

Patchy Image Structure Classification Using Multi-Orientation Region Transform

Xiaohan Yu^{1,2}, Yang Zhao^{1,2}, Yongsheng Gao^{1,*}, Shengwu Xiong^{2,*}, Xiaohui Yuan²

¹School of Engineering and Built Environment, Griffith University, Australia

²School of Computer Science and Technology, Wuhan University of Technology, China

{xiaohan.yu2, yang.zhao4}@griffithuni.edu.au, yongsheng.gao@griffith.edu.au, {xiongsw, yuanxiaohui}@whut.edu.cn

Abstract

Exterior contour and interior structure are both vital features for classifying objects. However, most of the existing methods consider exterior contour feature and internal structure feature separately, and thus fail to function when classifying patchy image structures that have similar contours and flexible structures. To address above limitations, this paper proposes a novel Multi-Orientation Region Transform (MORT), which can effectively characterize both contour and structure features simultaneously, for patchy image structure classification. MORT is performed over multiple orientation regions at multiple scales to effectively integrate patchy features, and thus enables a better description of the shape in a coarse-to-fine manner. Moreover, the proposed MORT can be extended to combine with the deep convolutional neural network techniques, for further enhancement of classification accuracy. Very encouraging experimental results on the challenging ultra-fine-grained cultivar recognition task, insect wing recognition task, and large variation butterfly recognition task are obtained, which demonstrate the effectiveness and superiority of the proposed MORT over the state-of-the-art methods in classifying patchy image structures. Our code and three patchy image structure datasets are available at: <https://github.com/XiaohanYu-GU/MReT2019>.

Introduction

Patchy image structure classification is a fundamental yet significant research topic in computer vision and artificial intelligence (AI) research communities. Recent years have witnessed great progress of patchy image structure classification in applications highly related to AI agriculture and smart farming, such as plant branch structure classification (Isokane et al. 2018), leaf vein classification (Larese et al. 2014), and insect wing vein classification (Martineau et al. 2017). Patchy image structures characterize topological information of both exterior contours and interior structures from a target shape, which are important for shape modeling and classification. An example of illustrating the patchy image structures is shown in Fig. 1. Unfortunately, classifying such patchy image structures remains as an open problem,

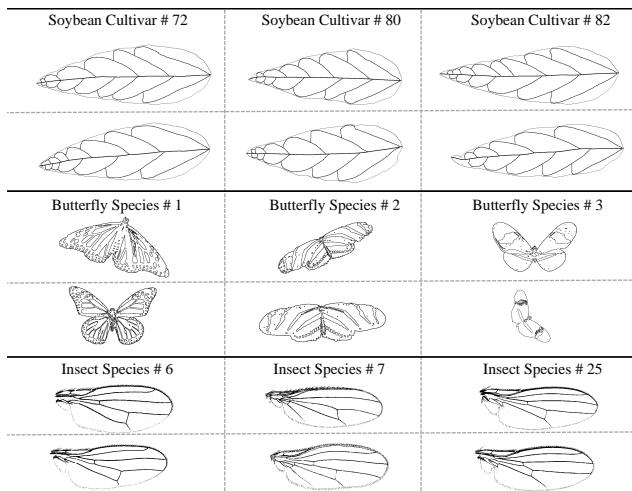


Figure 1: An example of illustrating the small inter-class differences (compare horizontally) in comparison to their intra-class variations (compare vertically) in patchy image structure classification.

due to the fact that target shapes from various categories may have highly similar contours and flexible interior structures.

In the past decades, many researchers have devoted their efforts to addressing this challenging problem. Existing approaches can be broadly classified into two categories: (1) handcrafted feature based methods and (2) deep convolutional neural network (ConvNets) techniques. Most approaches in the first category focus on developing effective feature representations of the shape contour (Ling and Jacobs 2007; Srivastava et al. 2011; Hu et al. 2012; Wang et al. 2012; Wang and Gao 2014; Demisse, Aouada, and Ottersten 2018). Although these methods achieved promising performance in general shape classification tasks, they may fail to function when the contour of objects among different classes are highly similar, such as the patchy image structures shown in Fig. 1.

Alternatively, one may consider applying ConvNets techniques, which are currently driving advances in image clas-

*Corresponding authors

sification tasks, to the patchy image structure classification. The most powerful and practical ConvNets methods are the so-called backbone networks, such as Alexnet (Krizhevsky, Sutskever, and Hinton 2012), VGG-16 (Simonyan and Zisserman 2014), and ResNet-50 (He et al. 2016), which have been proven to be very effective in large-scale image classification tasks. More recently, approaches (Fu, Zheng, and Mei 2017; Zheng et al. 2019; Yang et al. 2018; Chen et al. 2019) focusing on fine-grained classification may also contribute to the challenging patchy image structure classification. However, a major limitation is that deep learning techniques rely heavily on the amount of the training samples, *i.e.*, they may fail to perform when limited training samples are provided (Donahue et al. 2014).

In this paper, we address the patchy image structure classification problem by proposing a novel multi-orientation region transform method that can simultaneously characterize exterior contour, flexible interior structure, and inter-relationship between them. The contributions of our work are summarized as follows: (1) A Multi-Orientation Region Transform (MORT) is proposed to provide unified discriminative description of the flexible interior connection structure, shape of the exterior contour, and their inter-relationship in the patchy distance map of a patchy image structure with arbitrary orientation. (2) The proposed MORT is theoretically and experimentally proved to be rotation and translation invariant, which is important for effective and practical shape description and matching. (3) The proposed MORT can be extended to combine with ConvNets based features for further performance enhancement. (4) The encouraging experimental results demonstrate the effectiveness and efficiency of the proposed MORT over the state-of-the-art benchmarks on patchy image structure classification.

Related work

Existing methods for patchy image classification can be coarsely categorized into handcrafted feature based methods and deep convolutional neural network (ConvNets) techniques. A common strategy of handcrafted feature based methods is to present contour-based feature representations, which has yielded desirable performance for shape classification on various publicly available shape databases (Hu et al. 2012; Ling and Jacobs 2007; Wang and Gao 2014; Wang et al. 2012). Another line of work focuses on deformation-based shape analysis and classification. A recent progress on deformation-based shape analysis (Srivastava et al. 2011) enables simple physical interpretations of the resulting deformations. However, their proposed similarity metric is based on the differential quantities which are highly sensitive to noise and local perturbations. To overcome the sensitivity to local perturbations, a deformation based curved shape representation (DBCSR) (Demisse, Aouada, and Ottersten 2018) is proposed based on the relative transformation matrices between neighboring points, such that shapes are represented as elements of matrix Lie group. However, these methods may fail to function when classifying shapes with similar contours in ultra-fine granularity (e.g. leaf images from the same species but different cultivars).

Recently, ConvNets methods have been successfully applied in various visual tasks, thus may also contribute to the challenging patchy image structure classification. Most existing ConvNets methods are designed based on the so-called backbone networks, such as Alexnet (Krizhevsky, Sutskever, and Hinton 2012), VGG-16 (Simonyan and Zisserman 2014), and ResNet-50 (He et al. 2016). Given the effectiveness and practicality of these backbone networks in image classification tasks, they are naturally regarded as the first options for exploring new classification tasks (Wu et al. 2019).

Another promising alternative is to apply the fine-grained classification techniques, which are particular capable of classifying similar shapes. Lin and Maji (2017) introduced statistics normalization methods to improve an architecture that can capture second-order statistics of convolutional features in a translationally invariant manner. Li et al. (2018) proposed to utilize the second-order information to improve the fine-grained classification performance. Some researchers propose to first locate the discriminative regions and then classifies based on these regions (Huang et al. 2016). Such methods, however, require additional bounding box annotations on targeted regions. Another technique focuses on localizing discriminative regions by attention mechanism in an unsupervised manner, without requiring extra annotations. Unfortunately, these methods (Zheng et al. 2017; Fu, Zheng, and Mei 2017; Zheng et al. 2019) require additional network structure (e.g., attention mechanism), leading to extra computation overhead for both training and inference stages. To that end, Yang et al. (2018) proposed a method that is weak supervised without requiring the annotations of parts or key areas. Chen et al. (2019) introduced a very effective Destruction and Construction (DCL) method for fine-grained classification. In their work, the input images are partitioned into local regions and then shuffled by a region confusion mechanism, leading to a classification network focusing on discriminative regions for spotting the subtle differences. Given a large number of training data, ConvNets based methods can yield desirable performance on classification tasks. However, the ConvNets based methods with limited training data may dramatically overfit the training data (Donahue et al. 2014).

Aforementioned limitations motivate us to develop an effective and efficient method, without requiring large training sets for the challenging patchy image structure classification. Our work has the same assumption with the state-of-the-art handcrafted feature based methods (Belongie, Malik, and Puzicha 2002; Ling and Jacobs 2007; Hu et al. 2012; Wang et al. 2012; Wang and Gao 2014; Srivastava et al. 2011; Demisse, Aouada, and Ottersten 2018), *i.e.*, the contours of shapes are assumed to be provided.

Multi-orientation region transform method

In this section, a novel Multi-Orientation Region Transform (MORT) method is proposed to encode not only global shape features but also local structures within a shape for finer-level structure pattern analysis. To better capture the interior structure, shape of the exterior contour, and their

inter-relationship, a patchy distance map (PDM) is first developed. The proposed MORT (embedded with the PDM) and its discrete form is then described to construct the final feature descriptor. Finally, we present the rotation, translation and scale invariance analysis of the proposed feature descriptor, as well as the similarity measure using the transform coefficient matrices.

Patchy distance map

In set theory, when all sets under consideration are subsets of a given set U , the absolute complement of A is defined as the set of elements in U but not in A . Based on above notions, a patchy shape is defined to be composed of two subsets, interior set and its absolute complement. Specifically, given a shape S with patchy structure, the patches are categorized into two types: (1) interior patches I , defined as closed sub-regions inside the shape when the contour is filtered out (e.g. closed pattern patches in a butterfly image); and (2) complementary patches C , defined as the absolute complement of interior patch $C = S \setminus I$.

In order to integrate both geometrical and topological features of the patchy shape, we propose a patchy distance map (PDM) to enable comprehensive skeleton-based feature description, in which each patch is encoded with distance transform. Specifically, we apply Euclidean Distance Transform (Maurer, Qi, and Raghavan 2003) to each patch, such that each pixel is assigned with a value computed by the distance between the pixel and the nearest patch boundary. By normalizing the distance map in each patch, both small patch and large patch can provide equally detailed description based on the PDM. An example of visualizing PDM is given in Fig. 2. For the convenience of description, we denote $f^I(x, y)$ as the interior patchy distance map function, and $f^C(x, y)$ as the complementary patch distance map function.

Multi-orientation region transform

Region integral on a given contour point. Given a contour point $p(u)$ (u ranges from 0 to 1), the arc length $l(t)$, from $p(u)$ to its end point $p(e)$ on the contour along clockwise direction, is defined as $l(t) = N/2^t$, where N is the perimeter of contour (or the number of contour points when the contour is sampled in discrete form), and t is the scale index ($t = 0, 1, \dots, Q$ where $Q = \log_2 N$). Denote $p(u)$ and $p(e)$ by coordinates $(x(u), y(u))$ and $(x(e), y(e))$, a region integral on $p(u)$ at scale t is defined as:

$$R_f^z(t, \theta_u^t) = \int_{-\infty}^{+\infty} \int_{x(u) \cos \theta_u^t + y(u) \sin \theta_u^t}^{x(e) \cos \theta_u^t + y(e) \sin \theta_u^t} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f^z(x, y) \delta(x \cos \theta_u^t + y \sin \theta_u^t - \lambda_u^t, x \sin \theta_u^t + y \cos \theta_u^t - \rho_u^t) dx dy d\lambda_u^t d\rho_u^t, \quad (1)$$

where λ_u^t is the vertical distance from the origin to the boundary line defined by: $x \sin \theta_u^t + y \cos \theta_u^t = \rho_u^t$. ρ_u^t is the vertical distance from the origin to the base line defined by: $x \cos \theta_u^t + y \sin \theta_u^t = \lambda_u^t$. θ_u^t is the angle between the base line and y-axis. $z \in \{I, C\}$ and $\delta(\phi, \chi)$ is the 2D Dirac




Figure 2: An example of illustrating distance maps of patchy image structures. The patchy structure distance map of (a) a soybean leaf image, (b) a butterfly image, and (c) an insect wing image.

delta function defined as

$$\delta(\phi, \chi) = \begin{cases} 1 & \text{if } \phi = \chi = 0 \\ 0 & \text{otherwise} \end{cases}. \quad (2)$$

Here, $\delta(\phi, \chi)$ in Eq. (1) ensures that only the points on the intersection of base line $x \cos \theta_u^t + y \sin \theta_u^t = \lambda_u^t$ and boundary line $x \sin \theta_u^t + y \cos \theta_u^t = \rho_u^t$ are counted for integral calculation. Given a point indexed by u under scale t , variables θ_u^t , λ_u^t and ρ_u^t together determine the integral region, denoted as $\Omega_f(t, \theta_u^t)$, in Eq. (1). As $p(u)$ moves along the contour, the region integral is performed over different $\Omega_f(t, \theta_u^t)$ defined by the scanning range and orientation of the boundary line at all possible scales (see Fig. 3).

MORT coefficient matrix. The MORT on a given patchy distance map function $f^z(x, y)$ at point $p(u)$ is defined as

$$MT_f^z(\theta_u) = [R_f^z(t, \theta_u^0), \dots, R_f^z(\log_2 N, \theta_u^Q)]^T \quad (3)$$

The MORT at a single point $p(u)$ generates a vector of $(Q + 1)$ coefficients with each describing the region integral for one scale at point $p(u)$. An example of the proposed MORT on point $p(u)$ is illustrated in Fig. 3. By moving $p(u)$ along the contour for a complete loop, which results in a varying $\theta \in [0, 2\pi)$, the above vector grows into a MORT coefficient matrix of $(Q + 1)$ by N dimensions:

$$MT_f^z = \begin{bmatrix} R_f^z(0, \theta_1^0) & \dots & R_f^z(0, \theta_N^0) \\ \vdots & \ddots & \vdots \\ R_f^z(Q, \theta_1^Q) & \dots & R_f^z(Q, \theta_N^Q) \end{bmatrix}. \quad (4)$$

This matrix describes how the region integrals at different scales synchronously vary when the point $p(u)$ moves. When scale $t = 1, 2, \dots, p(u)$ and $p(e)$ cut a half, a quarter, \dots , off the contour to steer the region integral at different orientations (see Fig. 3(b)&(c)). The greater the t is, the smaller integral region is scanned, which provides finer descriptions on local details of the target (see Fig. 3(d)-(i)). When scale $t = 0$, $p(u)$ and $p(e)$ become the same point after a complete loop, the region integral is performed over the whole region of the given image, making the first row elements in matrix MT_f^z share the same value. For example, the 2nd, 3rd, 4th rows in matrix MT_f^z are visualized in Fig.

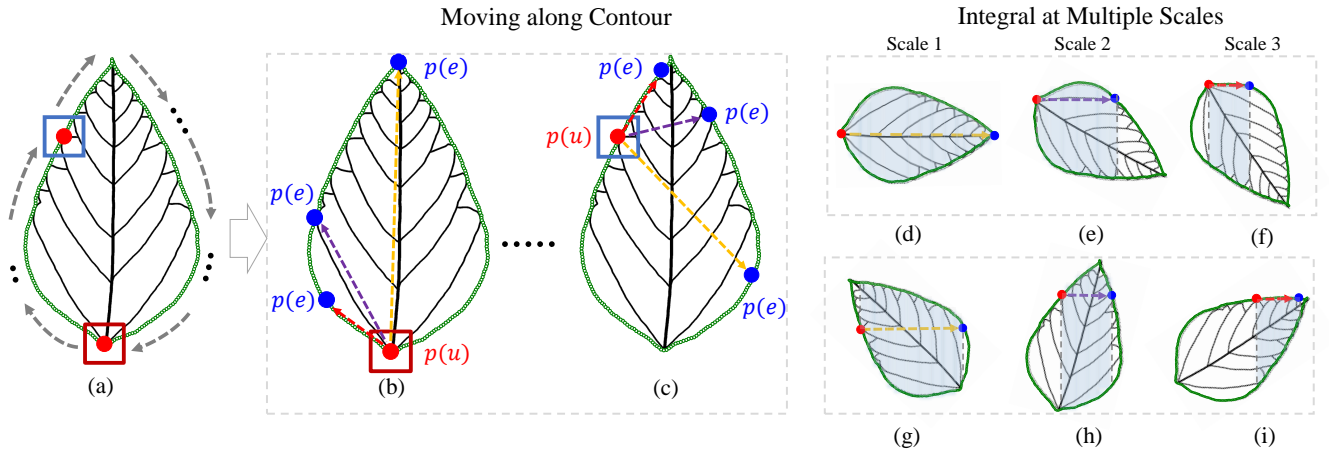


Figure 3: An example of illustrating the proposed Multi-Orientation Region Transform on a leaf image. (a) Input image with two example locations of a moving $p(u)$ point. (b) and (c): Base lines (indicating the orientations of the region integral) for scales 1, 2 and 3 (shown in yellow, purple, and red respectively) for the two example locations of $p(u)$ shown in (a). (d), (e) and (f) display the scanning ranges of the region integral for the $p(u)$ shown in (b) at scales 1, 2, and 3, respectively. (g), (h) and (i) display the scanning ranges of the region integral for the $p(u)$ shown in (c) at scales 1, 2, and 3, respectively.

3(d), Fig. 3(e), and Fig. 3(f) when point $p(u)$ moves to the position shown in Fig. 3(b).

Let $\{(x_k, y_k) | k = 1, \dots, N_f\}$ be the subset of pixels enclosed in the outer contour of the object shape, where N_f is the total number of pixels enclosed in the contour. The procedure of computing the discrete MORT coefficient matrix MT_f^z is presented in Algorithm 1, in which Eq. (1) becomes the sum of patchy distance map (PDM) values inside the moving integral region $\Omega_f(t, \theta_i^t)$.

As Steps 2, 3 and 4 take time $O(N)$, $O(\log_2 N)$ and $O(N_f)$ respectively, the algorithm 1 has a computational complexity of $O(N_f N \log_2 N)$.

Final MORT feature descriptor. By extracting features from various orientations and scales, the image region can be described in a coarse-to-fine manner to provide comprehensive description of the shape. More importantly, this allows encoding the inter-relationship of contour and interior context of a shape.

When the initial location of the contour point, that steers the Multi-Orientation Region Transform (i.e., the red point in Fig. 3, moves clockwise, the entire column of r_i^z in MT_f^z shifts to the right. The magnitudes of its 1D Fourier transform coefficients are calculated by

$$\tilde{R}^z(t, k) = \left(\frac{1}{N} \right) \left| \sum_{i=1}^N R_f^z(t, \theta_i^t) \exp\left(-\frac{2\pi j i k}{N}\right) \right|, \quad (5)$$

where $k = 1, \dots, N$ and $t = 0, \dots, \log_2 N$. To make the generated feature descriptor robust to noise and compact, the lowest M order coefficients are used to describe the target, where $M \ll N$. And $\tilde{R}^z(t, k)$ are used to construct the final

feature descriptor as

$$\Psi_f^z = \begin{bmatrix} \tilde{R}^z(0, 1) & \dots & \tilde{R}^z(0, M) \\ \vdots & \ddots & \vdots \\ \tilde{R}^z(\log_2 N, 1) & \dots & \tilde{R}^z(\log_2 N, M) \end{bmatrix}. \quad (6)$$

Invariance analysis and similarity measure

In this section, we present the rotation, translation and scale invariance analysis of the feature descriptor, which is important for effective and practical shape matching (Hong and Soatto 2015).

Invariance analysis. Lemma 1: feature descriptor Ψ_f^z is rotation invariant. It is not difficult to prove the rotation invariance. According to Eq. (5) and Eq. (6), every element in Ψ_f^z is invariant to the initial location of the contour point that steers the MORT, and thus enable Ψ_f^z to be invariant to rotation of the whole shape.

Lemma 2: A translation of $f^z(x, y)$ by a vector $\vec{u} = (x_0, y_0)$ (i.e. $f^z(x - x_0, y - y_0) = h^z(x, y)$) do not change any element in MT_f^z : $R_f^z(t, \theta_u^t) = R_h^z(t, \theta_u^t)$. This shows that every element in MT_f^z is translation invariant, and thus Ψ_f^z is translation invariant.

Lemma 3: A scaling of $f^z(x, y)$ by a factor γ (i.e. $f^z(\gamma x, \gamma y) = h^z(x, y)$) changes each element in MT_f^z by a factor $1/\gamma^2$: $R_h^z(t, \theta_u^t) = 1/\gamma^2 R_f^z(t, \theta_u^t)$. In summary, the feature descriptor Ψ_f^z is invariant to rotation and translation, and becomes scaling invariant if it is normalized by the area of the shape.

Given two sets of matrices $T = \cup_k \Psi_{(f(T).k)}^z = \{\Psi_{(f(T).1)}^z, \dots, \Psi_{(f(T).K)}^z\}$ (refer to Eq. (6)) and $M = \cup_k \Psi_{(f(M).k)}^z = \{\Psi_{(f(M).1)}^z, \dots, \Psi_{(f(M).K)}^z\}$ representing the MORT coefficient matrices of the test and model patchy

Algorithm 1 Calculating discrete Multi-Orientation Region Transform

Input:

- $f^I(x, y)$: interior patch distance map;
- $f^C(x, y)$: complementary patch distance map;
- N : number of sampling points on contour;

Output:

- MT_f^I : interior MORT coefficient matrix;
 - MT_f^C : complementary MORT coefficient matrix;
 - 1: Initialize the two MORT coefficient matrices as zero matrices: $MT_f^I = MT_f^C = \mathbf{0}_{(\log_2 N + 1) \times N}$;
 - 2: **for** $i = 1$ to N **do**
 - 3: **for** $t = 0$ to $\log_2 N$ **do**
 - 4: **for all** k such that $(x_k, y_k) \in \Omega_f(t, \theta_i^t)$ **do**
 - 5: $R_f^I(t, \theta_i^t) = R_f^I(t, \theta_i^t) + f^I(x_k, y_k)$;
 - 6: $R_f^C(t, \theta_i^t) = R_f^C(t, \theta_i^t) + f^C(x_k, y_k)$;
 - 7: **end for**
 - 8: **end for**
 - 9: $\mathbf{r}_i^I = [R_f^I(0, \theta_i^0), \dots, R_f^I(\log_2 N, \theta_i^{\log_2 N})]^T$;
 - 10: $\mathbf{r}_i^C = [R_f^C(0, \theta_i^0), \dots, R_f^C(\log_2 N, \theta_i^{\log_2 N})]^T$;
 - 11: **end for**
 - 12: $MT_f^I = [\mathbf{r}_1^I, \dots, \mathbf{r}_N^I]$;
 - 13: $MT_f^C = [\mathbf{r}_1^C, \dots, \mathbf{r}_N^C]$;
-

shapes respectively. Each matrix set is comprised of K matrix pairs of Ψ_f^I and Ψ_f^C and K is the number of types in each class. Note that $\Psi_{(f(T),k)}^z$ in T can only be matched against the matrix of the same type in M , that is $\Psi_{(f(M),k)}^z$ with the same type index k . The dissimilarity between two given samples can be measured using the fast L_1 Minkowski distance of the two matrices:

$$\begin{aligned} \text{dis}(\Psi_{(f(T),k)}^z, \Psi_{(f(M),k)}^z) &= \sum_{k=1}^K [|\Psi_{(f(T),k)}^I - \Psi_{(f(M),k)}^I| \\ &+ |\Psi_{(f(T),k)}^C - \Psi_{(f(M),k)}^C|]. \end{aligned} \quad (7)$$

It is worth noting that the proposed method (if needed) can be used as feature extractor together with a classifier, such as the classic support vector machine (SVM), for classification tasks (Dalal and Triggs 2005).

Experiments and discussions

In this section, we first introduce three publicly available patchy image structure datasets for classification evaluation. We then present the experimental results including both classification accuracy and computational time of the proposed MORT together with state-of-the-art methods. Finally, we present a fused MORT by integrating MORT with ConvNets methods via feature-level fusion, and report evaluation results on the three patchy image structure datasets.

Patchy image structure datasets

SoyCultivarVein dataset. The SoyCultivarVein dataset is a publicly available dataset, which comprises 100 categories (cultivars) with 6 samples (leaf images) in each cultivar




Figure 4: An example of illustrating (a), (d), (g) patchy image structures and their corresponding (b), (e), (h) interior patches (white patches), and (c), (f), (i) complementary patches (white patches).

and thus has a total number of $100 \times 6 = 600$ images (Yu et al. 2019). The leaves in the SoyCultivarVein dataset are highly similar due to the fact that they all belong to the same species, making it a new and challenging dataset for the artificial intelligence and pattern analysis research community.

BtffPIS dataset. The butterfly patchy image structure dataset (BtffPIS), is constructed by applying the Canny edge detection (Canny 1986) to the binarized images of the first 50 images in each class from the public available Leeds butterfly dataset (Wang, Markert, and Everingham 2009). There are $50 \times 10 = 500$ images in the BtffPIS dataset. The BtffPIS dataset comprises large intra-class variations including rotation and scaling variations in each class.

IwPIS dataset. The insect wing patchy image structure dataset (IwPIS) (Tofilski 2004), is adopted for evaluation, which comprises 25 classes of insect wings with 2 samples in each class. Although small, the IwPIS is very challenging for classification tasks due to the fact that all the samples are from the same order called Diptera. In other words, samples in IwPIS have very small inter-class variances that are difficult to be classified.

Comparisons with shape-based benchmarks

We compare the proposed MORT with ten state-of-the-art shape-based benchmarks. They are: (1) two versions of Shape Contexts, i.e., standard Shape Contexts (SC) and Shape Contexts with dynamic programming (SC-DP) (Belongie, Malik, and Puzicha 2002), (2) Inner Distance Shape Contexts (IDSC) and Inner Distance Shape Contexts with dynamic programming (IDSC-DP) (Ling and Jacobs 2007), (3) square-root velocity (SRV) method (Srivastava et al. 2011), (4) deformation-based curved shape representation (DBCSR) with uniform sampling and optimal sampling (Demisse, Aouada, and Ottersten 2018), (5) Hierarchical

Table 1: Accuracy & matching time on SoyCultivarVein dataset (K=3)

Algorithm	Accuracy (%)	Time (ms)
SC	37.58	2.38×10^0
SC+DP	36.90	1.17×10^{-1}
IDSC	49.07	2.26×10^0
IDSC+DP	46.93	1.14×10^{-1}
SRV	37.09	6.89×10^2
DBCSR (opt)	35.72	3.69×10^4
DBCSR (uni)	34.89	3.63×10^4
HSC	43.98	7.43×10^{-3}
MDM	39.39	3.93×10^{-2}
HF	40.46	1.02×10^1
Proposed MORT	53.43	6.62×10^{-2}

String Cuts (HSC) (Wang and Gao 2014), (6) Multiscale Distance Matrix (MDM) (Hu et al. 2012), and (7) Height Functions (HF) (Wang et al. 2012). The widely used Nearest Neighbor score (1NN) is employed for performance measurement, as used in the benchmark methods.

Note that, the MORT does not restrain how the interior and complementary patches shall be defined. In our experiments, for SoyCultivarVein dataset, the interior patches are defined as the patches enclosed only by vein points, while the complementary patches are those enclosed by a mixture of vein and contour points (see Figs. 4(a-c)). For BtfPIS dataset and IwPIS dataset, the interior patches are patches with bright pixels (i.e., their intensities are above or equal to the binarization threshold), while the complementary patches are those formed by dark pixels (i.e., their intensities are below the threshold), as illustrated in Figs. 4(d-f) and Figs. 4(g-i), respectively.

Evaluation on SoyCultivarVein dataset (ultra-fine-grained). Table 1 illustrates the average classification accuracies of MORT together with the state-of-the-art shape-based methods. We repeat the classification evaluation 1000 times by reselecting different three samples randomly to construct the model set and the remaining samples as the testing set. The average results are reported. Using the 1NN evaluation protocol, the proposed MORT achieves the highest average classification accuracy of 54.20% (10.22% higher than the HSC and 5.13% higher than the IDSC), demonstrating its superiority in the ultra-fine-grained cultivar classification task. Table 1 also lists the computational cost (the average time for each matching) of all the competing methods. The matching speed of MORT is among the most efficient methods (ranked the second), demonstrating the efficiency of the MORT method.

In order to verify the rotation invariance of the proposed method, we construct a rotated SoyCultivarVein dataset, by rotating each leaf image in the SoyCultivarVein dataset with a random angle (from 0° to 360°). We compare the proposed method with two rotation invariant shape classification methods, HSC and MDM on the rotated SoyCulti-

Table 2: Accuracy on rotated SoyCultivarVein dataset

Algorithm	Accuracy (%)
HSC	45.11
MDM	39.30
Proposed MORT	54.81

Table 3: Accuracy & matching time on BtfPIS dataset (K=1)

Algorithm	Accuracy (%)	Time (ms)
SC	42.57	3.00×10^0
SC+DP	41.93	1.08×10^{-1}
IDSC	57.45	2.10×10^1
IDSC+DP	54.32	1.13×10^{-1}
SRV	60.49	1.14×10^3
DBCSR (opt)	52.95	4.23×10^4
DBCSR (uni)	58.53	3.80×10^4
HSC	60.78	8.10×10^{-3}
MDM	54.84	5.00×10^{-2}
HF	20.64	8.50×10^1
Proposed MORT	75.02	5.45×10^{-2}

varVein dataset. The experimental results (see Table 2 in comparison with the Table 1) confirm the theoretical analysis on rotation invariance of MORT.

Evaluation on BtfPIS dataset (large intra-class variation). Table 3 illustrates the average classification accuracies of MORT together with the state-of-the-art shape-based methods. We repeat the classification evaluation 1000 times by reselecting different 25 samples randomly to construct the test set (and the remaining 25 samples are used as the model set) and their average results are reported. Using the 1NN evaluation protocol, MORT achieves the highest classification accuracy of 75.02%, which are significantly higher than the state-of-the-art benchmarks (14.24% higher than the HSC and 17.57% higher than the IDSC). The superior performance of the proposed method over the state-of-the-art benchmarks demonstrates the robustness of MORT in classifying images with large intra-class variances. The matching speed of MORT on BtfPIS dataset is also among the most efficient methods (ranked the second).

Evaluation on IwPIS dataset (small inter-class difference). Table 4 shows the average classification accuracies and average matching time of MORT and the state-of-the-art methods on the challenging IwPIS dataset. Using the same repeating strategy and 1NN evaluation (half images as model set and the remaining as testing set) as adopted in previous experiments, MORT achieves the highest classification accuracy of 44.79%, which is higher than other 10 benchmarks (4.1% higher than the SC and 13.87% higher than the IDSC). The average matching time (for each matching) of MORT is 1.56×10^{-2} ms, which is the lowest

among all the benchmarks. The results show the superiority of MORT against state-of-the-art benchmarks in classifying patchy image structures with small inter-class variances.

Comparisons with ConvNets methods

Fused MORT. Combining complementary feature representations may significantly improve the classification performance (Wang et al. 2015). However, a key question is that whether the proposed MORT and the state-of-the-art ConvNets methods can provide each other complementary features for further performance improvement. To that end, we propose a fused MORT by concatenating the original MORT feature matrices and feature vectors extracted from a state-of-the-art ConvNets method, DCL (Chen et al. 2019), and then investigate the comparative results of ConvNets methods and the fused MORT. To facilitate the comparison between the proposed method and ConvNets methods, we combine the fused MORT (as the feature encoder) with SVM (as the classifier) using the same protocol as adopted by Wu et al. (2019).

Competitors. We compare the proposed MORT with the following state-of-the-art ConvNets methods: (1) Three state-of-the-art ConvNets models: Alexnet (Krizhevsky, Sutskever, and Hinton 2012), VGG-16 (Simonyan and Zisserman 2014), and ResNet-50 (He et al. 2016); and (2) three fine-grained state-of-the-art methods: improved B-CNN (Lin and Maji 2017), fast-MPN-COV (Li et al. 2018), and DCL (Chen et al. 2019).

Implementation details. All the models are implemented in Pytorch 1.0.0 and are pretrained on the ImageNet (Deng et al. 2009) and then fine-tuned on each dataset. In all experiments, the input images are resized to 440×440 , and cropped to 384×384 randomly for training. Standard data augmentations are applied including random rotation within ± 15 degree and horizontal flip with 0.5 probability. For fast-MPN-COV and improved B-CNN, the models are trained with the best setting in the implementation of fast-MPN-COV (Li et al. 2018). For the remaining methods, the models are trained with the default settings in the implementation of DCL (Chen et al. 2019).

Evaluation on patchy image structure datasets. For SoyCultivarVein, BtPIS and IwPIS datasets, we select the first half images from each category as the training set, and the remaining images as the testing set. Table 5 lists the classification accuracies of all the competing ConvNets methods on the three datasets. The fused MORT achieves the best classification accuracy of 65% and 98.00% on SoyCultivarVein dataset and BtPIS dataset, respectively. Nevertheless, the fused MORT ranks third in classification accuracy on IwPIS, with 6% and 12% lower than the improved B-CNN and fast-MPN-COV, respectively. A possible reason is that the fused MORT has much lower feature dimension (1×2092) per image compared with improved B-CNN (1×262144) and fast-MPN-COV (1×32896). The overall results show that the fused MORT surpasses the original MORT in classification accuracy on all the three patchy image structure datasets, indicating that ConvNets methods

Table 4: Accuracy & matching time on IwPIS dataset (K=1)

Algorithm	Accuracy (%)	Time (ms)
SC	40.78	3.14×10^0
SC+DP	37.59	1.40×10^{-1}
IDSC	30.92	2.92×10^0
IDSC+DP	34.90	1.36×10^{-1}
SRV	36.40	8.44×10^2
DBCSR (opt)	33.75	4.12×10^4
DBCSR (uni)	27.72	4.20×10^4
HSC	35.15	1.72×10^{-2}
MDM	27.06	4.08×10^{-2}
HF	36.96	1.42×10^1
Proposed MORT	44.79	1.56×10^{-2}

Table 5: Accuracy on three patchy image structure datasets

Algorithm	Accuracy (%)		
	SoyCultivarVein	BtPIS	IwPIS
Alexnet	14	93.20	44
VGG-16	16	95.60	48
ResNet-50	28	96.80	52
improved B-CNN	63	73.60	74
fast-MPN-COV	51	77.20	80
DCL	39	97.20	68
fused MORT	65	98.00	68

and the proposed MORT can provide each other with complementary information. The superior results of the fused MORT over ConvNets methods on three datasets indicate that the proposed MORT can extract discriminative features even with very limited training samples, leading to the performance improvements of fused MORT over the sole use of DCL or original MORT.

Conclusion

In this paper, we presented a novel Multi-Orientation Region Transform (MORT) method, which is rotation, translation and scale invariant, for effective and efficient classification of patchy image structures. The proposed MORT can extract local structural features at various scales and orientations for comprehensive shape description. The encouraging experimental results on three patchy image structure databases demonstrate the effectiveness and efficiency of MORT for patchy image shape classification. Moreover, the results also indicate that the proposed MORT and ConvNets methods can provide each other with important complementary features to further improve the classification performance.

Acknowledgement

This work was supported in part by the Australian Research Council under Discovery Grant DP180100958 and Linkage Grant LP170100326.

References

- Belongie, S.; Malik, J.; and Puzicha, J. 2002. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Ana. Mach. Intell.* (4):509–522.
- Canny, J. 1986. A computational approach to edge detection. *IEEE Trans. Pattern Ana. Mach. Intell.* (6):679–698.
- Chen, Y.; Bai, Y.; Zhang, W.; and Mei, T. 2019. Destruction and construction learning for fine-grained image recognition. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 5157–5166.
- Dalal, N., and Triggs, B. 2005. Histograms of oriented gradients for human detection. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 886–893.
- Demisse, G. G.; Aouada, D.; and Ottersten, B. 2018. Deformation based curved shape representation. *IEEE Trans. Pattern Ana. Mach. Intell.* 40(6):1338–1351.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 248–255.
- Donahue, J.; Jia, Y.; Vinyals, O.; Hoffman, J.; Zhang, N.; Tzeng, E.; and Darrell, T. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *Proc. Int. Conf. Mach. Learn. (ICML)*, 647–655.
- Fu, J.; Zheng, H.; and Mei, T. 2017. Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 4438–4446.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 770–778.
- Hong, B.-W., and Soatto, S. 2015. Shape matching using multiscale integral invariants. *IEEE Trans. Pattern Ana. Mach. Intell.* 37(1):151–160.
- Hu, R.; Jia, W.; Ling, H.; and Huang, D. 2012. Multiscale distance matrix for fast plant leaf recognition. *IEEE Trans. Image Process.* 21(11):4667–4672.
- Huang, S.; Xu, Z.; Tao, D.; and Zhang, Y. 2016. Part-stacked cnn for fine-grained visual categorization. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 1173–1182.
- Isokane, T.; Okura, F.; Ide, A.; Matsushita, Y.; and Yagi, Y. 2018. Probabilistic plant modeling via multi-view image-to-image translation. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 2906–2915.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Proc. Advances in Neural Info. Process. Sys. (NeurIPS)*, 1097–1105.
- Larese, M. G.; Bayá, A. E.; Craviotto, R. M.; Arango, M. R.; Gallo, C.; and Granitto, P. M. 2014. Multiscale recognition of legume varieties based on leaf venation images. *Expert Syst. Appl.* 41(10):4638–4647.
- Li, P.; Xie, J.; Wang, Q.; and Gao, Z. 2018. Towards faster training of global covariance pooling networks by iterative matrix square root normalization. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 947–955.
- Lin, T.-Y., and Maji, S. 2017. Improved bilinear pooling with cnns. *Proc. British Mach. Vis. Conf. (BMVC)*.
- Ling, H., and Jacobs, D. W. 2007. Shape classification using the inner-distance. *IEEE Trans. Pattern Ana. Mach. Intell.* 29(2):286–299.
- Martineau, M.; Conte, D.; Raveaux, R.; Arnault, I.; Munier, D.; and Venturini, G. 2017. A survey on image-based insect classification. *Pattern Recognit.* 65:273–284.
- Maurer, C. R.; Qi, R.; and Raghavan, V. 2003. A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Trans. Pattern Ana. Mach. Intell.* 25(2):265–270.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Srivastava, A.; Klassen, E.; Joshi, S. H.; and Jermyn, I. H. 2011. Shape analysis of elastic curves in euclidean spaces. *IEEE Trans. Pattern Ana. Mach. Intell.* 33(7):1415–1428.
- Tofilski, A. 2004. Drawing, a program for numerical description of insect wings. *Journal of Insect Science* 4(1).
- Wang, B., and Gao, Y. 2014. Hierarchical string cuts: a translation, rotation, scale, and mirror invariant descriptor for fast shape retrieval. *IEEE Trans. Image Process.* 23(9):4101–4111.
- Wang, J.; Bai, X.; You, X.; Liu, W.; and Latecki, L. J. 2012. Shape matching and classification using height functions. *Pattern Recognit. Lett.* 33(2):134–143.
- Wang, L.; Zhang, J.; Zhou, L.; Tang, C.; and Li, W. 2015. Beyond covariance: Feature representation with nonlinear kernel matrices. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 4570–4578.
- Wang, J.; Markert, K.; and Everingham, M. 2009. Learning models for object recognition from natural language descriptions. In *Proc. British Mach. Vis. Conf. (BMVC)*.
- Wu, X.; Zhan, C.; Lai, Y.-K.; Cheng, M.-M.; and Yang, J. 2019. Ip102: A large-scale benchmark dataset for insect pest recognition. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 8787–8796.
- Yang, Z.; Luo, T.; Wang, D.; Hu, Z.; Gao, J.; and Wang, L. 2018. Learning to navigate for fine-grained classification. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*.
- Yu, X.; Gao, Y.; Xiong, S.; and Yuan, X. 2019. Multiscale contour steered region integral and its application for cultivar classification. *IEEE Access* 7:69087–69100.
- Zheng, H.; Fu, J.; Mei, T.; and Luo, J. 2017. Learning multi-attention convolutional neural network for fine-grained image recognition. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 5209–5217.
- Zheng, H.; Fu, J.; Zha, Z.-J.; and Luo, J. 2019. Looking for the devil in the details: Learning trilinear attention sampling network for fine-grained image recognition. In *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*.