On Recognizing Individual Salamanders

S. Ravela
Dept. of Earth, Atmospheric, and Planetary Sciences
Massachusetts Institute of Technology
ravela@mit.edu

Lloyd R. Gamble
Dept. of Natural Resources Conservation *
University of Massachusetts, Amherst
lgamble@forwild.umass.edu

Abstract

Recognizing individual salamanders is an important step in assessing their migratory patterns in ecological studies. In this paper, we develop an algorithm that uses distributions of multi-scale differential features to recognize individual salamanders from a database of salamander photographs. In the proposed method, images of salamanders are warped to have straight medial axes and differential features are extracted in windowed regions of interest in the straightened image. Feature distributions are computed in each region and concatenated together to form an efficient image representation. Results, at a 91% recognition rate, suggest that the technique is promising and can be scaled to much larger databases.

1 Introduction

In the field of wildlife ecology, researchers require the ability to uniquely identify individual animals in order to address a range of questions fundamental to the discipline. For example, to estimate population sizes, demographic rates or movement distances, biologists must be able to observe individuals at multiple points in time and distinguish newly observed individuals from repeated observations. Accomplishing these needs has typically required radio-telemetry and mark-recapture techniques in which individuals are physically marked or tagged (e.g., with ear tags, leg bands, shell notches or color marking). These methods are intrusive to varying degrees and may in fact affect the fate and/or behaviors of the animals being studied. In addition, some taxonomic groups like small amphibians may not effectively retain any marks or tags long enough to be useful. For these reasons, new techniques are needed which allow individual identification but minimize observer impacts on study animals and their behaviors.

Numerous efforts have been made to exploit unique visual traits in the identification of individuals, but these have largely been limited to marine mammals, using scars and/or fin abnormalities [8], or large terrestrial mammals like the African zebra [22]. While most of these efforts pre-date modern computer-based pattern recognition technologies, there are notable exceptions like Kelly's recent work [9] applying 3-D computer matching techniques in a long-term study of the Serengeti cheetah. Work with patterned amphibians has been mostly limited to manual techniques in which researchers use spot counts in predetermined regions (e.g., tail, left or right ventral section) [5, 14] or manually compare photographs to identify matches among small numbers of total captures [1, 5, 17].

The study organism in this research is the marbled salamander (*Ambystoma opacum*), a terrestrial salamander that is protected as a "Threatened Species" under the Endangered Species Act of Massachusetts. The ecological objectives of this research are to measure site fidelity (how often do individuals return to natal breeding sites?), dispersal rates and distances (how often do animals colonize new breeding sites and how far do they go?), and a range of demographic rates (e.g., survival, breeding frequency, age at maturity). By matching photographs of salamander taken at capture points across sites and across years, we will be able to document a chronology of movement and survival for individuals observed multiple times. This type of information is critically needed to understand the life history of this rare species and to better inform efforts directed at its conservation. Specifically, these data will be used to measure local- and metapopulation-level processes and to build population models which begin to address how much habitat is needed to protect viable populations of this species.

In this paper, we evaluate the use of multi-scale differential features and their distributions to identify salamanders in a study of population biology and movement ecology in western Massachusetts, USA. As such, the proposed technique is general; variations have hitherto been applied successfully in image retrieval tasks such as finding similar scenes, trademarks, binary shapes, textures and face recognition [27, 28, 30].

Currently, salamanders are photographed at capture-sites in the field by placing them in a box with relatively

*This material is based on work supported in part by the following sponsors: 1. Massachusetts Natural Heritage and Endangered Species Program, 2. Sweet Water Trust, 3. The NSF via the ITR program.

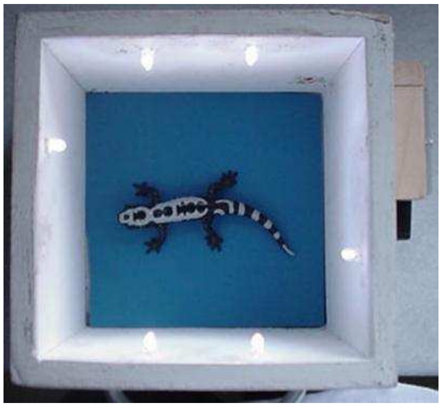

Figure 1: Acquiring Salamanders: A custom enclosure with lighting and a camera is used to capture photographs of salamanders at sites.

diffuse lighting (Figure 1). We also have data that predates this method, where images were taken in uncontrolled lighting. Therefore, our datasets exhibit illumination variations. The database images also exhibit scale variation(see Figures 3 and 4) because the salamander body changes with time and salamanders have varied dimensions.

The segmentation of current salamanders from the background is a relatively straightforward problem and for older data this was done manually. However, the body of the salamander is flexible, and rather than resort to deformable template matching (see Section 1.1), we *straighten* them first. This is done by extracting the medial axis, interpolating the axis with a spline and (un)warping small rectangular windows along the breadth of the salamander onto a straight line of length equal to the medial axis length. The result of a typically straightened salamander is shown in Figure 2. After pre-processing, the database consists of straightened salamanders taken over prolonged periods of time and the retrieval problem is to find a prior instance(s) of a presently photographed (and unwarped) salamander in the database.

The plan for recognition is to construct an appearance representation using distributions of local features of the image brightness surface. The claim is that the distinctive patterns on the backs of salamanders are good candidates for an appearance-based recognition technique. In particular, we use local features obtained by applying operators to the image that, equivalently, can be thought of as tunable spatial-frequency filters, statistical descriptors of the brightness surface, or approximations of the local shape of the image brightness surface. Specifically, multi-scale differential features are used [3, 10, 11, 15, 18, 27, 28, 30, 31, 23, 32, 33] and this choice is motivated by arguments [3, 11] that the local structure of an image can be represented in a stable and robust manner by the outputs of a set of multi-scale Gaussian derivative filters (MGDFs) applied to an image. In order to deduce similarity between two salamander images, multi-scale differential features are composed into histograms within overlapping windows along the length of the salamander. The *string* of multi-scale histograms is treated as a vector and correlated to deduce similarity (see Section 2).

1.1 Related Work

A simple approach to recognition may be to match the images of salamanders using normalized correlation. We obtained poor results in doing so because (a) anisotropic scale changes were present and (b) strong variability in local correlations could not be distinguished with modest, but unvarying, local correlation. Another technique could be to search across deformation parameters using brightness or some features of brightness [30] as the template, but this is expensive. Yet another possibility is the use of deformable templates on the unwarped images directly. This is not attempted because it is likely to be quite computationally expensive as well. Our motivation is statistically-based; continuum features and their transformations robust to scale and illumination are computed from the image and their statistics in windowed regions of the image are used as the representation. In the absence of knowledge of specific locations to compute features at, or regions to compute their statistics in, both the feature locations and windows are uniformly sampled from the appropriate space. The results are excellent for the salamander application presented here and this technique has been applied to other retrieval problems [24, 25, 27, 28].

From an appearance representation standpoint, principal component analysis (PCA) based techniques are more relevant. PCA has been pioneered by Kirby and Sirovich [13] as a representation for faces which was also developed into an effective face recognition system by Turk and Pentland [34]. Generalizations to multiple views [19, 21], illumination changes [19], other classes of objects [19] have been implemented. The approach presented in this paper is different and in fact, one of the conclusions drawn [24] is that a differential scale-space decomposition performs equally-well in several recognition and retrieval tasks.

The use of Gaussian multi-scale differential features and their representations for recognition and retrieval has also been studied in the literature. In particular, Rao and Ballard [23] develop an approach where multi-scale derivatives are extracted from the image at key locations and matched individually (using the so-called steerability property and scale-shifts) with those computed in an image of a scene. A vote count of valid matches is used to recognize the presence of the object within the scene. Schmid and Mohr [33] develop another approach where rotational

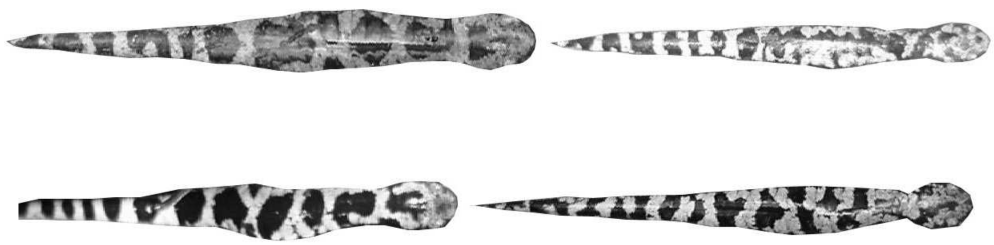

Figure 2: Unwarped and segmented Salamanders in the database. Illumination and scale variations are most prominently present in this database.

invariants at several scales are extracted at detected feature points. These features and their spatial (angle) relationships are then used for matching. Ravela and Manmatha [29] develop a technique for using rotational invariants and their spatial relationships, but the features are obtained at sampled locations in the image, and the spatial relationships are enforced online during matching. The technique presented here differs from these three approaches substantially. Rather than use features tabula-rasa, or with local constraints, their statistics over windows are computed. Such a method, we argue, explicitly represents the variability of continuum features. Other methods [20] that employ learning relationships between local features or their relationships are not considered in this implementation.

Closer to the approach presented here is work by Schiele and Crowley [32], who use multi-dimensional histograms of differential features for matching. The difference is that, here, multi-scale histograms are concatenated and hence form a much more efficient representation. While it remains to be seen whether certain recognition tasks will require the the use of multi-dimensional representation, we have found this representation to be successful at a number of retrieval tasks [24, 25, 27, 28]. Most closely related to the approach presented here is work by Nastar [18] and Dorai and Jain [2]; both use the shape-index [12] features we do, the former at a single scale and the latter for 3D-surfaces. Neither work uses the local orientation of the gradient.

The multi-scale differential features used here can be related to commonly used texture features. In the context of image retrieval Ma et. al. [16] use Gabor filters to retrieve images with similar texture. Gabor jets [35] have also been used for face recognition. A comparison between Gaussian and Gabor filters is instructive. Gabor filters are sine modulated Gaussian functions, which can be tuned to respond to a bandwidth around a certain center frequency. They exhibit compactness in space and frequency, are optimal in the sense of the uncertainty principle (time-bandwidth product) and are complete. Gabor filters are not equivariant with rotations, and separable implementations are expensive. In contrast, Gaussian derivatives exhibit the same time-bandwidth property and although they have infinite support, they can be safely truncated at around four standard deviations. While Gaussian derivatives have coupled bandwidth and center frequency, in practice separate tuning is not necessary. Rather, the derivatives provide a "natural" sampling of the frequency space, because they represent the orders of approximation in a Taylor series sense. The significant advantage of using the Gaussian derivatives is that they are equivariant with rotations [4], separable, and efficient implementations are possible. There are several other interesting properties and the reader is referred to [7, 27] for a more basic review.

2 Features and Matching Algorithm

The steps involved in deducing similarity between a query salamander image and a database image are as follows: Database images are filtered *a priori* with Gaussian derivatives, and then, at each pixel, the gradient orientation and surface curvature is computed. A query image is filtered the same way and multi-scale histograms of curvature and orientation are correlated to measure similarity. The top N ranks are presented to the user with the objective of locating the query in the database. Below, the use of differential features and the steps in the algorithm are discussed.

2.1 Differential features:

The simplest differential feature is a vector of spatial derivatives. For example, given an image I, and some point, p, the first two orders of spatial derivatives can be used as a feature (vector). This vector approximates the shape of the local intensity surface in the sense of a second order Taylor approximation. Including higher orders produces a more precise approximation. Derivatives capture useful statistical information about the image. The first derivatives represent the gradient or "edgeness" of the intensity and the second derivatives can be used to represent curvatures (e.g. bars, blobs).

However it is important that derivatives be computed in a stable manner. Derivatives will be stable if, instead of using just finite differences, they are computed by filtering an image with normalized Gaussian derivative filters (actually any C^∞ function will do [3]). In two dimensions, a Gaussian derivative is the derivative of the function

$$G(.,\sigma) = \frac{1}{2\pi\sigma^2} exp - \frac{x^2+y^2}{2\sigma^2}$$

In the frequency domain, a Gaussian derivative filter is a band-pass filter. Computing derivatives by filtering with a Gaussian derivative at a certain scale, therefore, implies that only a limited band of frequencies are being observed. Thus, in order to describe the original image more completely, a multi-scale representation is necessary. Sampling the scale-space of the image becomes essential.

2.2 Gaussian scale-space:

The necessity of a multi-scale representation described above can be concluded for *any* smooth band-limiting filter by using the commutativity of differentiation and convolution. The Gaussian happened to be a convenient function; it has natural scale parameterization, smoothness and self-similarity across scales. However, the Gaussian is more than just convenient. There are compelling theory and implementation related arguments for using multi-scale Gaussian derivatives to form appearance features. In particular, it has been shown by several authors [3, 10, 15, 31, 36] that under certain general constraints, the (isotropic) Gaussian filter forms a unique operator for representing an image across the space of scales. Structures (such as edges) observed at a coarser scale can be related to structures already present at a finer scale and not as an artifact of the filter. In general, the Gaussian (linear) scale-space serves as an unbiased (without using any other information) front end (pre-processor) for representing the image from which differential features may be computed. It is beyond the scope of this document to engage in a full discussion about the scale-space image representation and, instead, the reader is referred to the following papers [3, 15, 31, 36, 27]. Other reasons for choosing the Gaussian are presented in Section 1.1.

2.3 Multi-scale Histograms

The basic idea is to compute features at several scales, generate histograms of features at each scale and concatenate the histograms to generate a vector for matching. It is, therefore, necessary to extend the justification for multi-scale features to their histograms. For the purpose of this discussion, consider brightness as the feature and that we are representing an image by its brightness histogram. Under such a representation, it can be argued that any permutation of the image produces the same histogram and hence is admissible under matching. Obviously this is not desirable. However consider two images, one of a salamander and the other a random permutation of its pixels (assuming, in effect, each pixel is i.i.d). Upon diffusing both images, one finds that the histograms start to differ drastically, with the salamander image's histogram retaining structure and the permuted version's histogram approaching a normal distribution [24].

The argument being made here is that multi-scale histograms weakly represent image-structure because diffusion propagates pixel interactions. We contend that constructing histograms at several scales makes very dissimilar images inadmissible. This hypothesis seems to effectively be the reason for good results in experiments here (see Section 3 and elsewhere [24]). The utility of multi-scale feature histograms then is the same as for multi-scale features. They provide a far better representation of image structure than at a single scale. In contrast to using multi-scale templates or multi-dimensional histograms, they are more compact. Although they are global in nature, when it is necessary to develop a more *discriminating* representation, they can be used like local feature-vectors when computed in windowed neighborhoods of feature points or sampled locations. The advantage of doing so is to represent the variability of filter responses while retaining the advantages of a local representation. We make this choice in experiments here by simply partitioning the image into overlapping windows. The advantages of multi-scale distributions with brightness as the feature discussed in this section clearly extends to histograms of differential features because differentiation and convolution commute.

2.4 Curvature, Orientation and Matching

Several differential features can be constructed from derivatives and several representations and methods have been developed [23, 32, 30, 33, 28, 27, 26] for recognition and retrieval. The choice of these features depends on several factors, primary (among these) for this application is tolerance to illumination and scale changes.

The first feature is drawn from the curvatures of the isophote and flowlines curves, that is, the curvatures of iso-intensity contours and the gradient integral curves. These features are nothing more than the second order spatial derivatives expressed in a coordinate frame (see [3]) and determined by the orientation of the local intensity gradient. The isophote and flowline curvatures are invariant to image plane rotations, monotonic intensity variations and further, their ratios are, in practice, quite tolerant to scale variations of the entire image. The isophote (N) and flowline (T) curvatures are defined as [12, 3]:

$$N = A \star \left[2I_xI_yI_{xy} - I_x^2I_{yy} - I_y^2I_{xx}\right] \quad (1)$$

$$T = A * \left[I_{xy}(I_x^2 - I_y^2) + I_xI_y(I_{yy} - I_{xx})\right] \quad (2)$$

$$A = (I_x^2 + I_y^2)^{-\frac{3}{2}} \quad (3)$$

$I_x = I_x(p,\sigma)$ and $I_y = I_y(p,\sigma)$ are the first order partial spatial derivatives of image I around point p, computed using Gaussian derivative at scale σ. Similarly, I_{xx}, I_{xy}, and I_{yy} are the corresponding second derivatives. The isophote curvature N and flowline curvature T are then combined into a ratio called the shape index, expressed as follows [12, 2, 18]: $C = [0.5 - \frac{1}{\pi} * atan\frac{N+T}{N-T}]$. The shape index is in the range [0,1] and the index value C is undefined when either N and T are both zero, and is, therefore, not computed. This is interesting because very flat portions of an image (constant or constant slope in intensity) are eliminated.

The second feature used is local orientation. Local orientation is the direction of the local gradient. Orientation is independent of curvature and is stable with respect to scale and illumination changes. The orientation is simply defined as $P = atan2(I_y, I_x)$ Note that P is defined only at those locations where C is, and ignored elsewhere. As with the shape index, P is rescaled and shifted to lie between the interval [0,1].

In summary, the starting point for constructing a representation is the use of two non-linearly transformed first- and second-order differential features. Both these features exhibit tolerance to global illumination changes and isotropic scale changes.

Feature Histograms: Histograms of the shape index and orientation are used to represent the distributions of features over an image. Histograms form a global representation because they capture the distribution of local features and they are the simplest ways of estimating a non-parametric distribution. In this implementation, shape-index and orientation histograms are generated at several scales and represented as a one-dimensional record or vector. The representation of the image I is the vector $V_i =< H_c(\sigma_1)\dots H_c(\sigma_n), H_p(\sigma_1)\dots H_p(\sigma_n) >$. H_c and H_p are the curvature and orientation histograms respectively.

Matching feature histograms: Two representations are compared using normalized cross-covariance defined as $d_{ij} = \frac{V_i^{(m)} \cdot V_j^{(m)}}{||V_i^{(m)}||||V_j^{(m)}||}$, where $V_i^{(m)} = V_i - mean(V_i)$. The query histogram vector V_q is compared with each database histogram vector V_i. The corresponding images are ranked by their score. We call this algorithm the 1D curvature/orientation or CO-1 algorithm.

3 Salamander Recognition

The database used for obtaining the salamander was collected over a four year period (99-02), from fourteen different sites. Multiple images of salamanders may be taken at a single time and other instances of the same individual may be obtained at different sites and different times. The total number of database images used in this experiment was 370 and 69 queries were used to test the algorithm.

The algorithm was parametrized as follows. Each salamander image is divided into 5 overlapping windows along the length of the salamander with a 25% overlap. Eight scales were used in half-octave steps $(1\dots 8\sqrt{2})$. Ten bins each were used for shape-index and orientation features. The shape-index and orientation histograms were computed in each window at all eight scales and the histograms across the windows were concatenated to generate the image representation. Under this parameterization, each image requires approximately 3.2KB.

The parameters used here for bins, scales and windows follows extensive experimentation in another retrieval task [25]. It was shown that the CO-1 algorithm is insensitive to bin variations and hence only 10 bins were chosen. It was also shown that in the absence of particular scales to associate a pixel with, half-octave spacing of five to eight scales suffices for many retrieval scenarios.

The evaluation of the algorithm is fairly straightforward. With one exception, all the queries had one relevant image in the database. We simply count how many queries were correctly recognized in the N ranks, where N ranges to 5, the upper limit of how many retrievals the user is willing to peruse in the field. The performance of the algorithm is depicted in Table 1, where we see that approximately 91% recognition rate is achieved.

In Figure 3, five pairs of correct recognition instances are shown. The salamander images on the left are queries and the ones on the right are the retrievals at rank 1. These pictures depict the performance of the algorithm under interesting variabilities that include illumination changes and anisotropic scale changes. In one instance (third row) the salamander seems to have lost some weight. Note that all the pairs shown here come from different sites and therefore at different times.

In Figure 4, three examples of incorrect recognition are shown. Each row depicts a query on the left, the image retrieved at rank 1 in the middle and the relevant image on the right. The first image shows a large specularity in the query, an artifact of the imaging setup used for this acquisition. In this instance, the images were taken at night, by shining a bright light at the salamander. We have since modified the apparatus to be more uniform and diffuse, with an array of LED lights providing the illumination (see Figure 1). The second mismatch (second row) is obvious

Rate at Rank	1	2	3	4	5
	50/69	54/69	60/69	61/69	63/69

Table 1: The performance of MGDF method on Salamander Recognition

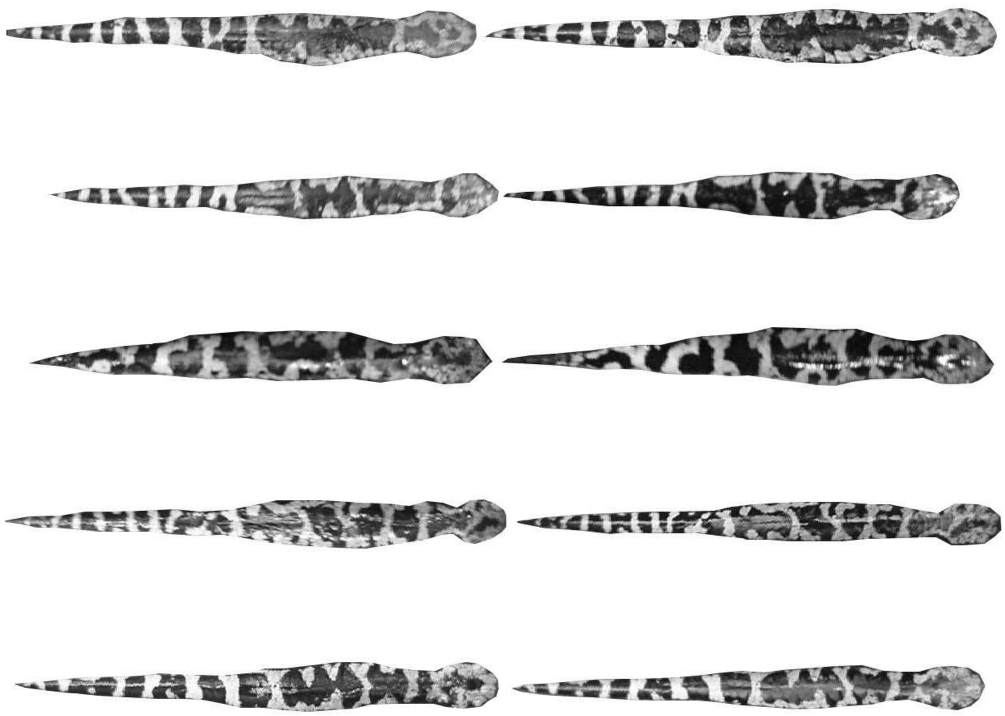

Figure 3: Correct salamander recognition examples. Each pair of images in a row depicts the query on the left and the recognized image on the right. All the recognized images were at rank 1.

Figure 4: Incorrect salamander recognition examples. Each triplet of images in a row depicts the query on the left and the recognized image in the middle and the relevant image in the database on the right.

to explain because it was imaged without the apparatus and the salamander was not segmented. In the third row, it is the specular reflection (see right image) that causes the mismatch and not the clipping of the tail in the query.

4 Summary and Conclusions

The results presented in this paper are very exciting for the following reasons. First, we started with the premise that the pattern on the back of a marbled salamander makes it amenable to appearance-based recognition. This hypothesis holds thus far. Second, the shape-index and orientation based method continues to demonstrate wide applicability. Both these features are insensitive to scale changes as are their histograms. The power of recognition substantially comes from the use of multi-scale histograms, which we argue narrows the admissible set of permutations of the image. To a lesser extent, the partitioning of the image into windows contributes to good results [24]. It is unclear if using networks of filter responses used at select feature points in the image will improve these results but that test has not yet been conducted. Our approach is to use distributions of filter outputs to represent the variability of the image or its locales. Given that this algorithm is fairly successful without any preprocssing of the brightness image (after unwarping), we are encouraged to scale this technique to a database of an order of magnitude larger and deploy it in the field.

The algorithm performs well with simple cross-covariance and no learning involved with respect to any of the parameters. We do not discount the necessity or potential for learning; however, we have observed that a representation based on the differential decomposition of the image at multiple scales is giving comparable performance to one based on learning a compact representation from the data, namely PCA [25].

This paper also indicates the next steps to be taken. The first of these is to develop an algorithm for removing specularity. Since the color of the specular reflection is the color of the illumination source, it may be possible to remove specularity using our new LEDs that produce light with a blue hue in contrast to the near black-white salamander images. The second improvement sought is scale-selection. Although the choice of scales presented here is a rather unbiased one, selecting scales [15] to associate with pixels may compact the representation even further and is an interesting research question. Finally, getting to 100% within the top 5 ranks is a pressing goal and we are exploring Bayesian and other statistical techniques to overlay on top of the image representations developed here.

Acknowledgment

The authors wish to thank Andrew Herst and Joshua Sarro for implementing the unwarping procedure, Daniel Pepin for assistance in design and construction of lightboxes, and Dana Brennan for her work in image-preprocessing.

References

[1] Doody, J. S., A photographic mark-recapture method for patterned amphibians. H Review 26(1): 19-21, 1995

[2] Chitra Dorai and Anil Jain, "COSMOS - A representation scheme for free form surfaces", ICCV 95, pp. 1024-1029, 1995.

[3] L M J Florack, The Syntactic Structure of Scalar Images, PhD Dissertation, University of Utrecht, 1993

[4] W. T. Freeman and E. H. Adelson, The design and use of steerable filters, IEEE Trans. Patt. Anal. and Mach. Intel., 13(9):891-906, 1991

[5] Gill, D. E. The metapopulation ecology of the red-spotted newt, Notophthalmus viridescens (Rafinesque). Ecological Monographs 48: 145-166, 1978

[6] Graham, R. E., Environmental effects on deme structure, dynamics and breeding strategy of Ambystoma opacum (Amphibia: Ambystomatidae), with an hypothesis on the probable origin of the marbled salamander life-style. Doctoral Dissertation, Rutgers University. 1971

[7] Gosta Granlund and Hans Knutsson, Signal Processing for Computer Vision, Kluwer Academic Publishers, 1995

[8] Hammond, P. S., S. A. Mizroch, and G. P. Donovan. Report of the workshop on individual recognition and estimation of cetacean population parameters. Reports of the International Whaling Commission, Cambridge, Massachusetts, 1990

[9] Kelly, M. J., Computer-aided photograph matching in studies using individual identification: an example from Serengeti cheetahs. Journal of Mammology 82(2): 440-449, 2001

[10] J. J. Koenderink, The Structure of Images, Biological Cybernetics, 50:363-396, 1984.

[11] J. J. Koenderink and A. J. van Doorn, Representation of Local Geometry in the Visual System, Biological Cybernetics, 55:367-375, 1987

[12] J. J. Koenderink and A. J. Van Doorn, Surface Shape and Curvature Scales, Image and Vision Computing, 10(8), 1992

[13] M. Kirby and L. Sirovich, Application of the Kruhnen-Loeve Procedure for the Characterization of Human Faces, IEEE Trans. Patt. Anal. and Mach. Intel., 12(1): 103-108, Jan. 1990.

[14] Loafman, P., Identifying individual spotted salamanders by spot pattern. Herp Review 22(3): 91-92, 1991

[15] T. Lindeberg, Scale-Space Theory in Computer Vision, Kluwer Academic Publishers, 1994

[16] W. Y. Ma and B. S. Manjunath, Texture-Based Pattern Retrieval from Image Databases, Multimedia Tools and Applications, 2(1):35-51, Jan. 1996

[17] Molines, K. and C. Swarth., The breeding migration of marbled salamanders (Ambystoma opacum) and spotted salamanders (A. maculatum) at the Jug Bay Wetlands Sanctuary on Maryland's coastal plain. Technical Report of Jug Bay Wetlands Sanctuary, 1999.

[18] C. Nastar, The image shape spectrum for image retrieval, Technical Report 3206, INRIA, June 1997.

[19] S. K. Nayar and H. Murase and S. A. Nene, Parametric Appearance Representation, Early Visual Learning, Oxford University Press, Feb. 1996.

[20] J. Piater, Visual Feature Learning, PhD Dissertation, University of Massachusetts at Amherst, 2001.

[21] A . Pentland, B. Moghaddam, and T. Starner, View-based Modular Eigenspaces for face recognition, Proc. Comp. Vision and Patt. Recognition, 1994.

[22] Peterson, J. C. B., An identification system for zebra (Equus burchelli Gray). East African Wildlife Journal 10: 59-63, 1972

[23] Rajesh Rao and Dana Ballard, Object Indexing Using an Iconic Sparse Distributed Memory, Proc. International Conference on Computer Vision, pp. 24-31, 1995.

[24] S. Ravela, On multi-scale differential features and their representations for recognition and retrieval, PhD Dissertation, University of Massachusetts at Amherst, 2003

[25] S. Ravela and A. Hanson, On multi-scale differential features for face recognition, In Proc. Vision Interface 01, Ottawa, June 2001.

[26] S. Ravela, S. and C. Luo, Appearance-based Global Similarity Retrieval of Images. In Advances in Information Retrieval, W. Bruce Croft (Ed), Kluwer Academic Publishers 2000.

[27] S. Ravela, R. Manmatha, Gaussian Filtered Representations of Images, Encyclopedia of Electrical and Electronic Engineering, John Webster (Editor), John Wiley, 1999

[28] S. Ravela, S. and R. Manmatha, Retrieving Images by Appearance, Proc. of the International Conf. on Computer Vision, (ICCV), Bombay, India, Jan 1998.

[29] S. Ravela and R. Manmatha, Image Retrieval by Appearance, Proc. of the 20th Intl. Conf. on Research and Development in Information Retrieval (SIGIR'97), July 1997, pp 278–285.

[30] S. Ravela, R. Manmatha, and E. M. Riseman, Retrieval from Image Databases using Scale Space Matching. Proc. of the European Conf. on Computer Vision ECCV '96, Cambridge, U.K., pages 273–282,Springer, April 1996

[31] B. M. ter Har Romeny, Geometry Driven Diffusion in Computer Vision, Kluwer Academic Publishers, 1994

[32] Bernt Schiele and James L. Crowley, Object Recognition Using Multidimensional Receptive Field Histograms, Proc. 4th European Conf. Computer Vision, Cambridge, U.K., April 96.

[33] Schmid, R. Mohr, Local Grayvalue Invariants for Image Retrieval, PAMI (19), No. 5, pp. 530-535, May 1997. 30.

[34] M. Turk and A. Pentland, Eigen Faces for Recognition, Jrnl. Cognitive Neuroscience, 3:71-86, 1991.

[35] L. Wiskott, J.-M. Fellous, N. Kruger and C. von der Malsburg, Face recognition by elastic bunch graph matching, IEEE Trans Patt. Anal. and Mach. Intell. 17(7):775-779, 1997.

[36] A. P. Witkin, Scale-Space Filtering, Proc. Intl. Joint Conf. Art. Intell., pp. 1019-1023, 1983