

Synthesis of Interest Point Detectors Through Genetic Programming

Leonardo Trujillo
Proyecto EvoVisión

Departamento de Ciencias de la Computación,
División de Física Aplicada
Centro de Investigación Científica y de
Educación Superior de Ensenada
Ensenada B.C. México
trujillo@cicese.mx

Gustavo Olague^{*}
Proyecto EvoVisión

Departamento de Ciencias de la Computación,
División de Física Aplicada
Centro de Investigación Científica y de
Educación Superior de Ensenada
Ensenada B.C. México
olague@cicese.mx

ABSTRACT

This contribution presents a novel approach for the automatic generation of a low-level feature extractor that is useful in higher-level computer vision tasks. Specifically, our work centers on the well-known computer vision problem of interest point detection. We pose interest point detection as an optimization problem, and are able to apply Genetic Programming to generate operators that exhibit human-competitive performance when compared with state-of-the-art designs. This work uses the repeatability rate that is applied as a benchmark metric in computer vision literature as part of the GP fitness function, together with a measure of the entropy related with the point distribution across the image. This two measures promote geometric stability and global separability under several types of image transformations. This paper introduces a Genetic Programming implementation that was able to discover a modified version of the DET operator [3], that shows a surprisingly high-level of performance. In this work emphasis was given to the balance between genetic programming and domain knowledge expertise to obtain results that are equal or better than human created solutions.

Categories and Subject Descriptors

I.4.7 [Image Processing and Computer Vision]: Feature Measurement—*feature representation, invariants*; I.2.2 [Artificial Intelligence]: Automatic Programming—*program synthesis*.

General Terms

Algorithms, Experimentation, Performance, Theory.

^{*}Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '06, July 8–12, 2006, Seattle, Washington, USA.
Copyright 2006 ACM 1-59593-186-4/06/0007 ...\$5.00.

Keywords

Evolutionary Computer Vision, Synthesis of Interest Points.

1. INTRODUCTION

Nowdays, interest point detectors are commonly used to approach low-level tasks as part of the pre-attentive stage that localize distinctive parts of the image that will be used in the attentive stage, in which relationships between these features and grouping takes place. Interest points are image pixels that exhibit a high level of variation with respect to a particular local measure. Common low-level features such as: edges, blobs, corners, and interest points simplify image analysis by reducing the overwhelming amount of information contained in images that higher level vision tasks need to process. Interest points in particular are used by state-of-the-art systems to perform high-level vision tasks, such as: object detection and recognition, matching, 3D reconstruction, tracking, and image registration, to name but a few. Different types of interest point operators¹ can be found in computer vision literature. Interest point operators can be grouped together according to the manner in which they model image information. From this we can identify two major groups. One group models images as three-dimensional surfaces. The idea is to extract measures directly related to the principal curvatures computed around each point [3, 4, 10]. A second more commonly used group of operators use the gradient distribution around each point captured by the local second moment matrix as their interest measure [8, 7, 20].

All current interest point detectors are hand coded designs, product of the analysis and interpretations of how the problem has been confronted by a human mind. Analysis is understood as the science which treats of the exact relations existing between quantities or magnitudes, and of the methods by which, in accordance with these relations, quantities sought are deducible from other quantities known or supposed; the science of spatial and quantitative relations. On the other hand, modern trends in computer science are developing new approaches that are based on the idea of automation of empirical learning through a process of synthesis. Synthesis is understood as the combination of separate

¹The terms *interest point operators* and *interest point detectors* are used interchangeably in computer vision literature.

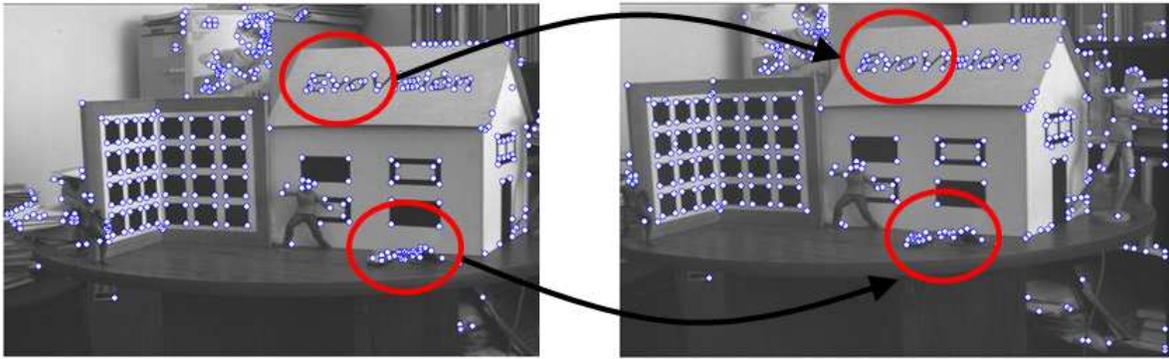


Figure 1: A stereo pair taken at the EvoVisión Laboratory, notice how interest points can be used to compute the correspondence between both images. The interest points depicted on both images were obtained with the IPGP2 presented in this work that outperforms past human designs.

elements of thought into a whole, as of simple into complex conceptions. Thus, synthesis refers to the art or process of making a compound by putting the ingredients together, as contrasted with analysis. In genetic programming it is common to find the idea of synthesizing a design from scratch [1]. However, we prefer the idea of using both approaches in our own research [18, 15, 17]. Analysis and synthesis, though commonly treated as two different methods, are, if properly understood, only the two necessary parts of the same method. Each is the relative and correlative of the other. This work presents how an appropriate analysis of the problem of interest point detection guide genetic programming in the search of a synthetic operator that exhibit human competitive results. The properties that are aimed to be fulfilled by an interest point detector are:

1. **Global separability** between extracted points.
2. **High information content** when compared to other pixels.
3. **Stability** under certain types of image transformations, *i.e.* image rotations.

Global separability of extracted points suggests that on an average-scene interest points should not be crowded together on isolated portions of the image, see Figure 1. This criterion is obviously image dependent, and requires a priori knowledge of the expected number of points and their position within the image. High information content refers to the uniqueness of the local neighborhood of the extracted points. This property would facilitate interest point matching based on local descriptors. Stability of detected interest points is probably the most important criterion and is the only one for which a widely accepted metric exists, the detectors repeatability rate [19]. This list is in no way exhaustive nor is it rigorous. However, it does express desired properties of interest point operators that are beneficial for high-level computer vision tasks that heavily rely on the extracted points.

Our approach poses interest point detection as an optimization problem, and presents a Genetic Programming based learning approach that constructs computer functions with stable and high performance, using a widely accepted measure. The evolved operators exhibit a high repeatability

rate and global separability. Furthermore, the Genetic Programming learning approach discovered a modified version of the DET interest point operator proposed by Beaudet in 1978. This last result exhibits the powerful capabilities of Genetic Programming for automatic generation of efficient and coherent problem solutions when an appropriate fitness function is implemented.

This paper is organized as follows: Section 2 presents a brief discussion on using machine intelligence to automatically construct computer vision applications. Section 3 is an overview of popular interest point operators. Section 4 outlines our proposed method and defines its implementation details. Section 5 shows preliminary experimental results. Section 6 is a discussion of the presented work and gives final conclusions, and finally in Section 7 we contemplate possible future work.

2. RELATED WORK

The development of machine learning algorithms that solve computer vision problems, is a relatively new and highly promising field of research [16]. These methods can be regarded as part of a broad class of algorithms that perform visual learning. Visual learning is the process in which an artificial system autonomously acquires knowledge from training images to solve a given visual task [11]. Most of the published work in this area is centered on solving mid-level and high-level vision tasks; these tasks include: feature selection, feature construction or synthesis, object detection, and image segmentation, to name but a few. On the other hand, learning low-level feature extraction has received less attention from the research community. A plausible explanation for this is the fact that common low-level operators have been extensively studied and are well understood by the computer vision research community. However, evolutionary computation has the capability of endowing a learning system with the ability to try new and uncommon image processing strategies that human designers might, or will not consider. If we learn image operators in this way, it is possible to provide researchers with deeper insights on the nature of the problem domain. For example, Bala *et al.* [2] use a Genetic Algorithm to perform feature selection for recognizing visual concepts. Sun *et al.* [21] perform PCA on two different object classes, and use a GA to select the best sub-

set of eigenimages for object recognition. Interestingly, the subset chosen by the learning algorithm shows that a high corresponding eigenvalue is not necessary nor sufficient for a given eigenimage to be useful for classification. Working at the feature synthesis level, machine learning algorithms are used to generate operators that extract specialized image features for a specific problem. Howard *et al.* [9] use Genetic Programming to generate image features for target detection in SAR images. Zhang *et al.* [22], use Genetic Programming to perform multiclass detection of small objects present in large images. This work uses domain independent pixel statistics as the GP terminal set, and shows how a single evolved program solves both the object detection and localization problem in a coupled manner.

Directly related to our work, and specifically addressing the problem of automatically learning an interest point detector, two main contributions exist. Ebner [5] posed interest point detection as an optimization problem, and attempted to evolve the Moravec interest point operator [13] using Genetic Programming. The author reports a 15% localization error between interest point detection of his evolved operator and that obtained using the Moravec detector. Nevertheless, this result cannot be taken as an acceptable evaluation criterion for the quality of the evolved operator due to the fact that the Moravec operator cannot be assumed to be a suitable performance metric. A second paper by Ebner *et al.* [6] presents an evolved operator that is optimized for computing the optical flow in a particular image sequence. Despite the fact that [6] showed promising results when computing optical flow estimation, the optimization criteria that was used does not guarantee the generality for other vision tasks where interest point detection is required.

The Genetic Programming learning approach used on both attempts by Ebner fails to capture the essence of the desired properties that a general interest point detector should attempt to fulfill. The shortcomings of these contributions are overcome in our present work by realizing a thorough analysis in order to define a suitable problem statement. We optimize interest point detectors with a measure of their stability given by the operators repeatability rate; as well as, a quantification of the operators global separability. Repeatability rate is a standard computer vision performance metric for interest point detectors [19], and global separability was estimated by including an entropy measure of the interest point localization histogram as part of the fitness evaluation function. As a result, a complete fitness measure was developed and applied in the genetic programming used to solve this task.

3. INTEREST POINT OPERATORS

This section covers only relevant aspects of interest point detection necessary to explain our work. For a more thorough review of interest point and corner detectors we recommend the work by Nobel [14], Schmid, *et al.* [19], and Olague and Hernández [17]. Interest point detection is a byproduct of research devoted to corner detection in images. Corner detectors are commonly classified in three main classes: *Countour based methods*, *Parametric model based methods* [17] and *Image intensity based methods* [13, 3, 10, 7, 8, 20]. The class of corner detectors that operate directly on the intensity image are more appropriately referred to as interest point detectors. These operators define a function that

operates on a local neighborhood and extracts a *cornerness* or *interest* measure from every pixel in the image. This operation produces a new image that can be referred to as the *interest image*, that is thresholded to extract the points with the highest *interest* measure. Conceptually, this type of operators were designed as corner detectors; however, their detection capabilities are not limited to points that conform to the geometric concept of "corner" [17]. Hence, they extract all points where image intensity variations are high with respect to a particular measure. Furthermore, the kind of image features they extract are better explained with the more general concept of *interesting* or *interest point*. In this context, interest points are image pixels that show a distinctive property that make them suitable for applications where specific scene points need to be tracked across multiple images.

Popular interest point detectors can be grouped according to the manner in which they model and extract image information. The first major group extracts an interest measure using the gradient distribution around each point captured by the local second moment matrix as their interest measure. The first method that established interest point detectors was proposed by Moravec in 1977. The work by Harris and Stephens [8], Forstner [7], and Shi and Tomasi [20], among others, extend the concept proposed by Moravec. From this group the Harris operator has emerged as the most popular detector used in vision applications. A second group of interest point operators extract measures directly related to the principal curvatures computed around each point. This includes detectors proposed by Baudet [3], Kitchen and Rosenfeld [10], and Dreschler and Nagel [4].

3.1 Interest Point Operator Performance

As previously mentioned, robust image features are necessary to solve computer vision problems related to: 3D reconstruction, image registration, matching, object detection, recognition, and optical flow estimation, to mention but a few. These applications require local image features that are simple to detect, provide useful information for post-processing, and show geometric stability under different types of image transformations. Such transformations include: translation, rotation, illumination change, scale change and affine transformations. Of the previous list of transformations, interest points are only suitable for robust detection in the presence of the first three. Region detectors robust to scale change and affine transformations require a more general operation beyond the scope of this study. Nevertheless, an efficient and reasonable evaluation criteria is required to estimate the reliability of interest point detectors. The very important work done by Schmid, *et al.* [19] established an effective measure that captures the essence of the desired characteristic of robustness and stability for low-level feature extractors. The authors establish the repeatability rate as this primary performance metric. Based on this performance metric that is obtained experimentally, they conclude that the Harris operator shows the greatest stability, outperforming every other detector included in the survey. The results reported in [19] have heavily contributed into making the Harris interest point detector the most widely used detector within the computer vision community. However, their list of compared detectors is not exhaustive, being that they leave out the detectors proposed by Baudet, and Shi and Tomasi, to name a couple. Moreover, according

to our results Beaudet’s detector shows competitive results compared with Harris. Also, we would like to mention that the comparison is made through experimentation due to the difficulty to make such a comparison with analytical approaches [17]. This is also a favorable point to justify the use of Genetic Programming.

3.2 Repeatability

The performance metric for evaluating the operator’s stability and robustness is the repeatability rate that measures how interest point detection is independent of imaging conditions [19]. A point x_1 detected in image I_1 is repeated in image I_i if the corresponding point x_i is detected in image I_i . In the case of planar scenes a relation between points x_1 and x_i can be established with the homography $H_{1,i}$ where:

$$x_i = H_{1,i}x_1 \quad (1)$$

The repeatability rate measures the number of repeated points between both images, with respect to the total number of detected points. However, parts of image I_1 may not appear on the transformed image I_i . In order to account for this, repeated and detected points are only counted if they lie in the common parts of image I_1 and image I_i . Furthermore, a small amount of detection error needs to be taken into account because exact localization is not required in most computer vision applications. This is in direct contrast with high precision applications common in the photogrammetric community that would require the more precise capabilities of parametric model based corner detectors [17]. Consequently, to compute the repeatability rate, a repeated point is said to be detected at pixel x_i if it lies within a given neighborhood of x_i of size ϵ , see figure 2.

The set of point pairs (x_1^c, x_i^c) that lie in the common part of both images and correspond within an error of size ϵ is defined by:

$$R_i(\epsilon) = \{(x_1^c, x_i^c) \mid \text{dist}(H_{1,i}x_1^c, x_i^c) < \epsilon\} \quad (2)$$

Thus the repeatability rate $r_i(\epsilon)$ of points extracted from image I_i with respect to points from image I_1 , is defined by the following equation:

$$r_i(\epsilon) = \frac{|R_i(\epsilon)|}{\min(\gamma_1, \gamma_i)} \quad (3)$$

where $\gamma_1 = |\{x_1^c\}|$ and $\gamma_i = |\{x_i^c\}|$ are the total number of points extracted from image I_1 and image I_i respectively [19].

4. OUTLINE OF OUR APPROACH

We present a visual learning approach to automatically construct an interest point detector using Genetic Programming. Each individual in the GP population represents a candidate interest point operator. To define an application of Genetic Programming it is necessary to define three concepts: *The Fitness Function*, *Function Set* and the *Terminal Set*.

4.1 Fitness Function

Our approach uses a fitness assignment that is proportional to its mean repeatability rate $r_J(\epsilon)$ computed for a set $J = \{I_i\}$ of n training images, where $i = 1 \dots n$. A base image I_i is used to compute the repeatability on all other images in J . However, the GP search could easily get lost

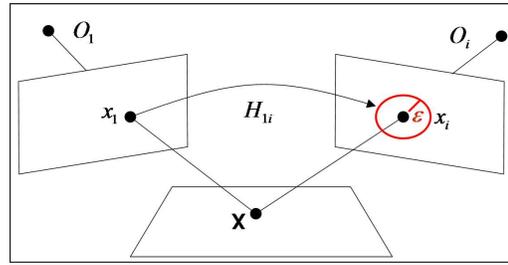


Figure 2: A 3D point X is projected onto points x_1 and x_i on images I_1 and I_i respectively. Point x_1 is said to be repeated by x_i , if a point is detected within a neighborhood of x_i of size ϵ . For planar scenes x_1 and x_i are related by the homography $H_{1,i}$.

in unwanted maxima if appropriate considerations are not taken into account when designing the fitness function. For example, one can imagine a degenerate case where the GP search could concentrate on individuals that extract useless points clustered together on textureless regions and still manage to have a high repeatability rate. Moreover, the training images present highly textured regions distributed across the image plane. Hence, a *good* detector should extract uniformly distributed points across the image plane. Consequently, three other terms were incorporated in the fitness function and combined in a multiplicative way:

$$f(x) = r_J(\epsilon) \cdot \phi_x^\alpha \cdot \phi_y^\beta \cdot N_{\%}^\delta \quad (4)$$

where the functions,

$$\phi_x = \frac{1}{1 + e^{-a(H_x - c)}} \quad (5)$$

$$\phi_y = \frac{1}{1 + e^{-a(H_y - c)}} \quad (6)$$

are sigmoidal functions used to promote point dispersion along the x and y directions. The terms H_x given by:

$$H_x = - \sum P(\cdot) \log_2 [P(\cdot)] \quad (7)$$

represent the entropy value of the spatial distribution of detected interest points along each direction. $P(\cdot)$ is approximated by the histogram of interest point localizations. Moreover, because of the logarithmic nature of the entropy function, ϕ_x and ϕ_y are set to promote entropy values lying within a very small range. The final term,

$$N_{\%} = \frac{\text{requestedpoints}}{\text{extractedpoints}} \quad (8)$$

is a penalizing factor that reduces the fitness value for detectors that return less than the total number of requested points. Finally, α , β and δ control the amount of influence that each term has on $f(x)$. Parameter setting was tuned experimentally. The appropriate value for parameters of (ϕ_x, ϕ_y) were obtained by estimating average point localization entropy produced by the Harris and DET operators.

4.2 Function Set

The function set F contains 6 unary functions and 5 binary functions. All functions, input and output, are data

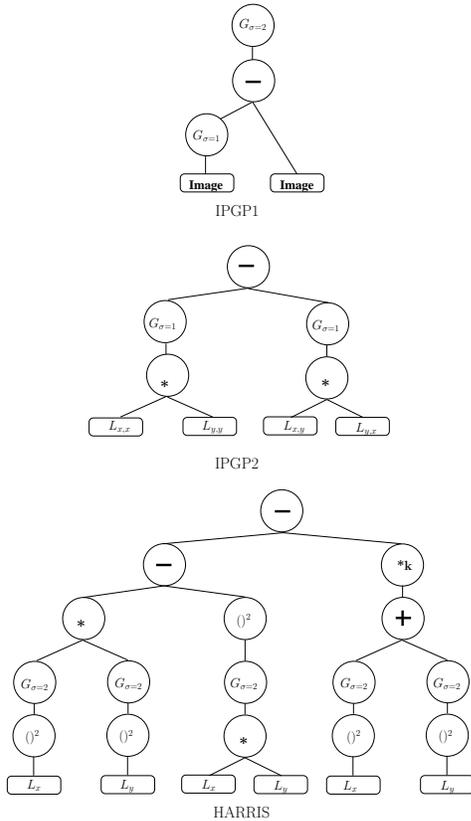


Figure 3: Tree representation for three different interest point detectors. From top to bottom: IPGP1, IPGP2 and Harris. It can be easily observed the simple structure found by the GP search when compared with the most popular interest point detector of computer vision.

matrices with the same size as images in J . The subset of binary and unary functions are:

$$F_{2ary} = \{+, -, | - |, *, /\} \quad (9)$$

$$F_{1ary} = \{A^2, \sqrt{A}, \log_2, EQ, G(\sigma = 1), G(\sigma = 2)\} \quad (10)$$

Where EQ is the histogram equalization, and $G(\sigma = x)$ are Gaussian filters with blur σ . The complete function set is:

$$F = F_{2ary} \cup F_{1ary} \quad (11)$$

4.3 Terminal Set

Great care was taken to design an appropriate terminal set. Thanks to previous understanding of the analytical properties of corner detectors described in [19, 17], we can conclude that an effective IP operator requires information pertaining to the rate of change in image intensity values. Consequently, the terminal set includes first and second order image derivatives. However, we do not claim that this set is necessary nor sufficient and further work will try to determine an optimal set of useful information for interest point detection. Furthermore, the terminal set is image dependent, which means that each image $I_i \in J$ has a corre-

sponding T_i defined by:

$$T_i = \{I_i, L_{i,x}, L_{i,x,x}, L_{i,x,y}, L_{i,y,y}, L_{i,y}, I_{i,\sigma=1}\} \quad (12)$$

Where $L_{i,w} = I_i * G_w(\sigma = 1)$ are image derivatives computed in the w direction using a convolution with Gaussian kernel derivatives, and $I_{i,\sigma=1}$ is the smoothed image computed by a convolution with a Gaussian smoothing function.

5. EXPERIMENTS

This section explains the implementation of our approach and summarizes our experimental results.

5.1 Implementation Details

The implementation of the previously described approach was programmed on Matlab, with the Genetic Programming toolbox GPLAB². The image sequence used for training was the VanGogh set of a planar scene with rotation transformations. For testing, two image sequences were used: Monet and Graph. The former is a sequence of a rotated planar scene and the latter is an image under different illumination conditions. All image sets were downloaded from the Visual Geometry Group website³, along with matlab source code for computing the repeatability rate and binary files for extracting Harris interest points. All experiments were made on a PC with AMD64 processor and 526MB of RAM running Linux OS. The following list specifies the GP runtime parameters for both experimental runs reported in the following section:

- **Population size and initialization:** We used a population size of 75 individuals initiated with the Ramped Half-and-Half method. The maximum size was set with a tree depth of 7 levels.
- **Crossover and Mutation:** The crossover probability was set to $p_c = 0.85$, and mutation probability was set to $p_\mu = 0.1$.
- **Selection Operator:** Selection for genetic operators was performed using a tournament selection method that uses lexicographic parsimony pressure [12]. In this kind of tournament selection if equally fit individuals are selected for competition, the smallest one is chosen. The tournament size was set to 4.
- **Survival method:** We use a keep-best survival strategy. In this method the best individual from both parents and children is retain for the new population. The remaining places in the new population are occupied by children only.
- **Fitness function parameter settings:** The parameters for our fitness function were set experimentally, and are the following: $a_x = 7$, $c_x = 5.05$, $a_y = 6$, $c_y = 4.3$, $\alpha = 10$, $\beta = 10$, $\delta = 2$.

5.2 Results

We present two different interest point operators generated with our approach: IPGP1 and IPGP2⁴ out of 30

²<http://gplab.sourceforge.net/index.html>, GPLAB A Genetic Programming Toolbox for MATLAB by Sara Silva

³<http://www.robots.ox.ac.uk/~vgg/research/>

⁴IPGP stands for "Interest Point detector with Genetic Programming"

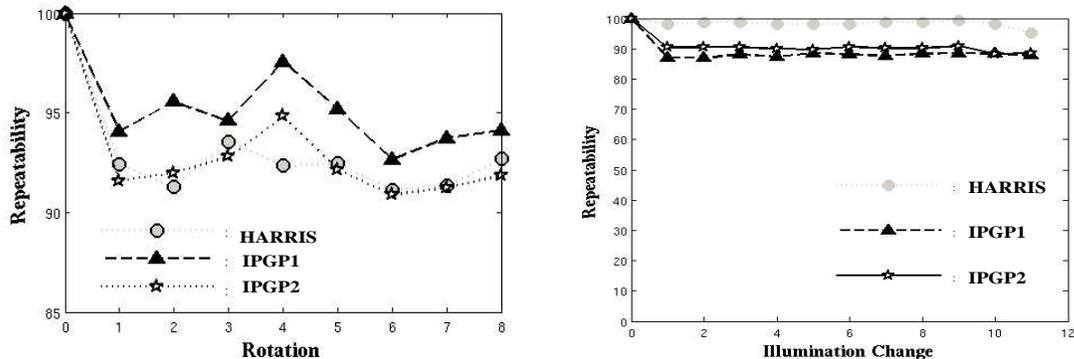


Figure 4: Performance measures: left) repeatability rate plotted against image rotation for the VanGogh images; right)repeatability rate plotted against illumination change for the Graph images.

IPGP1	$G(\sigma = 2) * [G(\sigma = 1) * I - I]$
IPGP2	$G(\sigma = 1) * [L_{xx} \cdot L_{yy}] - G(\sigma = 1) * [L_{xy} \cdot L_{yx}]$

Table 1: Mathematical expressions for IPGP1 and IPGP2.

attempts. Each detector was generated on a separate run of our algorithm. Even do more experiments have been carried out we believe that these operators suffice to exemplify the usefulness of our approach. The first operator IPGP1 has an extremely simple structure. The basic approach of IPGP1 for interest point extraction is a simple two step process: the first step is to extract high frequencies from the image by subtracting a smoothed version of the image from the original; the second step uses a Gaussian to smooth the extracted high frequencies. This is basically an operation that is known as difference-of-Gaussian. IGP1 does not use image derivatives to extract image regions with high signal variations. Despite its seemingly simple operation it achieves a 95% of average repeatability rate on the training image set that is the principal benchmark used by Schmid et al. in 2000. The second evolved detector IPGP2, represents a modified version of the DET operator proposed by Beaudet [3] in 1978. The DET operator is the determinant of the Hessian Matrix computed at each image pixel. Baudet defined the DET measure as follows:

$$DET = I_{x,x} \cdot I_{y,y} - I_{x,y} \cdot I_{y,x} \quad (13)$$

The function constructed by our learning approach for IPGP2 is expressed in the second row of Table 1. IPGP2 has the same basic structure as DET, with the added difference that the function is averaged around a local neighborhood with a smoothing Gaussian function. In this instance the GP was able to discover a similar human-made design. The average repeatability rate for IPGP2 on the training set was 92%.

Table 1 summarizes the mathematical expressions for each evolved operator, IPGP1 and IPGP2. Figure 3 shows the tree structure for IPGP1, IPGP2 and the Harris operator for comparison. Figure 4 shows two graphs characterizing the performance of both evolved operators, compared with the Harris detector. Finally, figure 5 shows extracted interest points from samples taken from each of the three image sequences.

6. DISCUSSION AND CONCLUSIONS

This paper presented a novel approach that performs learning to construct a low-level feature extraction function for computer vision applications. The type of low-level features that are extracted are image interest points. Our approach poses interest point detection as an optimization problem, and developed a learning methodology that allows a computer to automatically generate interest point detectors. Learning was conducted with Genetic Programming. The fitness evaluation function promoted performance that was optimized according to the repeatability rate and global separability of extracted points. Experimental results showed that the proposed approach generates reliable and compact operators. Our Genetic Programming implementation uses image derivatives as a terminal set, and simple arithmetic functions as a function set.

We present two learned operators IPGP1 and IPGP2 that exhibit competitive results when compared with the most popular interest point detector in computer vision proposed by Harris and Stephens. Furthermore, the construction of IPGP2 demonstrated the ability of Genetic Programming to rediscovered a modified version of a previous man made design, Beaudet’s DET operator, which is one of the earliest proposed interest point detectors. Nevertheless, the results presented in this paper, are not intended as a claim that the evolved operators are superior to any other. However, the results clearly demonstrate that learning techniques based on simulated evolution are capable of producing competitive and useful low-level feature extractors. Moreover, it is our belief that if using a similar methodology, it could be possible to design learning algorithms that generate appropriate feature extractors for different kinds of computer vision tasks.

This work provides an example in the area of computer vision of how evolutionary computation could achieve human-competitive results. This work fulfills six out of the eight criteria that are used in the evolutionary computation community to demonstrate if an automatically created result is considered human-competitive:

- (B) The result is equal to or better than a result that was accepted as a new scientific result at the time when it was published in a peer-reviewed scientific journal.
- (C) The result is equal to or better than a result that was placed into a database or archive of results maintained by

HARRIS

IPGP1

IPGP2

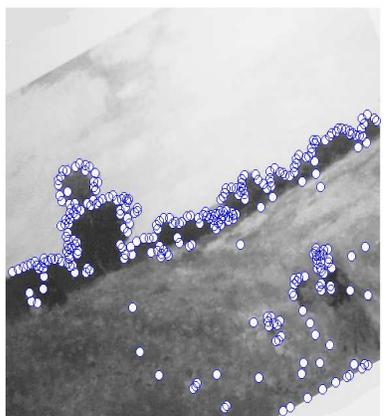
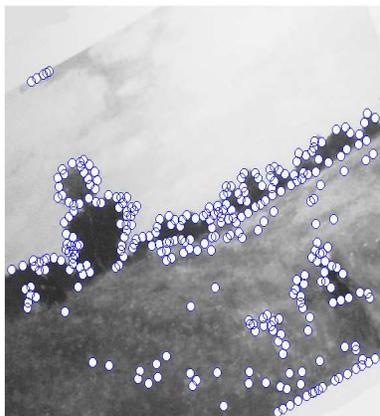
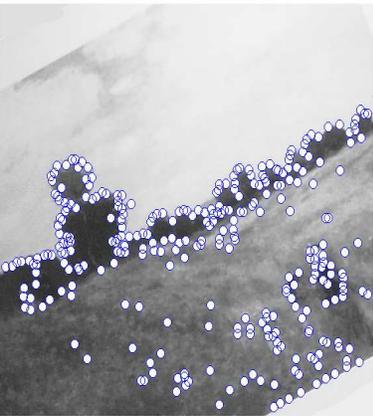
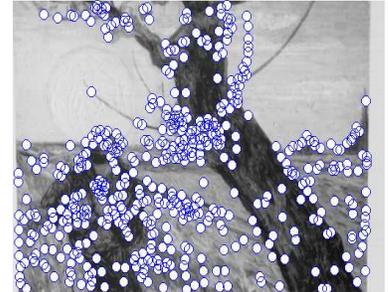
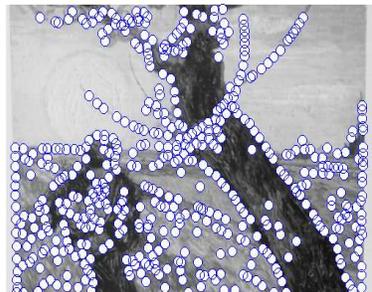
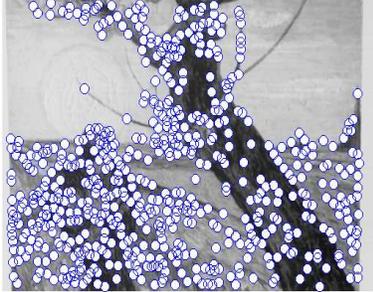


Figure 5: Sample images of extracted interest points. Columns correspond to Harris, IPGP1 and IPGP2 from left to right. Rows correspond VanGogh, Graph and Monet from top to bottom.

an internationally recognized panel of scientific experts.

(D) The result is publishable in its own right as a new scientific result independent of the fact that the result was mechanically created.

(E) The result is equal to or better than the most recent human-created solution to a long-standing problem for which there has been a succession of increasingly better human-created solutions.

(F) The result is equal to or better than a result that was considered an achievement in its field at the time it was first discovered.

(G) The result solves a problem of indisputable difficulty in its field.

7. FUTURE WORK

The current state of our study leaves some open issues, such as:

- Could a fitness function be designed that avoids parameter tuning? Could the incorporation of fitness function parameters in the evolutionary process be advantageous?
- Is the terminal set proposed in this paper sufficient or necessary? It is possible to augment the terminal set that we used. For example, we could add image responses to different types of filter banks or texture extraction methods.
- What is the effect of including more than one image sequence as a training set? Will this produce a more general operator?

Our work also opens the path to several promising research opportunities, that include:

- Explicitly adding a term to the fitness function that accounts for high information content around a local neighborhood, using the response to a local descriptor to estimate the usefulness of the extracted point.
- Construct specific interest point detectors that are optimized for a certain type of images, i.e. outdoor scenes, or indoor office environments.
- Extend this work to extract interest regions, a more general and far more interesting result for certain vision applications.

8. ACKNOWLEDGMENTS

This research was funded by UC MEXUS-CONACYT Collaborative Research Grant 2005; through the project "Intelligent Robots for the Exploration of Dynamic Environments". First author supported by scholarship 174785 from CONACyT. This research was also supported by the LAFMI project.

9. REFERENCES

- [1] S. H. Al-Sakran, J. R. Koza, and L. W. Jones. Automated re-invention of a previously patented optical lens system using genetic programming. In *EuroGP 2005*, volume LNCS 3447, pages 25–37, 2005.
- [2] J. Bala, K. DeJong, J. Huang, H. Vafaie, and H. Wechsler. Using learning to facilitate the evolution of features for recognizing visual concepts. *Evolutionary Computation*, 4(3):297–311, 1996.
- [3] P. R. Beaudet. Rotational invariant image operators. In *Proc. IAPR 1978*, pages 579–583, 1978.
- [4] L. Dreschler and H.-H. Nagel. Volumetric model and 3d trajectory of a moving car derived from monocular tv-frame sequence of a street scene. *Computer Vision, Graphics and Image Processing*, 20(3):199–228, 1981.
- [5] M. Ebner. On the evolution of interest operators using genetic programming. In *Proc. EuroGP'98*, pages 6–10, Paris, France, 14-15 1998.
- [6] M. Ebner and A. Zell. Evolving a task specific image operator. In *Proc. EvoIASP'99*, volume 1596, pages 74–89, Goteborg, Sweden, 28-29 1999. Springer-Verlag.
- [7] W. Forstner. A framework for low level feature extraction. In *Proc. ECCV 1994*, pages 383–394, Sweden, May 1994.
- [8] C. Harris and M. Stephens. A combined corner and edge detector. In *Proc. Fourth Alvey Vision Conf.*, volume 15, pages 147–151, 1988.
- [9] D. Howard, S. C. Roberts, and R. Brankin. Target detection in imagery by genetic programming. *Adv. in Eng. Software*, 30(5):303–311, 1999.
- [10] L. Kitchen and A. Rosenfeld. Gray-level corner detection. *Pattern Recognition Letters*, 1:95–102, 1982.
- [11] K. Krawiec and B. Bhanu. Visual learning by evolutionary feature synthesis. *IEEE Transactions on Systems, Man and Cybernetics, Part B, Special Issue on Learning in Computer Vision and Pattern Recognition*, 35(3):409–425, june 2005.
- [12] S. Luke and L. Panait. Lexicographic parsimony pressure. In *Proceedings of GECCO-2002*, pages 829–836, 2002.
- [13] H. P. Moravec. Towards automatic visual obstacle avoidance. In *IJCAI*, page 584, 1977.
- [14] J. Noble. *Descriptions of Image Surfaces*. D.Phil thesis, RoboticsResearch Group, Department of Engineering Science, Oxford University, 1989.
- [15] G. Olague. Automated photogrammetric network design using genetic algorithms. *Photogrammetric Engineering & Remote Sensing*, 68(5):423–431, 2002. Paper awarded the 2003 First Honorable Mention for the Talbert Abrams Award, by ASPRS.
- [16] G. Olague, S. Cagnoni, and E. Lutton. Introduction to the special issue on evolutionary computer vision and image understanding. *Pattern Recognition Letters*, to appear, 2006.
- [17] G. Olague and B. Hernández. A new accurate and flexible model based multi-corner detector for measurement and recognition. *Pattern Recognition Letters*, 26(1):27–41, 2005.
- [18] G. Olague and R. Mohr. Optimal camera placement for accurate reconstruction. *Pattern Recognition*, 35(4):927–944, 2002.
- [19] C. Schmid, R. Mohr, and C. Bauckhage. Evaluation of interest point detectors. *Int. J. Comput. Vision*, 37(2):151–172, 2000.
- [20] J. Shi and C. Tomasi. Good features to track. In *Proc. CVPR 1994*, Seattle, Jun 1994.
- [21] Z. Sun, G. Bebis, and R. Miller. Object detection using feature subset selection. *Pattern Recognition*, 37:2165–2176, 2004.
- [22] M. Zhang, V. B. Ciesielski, and P. Andreae. A domain-independent window approach to multiclass object detection using genetic programming. *EURASIP Journal on Applied Signal Processing*, 2003(8):841–859, jul 2003.