Pattern Recognition 50 (2016) 56–73

Contents lists available at ScienceDirect

Pattern Recognition

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

α-shapes for local feature detection

Christos Varytimidis*, Konstantinos Rapantzikos, Yannis Avrithis, Stefanos Kollias

National Technical University of Athens, Iroon Polytexneiou 9, 15780 Zografou, Greece

ARTICLE INFO

Article history:
Received 16 June 2014
Received in revised form
29 November 2014
Accepted 14 August 2015
Available online 28 August 2015

Keywords:
Local features
Point sampling
Triangulation
α-shapes

ABSTRACT

Local image features are routinely used in state-of-the-art methods to solve many computer vision problems like image retrieval, classification, or 3D registration. As the applications become more complex, the research for better visual features is still active. In this paper we present a feature detector that exploits the inherent geometry of sampled image edges using *α-shapes*. We propose a novel edge sampling scheme that exploits local shape and investigate different triangulations of sampled points. We also introduce a novel approach to represent the anisotropy in a triangulation along with different feature selection methods. Our detector provides a small number of distinctive features that is ideal for large scale applications, while achieving competitive performance in a series of matching and retrieval experiments.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Local features provide a balance between the sparseness of global representations and the density of features extracted on a fixed grid of locations. By ignoring non-salient image parts and focusing on distinctive regions they provide repeatability, discriminative power, computational efficiency and compactness. These properties boost computer vision applications including large-scale recognition, retrieval or 3D reconstruction.

State-of-the-art detectors like Hessian-Affine [1], MSER [2] and SURF [3] have been used in many computer vision applications and are quite mature and popular. However, the speed, stability and image coverage provided by those detectors are not ideal. The speed and image coverage of the Hessian-Affine detector are limited, while multiple detections often appear on nearby locations at different scales. The MSER detector is fast, but often extracts sparse regular regions that are not representative enough. SURF is also fast, but detections are often not stable enough. Recent publications compare state-of-the-art detectors not only by common statistics (e.g. repeatability/matching score), but also in diverse applications like image classification [4] or retrieval [5–7].

In an attempt to capture the dominant structural features in an image, we propose a detector based on a local shape representation rather than first- or second-order image derivatives. In this direction we employ α-shapes, a well known method in computational geometry introduced by Edelsbrunner et al. [8]. An α-shape is a subset of a triangulation of a point set in a Euclidean space, where scale-like parameter $\alpha \geq 0$ determines the *faces* of the triangulation (points, edges or triangles) that are included in the particular subset (see Section 3.3). Given a set of points, α-shapes involve a grouping process guided by α, and capture the shape of structures generated by this process. They can be thought of as a generalization of the *convex hull*, being parameterized by α. Starting from the point set for $\alpha = 0$, the subset of the triangulation expands to the convex hull at the other extreme $\alpha = \infty$ (see Fig. 1).

The set of all α-shapes (for all possible values of α) is a *filtration* over a triangulation of the point set, *i.e.* a partial ordering of simplices (edges and triangles in two dimensions) [9]. *Delaunay* triangulation is the most common choice, but since it is only based on point coordinates, it may be a poor representation depending on the application, *e.g.* a sampled function over an image may be more informative. *Weighted α-shapes* [10] on the other hand are based on a *regular* triangulation and provide a more flexible representation, by associating an additional scalar parameter per point. Teichmann and Capps introduce the *anisotropic* α-shapes [11], using an even richer representation per point. Anisotropic α-shapes are a generalization of weighted α-shapes, defined on non-Euclidean metric spaces.

In [12] we introduce WαSH, a detector based on weighted α-shapes that groups edge samples by exploiting location, gradient strength and local shape. To capture local shape, we devise an efficient way to overcome the main weakness of α-shapes, namely the automatic selection of α value that best represents the underlying point set. We also show how noisy points or groupings can be automatically filtered out by a shape-based stability measure. WαSH performs quite well and is controlled by a single and intuitive parameter.

In this paper we treat a local feature as a region delineated by a set of points sampled from its contour, as in WαSH. However, instead of using a uniform sampling scheme along the edges, we explore a non-uniform scheme whose sampling density is guided

* Corresponding author. Tel.: +30 2107722521; fax: +30 2107722492.
E-mail addresses: chrisvar@image.ntua.gr (C. Varytimidis), rap@image.ntua.gr (K. Rapantzikos), iavr@image.ntua.gr (Y. Avrithis), stefanos@cs.ntua.gr (S. Kollias).

http://dx.doi.org/10.1016/j.patcog.2015.08.016
0031-3203/© 2015 Elsevier Ltd. All rights reserved.

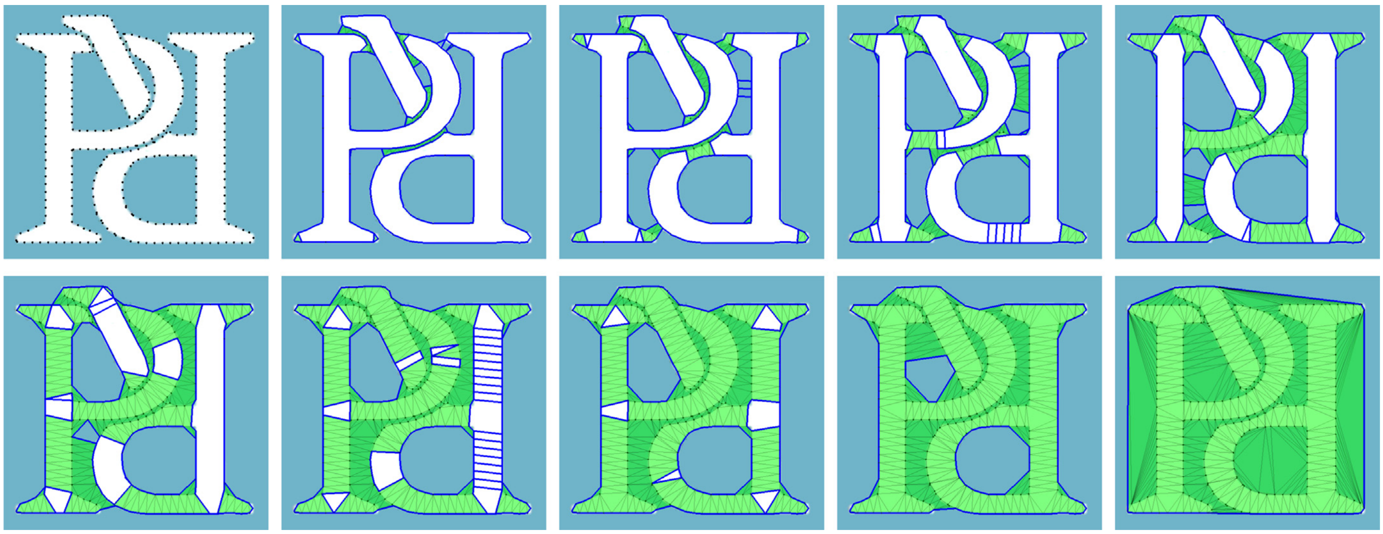

Fig. 1. Example of the *α-filtration*. Different instances of the filtration for different values of α: starting from the point set for $\alpha = 0$ (top left) and adding triangles and edges, we end up to the *convex hull* for $\alpha = \infty$ (bottom right). Observe how the cavities of the α-shapes correspond to cavities of the objects and blob-like regions.

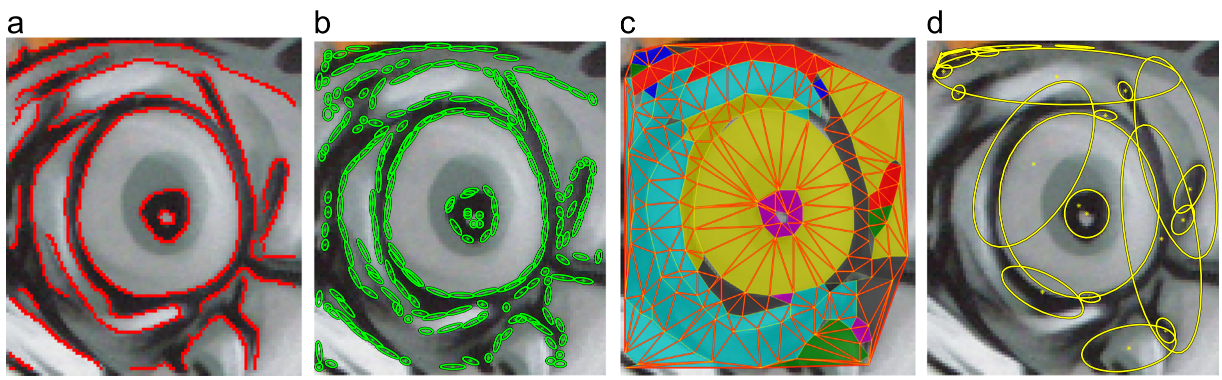

Fig. 2. Overview of our detector. (a) Edges of the input image are (b) non-uniformly sampled. (c) We create the α-filtration of a triangulation over the samples and track the evolution of connected components. (d) We extract features by selecting stable and prominent components.

by local edge shape. The major steps of the algorithm are summarized in Fig 2, namely (a) edge extraction, (b) non-uniform sampling based on local anisotropy, (c) triangulation of the samples and α-filtration construction, and (d) local feature extraction based on shape measures. Compared to WαSH, we improve and extend the method by

- applying non-uniform sampling based on image edges and local shape;
- introducing anisotropically weighted α-shapes, also adapted to local shape;
- comparing several triangulations to build the α-shapes; and
- proposing and evaluating different measures to select dominant components.

The remaining of the paper is organized as follows: In Section 2 we discuss related work, followed by the description of our method. In Section 3.1 we describe different sampling strategies and in Section 3.2 we provide an overview of the triangulations used. In Section 3.3 we introduce anisotropically weighted α-shapes. In Section 4.1 we describe the component trees used to track evolving shapes and in Section 4.2 we propose different measures to select dominant components as features, followed by an overview of our algorithm in Section 4.3. In Section 4.4 we show visual examples and discuss the qualities of our detector. The performance of our detector is experimentally evaluated and compared to the state-of-the-art in Section 5, followed by conclusions and discussion in Section 6.

2. Related work

Early region detectors were based on extending ideas found in corner detectors like Beaudet [13] and Harris and Stevens [14], which were based on the Hessian and the second moment matrix respectively. In his inspiring work, Lindeberg et al. [15] studied scale-invariant detectors and established the theoretical foundations for making them affine-invariant [16]. Based on these foundations, Lowe [17] introduced the *scale invariant feature transform* (SIFT), still one of the most popular detectors, which achieves invariance to scale and rotation based on the Difference-of-Gaussian (DoG) operator. Mikolajczyk et al. [1] extended the Harris–Laplace and Hessian–Laplace operators towards affine invariance using the Laplacian-of-Gaussian (LoG) operator in affine scale space.

More recently, Alcantarilla et al. [7] introduce the KAZE operator, which detects maxima of the Hessian in a nonlinear scale space built by diffusion filtering. Although the statistics are comparable to the state-of-the-art, the creation of the nonlinear scale-space is computationally expensive and the number of features is high. The fast variant of KAZE in [18] is still slower than the state-of-the-art, while not providing better performance.

The *maximally stable extremal regions* (MSER) of Matas et al. [2], one of the best performing region detectors in [19], detect regions of stable intensity and therefore avoid common problems of gradient-based methods like localization accuracy and noise. The idea is to compute a watershed-like segmentation and to select those regions that remain stable over a predefined set of thresholds. MSER are

Download English Version:

https://daneshyari.com/en/article/6940058

Download Persian Version:

https://daneshyari.com/article/6940058

Daneshyari.com