



ELSEVIER

Pattern Recognition Letters 23 (2002) 335–345

Pattern Recognition
Letters

www.elsevier.com/locate/patrec

Dynamic flies: a new pattern recognition tool applied to stereo sequence processing

Jean Louchet^{a,b,*}, Maud Guyon^a, Marie-Jeanne Lesot^a,
Amine Boumaza^b

^a *Lab. d'Electronique et Informatique, Ecole Nationale Supérieure de Techniques Avancées, 32 Boulevard Victor, 75739 Paris Cedex 15, France*

^b *INRIA Rocquencourt, BP75, 78153 Le Chesnay Cedex, France*

Received 27 March 2001; received in revised form 2 July 2001

Abstract

The “fly algorithm” is a fast artificial evolution-based technique devised for the exploration of parameter space in pattern recognition applications. In the application described, we evolve a population which constitutes a particle-based three-dimensional representation of the scene. Each individual represents a three-dimensional point in the scene and may be fitted with optional velocity parameters. Evolution is controlled by a fitness function which contains all pixel-level calculations, and uses classical evolutionary operators (sharing, mutation, crossover). The combined individual approach and low complexity fitness function allow fast processing. Test results and an application to mobile robotics are presented. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Artificial evolution; Pattern recognition; Computer vision; Image processing; Parameter space exploration

1. Introduction

Parameter-space vote approaches to computer vision and pattern recognition like the generalised Hough transform (Hough, 1962) are limited by the number of parameters and the size of the search space. Artificial evolution (Goldberg, 1989;

Holland, 1975; Rechenberg, 1994), seen as a class of parameter space exploration techniques, provides successful solutions to some image processing problems (Lutton and Martinez, 1994; Roth and Levine, 1992; Ser et al., 1999) but often suffer from the costly calculations needed if they have to manipulate populations where each individual is in itself a complex description of the scene. The “Parisian” approach (Collet et al., 1999; Louchet, 2000; Louchet, 2001) to artificial evolution introduces each individual as a part of the solution of the problem. Thus, the solution is represented by the total population or by a large part of it. The algorithm presented in this paper

* Corresponding author. Tel.: +33-1-4552 6075; fax: +33-1-4552 8327.

E-mail addresses: louchet@ensta.fr (<http://www.ensta.fr/~louchet>) (J. Louchet), guyon@ensta.fr (M. Guyon), lesot@ensta.fr (M.-J. Lesot), boumaza@inria.fr (<http://www-rocq.inria.fr/~boumaza>) (A. Boumaza).

uses 3-D points¹ ('flies') as the population's individuals, and evolves the population of flies using a pixel-based fitness function such that the flies concentrate onto the objects to be detected in the scene² (Section 2). Section 3 shows an extension to stereo image sequences through the introduction of velocity parameters. An application to mobile robot obstacle avoidance is shown in Section 4.

2. Evolving flies

2.1. Geometry and fitness function

A fly is defined as a 3-D point with coordinates (x, y, z) . The coordinates of the fly's projections are (x_L, y_L) in the image given by the left camera and (x_R, y_R) for the right camera. The cameras' calibration parameters are known and so x_L, y_L, x_R, y_R may be easily calculated from x, y, z using projective geometry (Haralick, 1980; Jain et al., 1994). If the fly is on the surface of an opaque object, then the corresponding pixels in the two images will normally have the same grey level.³ Conversely, if the fly is not on the surface of an object, thanks to the non-uniformity of objects and illumination, the

grey levels of its projections and their immediate neighbourhoods will be probably not identical. This is expressed as a fitness function used to control the evolution of the flies' population from random initial positions to converge onto the surfaces of the apparent objects.

The fitness function evaluates the degree of similarity of the pixel neighbourhoods of the projections of the fly onto each image. This ensures highest fitness values for the flies lying on the surface of an object:

$$\text{fitness}(\text{indiv}) = G \left/ \left(\sum_{\text{colours}} \sum_{(i,j) \in N} (L(x_L + i, y_L + j) - R(x_R + i, y_R + j))^2 \right) \right.$$

- (x_L, y_L) and (x_R, y_R) are the coordinates of the left and right projections of the current individual (see Fig. 1),
- $L(x_L + i, y_L + j)$ is the grey value of the left image at pixel $(x_L + i, y_L + j)$, and similarly with R for the right image.
- N is a neighbourhood introduced to obtain a more discriminant comparison of the fly's projections.

On colour images, square differences are calculated on each colour channel.

The normalizing factor is a local contrast measurement based on an image gradient calculation, giving higher fitness values to flies which project onto contrasted or textured areas of the images, and reducing the fitness of those over uniform regions. Experiments showed that the best results are obtained when G is defined as the square root of an image gradient norm such as




Fig. 1. Pixels b_1 and b_2 , projections of fly B, have identical grey levels, while pixels a_1 and a_2 , projections of fly A, which receive their illumination from two different physical points on the object's surface, have different grey levels.

¹ Particle swarms (Millonas, 1994; Eberhart and Kennedy, 1995) also use 3-D points as primitives. They were designed as optimisation tools using a metaphor of particles moving towards a target, using collective behaviours inspired by particle systems used in computer graphics. The essential difference is that our approach is based on evolutionary operators (selection, mutation, crossover) which are not used in particle swarm techniques. However one can outline some similarities between these two approaches: e.g. random speeds vs. mutations, or mutual avoidance vs. sharing.

² The tensor voting approach of Tang and Medioni, 1998 also uses a random distribution of 3-D points which do not evolve but are only used to initialise the calculation of a dense tensor field.

³ Or the same colour components if colour images are used. This is essentially true with matt (Lambertian) surfaces where scattering of incident light is isotropic. Usually, surfaces differ slightly from the Lambertian model, but this may be at least partly taken into account with the fitness function (see the remark on low spatial frequencies, Section 2.1). Reflections from glossy surfaces may give rise to virtual objects and wrong 3-D interpretation, as with other stereo analysis algorithms.

Sobel's (Jain et al., 1994). Thus, the highest fitness values are obtained for flies whose projections have similar *and significant* pixel surroundings. Additionally, we modified the denominator of the fitness function to suppress the continuous component and reduce its sensitivity to lower spatial frequencies (e.g. due to different camera responses).

2.2. Artificial evolution operators

An individual's chromosome is the triple (x, y, z) which contains the fly's coordinates. The population is initialised randomly inside the intersection of the cameras' fields of view, from a given distance (clipping line) to infinity (Fig. 2). The statistical distribution is chosen in order to obtain uniformly distributed projections in the left image. Their depth is chosen by uniformly distributing the values of z^{-1} between zero and $1/d_{\min}$.

Selection is elitist and deterministic. It ranks the flies according to their fitness values and retains the best individuals (typically 50%).

2-D sharing (Holland, 1975) reduces the fitness values of flies located in crowded areas to prevent them from getting concentrated into a small number of maxima. It reduces each fly's fitness by $K \times N$, where K is a "sharing coefficient" and N the number of flies which project into the left image within a distance R ("sharing radius") from the given fly. We determined an empirical rule to choose a satisfactory sharing radius R , considering that if the image is divided into squares of width $2R + 1$ pixels, there should ideally be one fly per square on average. Therefore $(2R + 1)^2 \times N_{\text{flies}} \approx N_{\text{pixels}}$ and




Fig. 2. The fly population is initialised inside the intersection of the cameras' 3-D fields of view.

$$R \approx \frac{1}{2} \left(\sqrt{\frac{N_{\text{pixels}}}{N_{\text{flies}}}} - 1 \right).$$

For example, on a 500×500 pixel image with 5000 flies, this gives $R = 3$.

Mutation allows an extensive exploration of the search space. It uses an approximation of a Gaussian random noise added to the flies' chromosome parameters (x, y, z) . We chose standard deviations $\sigma_x, \sigma_y, \sigma_z$ equal to R , so that they are of the same order of magnitude as the mean distance between neighbouring flies.

In real-world scenes, many objects contain straight lines or planar surfaces. We translated this geometrical property into a *barycentric crossover operator* which builds an offspring randomly located on the line segment between its parents: the offspring of two flies $F_1(x_1, y_1, z_1)$ and $F_2(x_2, y_2, z_2)$ is the fly $F_3(x_3, y_3, z_3)$ defined by $\text{OF}_3 = \lambda \text{OF}_1 + (1 - \lambda) \text{OF}_2$. The weight λ is chosen using a uniform random law in $[0, 1]$.⁴

2.3. Results

2.3.1. Synthetic images

Results are shown with the synthetic "Money" stereo image pair⁵ (Figs. 3 and 4) which allows easy readability of results, as most objects are vertical. Fig. 5 shows convergence results using a 5000 individual population, 50% mutation⁶ and 10% crossover rate⁷ at different stages of convergence. The scene is seen from above, showing the

⁴ It is generally accepted that such a crossover operator has contractive properties which may be avoided by using a larger interval. However the goal of this crossover operator is to fill in surfaces whose contours are easier to detect, rather than to extend them. It is therefore not desirable to use coefficients allowing the centre of gravity to lie outside the object's boundary. This is confirmed by most of our experiments showing that there is no benefit in using a wider interval, unlike in other applications of barycentric crossover.

⁵ "Money" colour image pair, 384×288 pixels, ©INRIA – Mirages project.

⁶ Unusually high mutation rates give best results in Parisian evolution.

⁷ Experiments showed that the best results are obtained with elitist selection, keeping the total mutation plus crossover rate around 60%.




Fig. 3. "Money" image, left.




Fig. 4. "Money" image, right.

axes x (right) and z (camera axis, pointing towards the top of the page).

2.3.2. Fine tuning the parameters

2.3.2.1. *Sharing.* Figs. 6–8 show how sharing affects the results on the same image data, showing the 1500 best individuals (30% of 5000) after 50 generations, with a mutation rate of 60%, no




Fig. 6. No sharing.

crossover and different sharing coefficients. Top images show the fly distributions seen from above (plan view); bottom images are front views (as seen from the camera) where darker pixels represent flies at shorter distances. Low sharing gives unwanted fly concentrations into high gradient regions; high sharing coefficients tend to degrade the results.

2.3.2.2. *Mutation and crossover rates.* Fig. 9 shows typical average fitness evolution curves with different combinations of mutation and crossover rates. The best results are obtained with mutation




Fig. 5. Convergence results after 10 (left), 50 (centre) and 100 (right) generations. The line of flies which appears at the background corresponds to the wooden panel in the image background; the four fly clusters in the foreground correspond to the four coins. Note that the coin on the right is not vertical and gives a less distinct cloud.

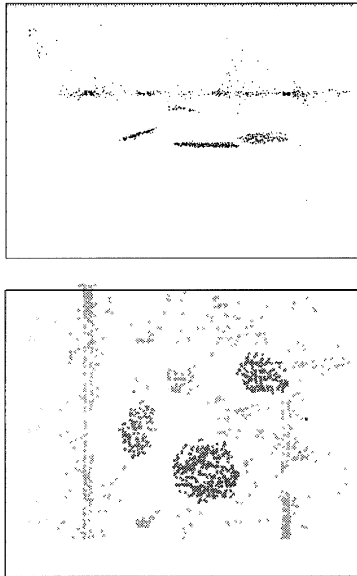


Fig. 7. Medium sharing (radius 2, coefficient 0.2).

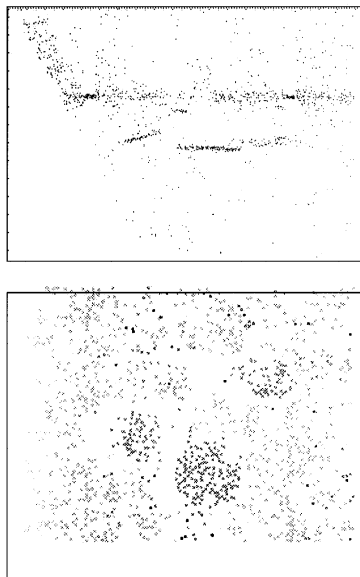


Fig. 8. High sharing (radius 3, coefficient 0.5).

and crossover rates around 50% and 10% on interior scenes. The bottom curve corresponds to a mutation-only evolution.




Fig. 9. Evolution of average fitness values for a 5000 individual population using three different combinations of mutation and crossover rates.

2.3.2.3. *Small populations.* When smaller populations are used (e.g. to reduce computation time) the sharing coefficient must be increased to ensure a fair repartition of the population, as seen in Section 2.2. In addition, higher crossover rates then allow more reliable object detection (see typical example in Fig. 10). A fair all-purpose compromise is around 10–20%, depending on the importance of planar objects in the scene.

Experiments on real data show that the choice of population size depends on processing time constraints and on the level of detail required in the application (size of objects to be detected; see details below).

2.3.3. Results on real images

We tested the algorithm on monochrome stereo 768×576 image pairs of interior scenes (Figs. 11 and 12). Typical results are shown on Fig. 13, using similar genetic parameters (5000 flies, 100 generations, 40% mutation, 20% crossover, sharing radius 2, sharing coefficient 0.3). One can see the two sides of the cabinet, the beginning of the wall on the right and the front half-circle of the stool in spite of its small number of pixels.

3. Detecting mobile obstacles: real time and dynamic flies

We have extended the algorithm to process stereo pair sequences taken by moving cameras.




Fig. 10. Using different crossover rates with a small population: 1000 flies, 100 generations.




Fig. 11. Left image.



Fig. 12. Right image.




Fig. 13. Results and plan view.

3.1. Random method

The random approach consists in keeping the same population evolving through frame changes. Thus, only the images (and therefore the param-

eters used by the fitness function) are modified while the fly population evolves. When motion is slow enough, using the results of the last step speeds up convergence significantly compared to using a totally new random population. This method may be seen as the introduction of a collective memory of space, exploiting the similarity between consecutive scenes. It does not alter much the simplicity of the static method and requires minimal algorithm changes. An extra mutation is applied to flies' positions at each frame change. To improve detection of new objects appearing in the field of view, we introduced an extra mutation operator, *immigration*, which creates new flies randomly in a way similar to the procedure already used to first initialise the population: used with a low probability (1–5%) this favours a convenient exploration of the whole search space.

3.2. Forward method

In order to process faster motion, we introduced an extended chromosome, which contains both position and speed coordinates in the camera's mobile coordinate system, and modified the fitness and genetic operators accordingly. To exploit the redundancy of data and the likely continuity of velocities in most applications, the most handy approach ("*forward approach*") is the following. Each fly's genome is the 6-tuple $(x, y, z, \dot{x}, \dot{y}, \dot{z})$. At the first frame of the sequence, the positions x, y, z are evolved in the usual way, and the velocities $\dot{x}, \dot{y}, \dot{z}$ are initialised randomly. Then for the following frames in the sequence, knowing the fly's genome $(x, y, z, \dot{x}, \dot{y}, \dot{z})$ at instant t , we are able to extrapolate its genome at instant

$t + \Delta t$ as $(x + \dot{x}\Delta t, y + \dot{y}\Delta t, z + \dot{z}\Delta t, \dot{x}, \dot{y}, \dot{z})$ and now evolve the speed components $(\dot{x}, \dot{y}, \dot{z})$. The two possible fitness functions we use are described in Section 3.5. At the end of this step, the flies' coordinates will be updated to $x + \dot{x}\Delta t$, $y + \dot{y}\Delta t$, $z + \dot{z}\Delta t$, so that the population is ready for the next step.

This method may be considered as the introduction of a speed memory in flies. It does not involve important changes in the program, except that evolution is applied to velocities rather than positions.

3.3. Backward method

To better exploit time redundancy, we have introduced the “backward” approach: to evaluate a fly with genome $(x, y, z, \dot{x}, \dot{y}, \dot{z})$, containing position and speed components at instant t , we calculate both the coherence of (x, y, z) with the current image pair and the coherence of $(x - \dot{x}\Delta t, y - \dot{y}\Delta t, z - \dot{z}\Delta t)$ with the previous images. This can be done two different ways, as detailed in Section 3.5.

The approach means that the set of pairs (old position, current position) and therefore all six parameters of the genome $(x, y, z, \dot{x}, \dot{y}, \dot{z})$ are explored more systematically than with the forward method. It implies that the algorithm has been thoroughly rewritten. The results given in Section 3.7 allow a qualitative comparison of these approaches.

3.4. Dynamic fitness functions: fitness consensus or grey level comparison?

The static fitness function described in Section 2.1 only involves the position parameters (x, y, z) of a fly. It is now necessary to define a “dynamic fitness function” to take into account two consecutive stereo pairs and evaluate a 6-component fly genome.

The first method (“fitness consensus”) consists in calculating the static fitness of the fly's old position $(x - \dot{x}\Delta t, y - \dot{y}\Delta t, z - \dot{z}\Delta t)$ on the old images, and calculating the static fitness of its new position (x, y, z) on the new images; then a fly will be given a high dynamic fitness only if both static fitness

values are high. For the results not to be biased by possible different sharing values, only the new position of the fly is used to calculate both fitness values, and the dynamic fitness is defined as the product of the new static fitness with a Gaussian function of the difference between old and new fitness:

dynamic fitness

$$= \text{fitness}_1 \exp \left(-\frac{(\text{fitness}_1 - \text{fitness}_2)^2}{2s^2} \right),$$

where the parameter s allows adjustment of tolerance on fitness dispersion. The flies' fitness values are easily stored from the preceding generation.

The second method consists in comparing the grey levels of all four images⁸ (“grey level comparison”) from two consecutive stereo pairs, and only gives high fitness values to flies whose six coordinates give similar grey levels on all four images. We built the dynamic fitness function as the inverse of the sum of the four grey level differences: old left vs. new left, old right vs. new right, old left vs. old right, new left vs. new right.

The first two terms will give a penalty to flies with wrong speeds, the last ones to flies with wrong positions: we had to introduce and adjust weight coefficients experimentally.

3.5. Genetic operators

In order to use the static version of the algorithm with image sequences, we introduced the “immigration” operator, which helps the algorithm take into account new objects that would appear in the scene but would not be useful in processing static scenes.

In the dynamic version, the genetic operators, immigration, mutation and barycentric crossover, are similar to those used in the static version.

⁸ It is also possible to apply this method using a single moving camera, but its accuracy may then be limited around the focus of expansion if the direction of motion is close to the camera axis.




Fig. 14. Random method: results on images 2, 4, 6 of the sequence.




Fig. 15. Forward approach, evaluation by fitness comparison: results on frames 2 (left) and 6 (right). Second and fourth images show velocities (magnified).

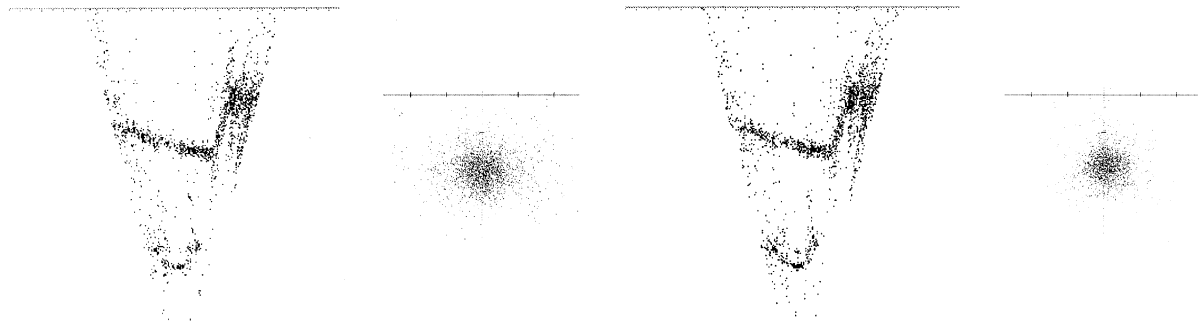


Fig. 16. Grey level comparison on four images: forward approach (left) and backward approach (right). With both methods, improvement on positions is obvious. The backward approach gives better velocities but needs longer calculations.

Immigration again creates totally new individuals located in the intersection of the two cameras' fields of view, and speeds randomly chosen in a given domain. Mutation now affects the norm and direction of the speed vector.

3.6. Initialising the population

The algorithm is initialised using its static version on the first frame. At the second frame it is necessary to initialise the velocities as well, which

is left to the user. In the examples below, the simulated robot's speed was known and used to help initialise the flies' velocities. Accurate a priori knowledge on speed allows reduction of the variance of the initial distribution, privileges local research and helps convergence.

3.7. Results

Results in Figs. 14–16 have been obtained with the same sequence of six image pairs of an interior scene. The order of magnitude V of the robot's speed was known, and used to initialise the flies' depth derivatives \dot{z} using a uniform distribution in the interval $[-3V/2, 3V/2]$. The other components \dot{x}, \dot{y} have been initialised with a Gaussian distribution around 0. Unless otherwise stated, figures show the results obtained at the sixth image, projected on the horizontal plane (x, z) , the camera pointing towards the top of the page. Velocity images (Figs. 15 and 16) show the projections of the velocity vectors on the horizontal plane (\dot{x}, \dot{z}) .

Concerning positions, results are more convincing when the fitness function is based on the comparison of grey levels on four images. The backward method gives a better convergence of speeds, and therefore possibly better results on long image sequences when motion is smooth enough, but uses more computing resources.

4. Application to robot obstacle avoidance

The unusual way the scene is described using the fly algorithm (compared to more classical stereovision algorithm used in mobile robotics) led us to adapt a classical robot navigation method in order to use the results of the fly algorithm as input data, and build a simulator (see Figs. 17 and 18).

To keep the potential speed benefit of the fly algorithm, we chose to implement a fast potential-based method (Koren and Borenstein, 1991; Zelek, 1998). We define a force derived from the sum of an attractive (the target) and a




Fig. 17. The scene seen by the robot.




Fig. 18. The robot facing a wall and a door aperture. White dots represent the flies memorised from the preceding positions.

repulsive (the flies) potential $\vec{F} = \vec{F}_a + \vec{F}_r$, acting as a steering command. The attractive force F_a points towards the target and has a constant amplitude except in the close neighbourhood of the target. The repulsive force F_r takes into account all the flies. Our internal 2-D representation of the robot's environment is inspired by Martin and Moravec's *robot evidence grids* (Martins and Moravec, 1996): each cell concentrates the information contained in the flies it contains and generates its contribution to the repulsive force, proportionally to the sum of the flies' fitness values. The complete simulator includes: a stereo camera pair simulator (using PovRay), the fly algorithm, the potential-based planner described above, and a simple kinetic robot simulator. Blockage situations corresponding to local potential minima are solved using two heuristics (Boumaza and Louchet, 2001) which create a new force attracting the robot out of the potential minimum (random walk and wall following methods) (Figs. 19–22).




Fig. 19. Blocked situation.




Fig. 20. Resolution using secondary targets.




Fig. 21. Another trajectory generated using secondary targets.




Fig. 22. A trajectory generated by wall following.

5. Conclusion

It is often considered that artificial evolution is slow and not suited to real-time applications. However:

- Evolution strategies are generally capable of adaptation, i.e. of coping with modifications of the fitness during the algorithm's execution (Salomon and Eggenberger, 1997). Real time refers to the ability of the system to exploit input data as soon as they become available, and provide the user with suitable updated results at any time.
- With conventional image processing algorithms (Jain et al., 1994; CUMULI project),⁹ a frame must have been fully scanned before the algorithm can calculate and deliver the result. The segmentation-free fly algorithm tolerates data updating during its execution and always uses freshly updated data. This feature can be exploited using CMOS camera technology allowing asynchronous data reading.
- The fly algorithm's results are updated continuously (at the rate of generations), which enables the user (e.g. a robot path planner) to react faster to external events.
- The main time-consuming task in artificial evolution is usually calculating the fitness function. Here, the "Parisian approach" allows one to spread the scene representation over a large number of extremely simple primitives; hence a simple fitness function and a fast algorithm.

Fairly unusual in image processing, the program's structure is largely application independent, as most of the problem-specific knowledge is expressed in the fitness function, enabling easy transfer to other applications. We are currently implementing the method described in this paper into ENSTA's mobile robot project, and are developing an application of the same pattern recognition technique to medical image reconstruction.

Average processing time is from 7 to 25 ms per generation for the random method, on an 800 MHz

⁹ Computational Understanding of Multiple Images, <http://www.inrialpes.fr/CUMULI>.

Pentium III PC, Linux, gcc, and may go up to about 50 ms with the backward method, depending on population size and genetic parameters. Detailed results may be found on the site: <http://www.ensta.fr/~louchet/FlyBenchmark.html>.

References

- Boumaza, A., Louchet, J., 2001. Using real-time evolution in robotics. In: EVOIASP 2001, Artificial Evolution in Image Analysis and Signal Processing, April 2001, Como, Italy.
- Collet, P., Lutton, E., Raynal, F., Schoenauer, M., 1999. Individual GP: an alternative viewpoint for the resolution of complex problems. In: Banzhaf, W., Daida, J., Eiben, A.E., Garzon, M.H., Honovar, V., Jakiela, M., Smith, R.E. (Eds.), Genetic and Evolutionary Computation Conf. GECCO99. Morgan Kaufmann, San Francisco, CA.
- Eberhart, R.C., Kennedy, J.A., 1995. New optimizer using particle swarm theory. In: Proc. Sixth Internat. Symp. on Micro Machine and Human Science, Nagoya. IEEE Service Center, Piscataway, NJ, pp. 39–43.
- Goldberg, D.E., 1989. Genetic Algorithms in Search, Optimization and Machine Learning. Addison-Wesley, Reading, MA.
- Haralick, R.M., 1980. Using perspective transformations in scene analysis. *Comput. Graphics Image Process.* 13, 191–221.
- Holland, J.H., 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Hough, P.V.C., 1962. Method and means of recognizing complex patterns. US Patent No. 3069 654, 18 December.
- Jain, R.C., Kasturi, R., Schunck, B.G., 1994. *Machine Vision*. McGraw-Hill, New York.
- Koren, Y., Borenstein, J., 1991. Vfh+: potential field methods and their inherent limitations for mobile robot navigation. In: Proc. IEEE Conf. on Robotics and Automation, ICRA'91, Sacramento, CA, April 7–12, 1991. pp. 1398–1404.
- Louchet, J., 2000. Stereo analysis using individual evolution strategy. In: Internat. Conf. on Pattern Recognition ICPR2000, Barcelona, Spain, September 2000.
- Louchet, J., 2001. Using an individual evolution strategy for stereovision. *Genetic Programming and Evolvable Machines* 2 (2), 101–109.
- Lutton, E., Martinez, P., 1994. A genetic algorithm for the detection of 2D geometric primitives in images. In: S. Ulman, S. Peleg (gen. co-chairs), 12th ICPR, Jerusalem, Israel, October 9–13, 1994. IEEE Computer Society, Los Alamitos, CA, pp. 526–528.
- Martins, M.C., Moravec, H.P., 1996. Robot evidence grid. Technical Report, The Robotics Institute, Carnegie Mellon University, March.
- Millonas, M., 1994. Swarms, phase transitions and collective intelligence. In: Langton, C.G. (Ed.), *Artificial Life III*. Santa Fe Institute Studies in the Sciences of Complexity, Vol. XVII. Addison-Wesley, Reading, MA.
- Rechenberg, I., 1994. Evolution strategy. In: Zurada, J.M., Marks II, R.J., Robinson, C.J. (Eds.), *Computational Intelligence Imitating Life*. IEEE Press, Piscataway, NJ, pp. 147–159.
- Roth, G., Levine, M.D., 1992. Geometric primitive extraction using a genetic algorithm. In: IEEE Computer Society Conf. on Computer Vision and Pattern Recognition, 1992. IEEE Press, Piscataway, NJ, pp. 640–644.
- Salomon, R., Eggenberger, P., 1997. Adaptation on the evolutionary time scale: a working hypothesis and basic experiments. In: Hao, J.-K., Lutton, E., Ronald, E., Schoenauer, M., Snyers, D. (Eds.), *Artificial Evolution '97*. Lecture Notes in Computer Science, Vol. 1363. Springer, Berlin, pp. 251–262.
- Ser, P.K., Clifford, S., Choy, T., Siu, W.C., 1999. Genetic algorithm for the extraction of nonanalytic objects from multiple dimensional parameter space. *Computer Vision and Image Understanding* 73 (1), 1–13.
- Tang, C.K., Medioni, G., 1998. Integrated surface, curve and junction inference from sparse 3-D data sets. In: Proc. ICCV 98, Bombay. IEEE Computer Society Press, Piscataway, NJ, pp. 818–823.
- Zelek, J.S., 1998. Complete real-time path planning during sensor-based discovery. In: IEEE RSJ Internat. Conf. on Intelligent Robots and Systems.