INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Efficient Dense Matching for Textured Scenes Using Region Growing

Maxime Lhuillier

N ° 3382

Mars 1998

THÈME 3

Rapport de recherche

ISSN 0249-6399

Efficient Dense Matching for Textured Scenes Using Region Growing

Maxime Lhuillier

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet Movi

Rapport de recherche n° 3382 — Mars 1998 — 22 pages

Abstract: We describe a simple and efficient dense matching method based on region growing techniques, which can be applied to a wide range of globally textured images. Our method can deal with non-rigid scenes and large camera motions. First a few highly distinctive features like points or areas are extracted and matched. These initial matches are then used in a correlation-based region growing step which propagates the matches in textured and more ambiguous regions of the images. The implementation of the algorithm is also given and is demonstrated on both synthetic and real image pairs.

Key-words: Dense Matching, Region Growing, Correlation

(Résumé : tsvp)

Unité de recherche INRIA Rhône-Alpes
655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN (France)
Téléphone : 04 76 61 52 00 - International: +33 4 76 61 52 00
Télécopie : 04 76 61 52 52 - International: +33 4 76 61 52 52

Appariement dense d'images texturées par croissance de régions

Résumé : Un algorithme simple et rapide d'appariement dense utilisant des techniques de croissances de régions est présenté. Il s'applique à une large classe d'images texturées: scènes non rigides et grands déplacements. L'algorithme se décompose en deux étapes. On extrait et apparie d'abord des primitives, points ou régions, se distinguant facilement des autres. Ces premiers appariements sont ensuites utilisés pour initialiser une étape de croissance de régions appariées. Cette propagation d'appariements dans les zones texturées permet de construire, par continuité, les zones difficiles à apparier à partir de zones voisines plus facilement appariables. On décrit l'implémentation de l'algorithme, puis des tests sur des images synthétiques et réelles sont proposés.

Mots-clé : Appariement dense, Croissance de régions, Corrélation

1 Introduction

Many algorithms have been proposed for dense matching. One popular approach is based on correlation, however this kind of algorithm is generally limited to relatively small disparity, hence small camera motion. For stereo images [DA89], [Kos93] whose epipolar geometry is known *a priori*, the search space can be reduced to a 1D along epipolar lines. Images rectification is usually used to accelerate the dense matching process, but does not allow zooming in/out of the camera.

Another approach is optical flow [BFB94], which handles non-rigid scenes but is limited to smaller displacements. Differential techniques give accurate estimation of displacements for smooth images, but fail for textured images and at depth discontinuities. Area-based matching techniques are fast, but do not perform well for sub-pixel displacements or dilations.

Occlusions are one of the major sources for wrong matches. Most of the recent stereo and optical flow work consists of incremental improvements to existing methods, to increase speed, accuracy or reliability. Only a few authors directly treat large occlusion stereo [IB94]. Hierarchical methods seem to be necessary to treat large disparity range. However, coarse-to-fine strategies might miss some texture details and fail at depth discontinuities.

Our algorithm uses mainly region growing techniques. Region growing is a classic approach for segmentation [HS85], [Mon87], [AB94], and finding shapes [Bra93]. It is also well adapted to semi-interactive image processing applications [AB94] (Interactive process are unavoidable for real applications, although we will only discuss fully automatic techniques here). In its simplest sense, region growing is the process of merging neighboring points (or collections of points) into larger regions based on homogeneity properties (cf. [HS85]).

Matching by region growing, or in other words dense matching propagation, is implicitly used in regularization (fill the gaps in disparity map from neighboring matches) and optical flow techniques [BFB94] (by iterative algorithms of global cost minimization). An explicit region growing method is introduced in the photogrammetry domain by [OC89] with the "Gotcha" (Gruen-Otto-Chau) ALSC (Adaptative Least Square Correlation) algorithm. It starts with approximate patch matches between two SPOT satellite images and refines them. Their recovered distortion parameters are used to predict approximate matches for new patches in the neighborhood of the first match. Then, these patch matches are refined and so on. Complements for building extraction are discussed in the same domain by [KM96]: a pyramidal algorithm to produce seed matches and extraction of linear elements to remove possible blunders are proposed. Like Iterated Closest Point (ICP) methods for registration of 3D shapes [BM92], dense matching based on region growing combines a best first searching strategy and propagation to neighborhood.

Our main assumption is that the scene is globally textured like many outdoor scenes. Non-rigid scenes and large displacements are allowed. Our computation time is fast without image rectification, therefore camera zooming in/out is allowed.

Our algorithm has two main steps. The first step extracts and matches a sparse set of highly distinctive features (unlike to [KM96]): *seed points* and *seed areas*. Seed points are points of interest and are matched by correlation. If the scene is rigid, a robust technique to match points of interest through the recovery of the unknown epipolar geometry [ZDFL94] could be used. Seed areas complete these matches in the most uniform colored areas. We extract and match them by simultaneously matching and region growing in the most uniform colored areas of the images.

The second step uses these initial matches to seed a dense matching propagation, using a best first matching strategy. This extends the matches to include the textured areas of the image. If the scene is rigid, we can use the epipolar geometry obtained in the first step to constrain the propagation in the second step. Our pixel-to-pixels propagation deals with fine texture details, and stops just at the occlusion borders if they are enough textured. The result is a pixelic matching, but it needs less calculations than patch-to-patches propagation and distortion parameter estimations like [OC89].

The report is organized as follows. The two steps of the dense matching algorithm are respectively described in Sections 2 and 3. Quantitative results on synthetic distortions and qualitative results on real image pairs are presented in Section 4. Section 5 summarizes the work, discusses its advantages and limitations and suggests some future directions.

2 Initial Matching

In this section, we show how to produce a set of initial candidate matches. We first justify the choice of seed points and seed areas. Secondly, we explain how to compare seed areas, and finally describe their matching and their region growing-based extraction.

2.1 Which features to choose ?

Matching points of interest is now a robust process, as demonstrated for example in [ZDFL94]. First, these points are extracted and matched with correlation. Because of noise and nearly repetitive patterns, a relaxation step and next a robust estimate of epipolar geometry seem to be necessary to produce reliable matches. However, a set of candidate point matches obtained by simple correlation could be sufficient to seed concurrent propagations.

Matching edge segments is well adapted for polyedric and low textured scenes. It is difficult to extract and match salient matches of edge segments in our case, because of our assumption of textured scenes.

Finally, it is known that segmentation is an instable process. Nevertheless, we need only to extract some initial seed area matches. A process will be described above which produces some reliable matches of the most uniform colored areas of the images. The simultaneous use of shape and mean color comparisons between some isolated and uniform colored regions in the globally textured images is sufficient to produce candidate seed area matches.

We use then seed areas and seed points. Such seed features could only be matched in weak distortion areas between the two images. Dense matching propagation will extend matching to more difficult and distorted regions to match. If the scene is rigid, the epipolar geometry is recovered while matching seed points.

2.2 How to Compare Two Seed Areas ?

First, seed areas will not usually be distinguishable if their area are too small (unless their colors are rare or their neighborhoods are very different, but this case is not considered here). On the other hand, areas which are too large are subject to significant perspective and segmentation distortions. So we limit the minimum and maximum sizes of our seed areas. In practice, it turns out that the same interval of allowed values is sufficient for many different types of images. The range interval is 100 to 2000 pixels for all our tests.

Two areas A and B are compared very simply by their mean color and their shape (see Figure 1):

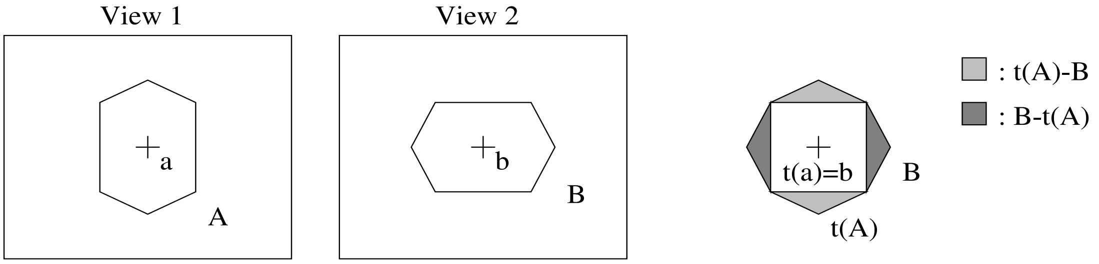

We want to compare areas A of view 1 and area B of view 2. Let a (resp. b) be the centroid of A (resp. B). Our simple shape-based criterion is the quotient of the gray surfaces by A's surface plus B's.

Figure 1: Comparison of two shapes.

$$c(A,B) = \frac{|t(A) - B| + |B - t(A)|}{|A| + |B|}$$

where $|A|$ is the area of A, "-" indicates set difference, and t is the translation from A's centroid to B's. Areas are easily discriminated by their colors and forms, and c allows for a little perspective distortion or initial segmentation error. Other measures could be used, such as the generalized Hausdorff measure [HJ95] or moments [BS97], but the two simple measures above have proved adequate in our experiments.

2.3 Extract Candidate Matches for Seed Areas

The first stage of our algorithm is an alternate sequence of region growing and matching steps for seed areas.

At the beginning, each pixel forms a separate region. During a growing step, for each connected 2×2 pixels block of pixels in the images, their regions are merged if their color difference is less than a threshold (see examples in Figure 2). The threshold is the same for all blocks, but increases between growing steps.

During a matching step, each region of the first image is compared to each region of the second using the above criteria: candidate matches are accepted if both of their areas are within the thresholds and their mean color and shape differences are small.

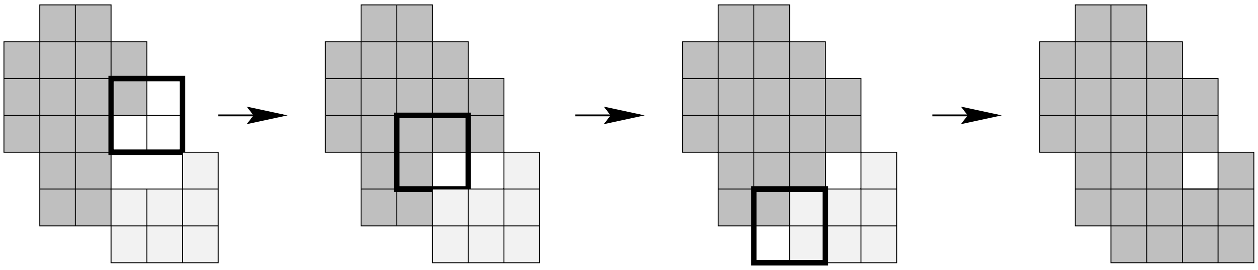

Each little square is a pixel. The big square is the selected 2*2 block of pixel for a merge process. The difference between the max and min colors values of its pixel is bounded by the upper threshold mentioned in the text. The last merge is the fusion of two different region and a individual pixel.

Figure 2: Tree successive pixel or region merges using most uniform colored 2×2 pixels block.

The algorithm is run several times at different color uniformity levels for two reasons. Firstly, region growing is not strictly identical in the two images because of noise and perspective distortion: successive tests are necessary to ensure good matches. Secondly, it allows the same interval of tested thresholds to handle many different types of views.

3 Dense Matching Propagation

We have described the first step of our algorithm in the previous section. A method to obtain seed point matches was cited. An algorithm was also proposed to obtain seed area matches. The second step is now described. We first justify the choice of the dense matching propagation strategy; then we give the principle of the algorithm; finally the algorithm and some implementation details are given. For clarity, the exact link between the first and the second step is explicited in the last part of this section.

3.1 Why Dense Matching Propagation ?

Our goal is dense matching for textured, non differentiable and noisy images. We choose a correlation-based method because it is simple and fast. Correlation is less sensitive to geometric distortion if small windows are used. However, matching with small windows can be ambiguous with nearly periodic textures such as those of outdoor scenes. Thus, a strength constraint is needed for reliable matching. We use the *continuity constraint*: except for some pixels on the object boundaries, the disparity must vary smoothly. Dense matching propagation is a simple and effective way to use this constraint: the propagation moves from less ambiguous matches to more ambiguous ones.

3.2 Principle

A disparity map Map stores the region of correct pixel matches. The algorithm consists of growing this region. Let $Start$ be a set of active pixel matches near its boundaries. At each step we remove the best match m from set $Start$. Match m is the seed for a local propagation: new matches in the neighborhood of m (see Figure 3) are added simultaneously to set $Start$ and map Map. New matches (a, b) are added only if neither pixel a nor b are already matched in Map.

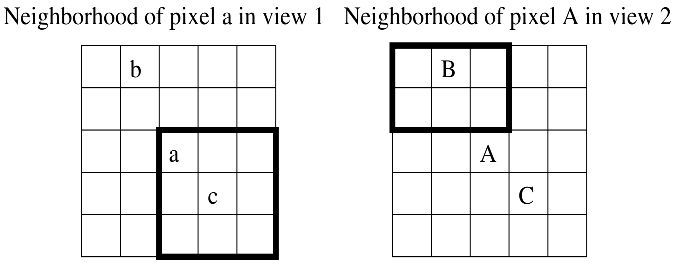

The neighborhood of a match (a,A) is a set of matchs included in the two 25-neighborhood of a and A. Possible correspondants of b (resp C) are in the black frame centered at B (resp c).

Figure 3: Definition of a match neighborhood.

Notice that:

- The set $Start$ is always included in the region of correct matches in Map and initial content of $Start$.
- The unicity constraint is guaranteed in Map by our choice of new matches. Thus, the number of steps and size of set $Start$ are bounded by the sum of the size of $Start$'s initial content and the area of the image.
- Choosing only one match in the neighborhood of m is inadequate. It does not produce a real 2D propagation, because size of set $Start$ can not increase and so can not contain the whole boundaries of region in Map in progress.
- The risk of bad propagation is reduced by the choice of the best match m of set $Start$.

Furthermore, the more textured the image, the lower the risk of bad propagation. We reduce the risk by forbidding local propagation in regions which are too smooth.

Propagation is begun by initializing set *Start* as mentioned in Section 3.4. Propagation is stopped by image borders, too smooth regions and already matched areas. Occlusion contours stop it too, if they separate two different textures. If yes, they are included in borders of a finished propagation in one of the image. Difficult occlusions for others algorithms could so easily be localized with our algorithm.

3.3 Implementation and Algorithm

Disparity map *Map* is bidirectional and injective. We use a heap [Mon87] for the set *Start* to store the potential seeds for local propagations and to select the best at each step. The complexity of propagation is then $Knlog(n)$, where n is the area of the image and K a constant. Notice that it is independent of any disparity bound, and that K is well known small. Let *Local* be a small auxiliary heap of pixel matches.

If a is a pixel, let $\mathcal{N}_x(a)$ be the $x \times x$ window centered at pixel a. Let $s(a)$ be some estimate of the color roughness in $\mathcal{N}_3(a)$ and s_0 be a lower threshold. We use $s()$ to forbid propagation into insufficiently textured areas $\{a, s(a) < s_0\}$.

Let $d(a,b)$ be a measure of the image intensity/color difference between $\mathcal{N}_3(a)$ and $\mathcal{N}_3(b)$, and value d_0 be an upper threshold. The ratio $r(a,b) = \frac{min(s(a),s(b))}{d(a,b)}$ is used as a measure of reliability for pixel match (a,b). Matches with the best (the uppest) reliabilities are first considered.

Let $(r_a, g_a, b_a), 0 \leq r_a, g_a, b_a < 1$ be the color of a pixel a. We use the following definitions:

$$
\begin{aligned}
n(a,b) &= .299|r_a - r_b| + .587|g_a - g_b| + .114|b_a - b_b| \\
d(a,b) &= \frac{1}{9} \sum_{\delta \in \{-1,0,1\} \times \{-1,0,1\}} n(a+\delta, b+\delta) \\
d_0 &= 0.07 \\
s(a) &= max\{n(a,b), b-a \in \{(1,0),(-1,0),(0,1),(0,-1)\}\} \\
s_0 &= 0.04 \\
r(a,b) &= \frac{min\{s(a), s(b)\}}{d(a,b)}
\end{aligned}
$$

The algorithm is

// First, initialize $Start$ as mentioned in the next paragraph.
//Next, propagate:
$Local \leftarrow \emptyset$
while $Start \neq \emptyset$ do
. pull from $Start$ the match (a, b) which maximizes reliability $r(a, b)$
. // Store in $Local$ the potential matches of local propagation:
. for each (c, d) in $\{(c, d), c \in \mathcal{N}_5(a), d \in \mathcal{N}_5(b), (d - b) - (c - a) \in \{-1, 0, 1\} \times \{-1, 0, 1\}\}$ do
. . if $s(c) > s_0$ and $s(d) > s_0$ and $d(c, d) < d_0$ // and possibles others constraints
. . then store match (c, d) in the heap $Local$
. . end if
. end do
. // Store in $Start$ and Map consistent matches of $Local$:
. while $Local \neq \emptyset$ do
. . pull from $Local$ the match (c, d) which maximizes $r(c, d)$
. . if c and d are not already stored in the disparity map Map then
. . store match (c, d) in the disparity map Map
. . store match (c, d) in the heap $Start$
. . end if
. end do
end do

If the scene is rigid, we add the epipolar constraint for match (c, d) in the line
if $s(c) > s_0$ *and* $s(d) > s_0$ *and* $d(c, d) < d_0$.

3.4 Link between the First and Second Step

The two steps were described in the previous sections. The first step produces candidate matches of seed points and seed areas, which are accurate up to a few pixels. The second step needs new candidate pixel matches with pixel accuracy.

If a seed point match is accurate about some pixels, it is a good trick to convert it to a set of concurrent, candidate pixel matches of its neighborhood. Best matches will be first selected, and a single good one is sufficient to provoke an avalanche of correct matches in the second step.

Seed area match (A, B) is converted to concurrent, candidate pixel matches in set $Start$ with the simple process below, where t_{AB} (resp. t_{BA}) is the translation vector which maps A's centroid to B's (resp. B's to A's).

for each pixel a of A's boundaries do
store candidate match $(a, t_{AB}(a))$ in the set $Start$
end do
for each pixel b of B's boundaries do

store candidate match $(t_{BA}(b), b)$ in the set $Start$
end do

These candidate matches of boundary pixels have proved adequate in our experiments to start a matching avalanche effect.

We combine the two steps using two possible strategies.

- All candidate matches start simultaneously concurrent propagations.
- Propagations are done one by one, best initial match first. A homogeneous reliability measure for seed point and seed area matches could be defined by the score of a first and bounded propagation.

Results are presented for the first strategy.

4 Experimental Results

Quantitative results on synthetic distortions and qualitative results on real image pairs are now discussed. It should be stressed that the same parameter values introduced in the various steps of the algorithm were used for all of the tests.

4.1 Visualizing Arbitrary Dense Matching with a Checker-Board

Since we test on non stereo pairs, disparity can not always be interpreted as depth, and it might also be large. Depth maps or displacement fields are not well adapted to display the results. We designed a global way to visualize dense matches for arbitrary images as follows. Pixels of the first image are colored with a gray-black checker-board. For each matched pixel of this image, we color the corresponding pixel of the second image with the same color. This makes it easy to visualize the match of each square and its distortion. A best way for color displays consists to blend a red-blue checked board with the original images.

4.2 Synthetic Distortions

Portions of variable textures from real outdoor images are collected to create image I_0 (see Figure 4). A second image I_1 is obtained from I_0 by translating the previous texture portions. Matchings between I_1 and $f(I_0)$ are now evaluated for some synthetic distortions f. The background is changed to forbid boundaries of textures portions to help propagation.

4.2.1 Measures

We evaluate coverage and accuracy of the dense matching propagation. For each texture and distortion f the following numbers are now evaluated:

- S is the maximum percentage of matching coverage surface defined by the lower threshold value s_0. Remember that propagation is only allowed for pixels a of area $\{a, s_0 < s(a)\}$.
- C is the percentage of matching coverage surface obtained (we have $C < S$).
- $E1$ (resp. $E2$ and $E3$) is the percentage of pixel matches, such that their matching error is less than 1 pixel (resp. 2 and 3 pixels).

The matching error of a pixel match (a, b) is defined by

$$err(a, b) = max(|b - f(a)|, |a - f^{-1}(b)|)$$

where $|.|$ is the Euclidean norm of the plane.

4.2.2 Evaluations

Propagation is initialized in each couple of texture portions by a single seed point match at their centers. Figure 6 (resp. 7) shows S, C, E1, E2, E3 values for each texture and for successive 5, 10, 20, 30 degree rotations (resp. 5, 10, 20, 30 % reductions). Figures 8 and 9 show the resulting distortions.

First, our visual matching checker-board suggests that the majority of the matches are good. One seed match suffices in majority of cases to propagate matching in the whole texture portion. Owing to the small correlation windows used, the most textured portions are matched for large distortions.

The larger distortion is, the lower matching coverage C is. Thus, the percentage of good matches is usually high. More than 90 % of pixel matches have less than 1 pixel error for 5, 10 degree rotations and 5, 10 % reductions, about 90 % of pixel matches have less than 2 pixels error for 20 degree rotation or 20% reduction.

Banana texture is too smooth for s_0 value and is never matched. Grass1 texture contains low and dense details. Then, Grass1 texture is subject to bad propagation (see 30% reduction). However, about 90 % of pixel matches have less than 1 pixel error for 5, 10 degree rotations and 5, 10 % reductions (with matching coverage $45 \leq C \leq 50$). Bough0 texture contains strong details. Then, bad propagations are limited (see 30 degree rotation and 30% reductions). Like some real image pairs, Bough0 texture shows matching periodic holes (see 5, 10 degree rotations and 5, 10 % reductions) dues to image discretization and uniform distortions.

4.2.3 Automatic Seeding

The goal of this paragraph is to show that correlation of points of interest is sufficient to produce good seed point matches. The test pair has very large disparities: images I_1 is matched with a 10 degree rotation of I_0 (see Figure 5). Each texture portion forms an isolated region: a portion is matched only if it contains a good seed point match. The results for matched portions are similar to previous ones.

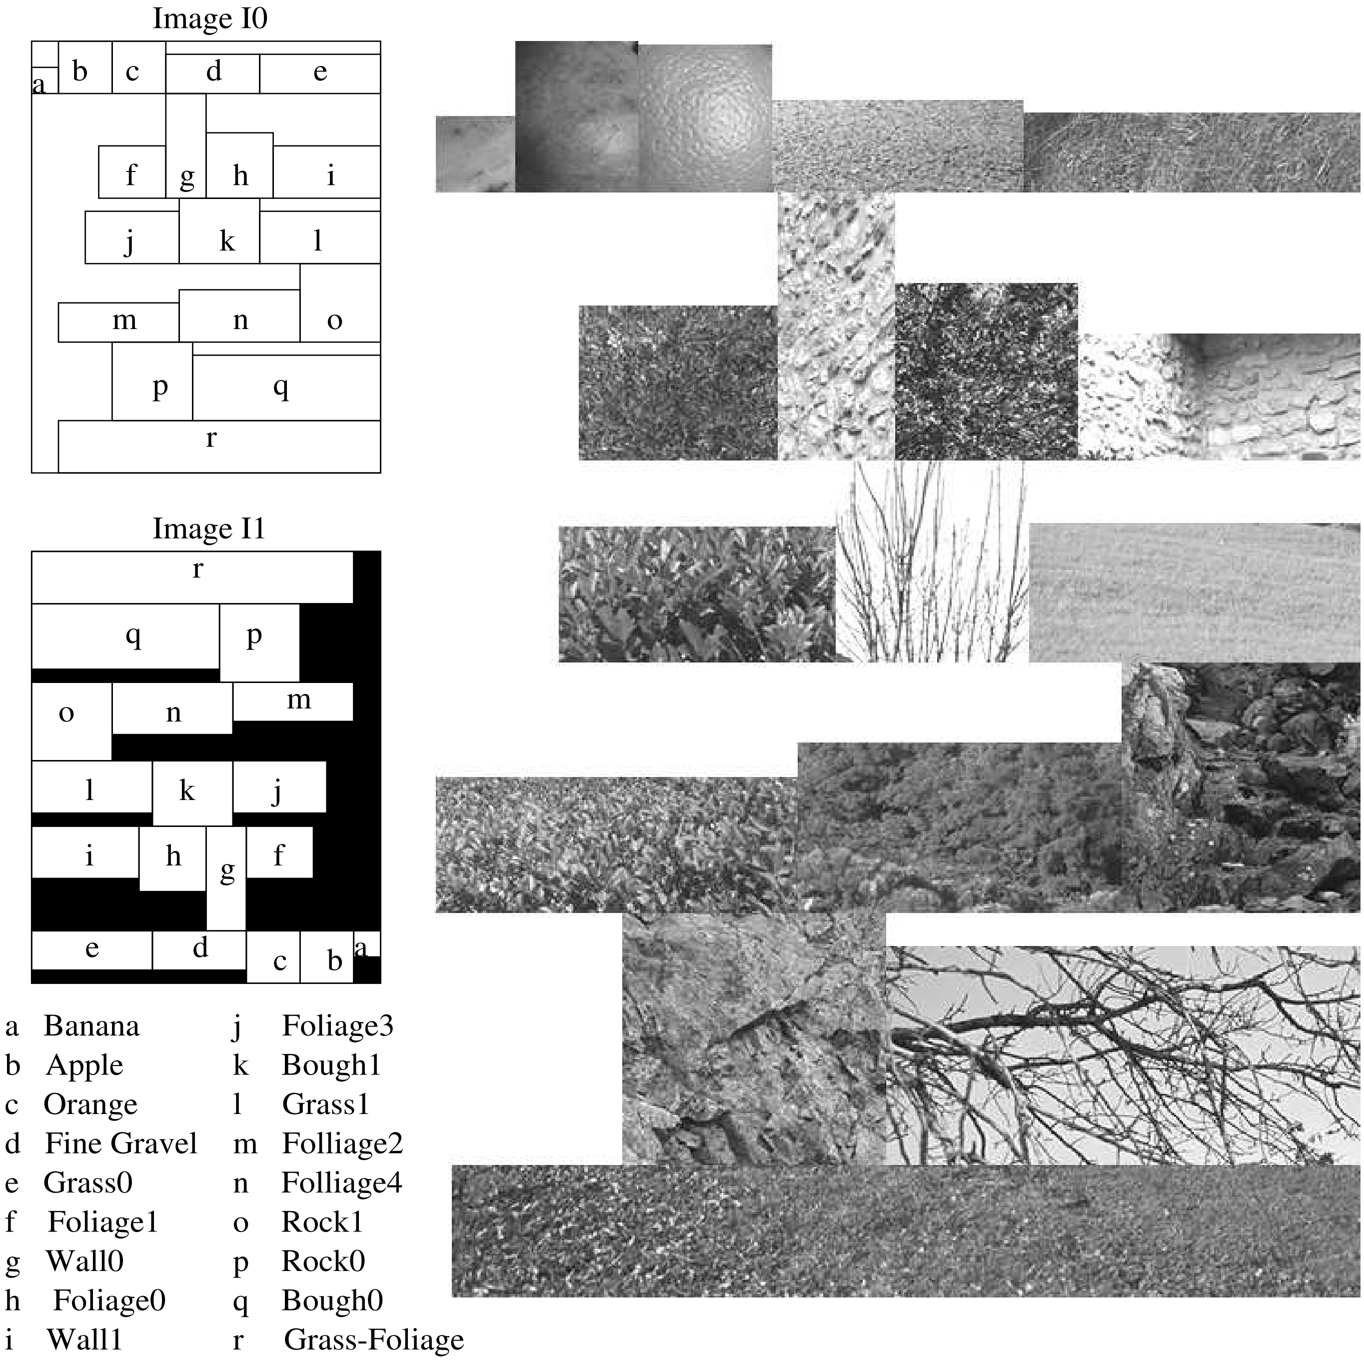

Figure 4: Textures for Tests.

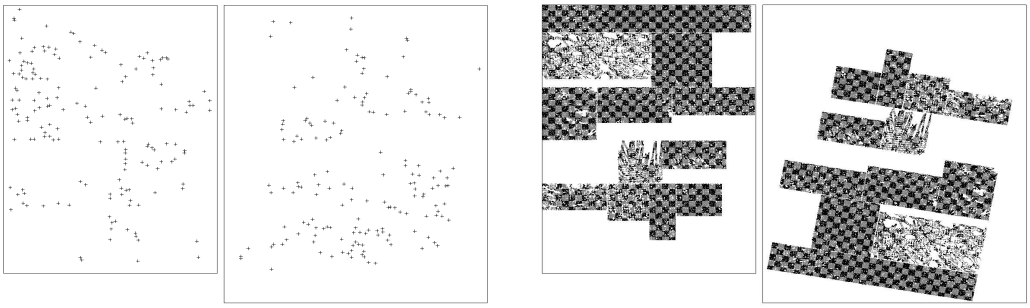

164 seed points matches (usually only one good seed point match per texture portion suffices).

The propagation result : texture portions without good seed point matches are not matched.

Figure 5: Automatic seeding.

4.3 Real image pairs

The first scene is a textured rivulet (see Figure 11). Colored Image dimensions are 512×768. The following table summarizes the process and times on a Ultra SPARC 1, 300Mhz:

Operations	Times	Results
Detect Harris points and correlation	5s	151 Seed point matches
Extract and match most uniform regions	14s	408 Seed area matches
Dense Matching Propagation	19s	238460 Pixel matches

The same detection and correlation as [ZDFL94] are used, but without relaxation or epipolar constraint. All seeds start concurrent propagation. Further, the dense matching propagation does not use epipolar constraint.

We show 6 images in Figure 11: the initial pair, the matched seed points and seed areas, and the dense matching propagation result. It should be noted that the initial pair is blurred on right down corner. Our visual “matching checker-board” suggests that the majority of the matches are good. Some seed matches are bad: the resulting propagations stop quickly thanks to good textures (see the small and isolated propagations near the border of the first image result).

Only a few seed matches are sufficient for highly textured scenes. We manually set 8 seed points matches and show propagation result in Figure 10. Manual correspondences are accurate to about 1-2 pixels. Only one seed match suffices to fill the majority of matching coverage (the background). However, more seed matches are necessary to obtain a similar result in the blurred region than the automatic seeding.

The second scene “street” is well textured too (cf. Figure 12). Results are obtained in the same conditions, and the majority of matches seems good without epipolar constraint. Tests on many others textured image pairs have also been successful.

	S	C	E1	E2	E3	S	C	E1	E2	E3	S	C	E1	E2	E3	S	C	E1	E2	E3
Bough0	83	46	89	97	99	83	40	86	96	98	83	0	0	0	0	83	0	0	1	11
Grass0	98	84	97	99	99	98	82	96	99	99	98	79	85	96	99	98	73	52	69	78
Foliage0	99	70	98	99	99	99	65	97	99	99	99	51	90	97	99	99	35	67	84	92
Orange	76	54	93	98	99	76	53	91	98	99	76	50	76	91	96	76	46	35	51	61
Apple	35	6	91	97	98	35	0	53	69	84	35	0	50	75	87	35	0	71	71	71
Banana	31	0	0	0	0	31	0	0	0	0	31	0	0	0	0	31	0	0	0	0
Grass-Foliage	98	86	97	99	99	98	83	96	99	99	98	79	87	97	99	98	72	63	80	89
Wall0	97	81	96	99	99	97	80	95	99	99	97	73	87	97	99	97	66	67	87	94
Wall1	79	64	94	99	99	79	61	91	98	99	79	58	80	95	98	79	51	54	77	87
Fine Gravel	99	86	97	99	99	99	84	95	97	98	99	81	86	95	98	99	73	54	69	79
Rock0	98	87	97	99	99	98	85	95	99	99	98	80	86	96	99	98	73	62	81	90
Rock1	91	78	95	99	99	91	76	93	99	99	91	72	83	96	99	91	66	58	80	90
Bough1	70	51	83	94	98	70	47	76	90	94	70	35	56	77	87	70	5	31	53	65
Foliage1	98	85	97	99	99	98	83	96	99	99	98	77	88	96	99	98	70	60	78	87
Foliage2	98	80	98	99	99	98	79	95	98	98	98	68	90	98	99	98	56	72	88	94
Foliage3	93	77	96	99	99	93	75	93	99	99	93	70	85	97	99	93	62	67	88	95
Grass1	81	50	94	98	99	81	48	92	98	99	81	44	75	87	93	81	35	16	24	29
Foliage4	94	82	96	99	99	94	81	93	98	99	94	79	84	96	99	94	76	62	82	92

Figure 6: Matching coverages and accuracies for 5,10,20,30 degree rotations

	S	C	E1	E2	E3	S	C	E1	E2	E3	S	C	E1	E2	E3	S	C	E1	E2	E3
Bough0	83	45	88	97	99	83	40	82	94	98	83	17	67	86	93	83	0	47	67	83
Grass0	98	77	96	99	99	98	69	91	98	99	98	55	73	91	98	98	42	35	58	74
Foliage0	99	66	97	99	99	99	57	95	99	99	99	41	87	97	99	99	24	60	79	89
Orange	76	52	93	99	99	76	46	87	97	99	76	40	56	79	90	76	30	7	16	23
Apple	35	8	88	96	98	35	6	85	95	99	35	8	60	84	94	35	7	8	16	27
Banana	31	1	0	0	0	31	0	0	0	0	31	0	0	0	0	31	0	0	0	0
Grass-Foliage	98	79	96	99	99	98	72	94	99	99	98	56	78	94	99	98	42	49	71	82
Wall0	97	77	95	99	99	97	69	93	99	99	97	55	80	96	99	97	41	57	84	94
Wall1	79	61	92	98	99	79	55	87	97	99	79	46	68	89	96	79	35	38	60	74
Fine Gravel	99	80	97	99	99	99	72	94	98	99	99	57	76	92	98	99	43	40	62	77
Rock0	98	81	96	99	99	98	73	92	98	99	98	58	77	95	99	98	44	49	76	89
Rock1	91	72	93	99	99	91	66	90	98	99	91	53	71	92	98	91	40	44	72	87
Bough1	70	50	76	89	95	70	45	72	89	96	70	32	50	69	79	70	10	13	23	31
Foliage1	98	79	96	99	99	98	71	95	99	99	98	55	80	95	99	98	41	49	72	85
Foliage2	98	76	96	99	99	98	68	95	99	99	98	53	85	96	98	98	36	68	90	96
Foliage3	93	73	94	99	99	93	66	91	98	99	93	53	77	95	99	93	39	53	81	92
Grass1	81	49	93	98	99	81	45	86	96	99	81	38	53	77	91	81	27	1	4	7
Foliage4	94	76	95	99	99	94	70	92	98	99	94	56	70	92	98	94	44	37	65	80

Figure 7: Matching coverages and accuracies for 5,10,20,30 % reductions

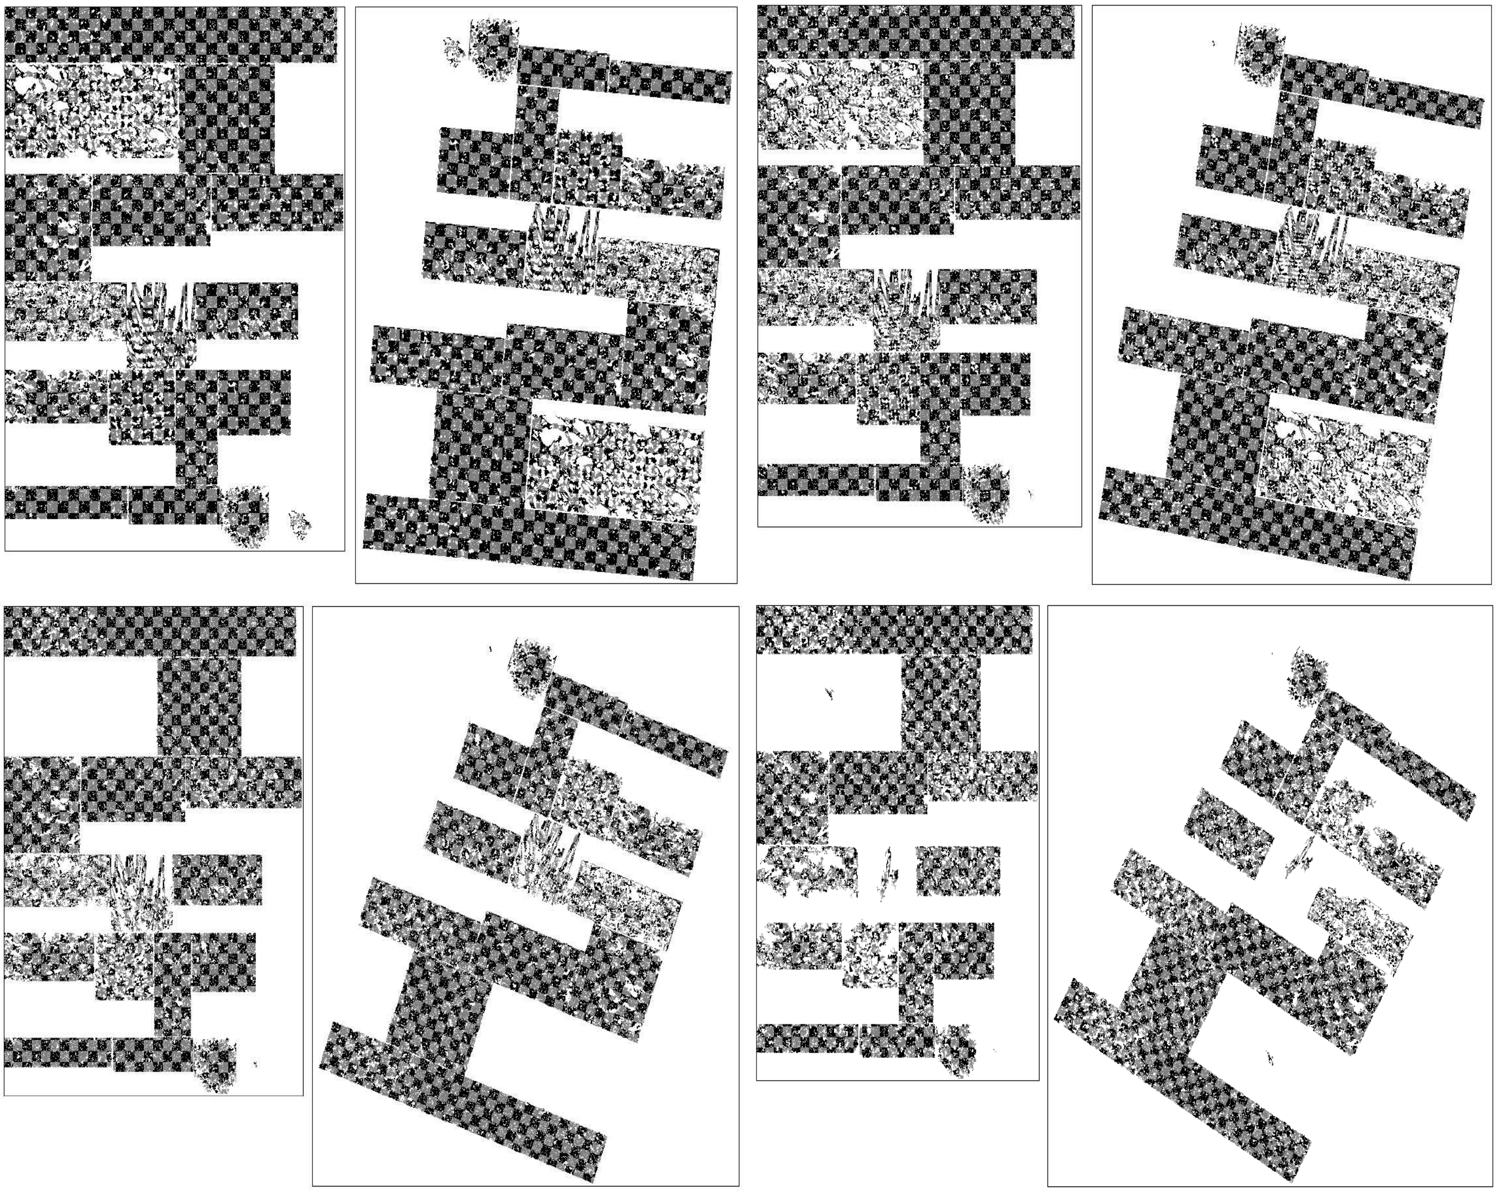

Figure 8: Matching for 5, 10, 20, 30, degree rotations.

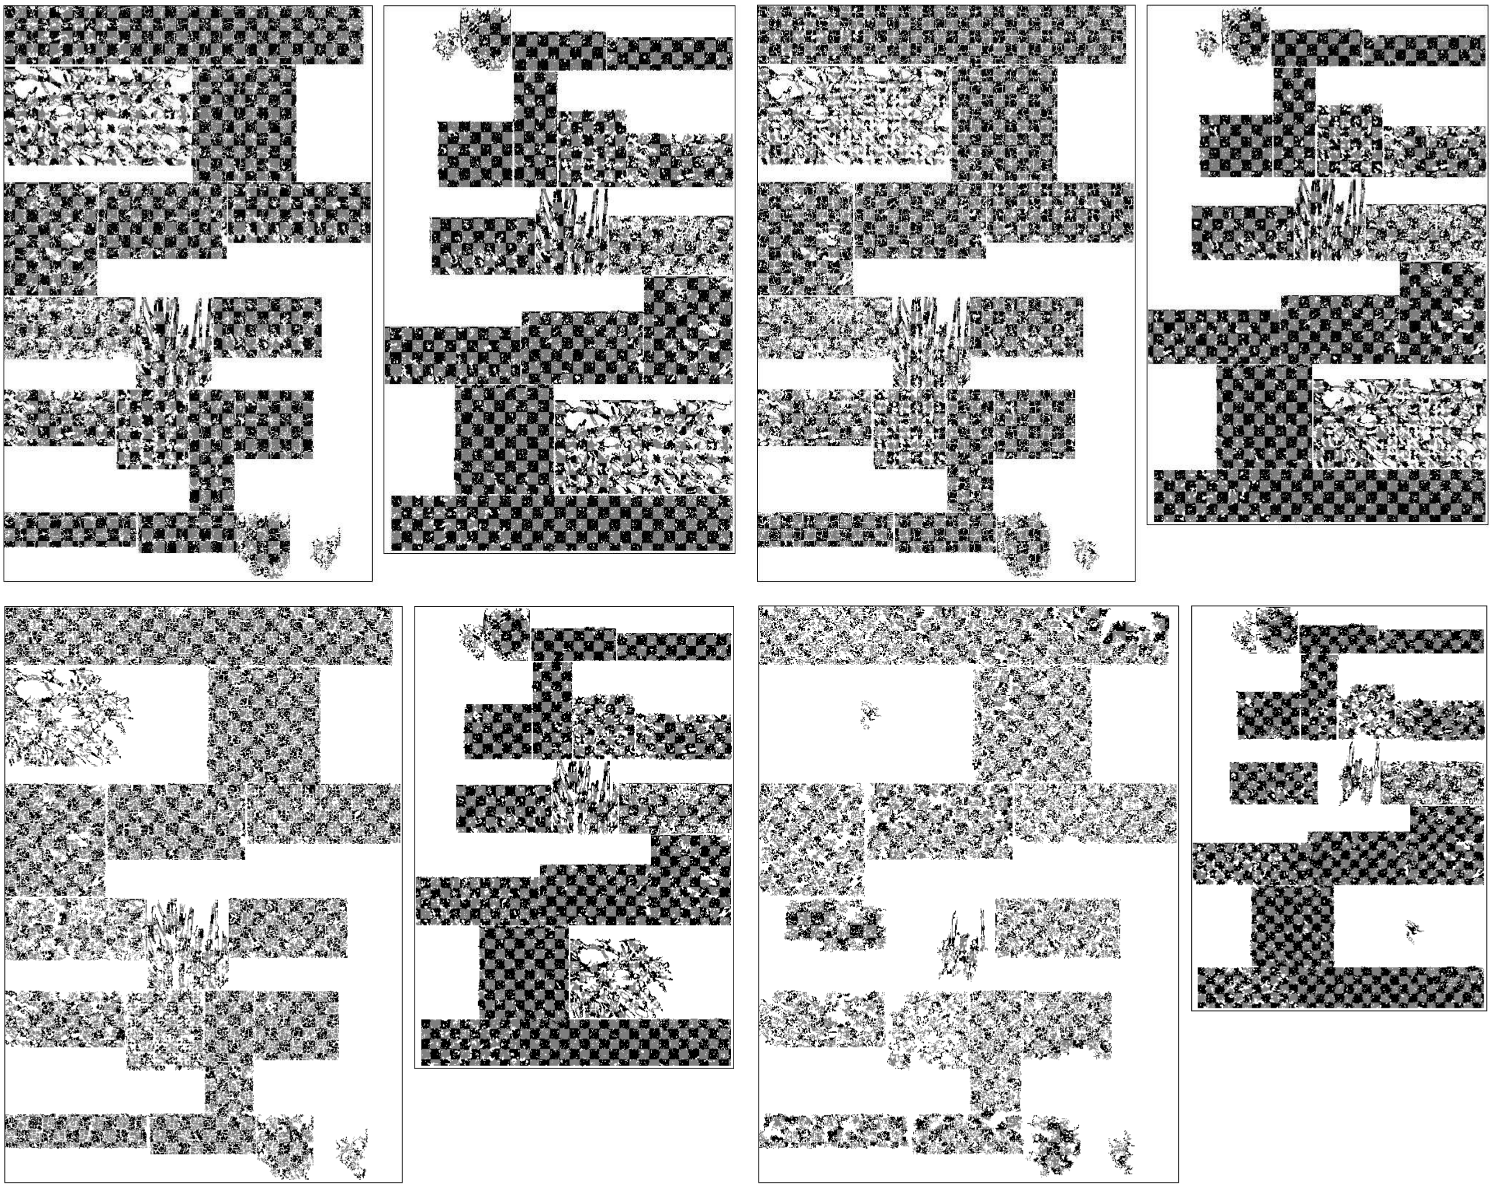

Figure 9: Matching for 5, 10, 20, 30 % reductions.

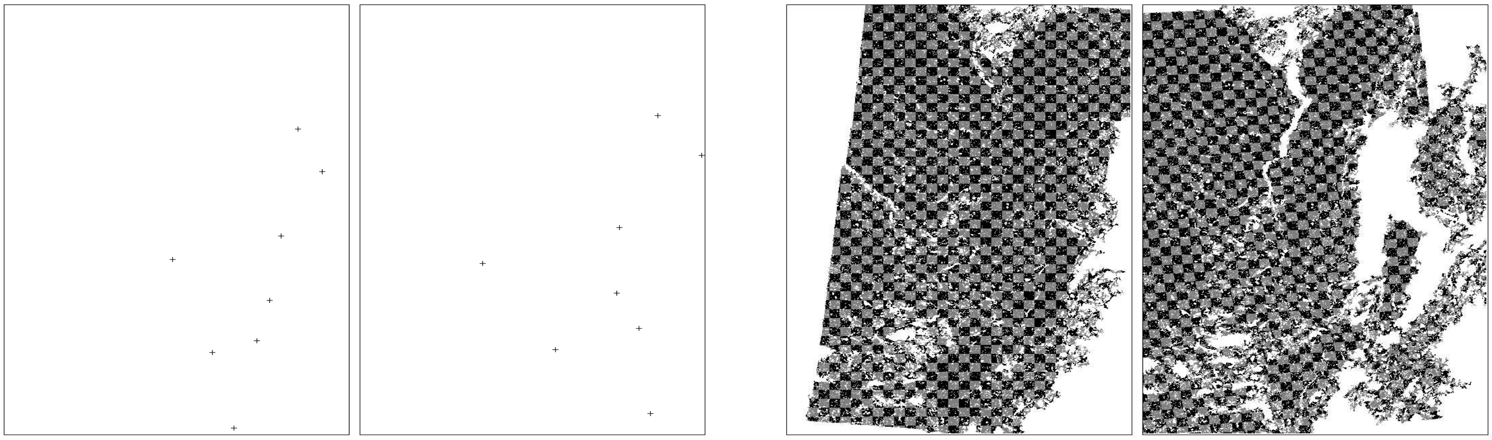

Figure 10: The result for “rivulet” image pair (some manual seed points).

We show finally the result for the less textured scene “bowl” (cf. Figure 13). Two propagations are then tested with lower roughness limits ($s_0 = 0.02$ and $s_0 = 0.01$) to obtain more matches. Since less texture information is available, the recovered epipolar constraint is needed to limit bad propagations. Only contours and well textured areas are matched for $s_0 = 0.02$. Good results are obtained for the whole bowl for $s_0 = 0.01$. Indeed, propagation comes from textured borders of fruits to less textured areas inside. However, the majority of matches are bad for the table.

5 Conclusion

A new method has been proposed for dense matching two textured images. The algorithm has two main steps. First, we extract and match the most uniform colored areas and interest points. Then, a correlation-based match propagation is started from these seeds, to produce a dense matching covering all sufficiently textured areas of the images. We have successfully tested the algorithm on textured images with large displacements and distortions. Rigidity constraint is not indispensable for sufficiently textured scenes. However, our method is not suitable for non-textured images like indoor scenes and manufactured objects: the dense matching propagation is immediately stopped.

If an initially matched seed is false and its neighborhood was little textured, the resulting propagation may not stop quickly. Such cases could be detected by limiting the global deformation allowed between the two images. Other constraints could also be introduced to avoid it too, such as matching compactness.

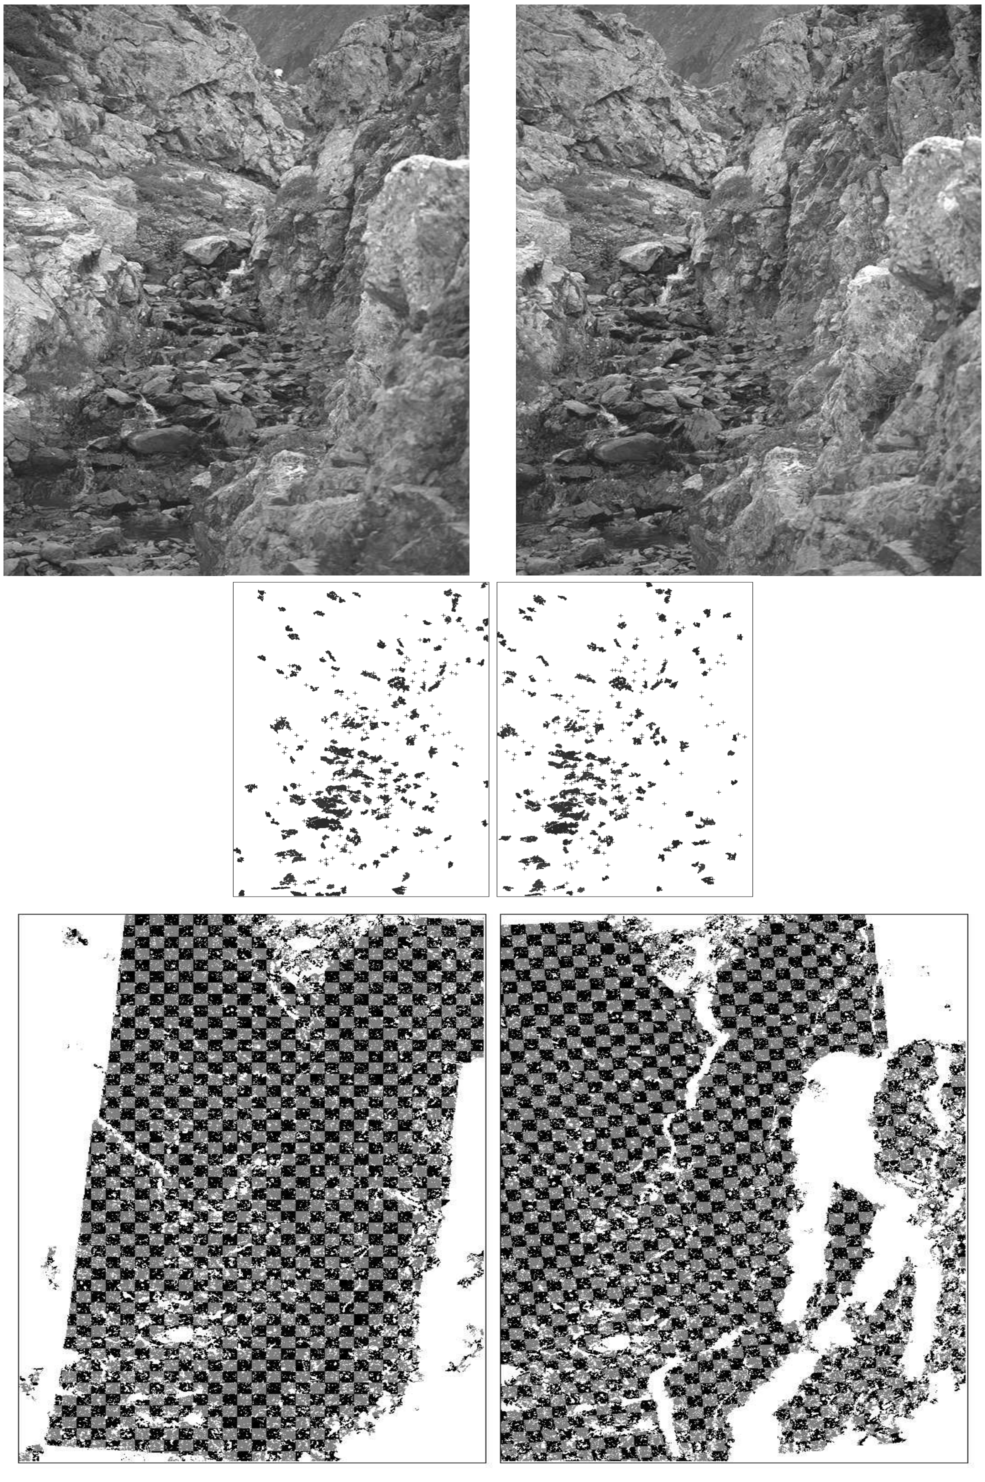

Figure 11: The result for “rivulet” image pair (automatic seeding).

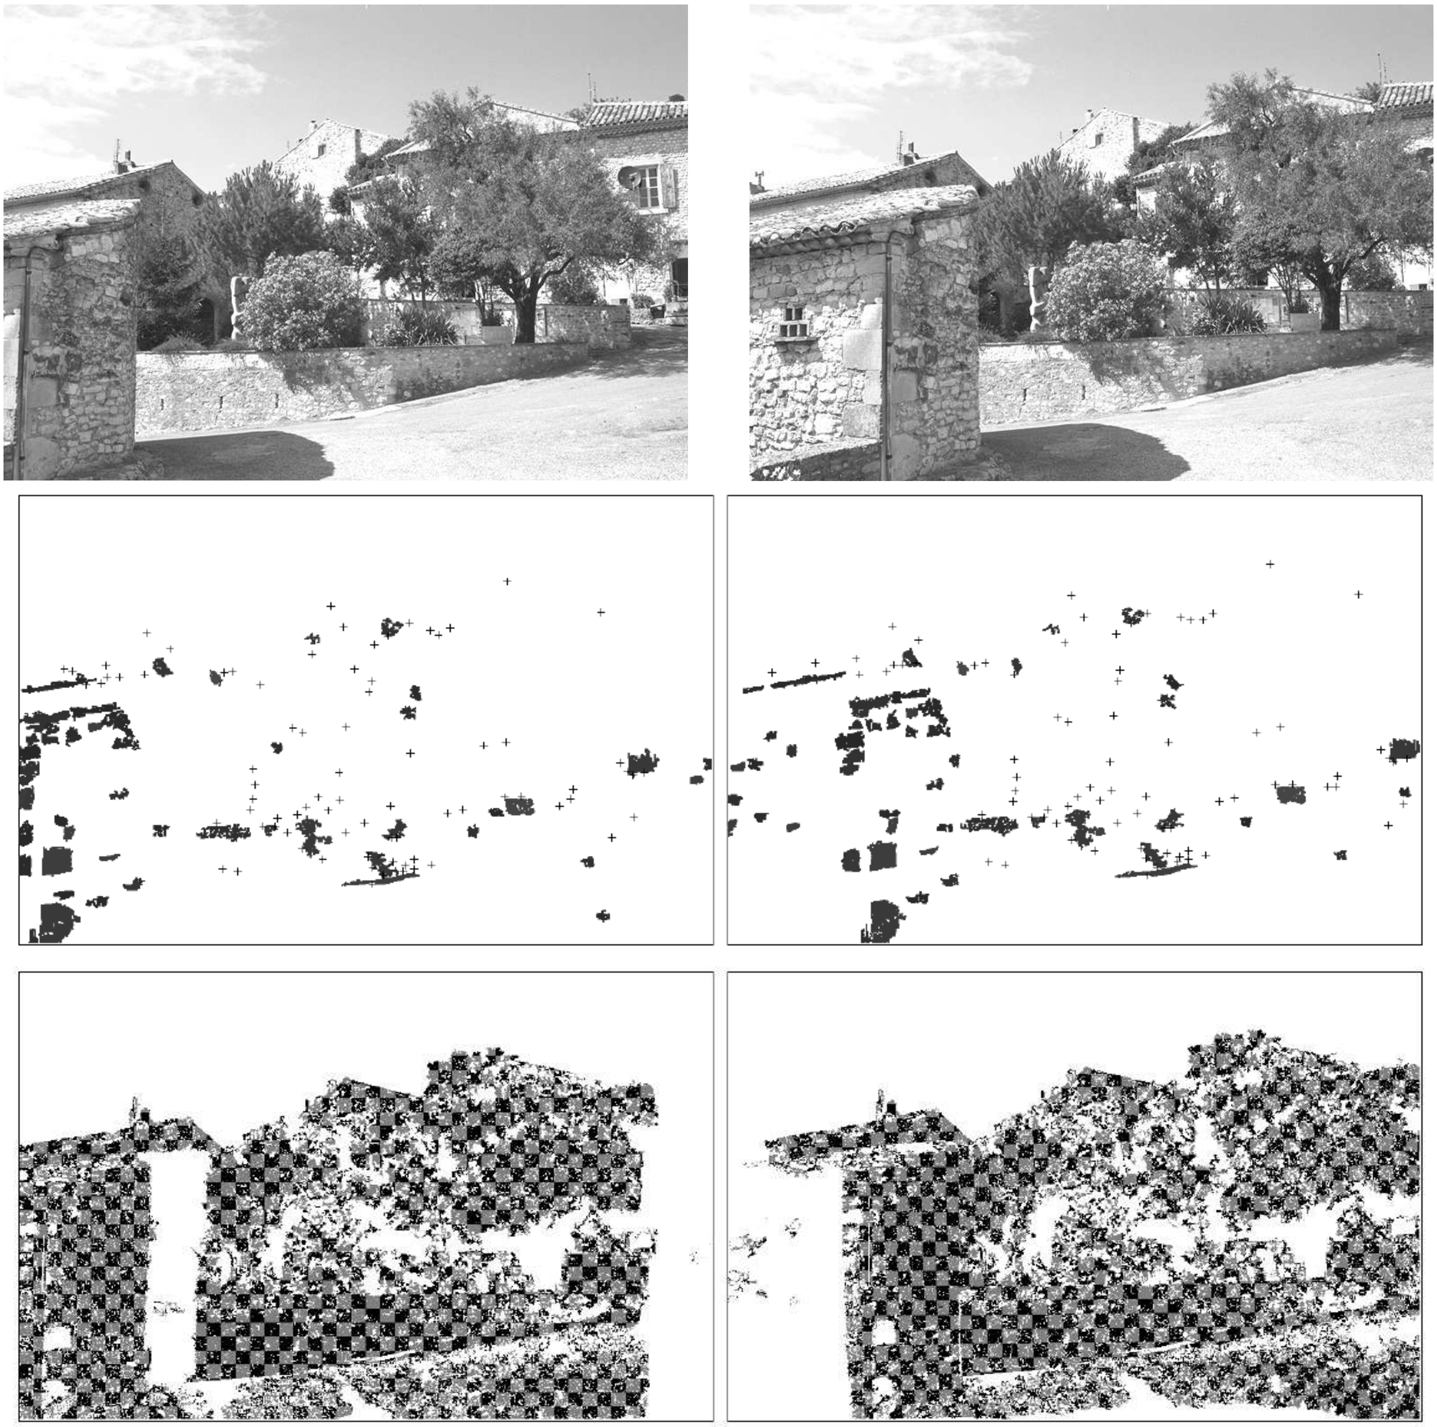

Figure 12: The result for "street" image pair.

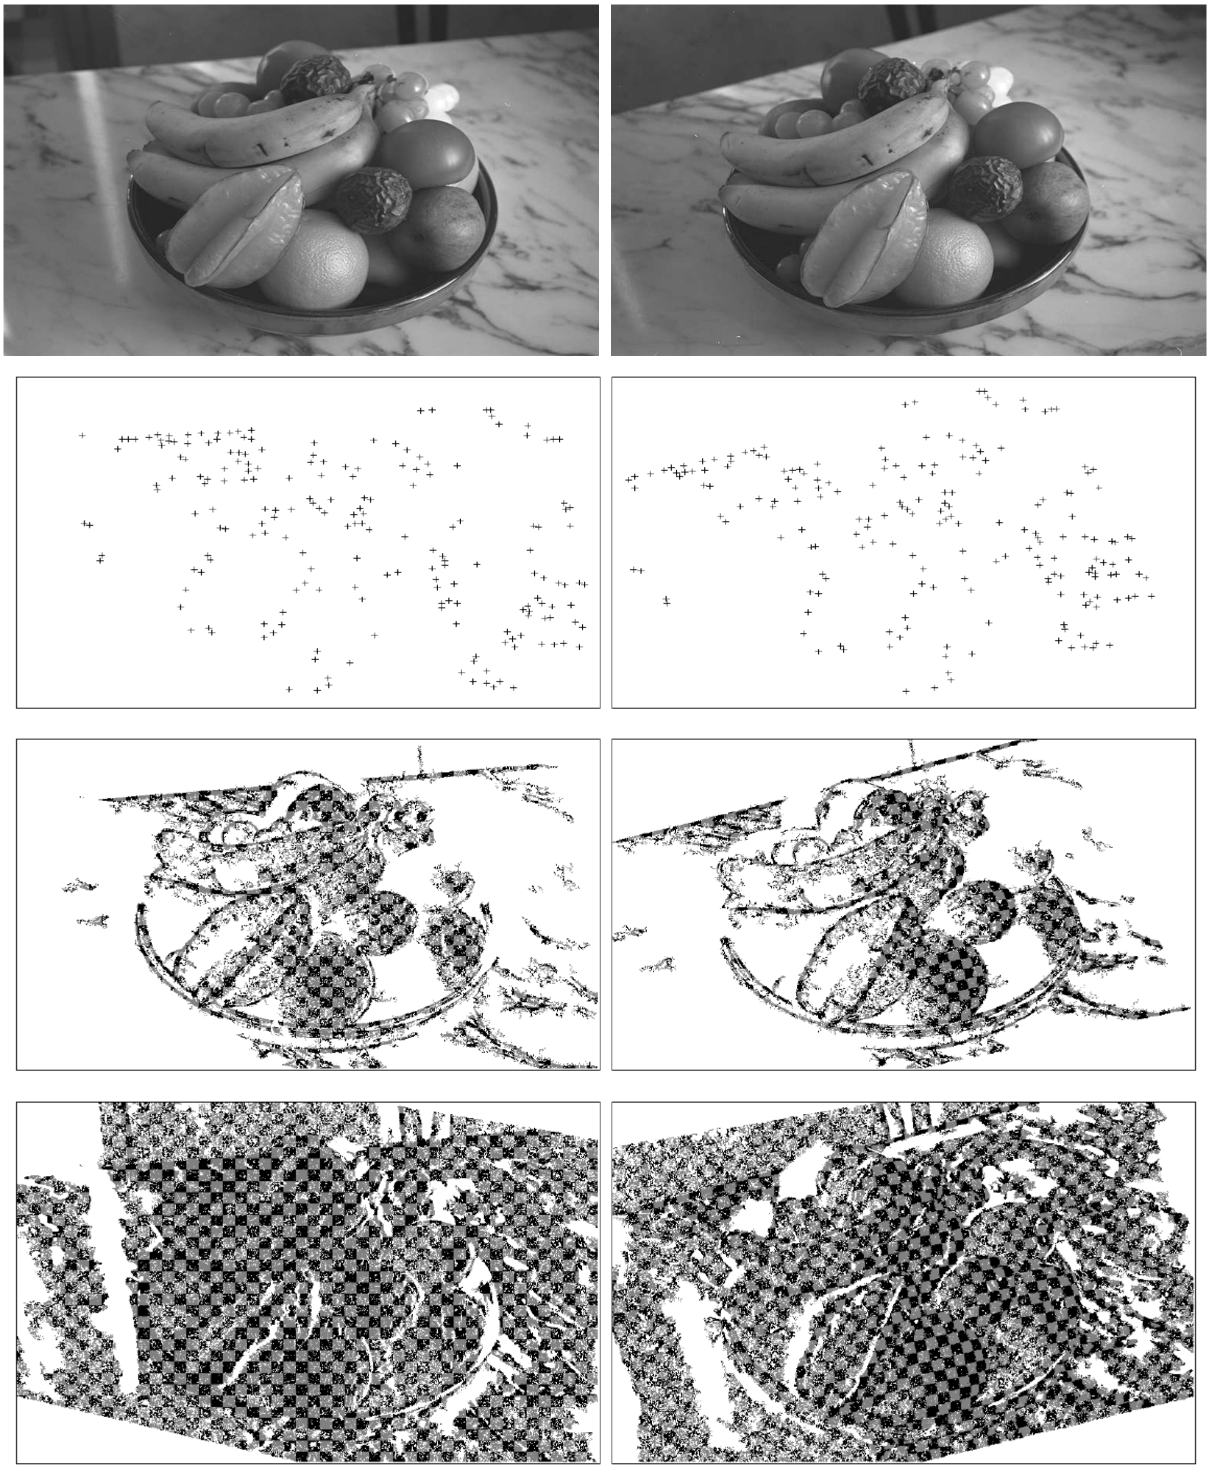

Figure 13: The result for “bowl” image pair.

References

[AB94] R. Adams and L. Bischof. Seeded region growing. IEEE *Transactions on Pattern Analysis and Machine Intelligence*, 16(6):641–647, June 1994.

[BFB94] J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, 1994.

[BM92] P.J. Besl and N.D. McKay. A method for registration of 3D shapes. IEEE *Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[Bra93] M. Brand. A short note on local region growing by pseudophysical simulation. *Proceedings of the Conference on Computer Vision and Pattern Recognition, New York, USA*, pages 782–783, 1993.

[BS97] D. Bhattacharya and S. Sinha. Invariance of stereo images vie the theory of complex moments. *Pattern Recognition*, 30(9):1373–1387, 1997.

[DA89] U.R. Dhond and J.K. Aggarwal. Structure from stereo – a review. IEEE *Transactions on Systems, Man and Cybernetics*, 19(6):1489–1510, November 1989.

[HJ95] D.P. Huttenlocher and E.W. Jaquith. Computing visual correspondence: Incorporating the probability of a false match. In *Proceedings of the 5th International Conference on Computer Vision, Cambridge, Massachusetts, USA*, pages 515–522, 1995.

[HS85] R.M. Haralick and L.G. Shapiro. Image segmentation techniques. *Computer Vision, Graphics, and Image Processing*, 29:100–132, 1985.

[IB94] S.S. Intille and A.F. Bobick. Disparity-space images and large occlusion stereo. In *Proceedings of the 3rd European Conference on Computer Vision, Stockholm, Sweden*, pages 179–186. Springer-Verlag, 1994.

[KM96] T. Kim and J.P. Muller. Automated urban area building extraction from high resolution stereo imagery. *Image and Vision Computing*, 14:115–130, 1996.

[Kos93] A. Koschan. What is new in computational stereo since 1989 : A survey on ereo papers. Technical report, Department of Computer Science, University of Berlin, august 1993.

[Mon87] O. Monga. An optimal region growing algorithm for image segmentation. *International Journal of Pattern Recognition and Artificial Intelligence*, 1(3):351–375, 1987.

[OC89] G.P. Otto and T.K. Chau. A region-growing algorithm for matching of terrain images. *Image and Vision Computing*, 7(2):83–94, 1989.

[ZDFL94] Z. Zhang, R. Deriche, O.D. Faugeras, and Q.T. Luong. A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry. *Artificial Intelligence*, 78(1-2):87–119, 1994. Appeared in October 1995, also INRIA Research Report No.2273, May 1994.

Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
http://www.inria.fr
ISSN 0249-6399