

Automatic individual identification of Saimaa ringed seals

ISSN 1751-9632
 Received on 28th February 2017
 Revised 22nd September 2017
 Accepted on 17th October 2017
 doi: 10.1049/iet-cvi.2017.0082
 www.ietdl.org

Tina Chehrsimin¹, Tuomas Eerola¹ ✉, Meeri Koivuniemi², Miina Auttila^{2,3}, Riikka Levänen², Marja Niemi², Mervi Kunnasranta^{2,4}, Heikki Kälviäinen¹

¹Machine Vision and Pattern Recognition Laboratory, School of Engineering Science, Lappeenranta University of Technology, Lappeenranta, Finland

²Department of Environmental and Biological Sciences, University of Eastern Finland, Joensuu, Finland

³Parks & Wildlife Finland, State Forest Enterprise (Metsähallitus), Savonlinna, Finland

⁴Natural Resources Institute Finland, FI-80100 Joensuu, Finland

✉ E-mail: tuomas.eerola@lut.fi

Abstract: In order to monitor an animal population and to track individual animals in a non-invasive way, identification of individual animals based on certain distinctive characteristics is necessary. In this study, automatic image-based individual identification of the endangered Saimaa ringed seal (*Phoca hispida saimensis*) is considered. Ringed seals have a distinctive permanent pelage pattern that is unique to each individual. This can be used as a basis for the identification process. The authors propose a framework that starts with segmentation of the seal from the background and proceeds to various post-processing steps to make the pelage pattern more visible and the identification easier. Finally, two existing species independent individual identification methods are compared with a challenging data set of Saimaa ringed seal images. The results show that the segmentation and proposed post-processing steps increase the identification performance.

1 Introduction

Wildlife photo-identification (Photo-ID) provides a tool to study and to monitor animal populations over time based on captured images of individuals. It has various applications in studying key aspects of the populations such as survival, dispersal, site fidelity, reproduction, health, and population size and density. Due to its non-invasive nature, the Photo-ID is a great alternative to more invasive techniques such as tagging that requires catching the animal and may cause stress to it, as well as may change its behaviour or increase mortality. The identification of individuals is based on distinctive permanent characteristics, such as fur patterns, pigmentations, scars, or shape. Traditionally, the identification has been performed manually by researchers. However, due to the rapid increase in the amount of image data, there is a demand for automated methods. Computer vision techniques provide an attractive tool to replace the laborious and time-consuming manual work.

This work focuses on the Saimaa ringed seal (*Phoca hispida saimensis*) which is an endangered subspecies of ringed seal only found in Lake Saimaa in Finland. At present, around 360 seals inhabit the lake, and on the average 60–85 pups are born annually. This small and fragmented population is threatened by various anthropogenic factors, especially by-catch and climate change [1]. The long-term and accurate assessment of the population is needed for conservation purposes, and camera trapping has been recently

launched as a new monitoring tool for the Saimaa ringed seal [2, 3]. Camera trapping is an especially effective method for the Saimaa ringed seal monitoring due to high site fidelity of the seals. The pelage of the ringed seal contains a distinctive patterning of dark spots surrounded by light grey rings (see Fig. 1). These patterns are unique to each seal, enabling the identification of individuals over their whole lifetime. Photo-ID data of the Saimaa ringed seal is mostly collected by game cameras during the moulting season in the spring [3]. The camera trapping produces a large number of seal images which have been identified manually until these days.

This work continues the study presented in [4] where the first steps towards automatic individual identification of the Saimaa ringed seal were taken. In the paper, a novel supervised segmentation method for the ringed seal and a simple identification method exploiting texture features were proposed. In this work, the segmentation method is further developed to decrease its computation time without sacrificing the performance. Moreover, a set of post-processing operations for segmented images are proposed to make the seals easier to identify. Finally, two animal identification methods are evaluated with a challenging data set of Saimaa ringed seal images to demonstrate the importance of the segmentation and post-processing operations.

2 Related work

Several approaches for automatic image-based animal identification can be found in the literature. Methods have been developed, for example, for polar bears [5], cattle [6], newts [7], giraffes [8], salamanders [9], snakes [10], insects [11], and marine mammals [12]. All of these methods use image processing and pattern recognition techniques to identify individuals. Most of the studies limit the individual identification to a certain animal species or species groups. For example, in [8], the effectiveness of wild-ID software in identifying individual Thornicroft's giraffes from a data set of 552 images was studied. The approach uses the scale invariant feature transform (SIFT) algorithm [13] to extract and match distinctive image features regardless of scale and orientation.



Fig. 1 Saimaa ringed seal

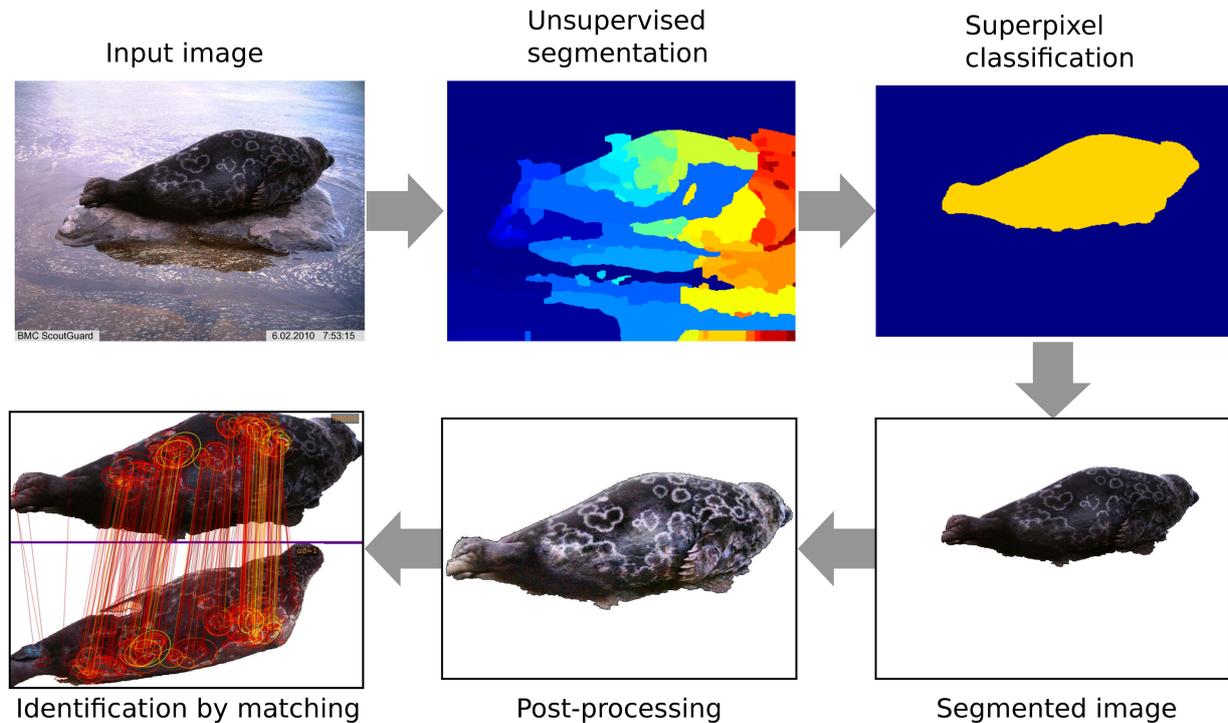


Fig. 2 Proposed method

In [14], a semiautomated method to match humpback whale flukes is presented. The method starts by detecting the fluke with the help of a user input and then computes various features such as spots, shape, texture, and damage areas that are used for the matching. In [15], another automatic method for matching of the humpback whale fluke patterns is presented. The method utilises morphological operators to obtain binary salient patterns that are matched using an iterative matching strategy and pattern similarity scores.

In [7], the suitability of biometric techniques for the identification of the great crested newts was investigated. Distinctive belly patterns were used to compare images of newts with an image database. Two different methods were used for the comparison: (i) the correlation coefficient of the greyscale pixel intensities, and (ii) the Hamming distance between the binary image segments.

All the above methods were developed for one species only and as such are not generalisable to Saimaa ringed seals. In [4], the first steps towards automatic individual identification of the Saimaa ringed seals were taken. The paper proposes a segmentation method for the Saimaa ringed seals using unsupervised segmentation and texture-based superpixel classification was proposed. Furthermore, a simple texture-based approach for the ringed seal identification was evaluated. However, the identification performance of the method is not good enough for most practical applications.

There have been also research efforts towards creating a unified approach applicable for identification purposes for several animal species. For example, in [16], the HotSpotter method to identify individual animals in a labelled database was presented. This algorithm is not species specific and has been applied to Grevy's and plain zebras, giraffes, leopards, and lionfish. HotSpotter uses viewpoint invariant descriptors and a scoring mechanism that emphasises the most distinctiveness keypoints and descriptors. In [17], a species recognition algorithm based on sparse coding spatial pyramid matching was proposed. It was shown that the proposed object recognition techniques can be successfully used to identify animals on sequences of images captured using camera traps in nature. One of the problems with the species independent individual identification methods is that they do not provide an automatic method to detect the animals in images. Therefore, either a manual detection or development of a detection method for the studied animal is needed. Furthermore, typically a higher

identification performance can be obtained by tuning the identification method for one species only.

3 Methodology

The proposed method for the identification of Saimaa ringed seals consists of three main parts: segmentation, post-processing, and identification (Fig. 2). First, the segmentation is applied to detect the seal and to eliminate the background that could complicate the identification process. Saimaa ringed seals tend to use same sites or areas inter-annually for moulting and hauling out, and the images are captured using static camera traps. Therefore, the same seal is often captured with the same background increasing the risk that a supervised identification algorithm learns to 'identify' the background instead of the actual seal if the full image or the bounding box around the seal is used. This may further lead to a system that is not able to identify the seal in a new environment making the separation of the animal from the background (segmentation) important. After the segmentation, the segmentation result is post-processed and the seal is identified based on a matching-based approach.

3.1 Segmentation phase

Automatic segmentation of animals is often difficult due to the camouflage colours of animals, i.e. the colouration and patterns are similar to the visual background of the animal. This makes it difficult to define a single criterion to distinguish the animal from the background. The proposed method to segment ringed seals starts with an unsupervised segmentation step to divide the image into segments or superpixels. After this, the superpixels are classified into two different classes: the seal or the background. A similar approach has been used, for example, to segment the optical disc in retinal images [18]. The first criterion is that ideally superpixels should be completely inside or outside the object, i.e. there should not be such superpixels that contain both seal and background pixels. Furthermore, the second criterion is that superpixels should be large enough to make the classification possible. Generally, the larger the superpixels are, the better, as long as they fulfil the first criterion.

In [4], the method proposed in [19, 20] was used for unsupervised segmentation of the Saimaa ringed seals. The method has been shown to produce the state-of-the-art performance in the Berkeley Segmentation Dataset [21]. It combines the globalised

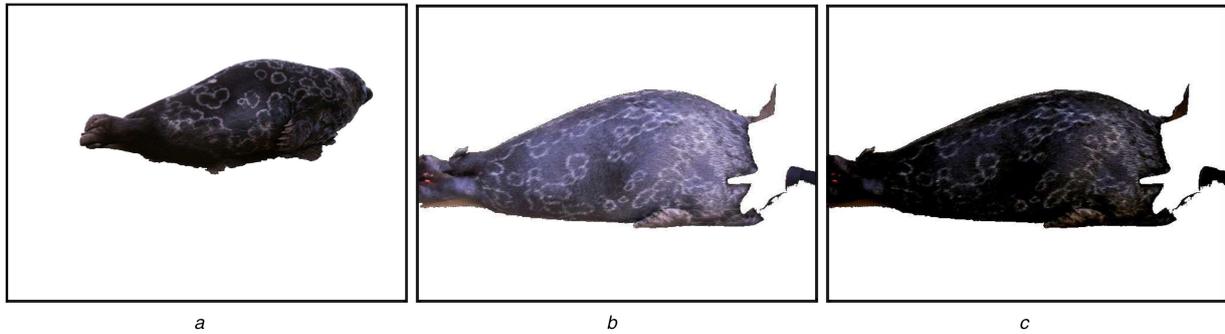


Fig. 3 Example of the histogram matching
 (a) Reference image, (b) Input image before normalisation, (c) Input image after normalisation

probability of boundary (gPb) detector, oriented watershed transform, and ultrametric contour map (UCM) to produce a weighted contour image that can be thresholded to produce superpixels. The size of the superpixels can be controlled by varying the threshold value. By selecting a proper value, the both criteria mentioned above can be fulfilled. However, the problem with the method is that it requires solving of a large, sparse eigensystem which takes a lot of time and computational efforts.

To keep the computation time low, we use multiscale combinatorial grouping (MCG) [22] to produce the weighted contour image. The method utilises an efficient normalised cuts algorithm to achieve a $20\times$ speed improvement without loss in performance. In practice, due to the large memory requirements of the gPb detector, large images need to be processed in pieces which increases the difference between the computation times even more. In our experiments, the computation time with the method used in [4] was 193 s for a 640×480 image while with the MCG-based method the computation time was 4.5 s. The experiment was run with a PC with i7-4600U (2.10 GHz) CPU and 8 GB of memory.

After the unsupervised segmentation step, feature descriptors are computed from the superpixels and each superpixel is classified into the seal or background class in a supervised manner. For this purpose, we utilise colour features (means of intensities of each colour component), local variation features, and two different texture features: blur invariant local phase quantisation (LPQ) [23] and local binary pattern histogram Fourier descriptors (LBP-HF) [24]. Local variation is computed from an intensity image as means of local range, local standard deviation, and local entropy of the segment with a 3×3 neighbourhood.

The LPQ descriptor uses locally computed phase information and decorrelated low-frequency coefficients to create a code word histogram that can be used for the texture classification. Since the low-frequency phase components are ideally blur invariant and the phase information is invariant to uniform illumination changes, the LPQ texture descriptor is a good choice for processing camera trap images which often are out of focus and suffer from large illumination changes between day and night time.

LBP-HF is a modified version of the well-known local binary pattern (LBP) descriptors [25]. The LBP descriptor is obtained by thresholding the neighbourhood of each pixel with the pixel value, presenting the result as a binary number, and computing histogram of binary numbers. LBP-HF is computed as a discrete Fourier transforms of the obtained non-invariant LBP histogram. One of the main advantages of this descriptor is that it is rotation invariant while it preserves the highly discriminative nature of LBP histograms.

Finally, the support vector machine (SVM) classifier is used to classify the superpixels. The comparison of different feature combinations and classifiers is presented in Section 4.2. After the superpixel classification step, all the superpixels that are classified to the seal class are combined to form the seal segment for identification purposes.

3.2 Segmentation post-processing

Due to errors in the segmentation process caused by non-uniform superpixels or classification errors, some pixels are incorrectly

marked as background or seal causing holes inside the seal region and multiple seal segments. Such errors have a negative effect on the identification performance. To avoid this, we apply morphological operations on the mask obtained by image segmentation. First, the closing operation is applied to fill the holes inside the seal segment. Next, the opening operation is applied to remove small isolated segments that can be considered as noise.

3.3 Colour normalisation

The same seal in different illumination conditions captured with different cameras appears in different colours which is a problem for identification algorithms. Since the colour of Saimaa ringed seals does not vary considerably, the seal segment should always have approximately similar colour histograms if the illumination is constant. Therefore, we utilise histogram matching to remove the effect of varying illumination conditions and colour differences caused by different cameras. Fig. 3 shows an example result of the histogram matching.

3.4 Contrast enhancement

The next step is to increase the contrast between dark spots surrounded by light grey rings to make the patterns essential for the identification more visible. For this we use the well-known contrast-limited adaptive histogram equalisation [26] algorithm. Fig. 4 shows an example result of the contrast enhancement.

3.5 Identification phase

For identification, we utilise two existing generic identification methods: Wild-ID [8] and HotSpotter [17]. The Wild-ID method consists of three main steps. The first step is to extract the SIFT features [14] for each image in a database. In this step, potential interest points are identified by utilising the scale and orientation invariant difference of a Gaussian function. Then, the local image gradients are measured at the selected scale in the regions near each interest point. In the second step, the candidate matched pairs of SIFT features are identified from two images by locating features of one image in the other image and minimising the Euclidean distance between the feature descriptors. In the third step, to choose reasonable matches among all potential matches the RANSAC algorithm [27] is employed to obtain a geometrically consistent subset of candidate matches.

HotSpotter [17] uses the Hessian-Hessian interest point detector, viewpoint invariant descriptors, and a scoring mechanism that emphasises the most distinctive keypoints and descriptors. For the identification, two matching algorithms are proposed: one-versus-one and one-versus-many matching. In the one-versus-one matching, the query image matches against each database image separately and it sorts out the database images by the similarity score to find the final rank of the result. In the one-versus-many matching, each descriptor is matched from the query image against all descriptors from the image database. In this method, a fast, approximate nearest neighbour search data structure is utilised. The scores are generated for each image according to these matches using the local naive Bayes nearest neighbour methods [28], and finally the scores are aggregated to generate the final similarity

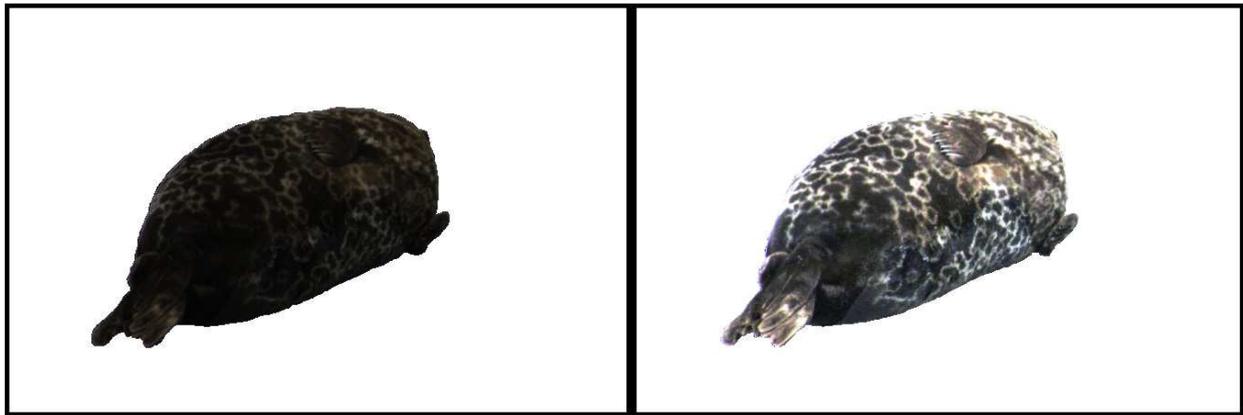


Fig. 4 Image before and after the contrast enhancement

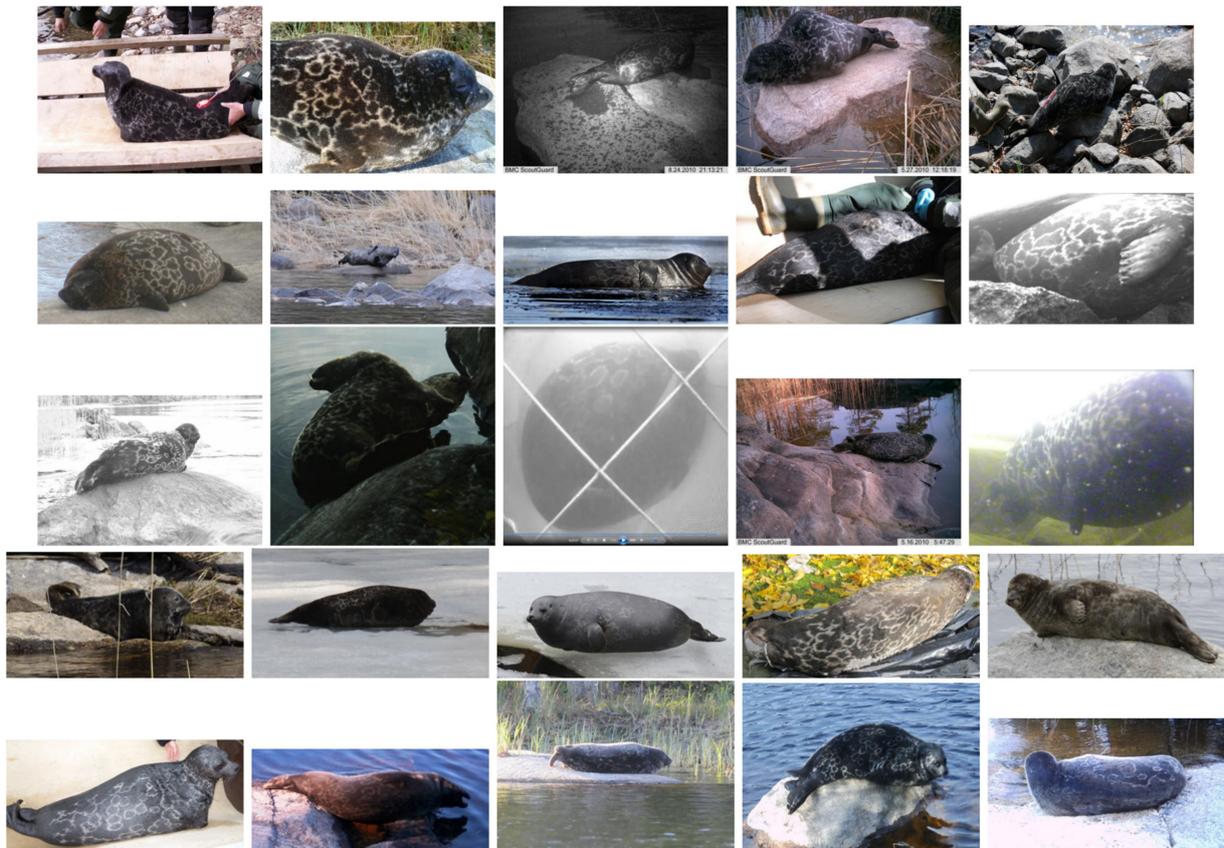


Fig. 5 Examples from the Saimaa ringed seals database

scores for each image. Based on the similarity scores all the images in the database are ranked from the best to the worst match.

4 Experiments

4.1 Data

To evaluate the method, a unique Photo-ID database of Saimaa ringed seal images collected by University of Eastern Finland (UEF) was used. The database contained total of 892 images of 147 individual seals. Most of the images contain one individual Saimaa ringed seal, and only few images contain two or more individuals. Example images from the database are shown in Fig. 5.

The database contained images of individual seals from the right and left sides because both flanks of the ringed seal have different pelage patterns. In addition, the database contained images from belly and back sides of the seals. However, most of the individual seals in the database had only one or few images. The selected identification methods require that the database contains at least one image of the same individual captured from

the same side as the query image. Therefore, the seals with only one image were omitted from the experiments. Also, images with multiple seals or too low quality for identification purposes were screened out from the data set. These images include underwater images, images with large obstacles, and images captured with low-end cameras from a long distance. It should be noted that, in practice, most of the image material collected for the monitoring purposes is obtained using automatic camera traps. These images have typically sufficient image quality for the identification purposes. The final data set contained 131 ringed seal individuals and total of 591 images with reasonable quality.

The seals were identified from the images by biologists from UEF to form the ground truth for identification. The segmentation ground truth was constructed by manually drawing the contour of the animals for 280 images (40 individuals). Moreover, to train and to evaluate the superpixel classifiers, manual labelling of superpixels was performed for all 591 images.

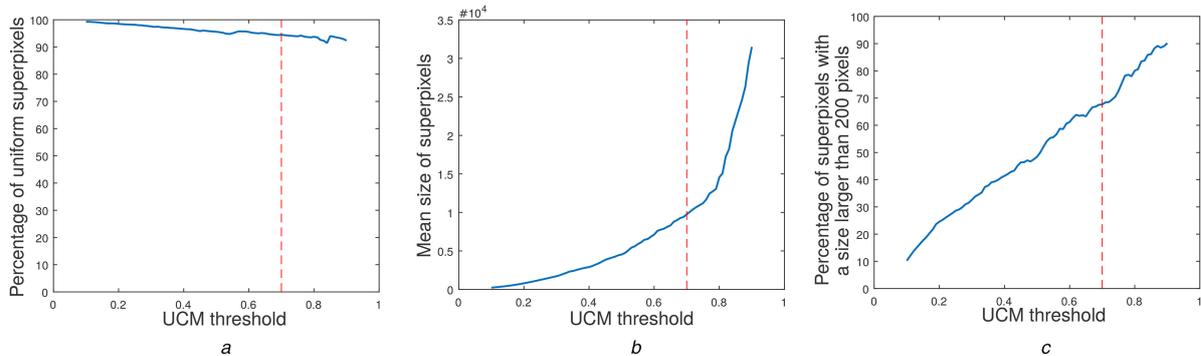


Fig. 6 Effect of the UCM threshold

(a) Percentage of uniform superpixels, (b) Mean size of superpixels, (c) Percentage of superpixels with a size >200 pixels

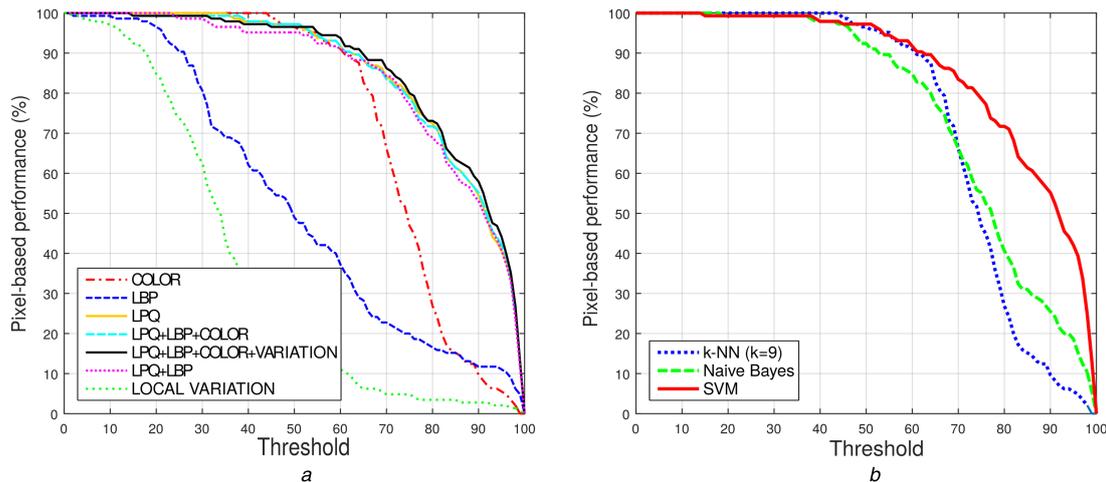


Fig. 7 Segmentation performance

(a) Comparison of feature descriptors with the SVM classifier, (b) Comparison of the classifiers (the best set of features selected for each classifier)

4.2 Segmentation experiments

The MCG segmentation algorithm [22] produces an UCM that is a weighted contour image. UCM is thresholded to produce an edge image that defines the superpixels. The size of the superpixels is controlled by varying the threshold value. When the threshold value is increased, similar neighbouring regions separated by an edge with a low weight are combined, causing a smaller number of larger superpixels. The success of the ringed seal segmentation depends highly on the selected threshold value. To select the value, two criteria should be considered: (i) superpixels that contain both the seal pixels and the background pixels should be minimised, and (ii) the size of the superpixels should be as large as possible without violating the first criteria to make the superpixels classification as robust as possible.

The evaluation of the UCM thresholds was carried out using the manually segmented 280 images. Three performance measures were selected to compare the threshold values: (i) the percentage of uniform superpixels, (ii) the mean size of superpixels, and (iii) the percentage of superpixels larger than 200 pixels. The third performance measure was selected based on an observation that superpixels smaller than 200 pixels do not usually contain enough texture information to be reliably classified. The performance measures were computed for all threshold values from 0.1 to 0.9

Table 1 Description of the features

Feature	Description	Length
mean colours	mean intensity of R, G, B colour channels	3
local variation	mean of local range, local standard deviation, and local entropy	3
LBP-HF	discrete Fourier transforms of local binary pattern histograms	38
LPQ	local phase quantisation	256

with a step size of 0.01. The results are presented in Fig. 6. It should be noted that the selection of the UCM threshold is always a trade-off between the size of the superpixels (larger superpixels are easier to classify) and the uniformity of the superpixels, and there is no unambiguous way to select the threshold value. In order to keep the number of non-uniform superpixels small while limiting the amount of small superpixels, the threshold value of 0.7 was selected for the further experiments.

In order to train and to test the superpixel classification methods, the unsupervised segmentation was performed to all images with the selected threshold value, and resulting superpixels were manually labelled into two classes: the seal and the background. The images were divided into the training set and the test set randomly. Twenty images were selected for the test set and the remaining 571 images were used as the training set. This was repeated 11 times, so in total the segmentation performance was tested with 220 images. Table 1 lists the evaluated features and Fig. 7a shows the results for different feature combinations with the SVM classifier with a quadratic kernel. The x-axis presents the threshold value for how many per cent of pixels needs to be correctly labelled (the seal or the background) that the image can be considered as correctly segmented and the y-axis presents the percentage of correctly segmented images. It can be seen that colour-based features and LBP-HF are alone not enough to obtain good superpixel classification accuracy, but both are useful when combined with the LPQ features and the best classification result is obtained by combining all the features. Fig. 7b presents the result for different classifiers. For the k -NN classifier k was experimentally set to 9 and for the naive Bayes classifier Gaussian distributions were used. The prior probabilities were obtained from the relative frequencies of the two classes in the training data. For each classifier, the best set of features was used. The results show that SVM outperforms the k -NN and naive Bayes classifiers.

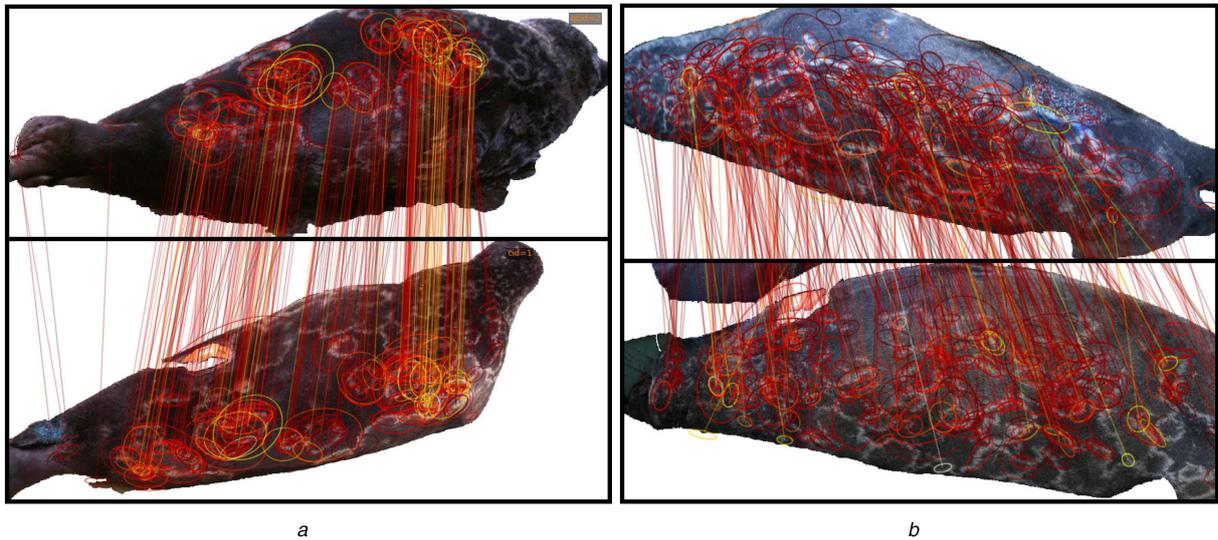


Fig. 8 Identification examples with query images (top row) and the top ranked images (bottom row)

(a) Correct identification (the top ranked seal is the same seal as in the query image), (b) Incorrect identification (the top ranked seal is different seal as in the query image)

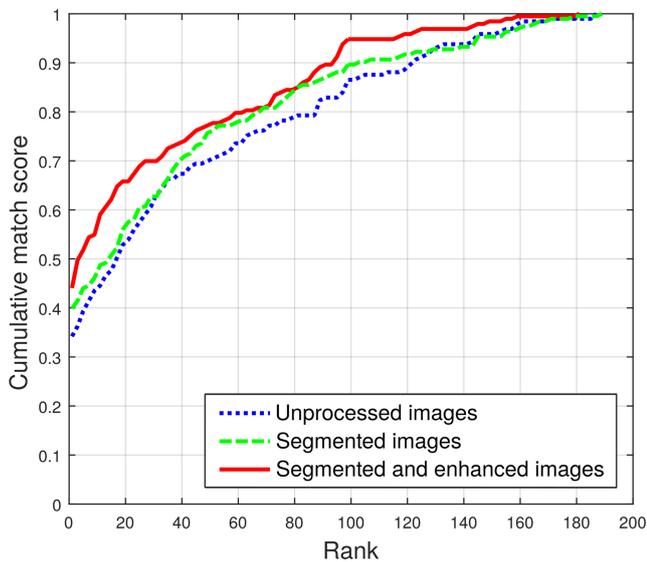


Fig. 9 Identification performance of the HotSpotter algorithm

4.3 Identification experiments

To evaluate the performance of the identification methods, 194 images were selected from the 220 images used to evaluate the segmentation method by removing the individuals that had only one image in the remaining set. Minimum of two images per individual are needed for evaluation purposes. To evaluate the performance of the identification methods, the database must contain at least one image of the same individual as in the query image. The remaining set contains 57 individuals with two to eight images per individual. On the average there are 3.4 images per individual (the median is 3).

The performance of the identification algorithms was evaluated separately for the original images, segmented images, and segmented and post-processed images. The identification methods were applied using the leave-one-out cross-validation, i.e. each of the 194 images was considered as a query image one by one and the remaining 193 images were used as a database.

With HotSpotter 44% of the segmented and post-processed images, 40% of segmented images, and 34% of the original images were correctly identified. This shows that both segmenting and post-processing of the images increase the performance compared with the original unprocessed images. With Wild-ID 13% of the segmented and post-processed images and 8% of the original images were correctly identified, which means that the HotSpotter outperformed it by a large margin. The results obtained using the

HotSpotter are also better than the preliminary results reported in [4] where only 10% of the seals were correctly identified, despite the fact that the number of seal individuals was smaller (40). Fig. 8 shows examples of correct and incorrect identifications with HotSpotter.

The performance of the HotSpotter algorithm was further measured using the cumulative match score (CMS) histogram commonly used in the face recognition research [29]. It measures how well the identification system ranks the identities in the database with respect to the input image. The N th bin in the CMS histograms tells the percentage of test images where the correct individual seal was in the set of the N best matches proposed by the identification algorithm. Fig. 9 presents the results for the HotSpotter method.

As can be seen from Fig. 9, with the segmented and post-processed images in 66% of the cases, the correct seal was within the 20 best matches. For the original and segmented images, the correct seal was within the best 20 matches in 57 and 53% of the cases, respectively. The Wild-ID identification software [8] does not produce the rank for the matches and therefore it was not possible to generate the CMS histogram.

5 Conclusion

In this paper, a segmentation method for Saimaa ringed seals using unsupervised segmentation and texture-based superpixel classification was proposed. Moreover, post-processing steps including morphological processing of the segmentation mask, colour normalisation, and contrast enhancement procedures were proposed to make seal individuals easier to identify. Finally, two existing individual identification methods, HotSpotter and Wild-ID were evaluated. Both segmentation and post-processing were shown to improve the identification results. The HotSpotter outperformed the Wild-ID and provided promising results with the very challenging data set of the Saimaa ringed seals. The future work will include developing a species-specific method for identification of ringed seals that outperforms the generic methods. Moreover, experiments will be carried out on a larger image database.

6 Acknowledgment

The authors thank the Wildlife Photo-ID Network funded by the Finnish Cultural Foundation.

7 References

- [1] Kovacs, K.M., Aguilar, A., Auriolos, D., *et al.*: 'Global threats to pinnipeds', *Mar. Mammal Sci.*, 2012, **28**, (2), pp. 414–436
- [2] Auttila, M., Niemi, M., Skrzypczak, T., *et al.*: 'Estimating and mitigating perinatal mortality in the endangered Saimaa ringed seal (*phoca hispida*)

- saimensis) in a changing climate', *Ann. Zoologici Fennici*, 2014, **51**, (6), pp. 526–534
- [3] Koivuniemi, M., Auttila, M., Niemi, M., *et al.*: 'Photo-id as a tool for studying and monitoring the endangered Saimaa ringed seal', *Endangered Species Res.*, 2016, **30**, pp. 29–36
- [4] Zhelezniakov, A., Eerola, T., Koivuniemi, M., *et al.*: 'Segmentation of Saimaa ringed seals for identification purposes'. Proc. of Int. Symp. on Advances in Visual Computing (ISVC), Las Vegas, USA, 2015, (LNCS, **9475**), pp. 227–236
- [5] Anderson, C.: 'Individual identification of polar bears by whisker spot patterns'. PhD thesis, University of Central Florida, Orlando, Florida, 2007
- [6] Tharwat, A., Gaber, T., Hassani, A., *et al.*: 'Cattle identification using muzzle print images based on texture features approach'. Proc. of the Fifth Int. Conf. on Innovations in Bio-Inspired Computing and Applications, 2014, pp. 217–227
- [7] Hoque, S., Hannan Bin Azhar, M.A., Deravi, F.: 'ZOMETRICS-biometric identification of wildlife using natural body marks', *Int. J. Bio-Science Bio-Technology*, 2011, **3**, (3), pp. 45–53
- [8] Halloran, K.M., Murdoch, J.D., Becker, M.S.: 'Applying computer aided photo-identification to messy datasets: a case study of Thornicroft's giraffe (*Giraffa camelopardalis thornicrofti*)', *Afr. J. Ecol.*, 2014, **53**, (2), pp. 147–155
- [9] Bendik, N.F., Morrison, T.A., Gluesenkamp, A.G., *et al.*: 'Computer-assisted photo identification outperforms visible implant elastomers in an endangered salamander, *Eurycea Tonkawae*', *PLoS One*, 2013, **8**, (3), p. e59424
- [10] Branzan Albu, A., Wiebe, G., Govindarajulu, P., *et al.*: 'Towards automatic model based identification of individual sharp-tailed snakes from natural body markings'. In Proc. of ICPR Workshop on Animal and Insect Behaviour, Tampa, FL, USA, 2008
- [11] Kayci, L., Kaya, Y., Tekin, R.: 'A computer vision system for the automatic identification of butterfly species via Gabor-filter-based texture features and extreme learning machine: GF + ELM', *TEM J.*, 2013, **2**, (1), pp. 13–20
- [12] Adams, J.D., Speakman, T., Zolman, E., *et al.*: 'Automating image matching, cataloging, and analysis for photo-identification research', *Aquatic Mammals*, 2006, **32**, (3), p. 374
- [13] Lowe, D.G.: 'Distinctive image features from scale-invariant keypoints', *Int. J. Comput. Vis.*, 2004, **60**, (2), pp. 91–110
- [14] Kniest, E., Burns, D., Harrison, P.: 'Fluke matcher: a computer-aided matching system for humpback whale (*Megaptera novaeangliae*) flukes', *Mar. Mammal Sci.*, 2010, **26**, (3), pp. 744–756
- [15] Rangelova, E., Pauwels, E.: 'Saliency detection and matching strategy for photoidentification of humpback whales'. Proc. of the 2005 ICGST Int. Conf. on Graphics, Vision and Image Processing, 2005, pp. 81–88
- [16] Crall, J.P., Stewart, C.V., Berger-Wolf, T.Y., *et al.*: 'Hotspotter - patterned species instance recognition'. Workshop on Applications of Computer Vision (WACV), January 2013, pp. 230–237
- [17] Yu, X., Wang, J., Kays, R., *et al.*: 'Automated identification of animal species in camera trap images', *EURASIP J. Image Video Process.*, 2013, **2013**, (1), p. 52
- [18] Cheng, J., Liu, J., Xu, Y., *et al.*: 'Superpixel classification based optic disc and optic cup segmentation for glaucoma screening', *IEEE Trans. Med. Imaging*, 2013, **32**, (6), pp. 1019–1032
- [19] Arbelaez, P., Maire, M., Fowlkes, C., *et al.*: 'Contour detection and hierarchical image segmentation', *IEEE Trans. Pattern Anal. Mach. Intell.*, 2011, **33**, (5), pp. 898–916
- [20] Arbelaez, P., Maire, M., Fowlkes, C., *et al.*: 'From contours to regions: an empirical evaluation'. IEEE Conf. on Computer Vision and Pattern Recognition, June 2009, pp. 2294–2301
- [21] Martin, D., Fowlkes, C., Tal, D., *et al.*: 'A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics'. Proc. of the 8th Int. Conf. on Computer Vision, July 2001, vol. **2**, pp. 416–423
- [22] Arbelaez, P., Pont-Tuset, J., Barron, J., *et al.*: 'Multiscale combinatorial grouping'. Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 328–335
- [23] Ojansivu, V., Heikkilä, J.: 'Blur insensitive texture classification using local phase quantization'. 3rd Int. Conf. on Image and Signal Processing, 2008, pp. 236–243
- [24] Ahonen, T., Matas, J., He, C., *et al.*: 'Rotation invariant image description with local binary pattern histogram Fourier features'. 16th Scandinavian Conf. on Image Analysis, 2009, pp. 61–70
- [25] Ojala, T., Pietikäinen, M., Harwood, D.: 'A comparative study of texture measures with classification based on featured distributions', *Pattern Recognit.*, 1996, **29**, (1), pp. 51–59
- [26] Pizer, S.M., Amburn, E.P., Austin, J.D., *et al.*: 'Adaptive histogram equalization and its variations', *Comput. Vis. Graph. Image Process.*, 1987, **39**, (3), pp. 355–368
- [27] Fischler, M.A., Bolles, R.C.: 'Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography' (Communications of the Association for Computing Machinery (ACM), 1981), pp. 381–395
- [28] McCann, S., Lowe, D.G.: 'Local naive Bayes nearest neighbor for image classification'. Conf. on Computer Vision and Pattern Recognition (CVPR), 2012, pp. 3650–3656
- [29] Jonathon Phillips, P., Moon, H., Rizvi, S.A., *et al.*: 'The FERET evaluation methodology for face recognition algorithms', *IEEE Trans. Pattern Anal. Mach. Intell.*, 2000, **22**, (10), pp. 1090–1104