Optimizing Random Patterns for Invariants-Based Identification

Maurizio Pilu
Personal Systems Laboratory
HP Laboratories Bristol
HPL-1999-45
March, 1999

E-mail: mp@hplb.hpl.hp.com

random point patterns, invariants, optimization

This paper addresses the optimization of pseudo-random planar point patterns for invariant-based identification or indexing. This is a novel problem and is formulated here as the maximization of the spacing of all the invariants when considered as points in a space. The task is of formidable complexity and a stochastic approximation strategy is proposed that yields interesting results.

Internal Accession Date Only

Optimizing Random Patterns for Invariants-Based Identification

Maurizio Pilu
Digital Media Department
Hewlett-Packard Laboratories
Bristol BS12 6QZ, UK
mp@hplb.hpl.hp.com

March 8, 1999

Abstract

This paper addresses the optimization of pseudo-random planar point patterns for invariant-based identification or indexing. This is a novel problem and is formulated here as the maximization of the spacing of all the invariants when considered as points in a space. The task is of formidable complexity and a stochastic approximation strategy is proposed that yields interesting results.

1 Introduction

Pseudo-random point configurations can be used for many vision tasks, both active and passive.

An example is the projection of a pseudo-random light pattern onto an arbitrary object to easily perform stereo matching [15] or depth estimation by triangulation [12]; random patterns have several advantages over time/color coded patterns in that they are much easier to generate and tends to be more robustly detected. Another example is robot navigation, where pseudo random robust landmarks are placed in the environment as beacons or reference points. The use of the arrangement of the random points for labeling is attractive in some situations because it can be performed more robustly than by proliferating the landmark types (such as different shapes, colors, etc). [12]. Random patterns are also applied to motion tracking from laser speckle, (which has recently found some exotic applications such as a digital

microphone [6]) and employed in some products such as the Logitech Marble mouse.

In most cases, pseudo-random patterns are used to identify and label a particular image region through their various, yet unique, local structures. Methods for matching random local structures include syntactic techniques (such as graph matching), cumulative (such as correlation), etc.

Alternatively, or in addition, to these methods, the use of projective invariants provides either an classification measure or an indexing tool which is insensitive to surface position and camera geometry; invariants have proven invaluable to curb the complexity of the search by order of magnitudes [7, 10].

This paper is concerned with a new problem, namely the generation of optimum pseudo random point patterns for invariant-based identification and investigate the problem for the particular case of the 5-point projective invariants [13].

The aim is to find configurations of planar points, each confined within a bound, that maximally space the distance between all the invariants computed with a given neighborhood structure.

Besides application to the aforementioned domains, this problem is *per se* of theoretical interest and its analytic solution would be extremely challenging, if not impossible. In this paper the optimization is achieved through stochastic approximation that iteratively adjusts the position of the points in the pattern while trying to maximize a spacing measure between the invariants.

In the rest of the paper, we first review some previous work in pseudo-random pattern optimization. Next we describe the five point invariants and explain the nature of the problem. Then the task is formalized and details to the solution are given. Finally, some results are presented for two specific cases that validate the method.

2 Previous works in pattern optimization

The use of projected pseudo-random patterns for active stereo or monocular triangulation has been subject to frequent investigations, of which we can cite just a few [15, 12, 14, 4].

Despite their intrinsic practical justifications, interesting theoretical problems were spun off by some authors on how to generate patterns that were optimal in some sense. In [12], for instance, a set of distinct landmarks are arranged as a regular grid but the types are shuffled such as to maximize the Hamming distance between neighboring windows. In [4] a methods is proposed to generate a pseudo-random colored pattern to make labeling local

and easier. Some other works are concerned with optimizing spatial-temporal pseudo-random patterns. For instance [5] uses Hilbert space filling curves to generate sub-optimal patterns layers whose local structures could be more robustly discriminated.

The present work is the first to address the optimization of random point patterns in order to make invariant based identification or indexing easier.

3 Planar projective invariants and their use

In general, an invariant is a property of a geometric structure that remains unchanged under some transformation. In this paper we are concerned with one of the most useful invariants for computer vision, the so called *five-point invariants*.

Given five coplanar points $\mathbf{p}_{i_1}, \mathbf{p}_{i_2}, \mathbf{p}_{i_3}, \mathbf{p}_{i_4}, \mathbf{p}_{i_5}$, with $\mathbf{p}_k = [x_k \ y_k \ 1]^T$, two functionally independent invariants can be derived from the properties of the cross-ratio as the ratios of the following determinants:

$$\iota' = \frac{|\mathbf{m}_{431}||\mathbf{m}_{521}|}{|\mathbf{m}_{421}||\mathbf{m}_{531}|} \qquad \iota'' = \frac{|\mathbf{m}_{421}||\mathbf{m}_{532}|}{|\mathbf{m}_{432}||\mathbf{m}_{521}|}$$

where $\mathbf{m}_{lmn} = [\mathbf{p}_{i_l} \mid \mathbf{p}_{i_m} \mid \mathbf{p}_{i_n}]$ and $|\mathbf{m}|$ is the determinant of $\mathbf{m}$.

The invariants of five coplanar points and their five-lines duals (henceforth called just invariants) have been used in numerous occasions in object recognition and indexing (e.g. [7]), texture analysis [2], robot navigation [11] and more.

4 Origin of the problem

Given a random point pattern and a particular way of computing invariants, one has to wonder whether these invariants are distinctive enough to allow safe discrimination of local unique configurations.

In Figure 1-top-left we show a random point pattern generated by perturbing the nodes of a planar, regular lattice with uniform isotropic noise. Let us use a 4-neighborhood structure (more in Section 5) to compute the invariants. We can map all the invariants ι' and ι'' thus calculated onto a *invariant plane* and have a visual feeling of how spaced they are (Figure 1-top-right).

It can be noticed that some lie very close to others and, given the particular chosen neighborhood structure (with the first point arranged at the center of the other four), the points are distributed in a butterfly-like cloud, with high density near the (1,1) point. However, they never coincide.

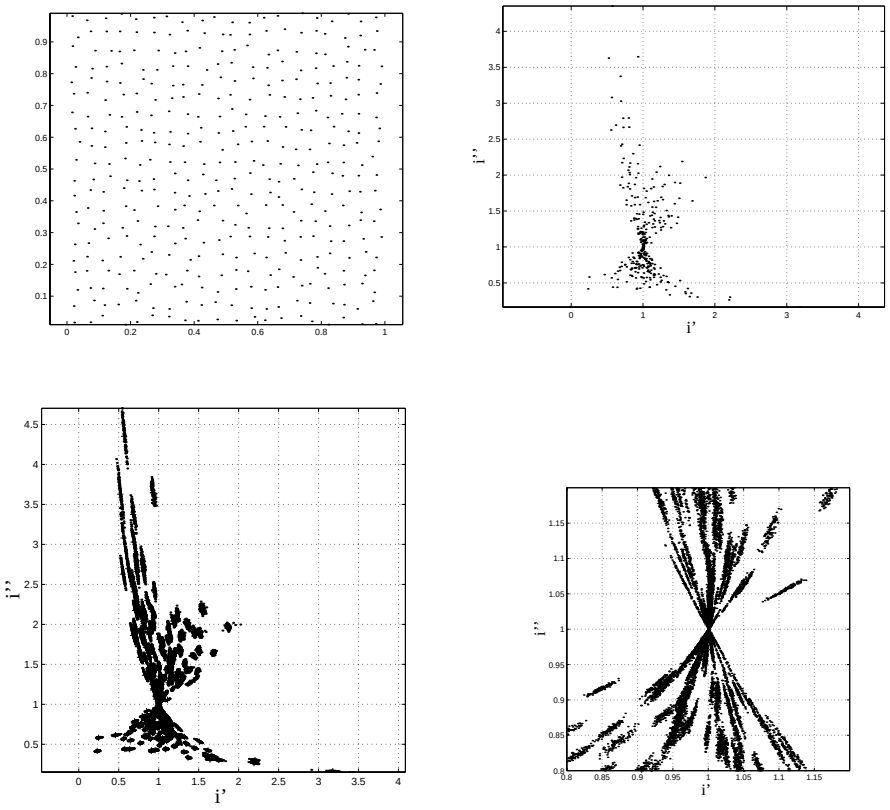

Figure 1: Top-Left: Random point pattern. Top-Right: Invariants. Bottom-Left: Noise propagated to the invariants. Bottom-Right: Blow-up around the (1,1) area.

The picture changes if we consider the points in the configuration with some positional uncertainty. By using a Monte Carlo technique, we can simulate what happens in the invariant plane to the invariants. In all experiments, the noise given to the points was equivalent to $\pm 1/4$ pixel in a 512x512 image. Figure 1-bottom-left shows that when the uncertainty envelopes are taken into account, the points cannot be considered separated any more and the classification could now be troublesome. Note that the noise envelopes are elongated and point to $(1, 1)$, reflecting, again, the particular neighborhood structure utilized for the experiment.

Not much work has been dedicated to the study of the noise properties of projective invariants. Forsyth *et al.* [3] proposed a simple linearized error propagation formula and performed stability testing for their recognition system. The classification problem in presence of noise has been well explored

by Meer *et al.* [10] who proposed several techniques reduce misclassification with uncertain data, and Maybank [9], where trade-offs in the use of the cross-ratio for model based visions were analyzed to avoid *degenerate* cases.

5 Problem definition

Let $P = \{\mathbf{p}_1 ... \mathbf{p}_n\}$ a set of n points in the real plane defined as $\mathbf{p}_k = \{x_k, y_k \in \Re\}$.

We can define a neighborhood structures of each point to compute the five-point invariants $\mathcal{N}$ defined as:

$$\mathcal{N}(\mathbf{p}_k) = \{\mathbf{p}_{i_1(k)=k} \ \cdots \ \mathbf{p}_{i_5(k)}\}$$

Note that for an arbitrary $\mathcal{N}$, in general for $j \neq k$ it is $\{i_1(j)...i_5(j)\} \cap \{i_1(k)...i_5(k)\} \neq \emptyset$, that is 5-tuples could have points in common or, in topology terms, $\mathcal{N}$ is not necessarily a Hausdorff neighborhood structure.

Let us now define the set of m distinct neighborhoods T over the set P defined recursively as

$$T = T \ \cup \ \{T_j = \mathcal{N}(\mathbf{p}_k) : T_j \neg \in T\}, \quad k = 1..n$$

Each neighborhood is used to compute the two five-point invariants ι' and ι'' which we map onto a invariant plane as a point that henceforth will be just called *invariant point* $\mathbf{i}_j = \{\iota'(T_j), \iota''(T_j)\}$. Let also define the set of all the m invariants $I = \{\mathbf{i}_1 \ ... \ \mathbf{i}_m\}$.

We should now define a measure E of I that reflects the "spacing" of the $\mathbf{i}_j$s in the invariant plane.

A straightforward solution would be to use the sum of distances or some form of system energy using a electrical particle analogy. However, simulations showed that these global metrics does not reflect the spacing between single particles in sub-optimal situations[1].

One of the possibilities is to use the sum of nearest-neighbor distances:

$$E(P) = \sum_{k=1}^{m} \| \ \mathbf{i}_k - \mathbf{i}_{c(k)} \ \|^2$$

where $c(k) \in \{1..m\}$ is the index of the closest invariant point to any given $\mathbf{i}_k$. Later we will refer also to the same metric applied to the $q\%$ of the invariant points that are the closest to others and will be referred as E_q.

[1] In particular, clusters of invariant points tend to form and act as a single entity while retaining similar global energy, as it happens in some real particle systems.

Once the neighborhood structure is chosen the measure E depends only on the point configuration P.

The optimization problem can then be expressed as

$$\hat{P} = \arg\max_{P} E(P)$$

subject to

$$\mathbf{p}_k \in B_k, \quad k = 1..N$$

that is we seek a configuration of points that maximizes the spacing in the dual invariant plane while keeping the points themselves within given bounds B_k.

6 The approximation method

In this section we describe the approximation method we have utilized to achieve a (sub-)optimum pattern in the E sense.

This is a multivariate, constrained optimization problem where the maximizer functional is in general a non-analytic function of the variables $\mathbf{p}_1 ... \mathbf{p}_n$ (due to the closest neighbor spacing measure).

After some experimentation with general purpose optimizers (such as simulated annealing) we have tailored a greedy approximation method that combines simultaneous stochastic perturbations of the pattern points and local gradient ascent.

The algorithm is summarized in the following. Starting from the initial point configuration, gradient ascent (Section 6.1) is applied until there is no improvement in E. At this point we simultaneously perturb the point configuration (Section 6.2) and repeat until the situation improves. If the E increases, then we start hill climbing again. If there is no improvement after 100 repetitions, the procedure stops.

In the implementation, a state is accepted only if E_{25} (Section 5) has improved *and* the minimum distance between two invariant points has not diminished. In fact, random perturbations that turn out good in "jammed" regions might move relatively isolated invariants closer to others which would cause a decrease in E and if this precaution was not taken many fruitful transitions would be wrongly rejected.

6.1 Local Gradient Ascent

Once a random perturbation has been accepted, gradient ascent is performed in order to improve the estimate.

The point configuration is updated according to the following rule:

$$\mathbf{p}_k \leftarrow \mathbf{p}_k + w_k \frac{\partial E}{\partial \mathbf{p}_k} \tag{1}$$

or more specifically for each coordinate:

$$x_k \leftarrow x_k + w_k \frac{\partial E}{\partial x_k}$$
$$y_k \leftarrow y_k + w_k \frac{\partial E}{\partial y_k}$$

where the w_k are weights that will be discussed later.

From Equation 5 we have

$$E = \sum_{k=1}^{m} (\iota'_k - \iota'_{c(k)})^2 + \sum_{k=1}^{m} (\iota''_k - \iota''_{c(k)})^2$$

and thus (equivalently for the y_k)

$$\frac{\partial E}{\partial x_k} = \sum_{\forall j : \mathbf{p}_k \in \mathcal{N}(\mathbf{p}_j)} 2\left(\iota'_j - \iota'_{c(j)}\right)\left(\frac{\partial \iota'_j}{\partial x_j} - \frac{\partial \iota'_{c(j)}}{\partial x_j}\right) + \sum_{\forall j : \mathbf{p}_k \in \mathcal{N}(\mathbf{p}_j)} 2\left(\iota''_j - \iota''_{c(j)}\right)\left(\frac{\partial \iota''_j}{\partial x_j} - \frac{\partial \iota''_{c(j)}}{\partial x_j}\right).$$

Note that the above sums are extended to all invariant points calculated with a neighborhood of which $\mathbf{p}_k$ was part.

The invariants ι' and ι'' are rational polynomial in ten variables and their derivatives are exceedingly cumbersome to extract so the math package Mathematica was used for their computation. They are omitted here for reason of space.

Given the non-convexity of these invariant functions, it is important that the w_k are chosen carefully in order to avoid over-relaxation. Classic optimization literature (see also [1]) suggests that w_k be equal to $w/\frac{\partial^2 E}{\partial \mathbf{p}_k^2}$, where $0 < w < 2$ to ensure convergence. In our case the second derivatives have such horrendous form that they are computed numerically using the derivatives from the previous iteration. At the first iteration of a hill climbing phase, all w_k are set to a very small value (10^{-6}).

Note that during this phase, the $c(i)$ are not changed, that is there is no flipping of the closest point to any given $\mathbf{i}_j$. This shrewdness allows a meaningful computation of the second derivatives and ensure continuity in the approximation.

In particular situations an application of Equation 2 could make one or more points fall outside the bounds B_k. If this happens, point movements are clipped at the boundaries.

6.2 Random perturbations

In order not to get away from local maxima, the coordinate of each $\mathbf{p}_k$ in the point plane are perturbed with random Gaussian noise of variance u_k^x and u_k^y.

We found that the success of random perturbations depend on the number n_s of points perturbed, which is normally rather small. The n_s points are randomly chosen such that points that are the closest to others are given higher probability of being selected.

The entity of the perturbation of a chosen point $\mathbf{p}_k$ is determined by how close to other invariant points the invariants computed with the point $\mathbf{p}_k$ are and by how many iterations have been carried out.

The empirical rule used is:

$$u_k^x = u_k^y = \exp\left(-\frac{t}{K}\right)\left(\alpha\frac{d_{min}}{\parallel \mathbf{i}_k - \mathbf{i}_{c(k)} \parallel^2}\right)$$

where t is the number of *successful* iterations, K is a cooling time constant, β controls the amount of initial noise and d_{min} is minimum Euclidean distance between two invariant points in I.

Hence, the noise is made decrease with t (or temperature), very simply simulating the behavior of multi particles systems. The dependence on the distance has been introduced to encourage invariants jammed in tight sub-configurations to jump out with higher probability than points in relatively sparser regions of the invariant plane. Note that care is taken so as to make the perturbation of the points fall inside their bounds B_k.

The "random walks" and the cooling are inspired by those of Simulated Annealing [8] but note that we never accept worse states and we also change the level of perturbation according to the point position and the temperature (t).

This method for generating random walks for this particular problem has been refined over several experiments but it is quite possible that there are, as for most optimization methods, other alternatives that are as effective if not better.

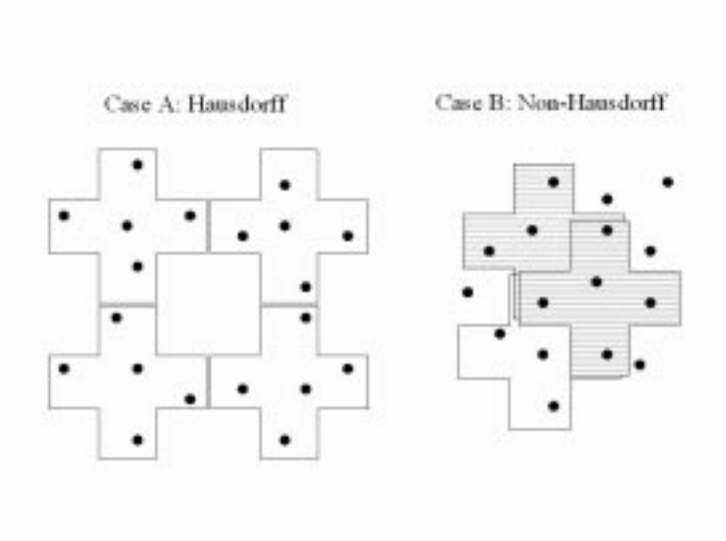

Figure 2: The two neighborhood structures analyzed. See text for details.

7 Cases analyzed

In this section experimental results are presented that show the effects of the optimization procedure.

Experiments have been carried out for two different neighborhood structures, which are shown in Figure 2.

Case A (left) is a Hausdorff structure $\mathcal{N}_a$ where the neighborhoods do not intersect; by definition, this case could be analyzed by having all the 5-tuples expressed with respect to canonical or Bookstein coordinates[2] but they are arranged on a plane for the sake of visual clarity. The second case B (right) is a non-Hausdorff structure $\mathcal{N}_b$ and five invariant points are computed for each point (as opposed to just one for $\mathcal{N}_a$). In both cases, points are free to move in square bounds but cannot trespass them.

It can be easily seen that the case of $\mathcal{N}_b$ is more complicated and interesting because a change in a point affects five invariants, which are now correlated; as we shall see, this makes the optimization problem all the harder.

Other structures could have been investigated and those illustrated here are by no mean intended to be the best ones for recognition tasks. However, once a neighborhood structure is chosen, the present method could be applied

[2]A convenient system of shape coordinates consisting of the coordinates of the points 3, 4 and 5 after the 5-tuple is rescaled and repositioned so that point 1 is fixed at (0,0) and point 2 is fixed at (1,0) in a Cartesian coordinate system.

straight away.

In the experiments that follow we used the following common values for these constants: $\beta = 0.01$, $K = 300$ and $w = 0.5$.

Note that in order to make the situation more realistic, the position of the pattern points were rounded off to simulated a 1/4 pixel sub-resolution of a 512x512 image.

Figure 3 shows an example of a random grid of 39 by 39 points with a $\mathcal{N}_a$ structure. The top row shows the initial and the optimized pattern, the second row the central densest region of the invariant plane before and after the optimization. The graphs show the evolution of the average of the bottom 25% d_{25} inter distances between invariants (left) and the minimum distance d_{min}(right). As it can also be seen visually, the d_{25} is about 2000% bigger and d_{min} has increased by 3 order of magnitudes. The final result was obtained in 1500 true iterations.

Figure 4 shows an example of another random grid of 20 by 20 points, this time with the more complex and interesting $\mathcal{N}_b$ structure. The optimization took much longer (5000 iterations) and predictably so, since local interactions cause good isolated changes to take several iterations to affect the global configuration [1]. Random perturbations of points are also less probable to be good, as they affect several invariants at the same time, and not always with concomitant increase in their inter-distance to the respective nearest neighbors. Having said so, the results are nonetheless good, with d_{25} up 900% and d_{min} up 2000%.

Several runs on the initial data confirmed that the results are always different and hence suboptimal, which could also be suspected by the non-asymptotic look of the graphs. The quality of the results is greatly affected by the choice of β as well as the cooling schedule and the average number of simultaneous random perturbations. In the first case n_s of 4-6 were perfectly acceptable, whereas for the second case, perturbations with more of 3 changes were rarely successful.

8 Conclusion

This paper proposes the problem of optimizing pseudo-random point patterns such that all invariants computed locally are maximally spaced and a stochastic approximation algorithm is described to do so.

The method could find applications wherever artificial patterns are projected or placed to uniquely identify a particular local neighborhood.

Future work could account for the uncertainty envelopes, rather than just distance of the invariants, as it is the uncertainty envelopes we would

like to space apart. Another interesting avenue would be to integrate the error measure based on invariant distance with other discriminant, such as correlation scores, in order to generate patterns that are (sub)optimal in several senses.

Acknowledgments

The author wishes to thank Adele Lorusso for the proof reading of this document..

References

[1] A. Blake and A. Zisserman. Visual reconstruction. In *MIT Press*, 1987.

[2] S. Chang, L.S. Davis, S.M. Dunn, J-O. Eklundh, and A. Rosenfeld. Texture discrimination by projective invariants. *Pattern Recognition Letters*, 5(5):337–342, 1987.

[3] D. Forsyth, J.L. Mundy, and A. Zisserman. Invariant descriptors for 3D object recognition and pose. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 13(10):971–990, 1991.

[4] P.M. Griffin, L.S. Narasimhan, and S.R. Yee. Generation of uniquely encoded light patterns for range data acquisition. *Pattern Recognition*, 25:609–616, 1992.

[5] E. Horn and N. Kiryati. Towards optimal structured light patterns. In *Conference on Recent Advances in 3-D Digital Imaging and Modeling*, pages 28–35, 1997.

[6] A. D. Houghton and P. A. Ivey. Motion sensor using laser speckle. Patent Application 9223915.1. University of Sheffield, UK.

[7] D. Jacobs. Matching 3D models to 2D images. *International Journal of Computer Vision*, 21:123–153, 1997.

[8] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.

[9] S. Maybank. Probabilistic analysis of the application of the cross ratio to model based vision. *International Journal of Computer Vision*, 14:199–210, 1995.

[10] P. Meer, R. Lenz, and S. Ramakrishna. Efficient invariant representations. *International Journal of Computer Vision*, 26(2):137–152, 1998.

[11] R. Mohr and L. Morin. Relative positioning from geometric invariants. In *IEEE Internation Conference on Computer Vision and Pattern Recognition*, pages 139–144, 1991.

[12] R.A. Morano, C. Ozturk, R. Conn, S. Dubin, S. Zietz, and J. Nissanov. Structured light using pseudorandom codes. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 20(3):322–327, March 1998.

[13] Joseph L. Mundy and A. Zisserman. *Geometric Invariance in Computer Vision*. The MIT Press, 1992.

[14] J.H. Nurre and E.L. Hall. Encoded Moire inspection based on a computer solid model. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 14(12):1214–1218, December 1992.

[15] J.P. Siebert and C.W. Urquhart. Active stereo: Texture enhanced reconstruction. *Electronics Letters*, 26(7):971–990, 1990.

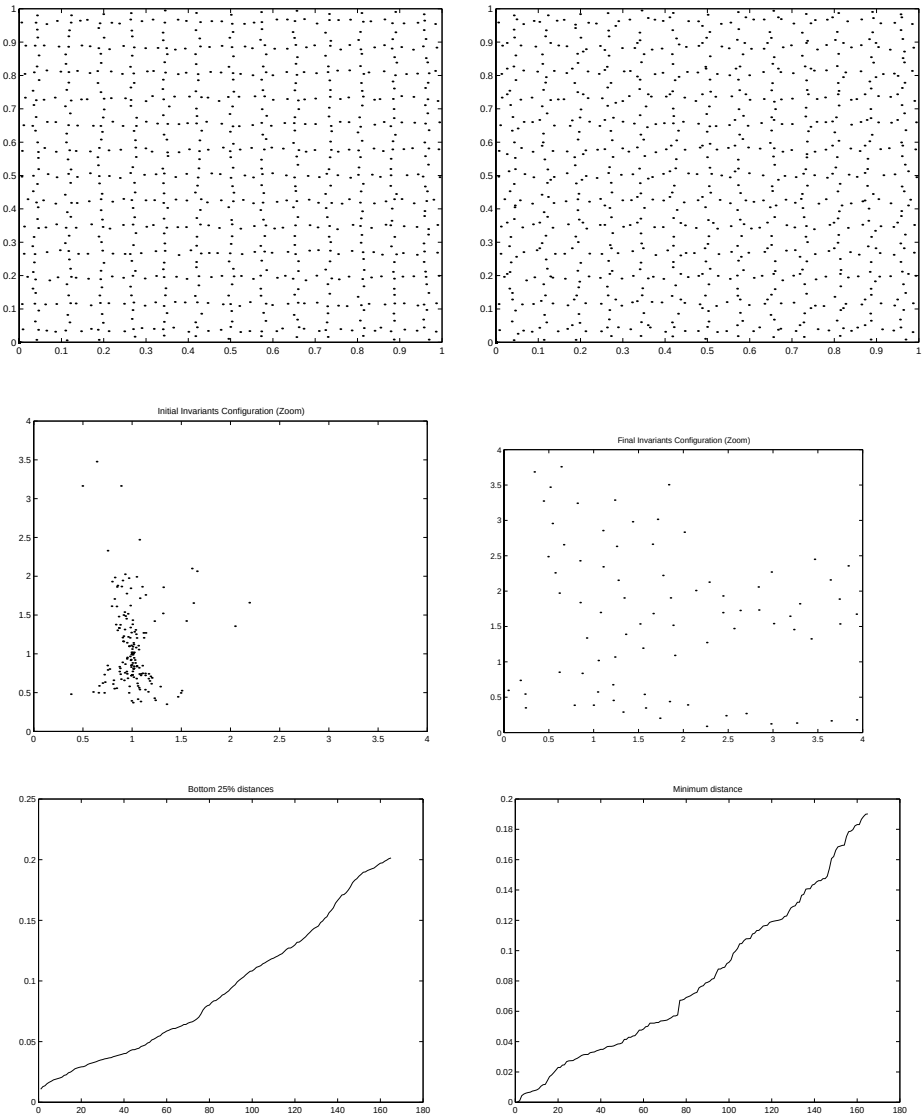

Figure 3: Case A: Top row: initial and optimised patterns; Mid row: initial and final blow up of the densest areas of the invariant plane displayed at the same scale; Bottom row: convergence graphs showing the bottom 25% and the minimum distances between invariants versus the number of succesful iterations.

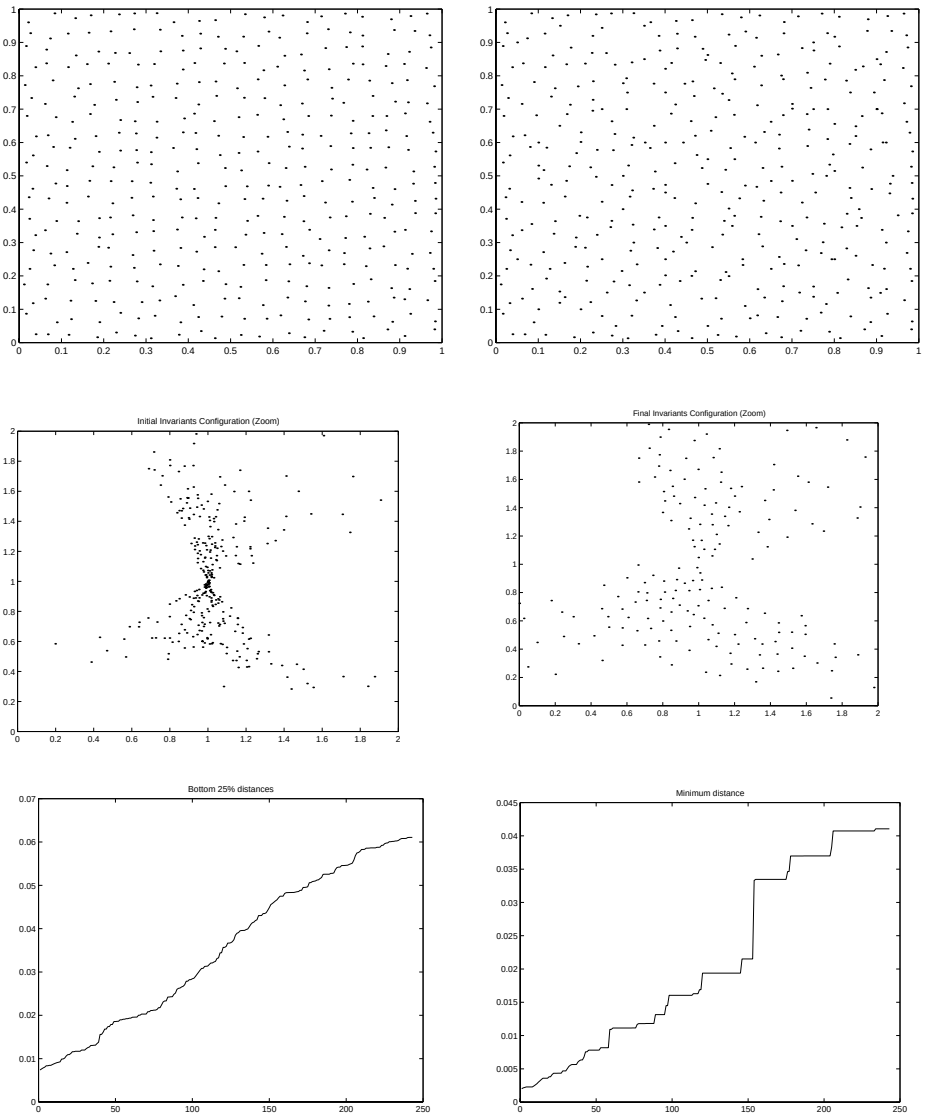

Figure 4: Case B: Top row: initial and final optimised patterns; Mid row: Initial and final blow up of the densest areas of the invariant plane displayed at the same scale; Bottom row: convergence graphs showing the bottom 25% and the minimum of distances between invariants versus the number of succesful iterations.