coupled with three possible methods for generating the histograms of visual words (VQ, SA, and SC). The six variants are used within the framework of the DR, SPM and CPM spatial structures. We then show that by fusing the two approaches (DCT-SA CPM and DCT-VQ CPM) via the MKL framework leads to an increase in recognition rates. Finally, we compare the MKL based system against three recently proposed systems in the literature. The various systems were implemented with the aid of the Armadillo C++ library [29].

¹ Here we have presented a binary classification problem. However, it can be easily generalised into a multi-class problem [25].

4.1. Datasets: ICPRContest and SNP HEP-2

For the experiments we used two publicly available datasets, briefly described below, in order to evaluate applicability of the various systems to differing assays and microscope parameters.

The ICPR HEP-2 Cell Classification Contest (ICPRContest) Dataset [11] contains 1,457 cells extracted from 28 specimen images². It contains six patterns: centromere, coarse speckled, cytoplasmic, fine speckled, homogeneous, and nucleolar. Each specimen image was acquired by means of fluorescence microscope (40-fold magnification) coupled with 50W mercury vapour lamp and with a CCD camera. The cell image masks were hand labelled. See Fig. 2 for examples. We followed the ICPR contest evaluation protocol for this dataset which only has one pair of train and test sets.

The SNP HEP-2 Cell (SNPHEP-2) Dataset³ [41] was obtained between January and February 2012 at Sullivan Nicolaides Pathology laboratory, Australia. This dataset has five patterns: centromere, coarse speckled, fine speckled, homogeneous and nucleolar. The 18-well slide of HEP-2000 IIF assay from Immuno Concepts N.A. Ltd. with screening dilution 1:80 was used to prepare 40 specimens. Each specimen image was captured using a monochrome high dynamic range cooled microscopy camera, which was fitted on a microscope with a plan-Apochromat 20x/0.8 objective lens and an LED illumination source. 4',6-diamidino-2-phenylindole (DAPI) image channel was used to automatically extract the cell image masks.

There are 1,884 cell images extracted from 40 specimen images. The specimen images are divided into training and testing sets with 20 images each (4 images for each pattern). In total there are 905 and 979 cell images extracted for training and testing. Five-fold validations of training and testing were created by randomly selecting the training and test images. Both training and testing in each fold contain around 900 cell images (approx. 450 cell images each). Examples are shown in Fig. 2.

Due to possible varying filtering effects caused by image capture equipment, tuning, operator bias, and/or environmental conditions (all of which can result in low-pass filtering), cell images with the same pattern can simply differ due to gross mismatches in frequency spectra. In turn this can lead to a degradation in recognition accuracy [43]. To counteract this undesirable effect, and to ensure a canonical image size is used, images from both datasets were downsampled by two to approximately 64×64 pixels.

4.2. Combinations of Local Features, Histogram Generation and Spatial Structures

We follow Lazebnik et al. [18] and Wiliem et al. [41] for SPM and DR implementations, respectively. The SVM classifier is used in all cases, with the kernels specified in Eqns. (7) and (10) for the SPM and DR methods, respectively. As noted in Section 3.4.3, a form of Eqn. (7) is used as the SVM kernel for the CPM method.

As there are three histogram encoding methods (ie. VQ, SA and SC) and two patch-level features (ie. SIFT and DCT), there are six variants of the BoW descriptor. For clarity, each variant is denoted by: *[patch-level features]-[histogram encoding method]*. For example, the variant using DCT as its patch-level features and VQ as its encoding method is called DCT-VQ.

The results, presented in Table 2, indicate that in most cases the proposed CPM system obtains the best performance, suggesting that it is taking advantage of both the specialised spatial layout for cells inherited from the DR approach, and the pyramid match kernel inherited from the SPM approach. The results also

² It is assumed that the cell images have been extracted from specimen images either via a manual or automated approach such as background subtraction [26, 27].

³ The SNPHEP-2 dataset is available for download at <http://staff.itee.uq.edu.au/lovell/snphep2/>

Table 2: Performance comparison of BoW descriptor variants on the ICPRContest and SNPHEp-2 datasets, using various spatial configurations (DR, SPM, CPM). The scores for SNPHEp-2 dataset shown as average correct classification rate. DR = dual region; SPM = Spatial Matching Pyramid; CPM = Cell Pyramid Matching.

Descriptor Variant	ICPRContest			SNPHEp-2		
	DR	SPM	CPM	DR	SPM	CPM
DCT-SA	64.9	64.3	65.9	79.5	80.3	81.2
DCT-VQ	54.5	57.1	61.2	80.7	77.9	80.8
DCT-SC	52.6	57.9	57.2	71.0	70.5	73.5
SIFT-SA	51.6	57.5	47.8	71.6	69.7	73.2
SIFT-VQ	55.6	53.8	59.0	64.9	74.4	75.0
SIFT-SC	60.8	59.9	62.1	76.2	73.6	76.3

Table 3: Performance of various systems fused via the MKL framework on the ICPRContest and SNPHEp-2 datasets. The “overall” column is the mean performance across the two datasets.

System Mixture	ICPRContest	SNPHEp-2	overall
(DCT-SA CPM) + (DCT-VQ CPM) + (SIFT-SC CPM)	66.9	82.5	74.70
(DCT-SA CPM) + (DCT-VQ CPM)	67.4	82.4	74.90
(DCT-SA CPM) + (SIFT-SC CPM)	66.3	81.2	73.75
(DCT-VQ CPM) + (SIFT-SC CPM)	64.0	79.7	71.85

show that in most cases the use of DCT based patch-level feature extraction leads to better performance than using SIFT based feature extraction. We conjecture that DCT obtains better performance as the SIFT descriptor needs a larger spatial support and is hence more likely to be affected by image deformations. Specifically, SIFT divides a given image patch into 4×4 subregions in which each has 4×4 bins, followed by extracting gradient information from each subregion [21]. Therefore, SIFT needs a spatial support of at least 16×16 pixels, which is relatively large when compared to the canonical cell image size of 64×64 . In contrast, standard DCT requires a much smaller spatial support of 8×8 pixels, making it less susceptible to image deformations.

4.3. Fusion via Multiple Kernel Learning

Based on the results obtained in the previous experiment, we have selected the overall top three systems (DCT-SA CPM, DCT-VQ CPM, SIFT-SC CPM) and evaluated fusing them via the MKL framework. The results for various mixtures of the three systems are shown in Table 3.

By using the mixture that obtains the best overall performance across both datasets, ie. DCT-SA CPM and DCT-VQ CPM, the recognition rate improves from 65.9% to 67.4% on the ICPRContest dataset, and from 81.2% to 82.4% on the SNPHEp-2 dataset.

Note that while it is possible to fuse information from more systems, there is no guarantee that this will always lead to better performance [12, 25, 31]. In further experiments (not shown here) we have found that combining more systems can decrease performance. As the main aim was to show the possible advantage of using MKL, we leave a detailed study for future work.

Figure 5: Performance comparison of various systems on the ICPRContest and SNPHEp-2 datasets.

4.4. Comparative Evaluation of Systems

In this section we compare the MKL based approach (where information from DCT-SA CPM and DCT-VQ CPM is fused) against three recently proposed systems in Wiliem et al. [41], Cordelli et al. [7] and Strandmark et al. [35].

The system in [41], denoted *Wiliem*, is based on the DCT-SA descriptor with the DR spatial structure and Nearest Convex Hull classifier. We denote the system in [35] by *Strandmark*, and used the code provided by the authors. The system employs various image statistics (eg., mean, standard deviation) and morphological features (eg., number of objects, area). The random forest classifier is used.

We implemented the best reported descriptor in [7], denoted by *Cordelli*, which is comprised of features such as image energy, mean and entropy, calculated from intensity and LBP channels. The LBP channel is computed by computing the local pattern code for each pixel in the intensity channel. We selected Logistic Boosting (LogitBoost) as the classifier instead of AdaBoost as the former obtained better performance.

The results are presented in Fig. 5. On the ICPRContest dataset, the Cordelli and Strandmark systems obtain comparable performance. However, the performance of Cordelli is considerably lower than Strandmark on the SNPHEp-2 dataset, indicating that the Cordelli system is not able to generalise to various recording conditions. The Wiliem system obtains better performance than Cordelli and Strandmark on both datasets, with a considerable advantage over Strandmark on SNPHEp-2. However, the proposed MKL based system obtains the best performance on both datasets, with a marked increase over Wiliem on the ICPRContest dataset.

4.5. Cell Level and Image Level Performance on the ICPRContest Dataset

Using the proposed MKL-based system from Section 4.4, Fig. 6 shows the confusion matrix for the classification results on the ICPRContest dataset. We also present the image level classification in Fig. 7. In image level classification we simply determine the label of an image based on the most frequent cell pattern. In this setting, the MKL-based system achieves 71.4%.

We also report Leave-One-Out validation results for ICPRContest in Table 4 as well as Figs. 8 and 9. In this setting, the validation constructs 28 splits of train and test images, where for each split cells belonging to a particular specimen image are used as the test images, and the rest as training images.

Actual class	Centromere	.70	.00	.09	.15	.05	.00
	Homogeneous	.03	.83	.02	.02	.07	.04
	Nucleolar	.12	.04	.65	.04	.10	.05
	Coarse speckled	.17	.02	.01	.55	.20	.05
	Fine speckled	.24	.11	.00	.21	.39	.05
	Cytoplasmic	.00	.00	.06	.04	.00	.90
		Predicted class					
		Centromere	Homogeneous	Nucleolar	Coarse speckled	Fine speckled	Cytoplasmic

Figure 6: Cell level confusion matrix of the proposed MKL-based system on the ICPRContest dataset. Each row and column represents instances of an actual class and predicted class, respectively. The elements in every row are normalised to one. The average accuracy is 67.4%. Note that as the number of instances in each actual class is different, the average accuracy cannot be obtained by averaging the diagonal elements of the matrix.

Actual class	Centromere	.67	.00	.00	.33	.00	.00
	Homogeneous	.00	1.0	.00	.00	.00	.00
	Nucleolar	.00	.00	1.0	.00	.00	.00
	Coarse speckled	.33	.00	.00	.33	.33	.00
	Fine speckled	.50	.00	.00	.00	.50	.00
	Cytoplasmic	.00	.00	.00	.00	.00	1.0
		Centromere	Homogeneous	Nucleolar	Coarse speckled	Fine speckled	Cytoplasmic
		Predicted class					

Figure 7: Image level confusion matrix of the proposed MKL-based system on the ICPRContest dataset. The average accuracy is 71.4%.

Actual class	Centromere	.70	.06	.04	.12	.08	.00
	Homogeneous	.03	.65	.04	.03	.21	.04
	Nucleolar	.22	.07	.61	.06	.02	.02
	Coarse speckled	.31	.05	.03	.32	.27	.02
	Fine speckled	.06	.26	.05	.29	.29	.04
	Cytoplasmic	.10	.01	.03	.01	.06	.79
		Centromere	Homogeneous	Nucleolar	Coarse speckled	Fine speckled	Cytoplasmic
		Predicted class					

Figure 8: Cell level confusion matrix of the proposed MKL-based system on the ICPRContest dataset using Leave-One-Out validation protocol. The average accuracy is 56.8%.

Figure 9: Image level confusion matrix of the proposed MKL-based system on the ICPRContest dataset using Leave-One-Out validation protocol. The average accuracy is 64.3%.

Table 4: Cell level classification performance for each cell image. Ce = Centromere; Ho = Homogeneous; Nu = Nucleolar; Co = Coarse speckled; Fi = Fine speckled; Cy = Cytoplasmic.

Specimen image number	True class	Cells assigned to each class											
		(in absolute number)						(in %)					
		Ce	Ho	Nu	Co	Fi	Cy	Ce	Ho	Nu	Co	Fi	Cy
1	Homogeneous	0	43	2	0	16	0	0.0	70.5	3.3	0.0	26.2	0.0
2	Fine speckled	4	8	10	19	6	1	8.3	16.7	20.8	39.6	12.5	2.1
3	Centromere	87	0	1	1	0	0	97.8	0.0	1.1	1.1	0.0	0.0
4	Nucleolar	23	3	32	0	5	3	34.8	4.5	48.5	0.0	7.6	4.5
5	Homogeneous	8	26	7	0	5	1	17.0	55.3	14.9	0.0	10.6	2.1
6	Coarse speckled	44	9	0	15	0	0	64.7	13.2	0.0	22.1	0.0	0.0
7	Centromere	54	1	1	0	0	0	96.4	1.8	1.8	0.0	0.0	0.0
8	Nucleolar	25	0	30	0	0	1	44.6	0.0	53.6	0.0	0.0	1.8
9	Fine speckled	0	20	1	2	23	0	0.0	43.5	2.2	4.3	50.0	0.0
10	Coarse speckled	16	0	0	7	10	0	48.5	0.0	0.0	21.2	30.3	0.0
11	Coarse speckled	1	0	1	7	32	0	2.4	0.0	2.4	17.1	78.0	0.0
12	Coarse speckled	4	0	5	35	0	5	8.2	0.0	10.2	71.4	0.0	10.2
13	Centromere	13	0	3	30	0	0	28.3	0.0	6.5	65.2	0.0	0.0
14	Centromere	4	19	1	9	30	0	6.3	30.2	1.6	14.3	47.6	0.0
15	Fine speckled	9	7	0	36	3	8	14.3	11.1	0.0	57.1	4.8	12.7
16	Centromere	36	0	2	0	0	0	94.7	0.0	5.3	0.0	0.0	0.0
17	Coarse speckled	0	2	0	3	14	0	0.0	10.5	0.0	15.8	73.7	0.0
18	Homogeneous	0	19	0	3	14	0	0.0	45.2	0.0	7.1	47.6	0.0
19	Centromere	56	0	7	2	0	0	86.2	0.0	10.8	3.1	0.0	0.0
20	Nucleolar	3	2	41	0	0	0	6.5	4.3	89.1	0.0	0.0	0.0
21	Homogeneous	1	50	3	0	7	0	1.6	82.0	4.9	0.0	11.5	0.0
22	Homogeneous	1	78	2	7	20	11	0.8	65.5	1.7	5.9	16.8	9.2
23	Fine speckled	0	20	0	3	28	0	0.0	39.2	0.0	5.9	54.9	0.0
24	Nucleolar	1	11	45	15	1	0	1.4	15.1	61.6	20.5	1.4	0.0
25	Cytoplasmic	10	0	0	0	6	8	41.7	0.0	0.0	0.0	25.0	33.3
26	Cytoplasmic	0	0	1	0	0	33	0.0	0.0	2.9	0.0	0.0	97.1
27	Cytoplasmic	1	1	0	1	0	35	2.6	2.6	0.0	2.6	0.0	92.1
28	Cytoplasmic	0	0	2	0	1	10	0.0	0.0	15.4	0.0	7.7	76.9

5. Conclusions

In this paper we have proposed a cell classification system comprised of a Cell Pyramid Matching (CPM) descriptor combined with Multiple Kernel Learning. The inspiration for the proposed CPM approach is drawn from Spatial Pyramid Matching (SPM) and Dual Region (DR) descriptors. The major contributions of this study are: (1) proposing a more effective adapted version of SPM for cell images; (2) an extensive study on Bag-of-Words descriptor variants and various spatial structures.

We evaluated numerous configurations on two publicly available datasets: ICPR HEp-2 cell classification contest dataset and the new SNPHEp-2 dataset. We found that DCT patch-level features in conjunction with soft-assignment/probabilistic encoding of histograms lead to the highest discrimination performance. We also found that the proposed CPM spatial layout is more effective than SPM and DR structures. The proposed CPM also has an advantage of not having heuristic parameters and leads to a much shorter descriptor length. The experiments show that the proposed system consistently delivered high performance and is more robust than three recent CAD systems presented in [7, 35, 41].

Acknowledgements

This research was partly funded by Sullivan Nicolaides Pathology, Australia and the Australian Research Council (ARC) Linkage Projects Grant LP130100230. NICTA is funded by the Australian Government as represented by the *Department of Broadband, Communications and the Digital Economy*, as well as the Australian Research Council through the *ICT Centre of Excellence* program. NUS-ZJU Sensor-Enhanced Social Media (SeSaMe) Centre is supported by the Singapore National Research Foundation under its International Research Centre @ Singapore Funding Initiative and administered by the Interactive Digital Media Programme Office.

References

- [1] M. Aharon, M. Elad, and A. Bruckstein. K-SVD: an algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Processing*, 54(11):4311–4322, 2006.
- [2] W. Bel Haj Ali, P. Piro, D. Giampaglia, T. Pourcher, and M. Barlaud. Biological cell classification using bio-inspired descriptor in a boosting k-NN framework. In *International Conference on Pattern Recognition (ICPR)*, 2012.
- [3] C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [4] N. Bizzaro, R. Tozzoli, E. Tonutti, A. Piazza, F. Manoni, A. Ghirardello, D. Bassetti, D. Villalta, M. Pradella, and P. Rizzotti. Variability between methods to determine ANA, anti-dsDNA and anti-ENA autoantibodies: a collaborative study with the biomedical industry. *Journal of Immunological Methods*, 219(1-2):99–107, 1998.
- [5] S. D. Cataldo, A. Bottino, E. Ficarra, and E. Macii. Applying textural features to the classification of hep-2 cell patterns in iif images. In *International Conference on Pattern Recognition (ICPR)*, 2012.
- [6] A. Coates and A. Y. Ng. The importance of encoding versus training with sparse coding and vector quantization. In *Int. Conf. Machine Learning*, 2011.
- [7] E. Cordelli and P. Soda. Color to grayscale staining pattern representation in IIF. In *International Symposium on Computer-Based Medical Systems*, pages 1–6, 2011.
- [8] P. Elbischger, S. Geerts, K. Sander, G. Ziervogel-Lukas, and P. Sinah. Algorithmic framework for HEp-2 fluorescence pattern classification to aid auto-immune diseases diagnosis. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 562–565, 2009.

- [9] I. Ersoy, F. Bunyak, J. Peng, and K. Palaniappan. HEp-2 cell classification in IIF images using shareboost. In *International Conference on Pattern Recognition (ICPR)*, 2012.
- [10] P. Foggia, G. Percannella, P. Soda, and M. Vento. Early experiences in mitotic cells recognition on HEp-2 slides. In *International Symposium on Computer-Based Medical Systems*, pages 38–43, 2010.
- [11] P. Foggia, G. Percannella, P. Soda, and M. Vento. Benchmarking HEp-2 cells classification methods. *IEEE Transactions on Medical Imaging*, 32(10):1878–1889, 2013.
- [12] P. Gehler and S. Nowozin. On feature combination for multiclass object classification. In *IEEE International Conference on Computer Vision*, 2009.
- [13] S. Ghosh and V. Chaudhary. Feature analysis for automatic classification of hep-2 florescence patterns: Computer-aided diagnosis of auto-immune diseases. In *International Conference on Pattern Recognition (ICPR)*, 2012.
- [14] M. N. Gurcan, L. E. Boucheron, A. Can, A. Madabhushi, N. M. Rajpoot, and B. Yener. Histopathological image analysis: A review. *IEEE Reviews in Biomedical Engineering*, 2:147–171, 2009.
- [15] R. Hiemann, T. Bttner, T. Krieger, D. Roggenbuck, U. Sack, and K. Conrad. Challenges of automated screening and differentiation of non-organ specific autoantibodies on HEp-2 cells. *Autoimmunity Reviews*, 9(1):17–22, 2009.
- [16] T. Hsieh, Y. Huang, C. Chung, and Y. Huang. HEp-2 cell classification in indirect immunofluorescence images. In *Int. Conf. Information, Communications and Signal Processing*, pages 1–4, 2009.
- [17] R. Khutlang, S. Krishnan, R. Dendere, A. Whitelaw, K. Veropoulos, G. Learmonth, and T. S. Douglas. Classification of mycobacterium tuberculosis in images of ZN-stained sputum smears. *IEEE Trans. Information Technology in Biomedicine*, 14(4):949–957, 2010.
- [18] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2169–2178, 2006.
- [19] K. Li and J. Yin. Multiclass boosting svm using different texture features in hep-2 cell staining pattern classification. In *International Conference on Pattern Recognition (ICPR)*, 2012.
- [20] C. Liu, J. Yuen, and A. Torralba. SIFT flow: Dense correspondence across scenes and its applications. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 33(5):978–994, 2011.
- [21] D. G. Lowe. Distinctive image features from Scale-Invariant keypoints. *Int. J. Comput. Vision*, 60:91–110, 2004.
- [22] P. L. Meroni and P. H. Schur. ANA screening: an old test with new recommendations. *Annals of the Rheumatic Diseases*, 69(8):1420–1422, 2010.
- [23] P. Perner, H. Perner, and B. Mller. Mining knowledge for HEp-2 cell image classification. *Artificial Intelligence in Medicine*, 26:161–173, 2002.
- [24] B. Pham, S. Albarede, A. Guyard, E. Burg, and P. Maisonneuve. Impact of external quality assessment on antinuclear antibody detection performance. *Lupus*, 14(2):113–119, 2005.
- [25] A. Rakotomamonjy, S. Bach, F. R. and Canu, and Y. Grandvalet. SimpleMKL. *Journal of Machine Learning Research*, 9:2491–2521, 2008.
- [26] V. Reddy, C. Sanderson, and B. C. Lovell. Adaptive patch-based background modelling for improved foreground object segmentation and tracking. In *International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 172–179, 2010.
- [27] V. Reddy, C. Sanderson, and B. C. Lovell. Improved foreground detection via block-based classifier cascade with probabilistic decision integration. *IEEE Transactions on Circuits and Systems for Video Technology*, 23(1):83–93, 2013.
- [28] R. Rubinstein, A. M. Bruckstein, and M. Elad. Dictionaries for sparse representation modeling. *Proceedings of the IEEE*, 98(6):1045–1057, 2010.
- [29] C. Sanderson. Armadillo: An open source C++ linear algebra library for fast prototyping and computationally intensive experiments. Technical report, NICTA, 2010.

- [30] C. Sanderson and B. C. Lovell. Multi-region probabilistic histograms for robust and scalable identity inference. In *Lecture Notes in Computer Science (LNCS)*, volume 5558, pages 199–208, 2009.
- [31] C. Sanderson and K. Paliwal. Identity verification using speech and face information. *Digital Signal Processing*, 14(5):449–480, 2004.
- [32] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [33] V. Snell, W. Christmas, and J. Kittler. Texture and shape in fluorescence pattern identification for auto-immune disease diagnosis. In *International Conference on Pattern Recognition (ICPR)*, 2012.
- [34] P. Soda and G. Iannello. Aggregation of classifiers for staining pattern recognition in antinuclear autoantibodies analysis. *IEEE Trans. Information Technology in Biomedicine*, 13(3):322–329, 2009.
- [35] P. Strandmark, J. Ulén, and F. Kahl. Hep-2 staining pattern classification. In *Int. Conf. Pattern Recognition*, 2012.
- [36] I. Theodorakopoulos, D. Kastaniotis, G. Economou, and S. Fotopoulos. Hep-2 cells classification via fusion of morphological and textural features. In *IEEE International Conference on Bioinformatics and Bioengineering (BIBE)*, 2012.
- [37] G. Thibault and J. Angulo. Efficient statistical/morphological cell texture characterization and classification. In *International Conference on Pattern Recognition (ICPR)*, 2012.
- [38] J. Tropp and S. Wright. Computational methods for sparse solution of linear inverse problems. *Proceedings of the IEEE*, 98(6):948–958, 2010.
- [39] J. van Gemert, C. Veenman, A. Smeulders, and J. Geusebroek. Visual word ambiguity. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 32(7):1271–1283, 2010.
- [40] A. S. Wiik, M. Hier-Madsen, J. Forslid, P. Charles, and J. Meyrowitsch. Antinuclear antibodies: A contemporary nomenclature using HEp-2 cells. *Journal of Autoimmunity*, 35:276–290, 2010.
- [41] A. Wiliem, Y. Wong, C. Sanderson, P. Hobson, S. Chen, and B. C. Lovell. Classification of human epithelial type 2 cell indirect immunofluorescence images via codebook based descriptors. In *IEEE Workshop on Applications of Computer Vision (WACV)*, 2013.
- [42] Y. Wong, M. T. Harandi, C. Sanderson, and B. C. Lovell. On robust biometric identity verification via sparse encoding of faces: Holistic vs local approaches. In *IEEE International Joint Conference on Neural Networks*, pages 1762–1769, 2012.
- [43] Y. Wong, C. Sanderson, S. Mau, and B. Lovell. Dynamic amelioration of resolution mismatches for local feature based identity inference. In *International Conference on Pattern Recognition*, pages 1200–1203, 2010.
- [44] J. Yang, K. Yu, Y. Gong, and T. Huang. Linear spatial pyramid matching using sparse coding for image classification. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 1794–1801, 2009.