

Eigencontours: Novel Contour Descriptors Based on Low-Rank Approximation

Wonhui Park
Korea University

whpark@mcl.korea.ac.kr

Dongkwon Jin
Korea University

dongkwonjin@mcl.korea.ac.kr

Chang-Su Kim
Korea University

changasukim@korea.ac.kr

Abstract

Novel contour descriptors, called *eigencontours*, based on low-rank approximation are proposed in this paper. First, we construct a contour matrix containing all object boundaries in a training set. Second, we decompose the contour matrix into eigencontours via the best rank- M approximation. Third, we represent an object boundary by a linear combination of the M eigencontours. We also incorporate the eigencontours into an instance segmentation framework. Experimental results demonstrate that the proposed eigencontours can represent object boundaries more effectively and more efficiently than existing descriptors in a low-dimensional space. Furthermore, the proposed algorithm yields meaningful performances on instance segmentation datasets.

1. Introduction

Contour is one of the most important object descriptors, along with texture and color. The boundary of an object in an image is encoded in contour description, which is useful in various applications, such as image retrieval [4, 31, 32], recognition [17, 25, 29], and segmentation [16, 19, 27, 28, 34]. It is desirable to represent object boundaries compactly, as well as faithfully, but it is challenging to design such contour descriptors due to the diversity and complexity of object shapes.

Early contour descriptors were developed mainly for image retrieval [4, 17, 31, 32]. An object contour can be simply represented based on the area, circularity, and/or eccentricity of the object [30]. For more precise description, there are several approaches, including shape signature [6, 18, 27], structural analysis [5, 7, 10, 21, 28], spectral analysis [4, 31], and curvature scale space (CSS) [8, 17].

Recently, contour descriptors have been incorporated into deep-learning-based object detection, tracking, and segmentation systems. In [35], bounding boxes are replaced by polygons to enclose objects more tightly. In [3], ellipse fitting is done to produce a rotated box of a target object to be tracked. For instance segmentation, contour-

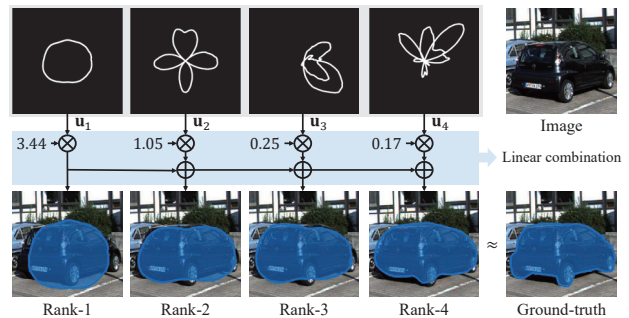


Figure 1. Illustration of the eigencontour representation. The boundary of a vehicle is represented by a linear combination of four eigencontours: u_1 , u_2 , u_3 and u_4 . First, u_1 approximates the object boundary roughly. Next, u_2 is used to refine the boundary by adjusting top and bottom parts, as well as front and rear ones. To reconstruct more complex parts, such as wheels and bumper, u_3 and u_4 should be used as well. These eigencontours were determined by analyzing the boundaries of all objects in the ‘car’ category in the KINS dataset [22].

based techniques have been proposed that represent pixel-wise masks by contour descriptors based on shape signature [27] or polynomial fitting [28]. Even though these descriptors can localize an object effectively, they may fail to reconstruct the object boundary faithfully. Also, they consider the structural information of an individual object only, without exploiting the shape correlation between different objects.

In this paper, we propose novel contour descriptors, called *eigencontours*, based on low-rank approximation. First, we construct a contour matrix containing all object boundaries in a training set. Second, we decompose the contour matrix into eigencontours, based on the best rank- M approximation of singular value decomposition (SVD) [2]. Then, each contour is represented by a linear combination of the M eigencontours, as illustrated in Figure 1. Also, we incorporate the eigencontours into an instance segmentation framework. Experimental results demonstrate that the proposed eigencontours can represent object boundaries more effectively and more efficiently than the existing contour descriptors [27, 28]. Moreover, utilizing the existing

framework of YOLOv3 [23], the proposed algorithm yields promising instance segmentation performances on various datasets — KINS [22], SBD [12], and COCO2017 [15].

This work has the following contributions:

- We propose the notion of eigencontours — data-driven contour descriptors based on SVD — to represent object boundaries as faithfully as possible with a limited number of coefficients.
- The proposed algorithm can represent object boundaries more effectively and more efficiently than the existing contour descriptors.
- The proposed algorithm outperforms conventional contour-based techniques in instance segmentation.

2. Related Work

The goal of contour description is to represent the boundary of an object in an image compactly and faithfully. Simple contour descriptors are based on the area, circularity, and/or eccentricity of an object [30], and basic geometric shapes, such as rectangles and ellipses, can be also used. However, these simple descriptors cannot preserve the original shape of an object faithfully [33], [24]. For more sophisticated description, there are four types of approaches: shape signature [6, 18, 27], structural analysis [5, 10, 21, 28], spectral analysis [4, 31], and CSS [8, 17]. First, a shape signature is a one-dimensional function derived from the boundary coordinates of an object. For example, a polar coordinate system is set up with respect to the centroid of an object. Then, the object boundary is represented by the (r, θ) graph, called the centroidal profile [6]. Also, an object shape can be represented by the angle between the tangent vector at each contour point and the x -axis [18]. Second, structural methods divide an object boundary into segments and approximate each segment to encode the whole boundary. In [10], the boundary is represented by a sequence of unit vectors with a few possible directions. In [21], polygonal approximation is performed to globally minimize the errors from an approximated polygon to the original boundary. In [5], segments of an object contour are represented by cubic polynomials. Third, in spectral methods, boundary coordinates are transformed to a spectral domain. In [4], a wavelet transform is used for contour description. In [31], the Fourier descriptors are derived from the Fourier series of centroidal profiles. Fourth, in CSS [17], a boundary is smoothed by a Gaussian filter with a varying standard deviation. Then, the boundary is represented by the curvature zero-crossing points of the smoothed curve at each standard deviation.

Recently, attempts have been made to improve the performances of deep-learning-based vision systems. In [35], a bounding box for object detection is replaced by an octagon to enclose an object more tightly via polygonal ap-

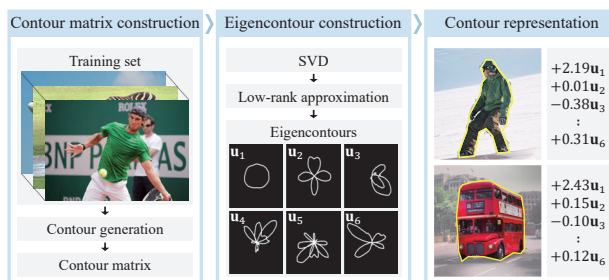


Figure 2. Overview of the proposed algorithm.

proximation. In [3], a rotated box for a target object is determined based on ellipse fitting, in order to cope with object deformation in a visual tracking system. For instance segmentation, contour-based approaches [27, 28] have been developed, which reformulate the pixelwise classification task as the boundary regression of an object. To this end, these methods encode segmentation masks into contour descriptors. In [27], centroidal profiles are used to describe object boundaries. In [28], each segment of a boundary is represented by a few coefficients based on polynomial fitting. Although these methods are computationally efficient for localizing object instances, they often fail to reconstruct the boundaries of the object shapes faithfully.

The proposed algorithm aims to represent an object boundary as faithfully as possible by employing as few coefficients as possible. To this end, we develop eigencontours based on the best low-rank approximation property of SVD.

3. Proposed Algorithm

Instead of deriving contour descriptors based on prior assumptions on object boundaries, such as rectangular, elliptical, or polynomial models, we develop eigencontours by analyzing boundary data in a training set. In this sense, the proposed eigencontours are data-driven descriptors. Figure 2 is an overview of the proposed algorithm. First, we compose a contour matrix, containing all object boundaries in a training set. Second, we approximate the matrix, by performing the best rank- M approximation, to determine M eigencontours. Third, we represent an object boundary by a linear combination of the M eigencontours.

3.1. Mathematical Formulation

SVD and principal component analysis (PCA) are used in various fields to achieve dimensionality reduction and represent data concisely [2, 13, 14]. In this paper, we use SVD to represent object boundaries compactly and reliably. More specifically, we adopt a data-driven approach to exploit the distribution of object contours in a training set, instead of performing curve fitting [5] or Fourier analysis [31], in order to represent object boundaries efficiently in a low-dimensional space.

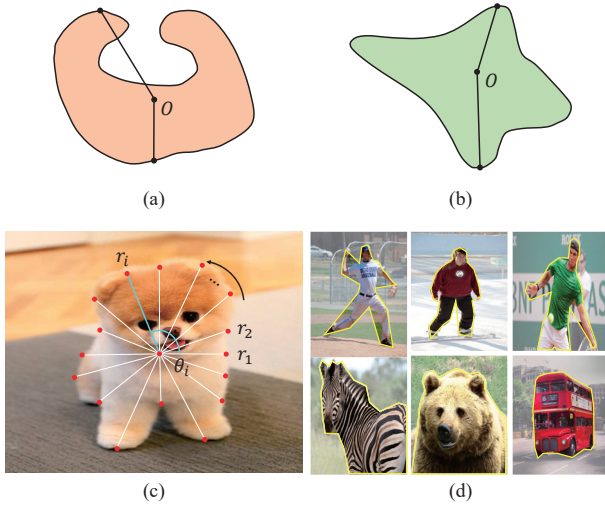


Figure 3. In (a), the set (or shape) is not star-convex because there are line segments not wholly contained in the set. In (b), it is star-convex because the segment from the center O to any point in the set is contained in the set. In (c), a star-convex contour is represented by polar coordinates. More examples of star-convex contours are in (d).

Star-convex contour generation: There is a tradeoff between accuracy and simplicity of a contour representation scheme: an accurate representation yields a high-dimensional feature vector, while too simple a representation cannot describe complicated boundaries precisely. To strike a good balance, we adopt the star-convexity assumption of object shapes. A regional set (or shape) is star-convex [26] if it contains a point such that the line segment from the point to any point in the set is contained in the set. Then, a star-convex contour is defined as the set of boundary points of a star-convex set. For example, Figure 3(a) is not a star-convex contour, but Figure 3(b) is a star-convex one.

To represent star-convex contours, we use centroidal profiles [6]. Given an object shape, we find the inner-center, which is the center of the circle of the maximum size wholly contained in the shape, as done in [28]. Then, with respect to the inner center, we describe the boundary using polar coordinates (r_i, θ_i) , $i = 1, 2, \dots, N$. The angular coordinates θ_i are sampled uniformly, so only the radial coordinates are recorded to represent the contour

$$\mathbf{r} = [r_1, r_2, \dots, r_N]^\top. \quad (1)$$

As in Figure 3(c), r_i is set to be the distance of the farthest object point from the center along the θ_i -axis. By construction, \mathbf{r} describes a star-convex contour.

Figure 3(d) shows more star-convex contours. With the infinite sampling $N = \infty$, a star-convex contour is guaranteed to enclose all object points, since it is the boundary of the star-convex hull of the object. However, with a finite

N , the star-convex contour may miss some object points, as well as include some non-object points. However, we see that the contours in Figure 3(d) represent object shapes quite faithfully.

Eigencontour space: In general, object shapes are well structured and thus highly correlated to one another, especially between objects in the same class. By exploiting this structural relationship using big data, we design effective contour descriptors. Specifically, we first construct a star-convex contour matrix $\mathbf{A} = [\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_L]$ from L training objects. Then, we perform SVD of the matrix \mathbf{A} ,

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \quad (2)$$

where $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N]$ and $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_L]$ are orthogonal matrices and $\mathbf{\Sigma}$ is a diagonal matrix, composed of singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$. It is known that

$$\mathbf{A}_M = [\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_L] = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^\top + \dots + \sigma_M \mathbf{u}_M \mathbf{v}_M^\top \quad (3)$$

is the best rank- M approximation of \mathbf{A} [2].

In (3), each approximate contour $\tilde{\mathbf{r}}_i$ is given by a linear combination of the first M left singular vectors $\mathbf{u}_1, \dots, \mathbf{u}_M$. In other words,

$$\tilde{\mathbf{r}}_i = \mathbf{U}_M \mathbf{c}_i = [\mathbf{u}_1, \dots, \mathbf{u}_M] \mathbf{c}_i. \quad (4)$$

We refer to these vectors $\mathbf{u}_1, \dots, \mathbf{u}_M$ as *eigencontours*, and the space spanned by $\{\mathbf{u}_1, \dots, \mathbf{u}_M\}$ as the *eigencontour space*.

Given a contour \mathbf{r} , we project it onto the eigencontour space to obtain the low-rank approximation

$$\tilde{\mathbf{r}} = \mathbf{U}_M \mathbf{c} \quad (5)$$

where the coefficient vector \mathbf{c} is given by

$$\mathbf{c} = \mathbf{U}_M^\top \mathbf{r}. \quad (6)$$

In (6), an N -dimensional contour \mathbf{r} is optimally approximated by an M -dimensional vector \mathbf{c} in the eigencontour space, where $M < N$. Also, the approximate $\tilde{\mathbf{r}}$ can be reconstructed from \mathbf{c} via (5). Note that eigencontours may have negative elements. Thus, in rare cases, the approximate $\tilde{\mathbf{r}}$ has negative elements. In such cases, we truncate the negative elements to 0 to ensure the star-convexity of $\tilde{\mathbf{r}}$.

Clustering in eigencontour space: To discover typical contour patterns in a dataset, contour clustering can be performed. Instead of the original contour space of dimension N , contours can be grouped more effectively and more efficiently in the eigencontour space of dimension M . This is because the original space is transformed to the eigencontour space by an isometry \mathbf{U}_M^\top . Specifically, let $\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_L$ be object contours, which are approximated via (4). Then, it can be easily shown that

$$\|\tilde{\mathbf{r}}_i - \tilde{\mathbf{r}}_j\| = \|\mathbf{c}_i - \mathbf{c}_j\|. \quad (7)$$

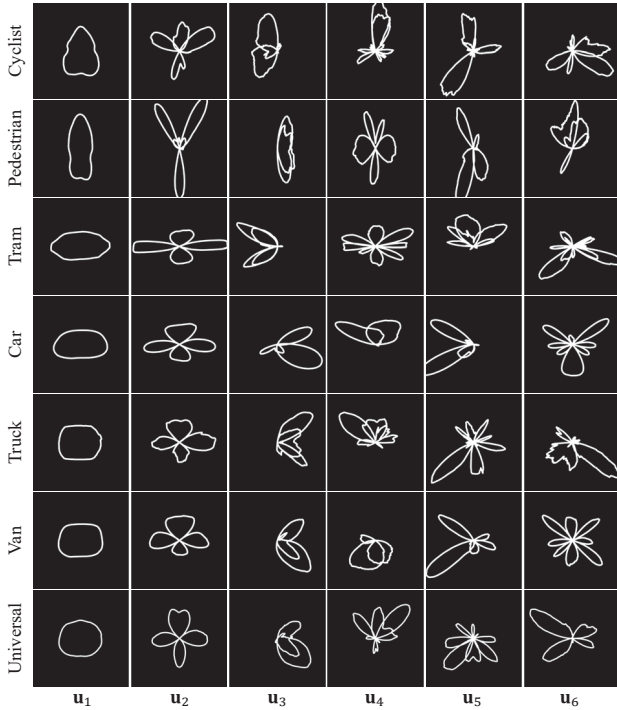


Figure 4. The first six eigencontours $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_6$ for the KINS dataset. The top six rows show the eigencontours for separate object categories in KINS, while the bottom one shows those for the universal set of all instances in the six categories.

In other words, the distances between contours in the original space are equal to those between the corresponding coefficient vectors in the eigencontour space. Hence, the clustering can be performed to yield the same results in both spaces, but it can be done more reliably and more efficiently in the eigencontour space because $M < N$. Note that, as the dimension of a space gets higher, clustering becomes more difficult because of the curse of dimensionality [1].

Regression in eigencontour space: Furthermore, it is also beneficial to find object contours in the eigencontour space. A contour regressor can be designed to detect object boundaries in images. To detect a star-convex contour in (1) in the original space, we should regress N variables. However, we can approximate all ground-truth contours of training objects using the first M eigencontours and train a network to regress M coefficients of \mathbf{c} in (6) in the eigencontour space. This approach requires the regression of fewer variables. Hence, the regression network also needs fewer parameters and is more efficient in both training and inference stages. The efficacy of the regression in the eigencontour space is demonstrated in Sections 4.2 and 4.3.

3.2. Examples and Analysis

Eigencontours: In this example, we use the KINS dataset [22], the instances of which are divided into seven cate-

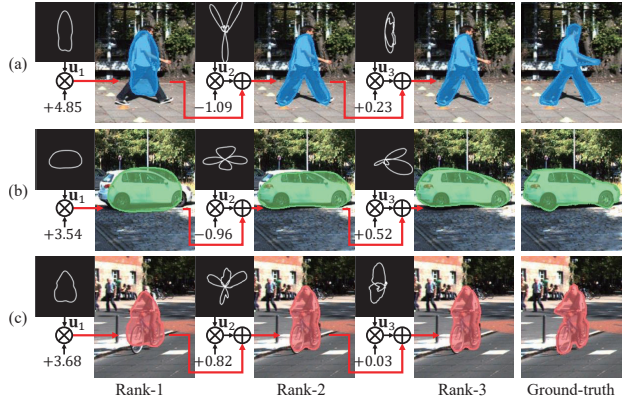


Figure 5. Illustration of linear combination of eigencontours.

gories. We determine the eigencontours for the six categories of ‘cyclist,’ ‘pedestrian,’ ‘tram,’ ‘car,’ ‘truck,’ and ‘van,’ respectively, except for ‘misc’ containing miscellaneous instances with unspecified classes. We also obtain the eigencontours for the universal set of all instances in the six categories. Each object boundary is represented by a 360-dimensional star-convex contour vector, by uniformly quantizing the 360-degree with an interval of 1° , *i.e.* $N = 360$.

Figure 4 shows the first six eigencontours $\mathbf{u}_1, \dots, \mathbf{u}_6$. For each category, the first eigencontour \mathbf{u}_1 describes rough outlines of typical instances. For example, most pedestrians stand or walk on sidewalks, as implied by the vertical shape of \mathbf{u}_1 for ‘pedestrian.’ By weighting \mathbf{u}_1 , the size of the shape can be controlled. Next, \mathbf{u}_2 is more complicated to represent detailed parts of instances. For ‘pedestrian,’ \mathbf{u}_2 is used to reconstruct a pair of legs, as shown in Figure 5(a). Also, \mathbf{u}_2 for ‘car’ generates a streamlined shape by refining the four sides of a car in Figure 5(b). The coefficient for \mathbf{u}_2 affects the horizontal and vertical sizes of the car. Similarly, \mathbf{u}_2 for ‘cyclist’ recovers bike wheels in Figure 5(c). In general, the coefficients for \mathbf{u}_1 and \mathbf{u}_2 are larger than those for the other eigencontours, and they are major factors for determining overall shapes. To represent those shapes more precisely, more eigencontours are required. Note that, for the three related categories of ‘car,’ ‘truck,’ and ‘van,’ \mathbf{u}_1 and \mathbf{u}_2 are similar to one another. Also, \mathbf{u}_1 for the universal set is a round shape to describe various instances in different categories.

Rank- M approximation: Figure 6 shows two object boundaries in the COCO2017 dataset [15] and their rank- M approximations. In this test, the eigencontours are determined for all training instances in all categories. The rank-1 approximations are not good enough; they represent the overall sizes of the objects only. The rank-2 approximations better reconstruct object shapes, but only roughly. As M gets larger, more faithful contours are restored. In this example, the objects have relatively complex shapes.

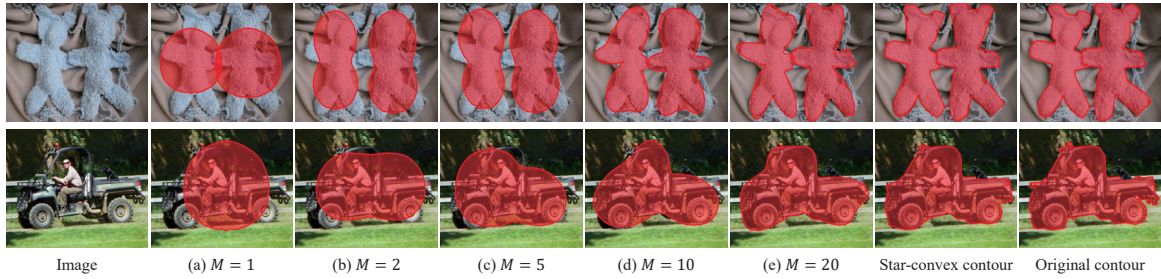


Figure 6. Object boundaries can be approximated using the first M eigencontours. As M gets larger, the rank- M approximations get closer to the star-convex conversions of the original contours.



Figure 7. Visualization of contour centroids in the 16-dimensional eigencontour space, according to object categories. Although the centroids are determined by grouping training data in the lower-dimensional space, each centroid represents the structure of the corresponding object category faithfully.

Hence, to represent their boundaries well, the rank-20 approximations are required, which are almost identical to the 360-dimensional star-convex contours. Although they cannot reconstruct the original contours perfectly, it is not because of the low-rank approximation, but because of the star-convex conversion. Note that, compared to the 360-dimensional star-convex contours, the rank-20 approximations reduce the dimensionality by a factor of 18.

Clustering in eigencontour space: For each of the six categories in the KINS dataset, we cluster the object boundaries in the 16-dimensional eigencontour space ($M = 16$) using the K -means algorithm, where K is set to 100. Figure 7 shows examples of contour centroids. We see that the centroids represent typical object shapes in the categories from different views. This indicates that eigencontours are

effective not only for representing individual contours faithfully, but also for clustering contours into typical patterns in a lower-dimensional space.

4. Experiments

4.1. Datasets

We use three datasets: KINS, SBD, and COCO2017. All these datasets were approved by institutional review boards.

KINS [22]: It is a dataset for amodal instance segmentation, built on the KITTI dataset [11]. It consists of 7,474 training and 7,517 test images. All instances are classified into seven categories, and an amodal segmentation mask is annotated for each instance.

SBD [12]: It is a semantic boundary dataset, re-annotated from the PASCAL VOC dataset [9]. Its 11,355 images are split to 5,623 training and 5,732 validation images. All instances are classified into 20 object categories. Each instance is annotated with its shape boundary without holes.

COCO2017 [15]: It is a large dataset for various tasks, such as object detection and segmentation. It contains 118K training images, 5K validation images, and 41K test images. The instance segmentation masks for objects in 80 categories are provided.

4.2. Comparative Assessment

Contour descriptors: It is desirable for contour descriptors to represent an object boundary compactly, as well as to reconstruct it faithfully. We compare the proposed eigencontours with the conventional contour descriptors [27,28]. For contour description, centroidal profiles are used in PolarMask [27], while polynomial fitting is performed to approximate the shape signature of a boundary in ESE-Seg [28]. In this test, the proposed eigencontours are determined for all instances in all categories in a training dataset.

For the quantitative assessment of contour descriptors, we employ the F-measure (\mathcal{F}) [20]. Specifically, bipartite matching is performed between the boundary points of a ground-truth contour and its approximated version. Then,

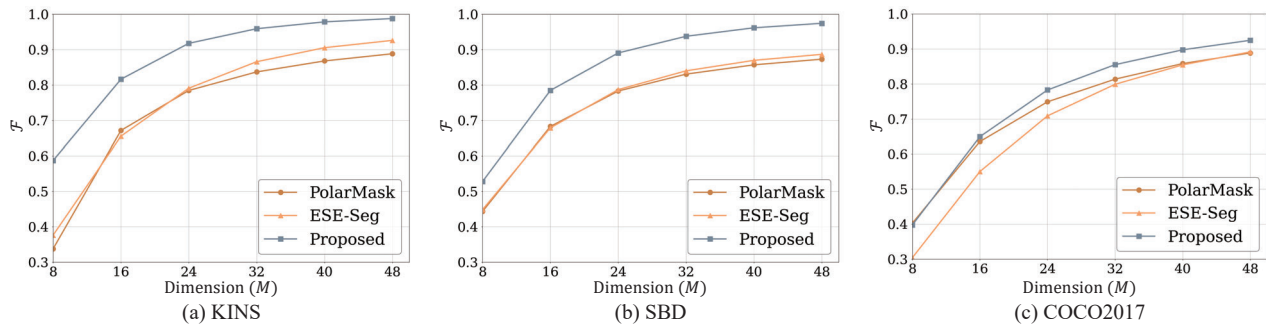


Figure 8. The \mathcal{F} score curves of the proposed eigencontours and the conventional contour descriptors in PolarMask [27] and ESE-Seg [28] according to the dimension M of the descriptors.



Figure 9. Qualitative comparison of boundary representations at $M = 16$. The left three images are from KINS, the middle three from SBD, and the remaining four from COCO2017.

the \mathcal{F} score is defined as the harmonic mean of the precision (\mathcal{P}) and the recall (\mathcal{R}) of the matching results.

Figure 8 compares the \mathcal{F} curves of the proposed eigencontours with those of the conventional descriptors according to the dimension M of the descriptors. In PolarMask, M radial coordinates in a centroidal profile are sampled to describe a contour. In ESE-Seg, M is the number of Chebyshev polynomial coefficients for approximating a contour. For all three datasets of KINS, SBD, and COCO2017, the proposed algorithm outperforms both PolarMask and ESE-Seg at every M . For KINS, the proposed algorithm achieves an \mathcal{F} score higher than 0.9 at $M = 24$, while the conventional ones need approximately double the dimension to yield a similar \mathcal{F} score. For SBD, similar tendencies are observed. For COCO2017, containing diverse instances with complicated shapes, the instances require higher-dimensional description than those in KINS and SBD. However, the proposed algorithm is still superior

Table 1. AUC- \mathcal{F} performances on KINS, SBD, and COCO2017.

	KINS	SBD	COCO2017
PolarMask [27]	75.47	76.23	74.05
ESE-Seg [28]	77.37	76.86	70.21
Proposed	89.17	86.51	76.92

to the conventional ones.

Table 1 compares the area under curve performances of the \mathcal{F} curves (AUC- \mathcal{F}) in Figure 8. The proposed algorithm outperforms the conventional algorithms by significant margins on all datasets. In other words, the proposed algorithm represents object boundaries more faithfully than the conventional algorithms, when the same number of coefficients are used for the contour description.

Figure 9 compares object boundaries approximated by the contour descriptors at $M = 16$. PolarMask fails to reconstruct curved parts. ESE-Seg provides better results,

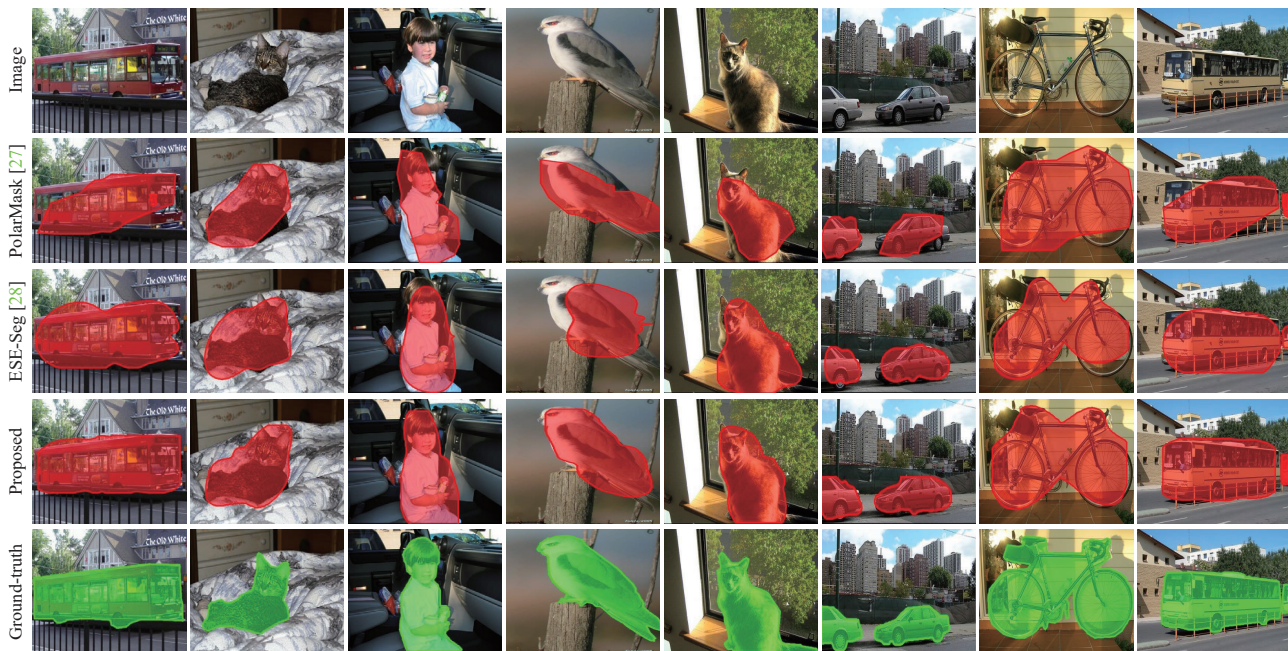


Figure 10. Comparison of instance segmentation results on the SBD dataset

Table 2. Comparison of the clustering performances on the COCO2017 dataset at $M = 16$ and $K = 500$.

	\mathcal{P}	\mathcal{R}	\mathcal{F}
PolarMask [27]	28.59	22.67	25.13
ESE-Seg [28]	30.31	24.31	26.82
Proposed	30.88	24.90	27.40

but it blurs complicated parts, especially the leg boundaries in the second and third columns. In contrast, the proposed eigencontour descriptors represent the object boundaries more accurately and more reliably.

Clustering in low-dimensional space: As mentioned in Section 3.1, it is possible to cluster object contours in a lower-dimensional descriptor space and obtain contour centroids there. To validate the effectiveness of the clustering in the proposed eigencontour space, we compare the clustering performances of the proposed algorithm on the COCO2017 dataset with those of PolarMask and ESE-Seg. To this end, we employ each algorithm to approximate all training boundaries into M -dimensional descriptors and obtain K centroids via K -means. Then, each contour in the dataset is matched with the nearest centroid, and the matching performance is computed in terms of \mathcal{P} , \mathcal{R} , and \mathcal{F} .

Table 2 compares the performances at $M = 16$ and $K = 500$. The proposed algorithm yields the best results in all three metrics, which indicates that the proposed algorithm can process object contours more reliably in a low-dimensional space. Qualitative comparison results of the

Table 3. Comparison of the AP_{50} , AP_{75} , and $AP_{\mathcal{F}}$ performances on the SBD validation dataset.

	AP_{50}	AP_{75}	$AP_{\mathcal{F}}$
PolarMask [27]	50.11	14.50	25.78
ESE-Seg [28]	52.14	20.48	27.37
Proposed	56.47	29.35	35.30

clustering are available in the supplemental document.

Instance segmentation: Both PolarMask and ESE-Seg were proposed for instance segmentation. To localize each instance, these methods reformulate the pixelwise classification as the regression of an object contour. The proposed eigencontours are more effective for this instance segmentation task as well. To demonstrate this, as done in ESE-Seg, we adopt YOLOv3 [23] as an object detector and modify its components. Given an input image, we predict an output map, in which each element contains an M -dimensional coefficient vector as well as the original YOLOv3 vector for bounding box regression and object classification. Then, we use the coefficient vector to linearly combine eigencontours to reconstruct the contour and shape mask of an object. The supplemental document describes the implementation details and the training procedure.

Table 3 compares the instance segmentation results on the SBD validation dataset at $M = 20$. The average precision (AP) performances, based on two intersection-over-union (IoU) thresholds of 0.5 and 0.75 and an \mathcal{F} score threshold of 0.3, are reported. The proposed algorithm per-



Figure 11. Instance segmentation results of the proposed algorithm on the SBD dataset.

forms better than PolarMask and ESE-Seg in terms of all three metrics. Figure 10 shows boundary regression results. PolarMask and ESE-Seg fail to reconstruct object boundaries reliably. In contrast, the proposed algorithm represents the boundaries more faithfully. Figure 11 shows more instance segmentation results.

4.3. Analysis

Dimension of eigencontour space (M): Table 4 lists the AUC- \mathcal{F} performances of the proposed algorithm on the SBD validation dataset according to the dimension, M , of the eigencontour space. At $M = 10$, the proposed algorithm yields poor scores, since object boundaries are too simplified and not sufficiently accurate. At $M = 20$, it provides the best results. At $M = 30$, it yields similarly good results. However, at $M = 40$, the performances are degraded further, which indicates that a high-dimensional space does not always lead to better results. It is more challenging to regress more variables reliably. There is a trade-off between accuracy and reliability. In this test, $M = 20$ achieves a good tradeoff.

Categorical eigencontour space: The proposed eigencontours are data-driven descriptors, which depend on the distribution of object contours in a dataset. Thus, different eigencontours are obtained for different data. Let us consider two options for constructing eigencontour spaces: categorical construction and universal construction. In the categorical construction, eigencontours are determined for each category in a dataset. In the universal construction, they are determined for all instances in all categories.

For the two options, \mathcal{F} score curves are presented according to the dimension M in the supplemental document. Table 5 compares the area under curve performances of the \mathcal{F} curves up to $M = 18$. The categorical construction pro-

Table 4. The instance segmentation performances of the proposed algorithm on the SBD validation dataset according to the dimension of the eigencontour space (M).

M	AP ₅₀	AP ₇₅	AP \mathcal{F}
10	49.96	26.19	30.27
20	56.47	29.35	35.30
30	55.85	28.89	36.15
40	54.45	22.42	32.21

Table 5. Comparison of the AUC- \mathcal{F} performances of categorical and universal eigencontours, $M \in [3, 18]$.

	KINS	SBD	COCO2017
Universal	64.65	60.22	47.19
Categorical	67.67	62.37	48.77

vides better performances than the universal construction, because it considers similar shapes in the same category only. In COCO2017, the gap between the two options is the smallest. This is because some object shapes are not properly represented due to occlusions and thus COCO2017 objects exhibit low intra-category correlation. In contrast, in KINS, whole contours are well represented because occluded regions are also annotated. Hence, the gap between the two options is the largest.

Limitations: The proposed eigencontours represent typical contour patterns in a dataset. Thus, if object contour patterns differ among datasets, the eigencontours for a dataset may be effective for that particular dataset only. To assess the dependency of eigencontours on a dataset, we conduct cross-validation tests between datasets in the supplemental document.

5. Conclusions

We proposed novel contour descriptors, called eigencontours, based on low-rank approximation. First, we constructed a contour matrix containing all contours in a training set. Second, we approximated the contour matrix, by performing the best rank- M approximation. Third, we represent an object boundary by a linear combination of the M eigencontours. Experimental results demonstrated that the proposed eigencontours can represent object boundaries more effectively and more faithfully than the existing methods. Moreover, the proposed algorithm yields meaningful instance segmentation performances.

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grants funded by the Korea government (MSIT) (No. NRF-2021R1A4A1031864 and No. NRF-2022R1A2B5B03002310).

References

- [1] Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966. 4
- [2] Avrim Blum, John Hopcroft, and Ravindran Kannan. *Foundations of Data Science*. 2015. 1, 2, 3
- [3] Bao X. Chen and John K. Tsotsos. Fast visual object tracking using ellipse fitting for rotated bounding boxes. In *Proc. ICCV Workshops*, 2019. 1, 2
- [4] Gene C.-H. Chuang and Chung-Chieh J. Kuo. Wavelet descriptor of planar curves: Theory and applications. *IEEE Trans. Image Process.*, 5(1):56–70, 1996. 1, 2
- [5] Luigi Cinque, Stefano Levialdi, and Alessio Malizia. Shape description using cubic polynomial Bezier curves. *Pattern Recog. Lett.*, 19(9):821–828, 1998. 1, 2
- [6] Edward R. Davies. *Machine Vision: Theory, Algorithms, Practicalities*. Elsevier, 2004. 1, 2, 3
- [7] Paul Dierckx. *Curve and Surface Fitting with Splines*. Oxford University Press, 1995. 1
- [8] Gregory Dudek and John K. Tsotsos. Shape representation and recognition from multiscale curvature. *Comput. Vis. Image Understand.*, 68(2):170–189, 1997. 1, 2
- [9] Mark Everingham, Luc V. Gool, Christopher K.I. Williams, John Winn, and Andrew Zisserman. The Pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010. 5
- [10] Herbert Freeman and John A. Saghri. Generalized chain codes for planar curves. In *Proceedings of the 4th International Joint conference on Pattern Recognition*, 1978. 1, 2
- [11] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *Proc. CVPR*, 2012. 5
- [12] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proc. ICCV*, 2011. 2, 5
- [13] Dongkwon Jin, Wonhui Park, Seong-Gyun Jeong, Heeyeon Kwon, and Chang-Su Kim. Eigenlanes: Data-driven lane descriptors for structurally diverse lanes. In *Proc. CVPR*, 2022. 2
- [14] David C. Lay. *Linear Algebra and Its Applications*. Pearson, 2007. 2
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and Charles L. Zitnick. Microsoft COCO: Common objects in context. In *Proc. ECCV*, 2014. 2, 4, 5
- [16] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Arbeláez, and Luc V. Gool. Convolutional oriented boundaries: From image segmentation to high-level tasks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017. 1
- [17] Farzin Mokhtarian and Alan K. Mackworth. A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(8):789–805, 1992. 1, 2
- [18] Peter J.V. Otterloo. *A Contour-Oriented Approach to Shape Analysis*. Prentice Hall International (UK) Ltd., 1991. 1, 2
- [19] Sida Peng, Wen Jiang, Huaijin Pi, Xiuli Li, Hujun Bao, and Xiaowei Zhou. Deep snake for real-time instance segmentation. In *Proc. CVPR*, 2020. 1
- [20] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc V. Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proc. CVPR*, 2016. 5
- [21] Juan-Carlos Perez and Enrique Vidal. Optimum polygonal approximation of digitized curves. *Pattern Recognition Letters*, 15(8):743–750, 1994. 1, 2
- [22] Lu Qi, Li Jiang, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Amodal instance segmentation with KINS dataset. In *Proc. CVPR*, 2019. 1, 2, 4, 5
- [23] Joseph Redmon and Ali Farhadi. YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 1, 7
- [24] Xing Shen, Jirui Yang, Chunbo Wei, Bing Deng, Jianqiang Huang, Xian-Sheng Hua, Xiaoliang Cheng, and Kewei Liang. Dct-mask: Discrete cosine transform mask representation for instance segmentation. In *Proc. CVPR*, 2021. 2
- [25] Jamie Shotton, Andrew Blake, and Roberto Cipolla. Multi-scale categorical object recognition using contour fragments. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2008. 1
- [26] Jean C. Stanek. A characterization of starshaped sets. *Canadian Journal of Mathematics*, 29(4):673–680, 1977. 3
- [27] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. PolarMask: Single shot instance segmentation with polar representation. In *Proc. CVPR*, 2020. 1, 2, 5, 6, 7
- [28] Wenqiang Xu, Haiyang Wang, Fubo Qi, and Cewu Lu. Explicit shape encoding for real-time instance segmentation. In *Proc. ICCV*, 2019. 1, 2, 3, 5, 6, 7
- [29] Yong Xu, Yuhui Quan, Zhuming Zhang, Hui Ji, Cornelia Fermüller, Morimichi Nishigaki, and Daniel Dementhon. Contour-based recognition. In *Proc. CVPR*, 2012. 1
- [30] Ian T. Young, Joseph E. Walker, and Jack E. Bowie. An analysis technique for biological shape. *Information and control*, 25(4):357–370, 1974. 1, 2
- [31] Dengsheng Zhang and Guojun Lu. A comparative study of Fourier descriptors for shape representation and retrieval. In *Proc. ACCV*, 2002. 1, 2
- [32] Dengsheng Zhang and Guojun Lu. Review of shape representation and description techniques. *Pattern Recognition*, 37(1):1–19, 2004. 1
- [33] Rufeng Zhang, Zhi Tian, Chunhua Shen, Mingyu You, and Youliang Yan. Mask encoding for single shot instance segmentation. In *Proc. CVPR*, 2020. 2
- [34] Mingmin Zhen, Jinglu Wang, Lei Zhou, Shiwei Li, Tianwei Shen, Jiaxiang Shang, Tian Fang, and Long Quan. Joint semantic segmentation and boundary detection using iterative pyramid contexts. In *Proc. CVPR*, 2020. 1
- [35] Xingyi Zhou, Jiacheng Zhuo, and Philipp Krahenbuhl. Bottom-up object detection by grouping extreme and center points. In *Proc. CVPR*, 2019. 1, 2