

Evolutionary feature selection for probabilistic object recognition, novel object detection and object saliency estimation using GMMs.

Leonardo Trujillo ^a, Gustavo Olague ^a,
Francisco Fernández de Vega ^b and Evelyne Lutton ^c

^a EvoVisión Project, Applied Physics Division, CICESE Research Center,
Km. 107 carretera Tijuana-Ensenada 22860, Ensenada, B.C. México.

^b Grupo de Evolución Artificial, Universidad de Extremadura, Centro Universitario
de Mérida C/Sta Teresa de Jornet, 38, 06800 Merida, Spain.

^c Complex Team, INRIA Roquencourt,
Domaine de Voluceau, BP 105, 78153 Le Chesnay Cedex, France.

{trujillo,olague}@cicese.mx, fcofdez@unex.es, evelyne.lutton@inria.fr

Abstract

This paper presents a method for object recognition, novel object detection, and estimation of the most salient object within a set. Objects are sampled using a scale invariant region detector, and each region is characterized by the subset of texture and color descriptors selected by a Genetic Algorithm (GA). Using multiple views of an object, and multiple regions per view, objects are modeled using mixtures of Gaussians, where each object represents a possible class for a particular image region. Given a set of objects, the GA learns a corresponding Gaussian Mixture Models (GMM) for each object in the set employing a *one vs. all* training scheme. Thence, given an input image where interest regions are detected, if a large majority of the regions are classified as regions of object *O* then it is assumed that said object appears in the imaged scene. The GA's fitness function promotes: 1) a high classification accuracy, 2) the selection of a minimal subset of descriptors, and 3) a high separation among models. The separation between two GMMs is computed using a weighted version of Fisher's linear discriminant, which is also used to estimate the most "salient" object among the set of modeled objects. Object recognition and *novel* object detection are done using confidence-based classification. Hence, when a non-modeled object is sampled, the detected regions are thereby identified as belonging to an unseen object and a new GMM is trained accordingly. Experimental results on the COIL-100 data set confirm the soundness of the approach.

1 Introduction

Currently, many computer vision systems address the problems of object detection and/or recognition using a sparse representation of image information through locally prominent




Figure 1: Abstract view of common object recognition vision systems.

image regions [8, 3], see Figure 1. A training phase consists on detecting stable image regions on an object using interest region detectors, and characterizing said regions using discriminative local descriptors [5, 3, 11, 10]. In this way, by relying on sparse local information the method is robust to partial object occlusions. During testing, an image is taken as input and the same region detection/description process is repeated. However, the extracted local information is now compared with stored object models and if appropriate matching criteria are met it is possible to identify known objects within the scene. This approach relies on the assumption that different local regions on an object will be highly separated in descriptor space, and thus requires highly discriminative region descriptors. This assumption will not hold true for objects with regular or repetitive patterns across their surface, i.e. a football or tomato. Furthermore, if object representations are learned in this manner, an intuitive comparison between two object models is not evident. For instance, if three object representations are learned, how can a measure of similarity be computed? These considerations are pertinent for a system that automatically identifies the “most salient” object, or image, from a given set. Automatic novelty detection is a line of research where these questions are essential [4]. Another application area relates to the automatic identification of visual landmarks; in robot navigation, for example, the norm is to use artificial or human selected landmarks.

This paper presents an approach where every region λ detected on an object O is taken as an instance of the same class, and is characterized with a feature vector of statistical descriptors computed in a feature space Φ of texture and color information. A GA searches within Φ for the smallest subspace $F \subseteq \Phi$ of statistical descriptors, of both texture and color, that yield the highest classification accuracy using a *one vs. all* scheme of maximum likelihood classification. The GA also searches for the best possible between-class separation of learned models. Therefore, the proposed approach does not require highly discriminative features because it uses a robust classifier, a known trade-off between descriptor design and classifier training. A GMM representation is used for each class (object), and a heuristic extension of Fisher’s linear discriminant is used to estimate an “*apparent*” measure of class separation among models with more than one component. Based on this measure of model separation the most *salient* object is identified by selecting the object with the highest between-class separation using a min-max operation. A further advantage of using a GMM based classifier is the ability to use confidence estimation to identify regions extracted from unknown objects as outliers and label them as

samples of a new class. Hence, it is possible to automatically train a new *one vs. all* classifier for the newly identified object. Experimental results in this paper only deal with objects in scenes with simple backgrounds. Nevertheless, the use of a multimodal models should allow the approach to extend to real world scenes where more within class variation is likely to occur. Recently, Markou and Singh [4] propose a similar system that carries out both novelty detection and classification, however several differences exist:

1. The current work is concerned with object recognition, on the other hand, the work in [4] only addresses ROI classification.
2. The work in [4] relies on prior segmentation, a drawback because segmentation is an ill-posed problem; this is avoided by using locally salient image regions.
3. The proposed feature space Φ is more compact than the one used in [4], with less redundant information. Furthermore, the GA used for feature selection maximizes accurate classification, minimizes the set of descriptors used, and maximizes the between-class separation of learned models. The authors in [4] use the sequential floating forward selection algorithm and do not consider between-class separation.
4. The proposed measure for class separation is based on Fisher’s linear discriminant which gives a closed form estimation computed directly from the learned GMMs; the Bhattacharya distance is employed in [4] along with NNet classifiers.
5. Novelty detection in the present work utilizes confidence-based classification of region descriptors, whereas [4] uses an heuristic criteria based on NNet output.
6. Finally, the COIL-100 data set used in the present work includes objects with information in feature space that tends to overlap, such as two toy cars with similar texture or two objects with the same color. On the other hand, [4] uses classes with marked differences among them, such as sky and chair classes.

2 Background

This section will give a brief review on some of the main concepts used throughout this work: scale invariant region detection, genetic algorithms, Gaussian mixture models, Fisher’s linear discriminant, and the texture and color feature space employed.

Scale Invariant Region Detection. Selecting a characteristic scale for local image features is a process in which local extrema of a function response, embedded into a linear scale-space, are found over different scales. The interest operator applied in the current work was synthesized with Genetic Programming, optimized for high repeatability and global region separability [9, 10], named K_{IPGP1^*} which is based on DoG filtering,

$$K_{IPGP1^*}(\mathbf{x}; t_j) = G_{t_j} * |G_{t_j} * I(\mathbf{x}) - I(\mathbf{x})|, \quad (1)$$

where $j = 0, 1, \dots, k$, and k is the number of scales to be analyzed, here it is set to $k = 15$. The size of a region is proportional to the scale at which it obtained its extrema value. For the sake of uniformity, all regions are scaled to a size of 41×41 pixels using bicubic interpolation before region descriptors are computed. Figure 2 shows sample interest regions extracted with the aforementioned detector.




Figure 2: Detected regions on three images from the COIL-100 data set.

Features	Description
Gradient information	<i>Gradient</i> , <i>Gradient magnitude</i> and <i>Gradient Orientation</i> ($\nabla, \ \nabla\ , \nabla_\phi$).
Gabor filter response	The sum of <i>Gabor filters</i> with 8 different orientations (<i>gab</i>).
Interest operators †	The response to 3 stable interest operators: <i>Harris</i> , <i>IPGP1</i> and <i>IPGP2</i> ($K_{Harris}, K_{IPGP1}, K_{IPGP2}$).
Color information	All the channels of 4 color spaces: <i>RGB</i> , <i>YIQ</i> , <i>Cie Lab</i> , and <i>rg chromaticity</i> ($R, G, B, Y, I, Q, L, a, b, r, g$).

† K_{IPGP1} is proportional to a DoG filter, and K_{IPGP2} is based on the determinant of the Hessian [9, 10].

Table 1: The complete feature space Φ .

Texture and Color Features. In order to appropriately describe each image region the search space Φ of possible features includes 18 different types of color and texture related information, see Table 1. To characterize the information contained along different channels, six statistical descriptors are computed: *mean* μ , *standard deviation* σ , *skewness* γ_1 , *kurtosis* γ_2 , *entropy* H and *log energy* E . This yields a total of 108 possible descriptor values for the multivariate GMMs. Because general statistical information is used, the descriptors will mostly be rotationally invariant.

Genetic Algorithms (GA) are stochastic heuristic search techniques that model, in an abstract manner, the principles of natural evolution [2]. The basic principles that a canonical GA follows are survival of the fittest (selection), recombination and replication of fit genetic material (crossover), and the introduction of novel genetic information (mutation), all of which are modeled as stochastic processes. These techniques operate over a set of parameterized solutions using population-based metaheuristics. GAs can manage a number of constraints and design decisions, and carry out a search in an intrinsic parallel manner; thence, GAs can be considered as a global optimization and search method. In the current work, the canonical GA with a binary string chromosome is employed.

Gaussian Mixture Models are a useful tool when it is necessary to model multimodal data, or as an approximation to different types of more complex distributions. The GMM pdf is defined as a weighted sum of Gaussian pdfs,

$$p(\mathbf{x}; \Theta) = \sum_{c=1}^C \alpha_c \mathcal{N}(\mathbf{x}; \mu_c, \Sigma_c), \quad (2)$$

where $\mathcal{N}(\mathbf{x}; \mu_c, \Sigma_c)$ is the c th multivariate Gaussian component with mean μ_c , covariance matrix Σ_c , and an associated weight α_c . Estimation of the mixture model parameters is

done using the EM algorithm when a fixed number of components is assumed. Alternatively, if a variable number of component is desired, with a maximum bound, it is possible to use the Greedy-EM [7]. Classification with GMMs can be done through Bayes rule, or using confidence-based classification [7]. A confidence value $\kappa \in [0, 1]$ and confidence region $\mathcal{R} \subseteq \Phi$ for a pdf are $0 \leq p(\mathbf{x}) < \infty, \forall \mathbf{x} \in \Phi$. κ is a confidence value related to a non-unique confidence region \mathcal{R} such that

$$\int_{\Phi \setminus \mathcal{R}} p(\mathbf{x}) d\mathbf{x} = \kappa. \quad (3)$$

A sample \mathbf{x} that lies within \mathcal{R} is considered a true member of the class modeled by p , otherwise it is classified as an outlier.

Fisher's Linear Discriminant. Fisher defined the separation between two distributions \mathcal{N}_i and \mathcal{N}_j as the following ratio

$$S_{i,j} = \frac{(\mathbf{w}(\mu_i - \mu_j))^2}{(\mathbf{w}^T(\Sigma_i + \Sigma_j)\mathbf{w})}, \quad (4)$$

where $\mathbf{w} = (\Sigma_i + \Sigma_j)^{-1}(\mu_i - \mu_j)$ [1]. Note that S is defined for unimodal pdfs, hence a weighted version \hat{S} that accounts for the weight α_i and α_j of the associated Gaussian components in a GMM is proposed, such that

$$\hat{S}_{i,j} = \frac{S_{i,j}}{1 + \alpha_i + \alpha_j}. \quad (5)$$

Hence, the separation between components with a small combined weight (they have less influence over their associated models) will *appear* to be larger with respect to the separation between components with larger weights. Therefore, let C_a and C_b represent the number of components of $p_a(\mathbf{x}; \Theta_a)$ and $p_b(\mathbf{x}; \Theta_b)$ respectively, then $S^{a,b}$ represents the *apparent* separation matrix of size $C_a \times C_b$ that contains the weighted separation $\hat{S}_{i,j}$ of every component of p_a with respect to every component of p_b . The final *apparent* separation measure \mathcal{S} between p_a and p_b is given by

$$\mathcal{S}^{a,b} = \inf(S^{a,b}). \quad (6)$$

3 Proposed Approach

This section describes the details of the proposed approach to object recognition, novel object detection, and salient object estimation; a flowchart view is depicted in Figure 3.

3.1 Learn Object Models

First, there is an initial off-line step in which interest regions from every object $O \in M$ are extracted and labeled accordingly; moreover, all 108 descriptor values are computed for each region. Afterwards, the GA performs feature selection, and learns appropriate GMMs for a subset N of the objects in M . Figure 3a shows the basic flow chart of a canonical GA, the two main aspects to discuss is how candidate solutions are represented and how fitness assignment is done. The other processes in the GA are standard: fitness proportional selection, mask crossover, single bit mutation and elitist survival strategy.




Figure 3: An overview of the proposed approach, a) Genetic Algorithm, b) Learn object models, c) Novel object detection.

Solution Representation: Each individual in the population is coded as a binary string $B = (b_1, b_2, \dots, b_{108})$ of 108 bits. Each bit is associated with one of the statistical descriptors in Φ . Therefore, if bit b_i is set to 1 its associated descriptor will be selected, with the opposite being true if $b_i = 0$. The feature vector \mathbf{x}_λ for each region λ is thereby given by the concatenation of the set of selected descriptors $F \subseteq \Phi$.

Fitness Evaluation: Here is where object models are learned and fitness is assigned to each individual in the population. For every object $O_j \in N$ a corresponding GMM $p_j(\mathbf{x}; \Theta_j)$ is trained with a *one vs. all* strategy with 70% of the regions, using the descriptor values selected by B . The GMM classifiers are trained with the EM algorithm. After training, a set $\mathcal{P} = \{p_i(\mathbf{x}; \Theta_i)\}$ of $|N|$ GMMs, on each $\forall O_i \in N$. Afterwards, the remaining 30% of image regions are used for testing and a corresponding accuracy score \mathcal{A}_i is computed using Bayes rule. Optimization is posed as a minimization problem, hence fitness is assigned by

$$f(B) = \begin{cases} \frac{B_{ones} + 1}{\mathcal{A}' \cdot \inf(\mathcal{S}^{p_i, p_j})} & \forall p_i, p_j \in \mathcal{P}, i \neq j, \text{ when } \forall \mathcal{A}_i > 0, \\ \frac{K \cdot B_{ones} + 1}{\mathcal{A}' + \varepsilon} & \text{otherwise.} \end{cases} \quad (7)$$

In the above equation, B_{ones} is the number of ones in string B , \mathcal{A}' is the average accuracy score of all the GMMs in \mathcal{P} , a penalization term set to $K = 2$, and $\varepsilon = 0.01$; hence, fitness depends upon testing and not training accuracy. The first case in Eq. 7 is applied when all of the classifiers were able to obtain an accuracy score, fitter individuals will minimize the number of selected descriptors and maximize the average testing accuracy \mathcal{A}' . Furthermore, the term $\inf(\mathcal{S}^{p_i, p_j})$ promotes between-class model separation by selecting the infimum of all the apparent separation measures computed for every object in N . On the other hand, the second case in Eq. 7 is applied when the EM algorithm fails to produce a valid GMM for one of the objects in N .

After a fixed number of iterations the GA stops and returns the fittest individual B^o

found so far. The best individual B^o is re-trained using the Greedy-EM instead of the basic EM, this is done for two reasons. First, the Greedy EM did not prove to be appropriate during evolution because it required more computation time and produced more runs that failed to converge. Secondly, once the GA has produced a valid high performance solution, the associated object models can be further enhanced by using the Greedy EM on B^o . Therefore, the GA returns the selected subset of descriptors F that characterize the objects in N , and a set of trained GMMs \mathcal{P}^o . Finally, the most salient object O^o in N is said to be modeled by the GMM p^o that satisfies the following,

$$p^o \leftarrow \arg \max_{p_i} (\mathcal{S}^{p_i, p_j}) \quad \forall p_i, p_j \in \mathcal{P}^o \quad \text{with } i \neq j. \quad (8)$$

3.2 Object Recognition and Novel Object Detection

In order to test the ability of the described approach to recognize known objects and detect novel objects (those without a corresponding $p_i \in \mathcal{P}^o$) the process in Figure 3c is followed. Given an image of an object $O_i \in M$, interest regions are detected and their corresponding descriptors, specified in F , are computed. The extracted regions are classified using confidence estimation with the models in \mathcal{P}^o . A confidence region within each GMM in \mathcal{P}^o is defined, with the confidence threshold set to $\kappa = 0.95$. Therefore, if a large majority, over 60%, of the regions lie within the confidence region of a given $p_j \in \mathcal{P}^o$ then it is said that object $O_i = O_j$, thereby accounting for a successful recognition. Otherwise, if regions are classified as outliers from all known classes, it is possible to tag them as belonging to an object not modeled in \mathcal{P}^o . Hence, if the percentage of regions classified as outliers is $\mathcal{A}_{out} > 60\%$, then the sampled object O_i is labeled as a new object, and a corresponding GMM is learned and added to \mathcal{P}^o .

4 Experimental Results

This section presents three different experiments to test the proposed object recognition system. The code was written mostly in MATLAB, the GMMBAYES Toolbox¹ was used for GMM training, and the Genetic Algorithms for Optimization Toolbox² was used as part of the GA code. The images used for testing are taken from the COIL-100 data set, Figure 4 shows the first 40 objects in the data set [6]. Every object is seen from 72 different views, interest regions are extracted from all of the views and tagged accordingly as the ground truth for each object. The basic parameters of the algorithm are the same in every run, only modifying the number of different objects used, the size of sets (M, N) . Three experiments are presented: Exp. 1 (10,5) with objects 1 - 10 from the data set; Exp. 2 (20,10) with objects 20 - 40; and Exp. 3 (40,25) with objects 1 - 40. The GMM classifiers were trained using EM with one Gaussian component, and if a solution was not found, the algorithm is restarted with 2 components, and so on. The results presented for each experiment are shown for object recognition and novel object detection. Table 2 shows the average accuracy score obtained after the initial object models are generated (Figure 3b), along with the fitness value, the number of features, the set of selected features F , errors in object recognition, and the salient object within the set. Table 3 presents the accuracy

¹GMMBAYES Matlab Toolbox <http://www.it.lut/project/gmmbayes>

²Genetic Algorithms for Optimization Toolbox by Andrey Popov <http://automatics.hit.bg>



Figure 4: These are the first 40 objects in the COIL-100 data set used in the reported experimental runs. The images used with the first two experiments are marked, while all 40 are used in the third. Salient objects selected by the separation criterion are circled. Object 32 is the only one for which novel object detection failed with $h = 60\%$.

Exp.	\mathcal{A}'	$f(\mathbf{B}^0)$	$\mathbf{B}^0_{\text{ones}}$	Features	Error	\mathbf{O}^0
1)	99.6	0.5	27	$\nabla_{(\gamma_2, H)}, \ \nabla\ _{(\sigma, \gamma_2)}, \nabla\phi_{(\gamma_2)}, K_{Harris}(E),$ $K_{IPGP1}(\sigma), R_{(\mu, H)}, G_{(\sigma, \gamma_1)}, B_{(\mu, \sigma, \gamma_1, H)},$ $Y_{(\mu, \gamma_2, H, E)}, I_{(\sigma, H)}, L_{(\sigma, E)},$ $a_{(\mu, \sigma)}, b_{(\sigma, E)}, g_{(\mu)}$	none	4
2)	99.2	1.5	43	$\nabla_{(\mu, \sigma, \gamma_2, E)}, \ \nabla\ _{(\sigma, \gamma_2, H)}, \nabla\phi_{(\mu, \sigma, \gamma_2)}$ $K_{Harris}(\gamma_1, H), K_{IPGP1}(\mu, E), K_{IPGP2}(\gamma_1, E),$ $gab_{(\gamma_1)}, R_{(\mu, \sigma, E)}, G_{(\mu)}, B_{(\sigma, \gamma_1, \gamma_2, H)}, Y_{(\mu, \gamma_2, H, E)},$ $I_{(\sigma, H)}, Q_{(\gamma_2, H)}, L_{(\mu)}, a_{(\mu, \sigma, H)},$ $b_{(\mu, \sigma, E)}, r_{(\mu, \sigma)}, g_{(E)}$	none	25
3)	98.7	6.4	37	$\nabla_{(\mu, \sigma, \gamma_2)}, \ \nabla\ _{(\mu, \sigma, \gamma_1, \gamma_2, H, E)}, \nabla\phi_{(\gamma_1, \gamma_2, H)},$ $K_{Harris}(\gamma_2, E), K_{IPGP1}(H), K_{IPGP2}(\gamma_2, H),$ $gab_{(\mu)}, R_{(\mu, \gamma_1, E)}, G_{(\mu)}, B_{(\gamma_2, H)}, Y_{(\mu, \sigma, \gamma_2, H, E)},$ $I_{(E)}, Q_{(\mu, \gamma_1)}, a_{(\gamma_1)}, b_{(\sigma)}, r_{(\sigma, H)}, g_{(\mu, \sigma)}$	none	4

Table 2: Performance when initial class models are learned; see text for further details.

	\mathcal{A}'_M	Errors	Salient Objects
Exp.1	99.72	none	objects 4, 7, 3
Exp.2	99.04	object 32	objects 36, 38, 25
Exp.3	98.68	object 32	objects 36, 28, 4

Table 3: Performance for novel object detection. Note that \mathcal{A}'_M represents the accuracy of region classification after a corresponding model is learned for every object $O \in M$.

score once a corresponding model is learned for every object $O \in M$, the incorrectly classified objects, and the three most salient objects found in each case. Given the high level of accuracy in both sets of results, it can be concluded that the problem of object recognition is almost perfectly solved for the set of images employed. Figure 5 shows the convergence graphs of each GA run, plotting the fitness of the best individual B^o found so far. The experiments were executed with 30, 30 and 40 iterations respectively.




Figure 5: Convergence plots that show the $\log(f(B^o))$ of the best individual found thus far by the GA in each of the experimental runs.

5 Discussion and Conclusions

The results presented in the previous section exhibit promising performance patterns. For all three experiments the algorithm was able to train extremely accurate classifiers using a fraction of the available descriptors. It is important to note that in Table 2 even do all experiments produce similar values for accuracy and number of descriptors, their associated fitness scores are different. This is due to the model separation measure $\inf(\mathcal{S}^{P_i, P_j})$ in the fitness function, because with more objects the space of possible objects models becomes crowded. All the classifiers trained in each experiment finished with a single Gaussian component, an unexpected outcome that can nevertheless be explained. Every object is small and tends to exhibit regular patterns across their surface; therefore, it was possible to characterize them with a single component in feature space. This suggests that GMMs would be more appropriate dealing with images that have a larger variations in descriptor space. Additionally, the convergence graphs in Figure 5 show two different patterns. First, starting from the random population the initial iterations produce very poor results, individuals in these generations are evaluated using the second case of the fitness function because the EM fails to find a valid model for at least one of the objects. Therefore, initial iterations attempt to find solutions B that are able to produce a classifier for every object in N . Once a good solution is found, and its genetic material begins to propagate throughout the population, the GA begins to optimize using the first case of the fitness function. With a valid classifier for every object it is then possible for the GA to explore the pruning of the feature space. Regarding novel object detection, the approach produced nearly perfect results with only one false negative, object 32. However, object 32 is almost identical to object 29, they only differ slightly in color space. Perhaps an interest operator that uses color information explicitly could help avoid ambiguous situations such as this. Finally, regarding the estimation of the most salient object within a

set, the algorithm also produced coherent selections. The objects selected as most salient, shown in Figure 4, are appreciably different than the rest, these objects tend to lack texture and exhibit small color variations. Furthermore, all of the other objects in the data set tend to have at least one similar counterpart, i.e. more than one toy car, and various small boxes. In conclusion, the proposed approach produced promising initial results for object recognition, novel object detection and salient object estimation. Future work concentrates on using images with complex backgrounds, in order to perform scene classification of real world images where the benefits of a multimodal model are expected to become evident.

Acknowledgements

Research funded by UC MEXUS-CONACyT Collaborative Research Grant 2005 through the project “Intelligent Robots for the Exploration of Dynamic Environments”, the Ministerio de Educación y Ciencia through the project Oplink - TIN2005-08818-C04, and Junta de Extremadura Spain. First author supported by scholarship 174785 from CONACyT México. This research was also supported by the LAFMI project, and special thanks is given to the Complex Team - INRIA Roquencourt and Grupo de Evolución Artificial - Universidad de Extremadura campus Mérida.

References

- [1] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7:179–188, 1936.
- [2] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press, Cambridge, MA, USA, 1992.
- [3] David G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision, 20-25 September, 1999, Kerkyra, Corfu, Greece*, volume 2, pages 1150–1157. IEEE Computer Society, 1999.
- [4] M. Markou and S. Singh. A neural network-based novelty detector for image sequence analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(10):1664–1677, 2006.
- [5] Krystian Mikolajczyk and Cordelia Schmid. A performance evaluation of local descriptors. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(10):1615–1630, 2005.
- [6] S.A. Nene, Nayar, and H. S.K., Murase. Columbia object image library (coil-100). Technical report, Department of Computer Science. Columbia University, 1996.
- [7] Pekka Paalanen, Joni-Kristian Kamarainen, Jarmo Ilonen, and Heikki Kälviäinen. Feature representation and discrimination based on gaussian mixture model probability densities-practices and algorithms. *Pattern Recognition*, 39(7):1346–1358, 2006.
- [8] Cordelia Schmid and Roger Mohr. Local grayvalue invariants for image retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(5):530–534, May 1997.
- [9] Leonardo Trujillo and Gustavo Olague. Synthesis of interest point detectors through genetic programming. In Mike Cattolico, editor, *Proceedings of GECCO 2006*, volume 1, pages 887–894. ACM, 2006.
- [10] Leonardo Trujillo and Gustavo Olague. Scale invariance for evolved interest operators. In Mario Giacobini et al., editor, *Proceedings of EvoWorkshops 2007*, volume 4448 of *Lecture Notes in Computer Science*, pages 423–430. Springer, 2007.
- [11] Leonardo Trujillo, Gustavo Olague, Pierrick Legrand, and Evelyne Lutton. Regularity based descriptor computed from local image oscillations. *Optics Express*, 15:6140–6145, 2007.