

Attributed Graph Matching for Image-Features Association Using SIFT Descriptors

Gerard Sanromà¹, René Alquézar², and Francesc Serratosa¹

¹ Departament d'Enginyeria Informàtica i Matemàtiques,
Universitat Rovira i Virgili, Tarragona, Spain
`gerard.sanroma, francesc.serratosa@urv.cat`

² Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain
`ralquezar@iri.upc.edu`

Abstract. Image-features matching based on SIFT descriptors is subject to the misplacement of certain matches due to the local nature of the SIFT representations. Some well-known outlier rejectors aim to remove those misplaced matches by imposing geometrical consistency. We present two graph matching approaches (one continuous and one discrete) aimed at the matching of SIFT features in a geometrically consistent way. The two main novelties are that, both local and contextual coherence are imposed during the optimization process and, a model of structural consistency is presented that accounts for the quality rather than the quantity of the surrounding matches. Experimental results show that our methods achieve good results under various types of noise.

Keywords: attributed graph matching, SIFT, image registration, discrete labeling, softassign.

1 Introduction

Image-features matching based on Local Invariant Features Extraction (LIFE) methods has become a topic of increasing interest over the last decade. LIFE methods extract stable representations from a selected set of characteristic regions (features) of the image. These local representations are aimed to be invariant at a certain extent to image deformations such as changes in illumination, position of the camera, ... Mikolajczyk and Schmid [1] identified Lowe's SIFT descriptors [2] as the most stable representations among a number of approaches.

SIFT features are located at the salient points of the scale-space. Each SIFT feature retains the magnitudes and orientations of the image gradient at its neighboring pixels. This information is represented in a 128-length vector.

Despite its efficiency, image-features matching based on local information is still subject to the misplacement of certain matches. A number of approaches have been presented aimed at fixing these misplacements by removing incorrect associations with the use of higher-level information. To cite some examples, RANSAC [3] has been successfully applied to outlier rejection by fitting a geometrical model. More in the topic of the present paper, Aguilar et al. [4] have

recently presented an approach that use a graph transformation to select a subset of geometrically consistent associations.

The main aim of our work is to fix the misplaced matches by relocating them when possible, so as to obtain a higher amount of useful matches than the outlier rejectors. We face this relocation as an attributed graph matching problem where we seek for the set of assignments that best fit the constraints imposed by both the local descriptors (attributes of the nodes) and the geometrical arrangement of the features (structural relations). One of the main contributions of this work is the development of a structural model of consistency that accounts, not only for edge-consistency, but also for the matching quality of the surrounding assignments.

We present two graph matching approaches. The first, presented in section 3, is cast in the continuous assignment framework provided by Softassign [5]. The second, presented in section 4, is a probabilistic model cast in a discrete labeling scheme [6]. We have evaluated the matching precision and recall of both methods under different sources of noise.

In section 2 some preliminary concepts are given. We present in section 5 comparative results with Aguilar et al.'s outlier rejector [4], RANSAC used to fit a fundamental matrix [7], and Luo and Hancock's structural graph matching approach [8]. Finally, in section 6 some conclusions are drawn.

2 Preliminaries

Consider an image I_M showing a certain scene. Consider another image I_D showing the same scene as I_M but with some random variations such as viewpoint change, illumination variation, nonrigid deformations in the objects of the scene, etc ... Consider two sets of SIFT features (keypoints) X, Y from the images I_D and I_M .

Definition 1. According to SIFT matching, a keypoint i from I_D with descriptor (column) vector x_i and position inside the image $(p_1^{(i)}, p_2^{(i)})$ is matched to a keypoint j from I_M iff:

$$\frac{\|x_i - y_j\|}{\|x_i - y_{i,2min}\|} < \rho \quad (1)$$

where $\|x\| = \sqrt{x^\top x}$ is the Euclidean length (L^2 norm), $y_{i,2min} \in Y$ is the descriptor with the second smallest distance from $x_i \in X$, and $0 < \rho \leq 1$ is a ratio defining the tolerance to false positives.

This is, a keypoint i from I_D is matched to the closest (in the descriptor-vector space) keypoint j from I_M if the ratio of their distance to the second smallest distance from i is below a certain value $0 < \rho \leq 1$. If this condition is not met, then keypoint i is leaved unmatched.

Definition 2. We define a graph G_M representing a set of SIFT keypoints from the image I_M as a three tuple $G_M = (V_M, M, Y)$ where $v_\alpha \in V_M$ is a node

associated to a SIFT keypoint with position $(p_1^{(\alpha)}, p_2^{(\alpha)})$, $y_\alpha \in Y$ is the SIFT descriptor associated to node v_α and, M is the adjacency matrix (thus, $M_{\alpha\beta} = 1$ indicates that nodes v_α and v_β are adjacent and $M_{\alpha\beta} = 0$ otherwise).

Consider also the graph $G_D = (V_D, D, X)$ that represent a set of keypoints from I_D .

Definition 3. We define the probability of matching nodes $v_a \in V_D$ with $v_\alpha \in V_M$ with regards to the nodes' attributes with the following quantity

$$P_{a\alpha} = \frac{\frac{1}{\|x_a - y_\alpha\|}}{\sum_{\alpha'} \frac{1}{\|x_a - y_{\alpha'}\|}} = \frac{1}{\|x_a - y_\alpha\| \sum_{\alpha'} \frac{1}{\|x_a - y_{\alpha'}\|}} \tag{2}$$

which is a quantity proportional to the inverse of the distance between their descriptors (normalized to sum up to one).

Definition 4. We define the threshold probability for sending a node $v_a \in V_D$ to null (i.e., leaving it unmatched) as

$$P_{a\emptyset} = \frac{\frac{1}{\rho \|x_a - y_{a,2min}\|}}{\sum_{\alpha'} \frac{1}{\|x_a - y_{\alpha'}\|}} = \frac{1}{\rho \|x_a - y_{a,2min}\| \sum_{\alpha'} \frac{1}{\|x_a - y_{\alpha'}\|}} \tag{3}$$

which places the threshold probability for $v_a \rightarrow \emptyset$ at the distance $\rho \|x_a - y_{a,2min}\|$ which is the maximum distance permitted for an v_α to satisfy $\frac{\|x_a - y_\alpha\|}{\|x_a - y_{a,2min}\|} < \rho$.

Note that the matching probabilities of equations (2) and (3) define the same matching criterion as definition 1.

It is a well-known strategy to state that a match from a node $v_a \in V_D$ to a node $v_\alpha \in V_M$ is more likely to occur as more nodes adjacent to v_a are assigned to nodes adjacent to v_α [5][8].

Definition 5. We define a hit as a node $v_b \in V_D$ adjacent to v_a that is matched to a node $v_\beta \in V_M$ adjacent to v_α .

In sections 3 and 4 we develop measures for gauging the structural consistency of a given match $v_a \rightarrow v_\alpha$. The novelty of the proposed measures lies on the fact that they do not rely on the quantity but on the quality of those hits.

Our approaches to attributed graph matching aim to estimate an assignment function $f : V_D \rightarrow V_M$ that best fits the criteria imposed by both the SIFT attributes (local constraints) and the structural relations of the graphs (contextual constraints). Accordingly, $f(a) = \alpha$ means that node $v_a \in V_D$ is matched to node $v_\alpha \in V_M$, and $f(a) = \emptyset$ means that it is not matched to any node.

Definition 6. We define the assignment variable S such that $s_{a\alpha} \in S$ and $s_{a\alpha} = 1$ if $f(a) = \alpha$ and $s_{a\alpha} = 0$ otherwise. The assignment variable is subject to the constraints $\forall a, \sum_\alpha s_{a\alpha} = \{0, 1\}$ and $\forall \alpha, \sum_a s_{a\alpha} = \{0, 1\}$.

This is, each node $v_a \in V_D$ can be assigned only to one node $v_\alpha \in V_M$.

In the following two sections we present a continuous and a discrete labeling approach to attributed graph matching.

3 A Continuous Labeling Approach

Graduated Assignment (Softassign) [5] is a well-known optimization algorithm that has been widely used to find suboptimal solutions to the graph matching problem. It estimates the assignment variable S that minimizes the following objective function:

$$\mathcal{F}(S) = -\frac{1}{2} \sum_a^{|V_D|} \sum_\alpha^{|V_M|} \sum_b^{|V_D|} \sum_\beta^{|V_M|} s_{a\alpha} s_{b\beta} C_{a\alpha b\beta} \tag{4}$$

where s_{ij} are the components of the assignment variable S (definition 6) and $C_{a\alpha b\beta}$ are the compatibility coefficients for the simultaneous associations $v_a \rightarrow v_\alpha$ and $v_b \rightarrow v_\beta$.

This measure originates from the relaxation labeling processes [9]. Gold and Rangarajan [5] have turned this minimization into an iterative assignment problem where a double stochastic matrix of continuous assignments \tilde{S} is updated at iteration $(n + 1)$ according to the following expression

$$\tilde{S}^{(n+1)} = \arg \max_{\tilde{S}} \sum_a^{|V_D|} \sum_\alpha^{|V_M|} Q_{a\alpha} \tilde{S} \tag{5}$$

where $Q_{a\alpha}$ is a quantity depending on the continuous assignment matrix of the current iteration $\tilde{S}^{(n)}$, and corresponds to the derivative of the objective function

$$Q_{a\alpha} = -\frac{\delta \mathcal{F}}{\delta \tilde{s}_{a\alpha}} = + \sum_b^{|V_D|} \sum_\beta^{|V_M|} \tilde{s}_{b\beta}^{(n)} C_{a\alpha b\beta} \tag{6}$$

where $0 \leq \tilde{s}_{ij} \leq 1$ is the (i, j) component of the continuous assignment matrix \tilde{S} . The assignment problem of equation (5) is solved in a continuous (soft) way using a continuation method controlled by a parameter to gradually push from continuous to discrete solutions (see reference [5] for more details about the algorithm).

We consider that a candidate association $v_a \rightarrow v_\alpha$ with a high probability regarding the local information but with low support from its surrounding matches is likely to be an outlier (i.e., a geometrically inconsistent association). In this case v_a should not be matched to v_α . On the other hand, a candidate association with a not-enough-high local probability (i.e., $P_{a\alpha} < P_{a\emptyset}$) but with high support from the surrounding matches, is likely to be an inlier. In that case v_a should be matched to v_α . We propose the following expression as it reflects this desired behaviour

$$T_{a\alpha} = \frac{P_{a\alpha}}{P_{a\emptyset}} + \frac{P_{a\alpha}}{P_{a\emptyset}} \left[\sum_b^{|V_D|} \sum_\beta^{|V_M|} \left(\frac{P_{b\beta}}{P_{b\emptyset}} D_{ab} M_{\alpha\beta} \tilde{s}_{b\beta} \right) - K_{\emptyset_1} \right] \tag{7}$$

where P_{ij} , $P_{i\emptyset}$ are the probabilities for matching $v_i \rightarrow v_j$, $v_i \rightarrow \emptyset$ regarding the nodes' attributes (equations (2), (3)) and; D and M are the adjacency matrices of G_D and G_M .

This measure is composed by a sum of two parts. The first part contributes with the matching quality regarding the nodes' local information, $\frac{P_{a\alpha}}{P_{a\emptyset}}$. This quotient is > 1 if $v_a \rightarrow v_\alpha$ is more likely than $v_a \rightarrow \emptyset$ in terms of local consistency and, ≤ 1 otherwise. The second part contributes with a quantity proportional to the sum of the quality of the *hits*. The *hits* are by definition the only terms of the double summatory different to zero (i.e., $\{(b, \beta) \mid D_{ab} = 1, M_{\alpha\beta} = 1, \tilde{s}_{b\beta} \neq 0\}$). The constant K_{\emptyset_1} represents the threshold contribution required from the *hits* in order to boost the overall measure $T_{a\alpha}$. This second part can be interpreted in the following way: it is > 0 if $v_a \rightarrow v_\alpha$ is more likely than $v_a \rightarrow \emptyset$ in terms of contextual consistency and, it is ≤ 0 otherwise.

Consider the case of a candidate association with a high local and a low contextual consistency. Despite of the high quantity of the first part of $T_{a\alpha}$, the negative contribution of the second part would smooth the overall measure. In the case of a candidate association with a not-enough-high local and a high contextual consistency, the positive contribution of the second part would boost the overall measure.

We have to express $T_{a\alpha}$ in the same terms of $Q_{a\alpha}$ (equation (6)) so that we can use it under the framework of Softassign. In the following expression we rearrange the terms in order to express our measure in terms of the compatibility coefficients $C_{a\alpha b\beta}$

$$Q_{a\alpha} = \sum_b^{|V_D|} \sum_\beta^{|V_M|} \tilde{s}_{b\beta}^{(n)} \left[\frac{P_{a\alpha}}{P_{a\emptyset}} \left(\frac{P_{b\beta}}{P_{b\emptyset}} D_{ab} M_{\alpha\beta} + \frac{1 - K_{\emptyset_1}}{N} \right) \right] \quad (8)$$

where the expression between brackets corresponds to $C_{a\alpha b\beta}$ of equation (6) and $N = \sum_b \sum_\beta \tilde{s}_{b\beta}$ that is a number approximately equal to the number of nodes of the graphs (due to the double stochastic nature of the assignment variable used in Softassign).

At the end of the algorithm the continuous assignment matrix \tilde{S} is turned into a (discrete) assignment variable S such that all the nodes $v_a \in V_D$ are assigned to some $v_\alpha \in V_M$. Finally, we remove the assignments $s_{a\alpha} = 1$ with coefficients $Q_{a\alpha} < 1$ since they do not satisfy the combined constraints.

4 A Discrete Labeling Approach

The idea of discrete labelling [6] is to visit each node and update f in order to gain the maximum improvement in our matching criterion. The difference with other approaches such as softassign [5] or probabilistic relaxation [9] is that the assignment variable is discretely updated, not allowing for soft assignments.

We want to maximise the joint probability of a graph given the assignment function f . To do so, our iterative algorithm visits all the nodes of the graph at each iteration, and updates f in order to increase this joint probability

$$P(G_D | G_M, f) = \prod_{v_a \in V_D} P(v_a \rightarrow v_{f(a)} | f) \quad (9)$$

The update equation of the assignment function is

$$f(u) = \arg \max_{\{\alpha=1\dots|V_M|\} \cup \{\emptyset\}} P(v_a \rightarrow v_\alpha | f) \tag{10}$$

(we use a cleaning heuristic in order to guarantee that f is an injective function).

We have designed our matching criterion as a product of the following two quantities:

$$P(v_a \rightarrow v_\alpha | f) = P_{a\alpha} R_{a\alpha} \tag{11}$$

where $P_{a\alpha}$ and $R_{a\alpha}$ stand for the matching probability according to the current node attributes and the structural relations, respectively. We use the multiplication to combine both quantities as we find it a natural way to combine probability measures as well as it doesn't need further parameters (as opposed to other operators such as linear weighting).

We use the expression presented in equation (2) (and equation (3)) to gauge the likelihood, regarding the node's attributes, of the putative match $v_a \rightarrow v_\alpha$ (and $v_a \rightarrow \emptyset$, when appropriate).

In the remainder of this section we develop the matching likelihood for the association $v_a \rightarrow v_\alpha$ regarding the structural relations ($R_{a\alpha}$). Our aim is to define a model that takes into account the quality of the surrounding matches.

Luo and Hancock [8] showed how to factorise, using the Bayesian theory, the hard-to-model matching probability given the entire state of the assignments S into easy-to-model unary assignment probability terms:

$$P(v_a \rightarrow v_\alpha | S) = g_a \prod_{v_b \in V_D} \prod_{v_\beta \in V_M} p(v_a \rightarrow v_\alpha | s_{b\beta}) \tag{12}$$

where $g_a = [1/p(v_a)]^{|V_D||V_M|-1}$ is a constant only depending on node v_a .

The model for the unary assignment probabilities presented in [8] used the Bernoulli distribution in order to accomodate *hits* and no *hits* (definition 5) with fixed probabilities $(1 - P_e)$ and P_e (being P_e the probability of error). We present a new model aimed at giving a more fine-grained measure by assessing the *hits* according with their quality, while giving room for possible structural errors in the case of no *hit*. The proposed expression is

$$p(v_a \rightarrow v_\alpha | s_{b\beta}) = P_{b\beta}^{D_{ab} M_{\alpha\beta} s_{b\beta}} [\xi P_{b\emptyset}]^{(1 - D_{ab} M_{\alpha\beta} s_{b\beta})} \tag{13}$$

where D and M are the adjacency matrices of G_D and G_M , respectively; $P_{b\beta}$ is the quality term of the association $v_b \rightarrow v_\beta$ (eq. (2)) and, $[\xi P_{b\emptyset}]$ is the ground-level contribution in the case of *no hit* expressed in reference to $P_{b\emptyset}$ (eq. (3)). The parameter $0 < \xi \leq 1$ regulates the ground-level contribution. When $\xi \rightarrow 0$, there is small room for structural errors and then, the update equation (11) relies mostly on the structural model. On the other hand, when $\xi \rightarrow 1$, the ground-level approaches the quality term and the structural model becomes ambiguous.

In a similar manner that is done in [8], we state equations (12) and (13) in the exponential form, obtaining the following expression

$$R_{a\alpha} = h_a \exp \left[\sum_b \sum_\beta \log \left(\frac{P_{b\beta}}{\xi P_{b\emptyset}} \right) D_{ab} M_{\alpha\beta} s_{b\beta} \right] \quad (14)$$

where $h_a = \exp[\sum_b \sum_\beta \log(\xi P_{b\emptyset})] g_a$ is a constant that does not depend on either the graph structure or the state of the correspondences.

Finally, we define the threshold probability for sending a node $v_a \in V_D$ to *null* according to the structural relations as

$$R_{a\emptyset} = h_a \exp \left[K_{\emptyset_2} \log \left(\frac{1}{\xi} \right) \right] = h_a \exp [-K_{\emptyset_2} \log(\xi)] \quad (15)$$

where $K_{\emptyset_2} \geq 0$, $K_{\emptyset_2} \in \Re$ is a parameter defining the minimum number of *hits* with quality term $P_{b\beta} \geq P_{b\emptyset}$ required for the match $v_a \rightarrow v_\alpha$ to be more structurally likely than $v_a \rightarrow \emptyset$.

The algorithm operates updating the assignment function f as stated in equation (10) with the new combined constraints provided in equation (11). This is, at each iteration, each node $v_a \in V_D$ is assigned to the node $v_\alpha \in V_M$ with the highest probability. If the target with the highest probability is \emptyset , then v_a is leaved unmatched.

5 Experiments

We have compared both the continuous and the discrete attributed graph matching approaches of the present work (C-AGM and D-AGM) to the following approaches: Graph Transformation Matching (GTM) [4], RANSAC used to fit a fundamental matrix [7] and, Structural Graph Matching with the EM Algorithm (SGM-EM) [8]. We have evaluated the matching *Precision* and *Recall* scores of each method under the following types of perturbations: *image distortions*, *geometrical noise* and *clutter* (point contamination). We have used the *F-measure* to plot the results. F-measure is defined as the weighted harmonic mean of Precision (P) and Recall (R) and its expression is $F = (2 \times P \times R) / (P + R)$.

The graphs used in our methods (C-AGM and D-AGM) have been generated as described in definition 2. Graph structures for all the methods using graphs (i.e., C-AGM, D-AGM, GTM and SGM-EM) have been generated using a K-nearest-neighbours approach with $K = 4$ (i.e., edges are placed joining a keypoint with its K nearest neighbours in space). All the methods have been initialized with the configuration of matches returned by a classical SIFT matching using a ratio $\rho = 1$ (the best value for the outlier rejectors). As the C-AGM method permits the use of continuous assignments we have initialized them with the probabilities due to local information (i.e., $s_{a\alpha} = P_{a\alpha}$). The keypoint-sets size used in the experiments has been $N = 20$. Our methods (C-AGM and D-AGM) have done 20 iterations, and we have used $\xi = 0.5$ (D-AGM) and $\rho = 1$ (C-AGM and D-AGM). We have empirically set $K_{\emptyset_1} = 0.6$ and $K_{\emptyset_2} = 2.3$ in the

clutter experiments, and $K_{\theta_1} = K_{\theta_2} = 0$ in the others. The tolerance threshold for RANSAC has been set to 0.01, and the number of iterations to 1000 (as suggested in [7]). The probability of error P_e for the SGM-EM method has been set to 0.0003, and the number of iterations to 100.

For each experiment we have arbitrarily chosen a grayscale image I_0 from the Camera Movements and Deformable Objects' databases used in [4].

In the image distortion experiments, we generate I_1 by simultaneously applying the following types of perturbations to I_0 : image resizing, to simulate changes in the distance from the objects in the image; image rotation, to simulate changes in viewpoint; image intensity adjustment, to simulate illumination changes and; gaussian white noise addition to pixel intensity values, to simulate deterioration in the viewing conditions.

We extract the SIFT keypoints from images I_1 and I_0 , obtaining coordinate vector-sets \mathbf{P} and \mathbf{Q} , and SIFT descriptor-sets X and Y , respectively. We define $\tilde{\mathbf{P}}$ as the result of the mapping from points in \mathbf{P} back to the reference of I_0 . We compute $\tilde{\mathbf{P}}$ by applying to \mathbf{P} the inverse resizing and rotation from the perturbation. We set the ground truth assignments on the basis of the proximity between the points in \mathbf{Q} and $\tilde{\mathbf{P}}$. Then, for a given $\mathbf{q}_i \in \mathbf{Q}$, we select as its ground truth assignment the most salient $\tilde{\mathbf{p}}_j \in \tilde{\mathbf{P}}$ among the ones falling inside a certain radius r from \mathbf{q}_i . Saliency is decided according to the gradient magnitude of the SIFT features [2]. The proximity radius has been set to $r = 0.03 \times l$, where l is the diagonal-length of the image. The keypoints that are not involved in any ground truth assignment are discarded. So, at the end of this step we end up with keypoint-sets $\mathbf{Q}' = (\mathbf{q}'_1, \dots, \mathbf{q}'_N)$ and $\mathbf{P}' = (\mathbf{p}'_1, \dots, \mathbf{p}'_N)$, and a bijective mapping $f_{gtr} : \mathbf{P}' \rightarrow \mathbf{Q}'$ of ground truth assignments.

Once the N ground truth assignments have been established, we implement the clutter by adding a certain amount of the remaining points in both \mathbf{P} and \mathbf{Q} to \mathbf{P}' and \mathbf{Q}' . Clutter points are carefully selected not to fall inside the radius of proximity r of any pre-existent point. Thus, we can safely assume that they have no correspondence in the other point-set.

Finally, geometrical noise consists on adding random gaussian noise with zero mean and a certain standard deviation σ_g to the point positions $\mathbf{p}_i = (p_x, p_y)$. This type of noise simulates nonrigid deformations in the position of the features.

Each plot is the average of the experiments on 10 images. Due to the random nature of the noise, we have run 10 experiments for each image.

Figure 1 shows the F-measure plots for an increasing amount of image distortions. Both geometrical noise and clutter have been set to zero.

Figure 2 shows the results for an increasing number of clutter points. The amount of point contamination has ranged from 0% to 80% of the total N points. Neither background geometrical noise nor image distortions have been introduced.

Figure 3 shows the results for geometrical noise with σ_g ranging from 0% to 50% of μ_d (where μ_d is the mean of the pairwise distances between the points). Neither image distortions nor clutter have been introduced.




Fig. 1. Image distortions




Fig. 2. Point contamination

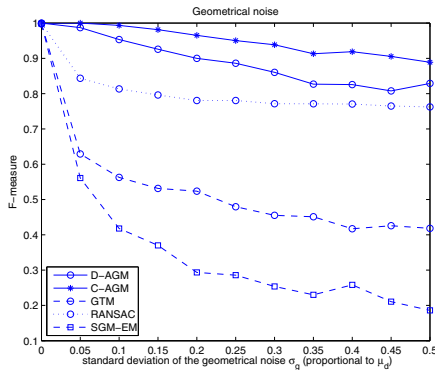


Fig. 3. Geometrical noise

6 Conclusions

We have presented a continuous and a discrete graph matching approach aimed at the matching of SIFT features in a structurally consistent way. They present two main novelties. On one hand, they force local and contextual consistency during the optimization process. On the other hand, they present a model of structural consistency based on the quality, rather than the quantity, of the surrounding matches. These features make them flexible and robust in front of various types of noise as seen in the experiments.

In the image distortion experiments, the methods that are not based on outlier rejection (C-AGM, D-AGM, SGM-EM) recover better than the others from matching misplacements. Specifically our attributed approaches (C-AGM, D-AGM) perform better than a purely structural one (SGM-EM). In the experiments with geometrical noise, the methods that only use structural information (GTM, SGM-EM) experience a considerable decreasing in performance. Our

approaches (C-AGM, D-AGM) remain the most stable even under severe noise conditions. In the point contamination experiments, outlier rejectors (GTM, RANSAC) show the best performance. The continuous approach (C-AGM) performs better than the discrete one (D-AGM) in the image distortions and geometrical noise experiments. Results suggest us to work towards the achievement of a better stability in front of point contamination.

References

1. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27(10), 1615–1630 (2005)
2. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60(2) (January 2004)
3. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24(6), 381–395 (1981)
4. Aguilar, W., Frauel, Y., Escolano, F., Martinez-Perez, M.E.: A robust graph transformation matching for non-rigid registration. *Image and Vision Computing* 27, 897–910 (2009)
5. Gold, S., Rangarajan, A.: A graduated assignment algorithm for graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18(4) (April 1996)
6. Waltz, D.: Understanding line drawings of scenes with shadows. In: *The Psychology of Computer Vision*, McGraw-Hill, New York (1975)
7. <http://www.csse.uwa.edu.au/~pk/research/matlabfns/>
8. Luo, B., Hancock, E.R.: Structural graph matching using the em algorithm and singular value decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(10) (October 2001)
9. Rosenfeld, A., Hummel, R.A., Zucker, S.W.: Scene labelling by relaxation operations. *IEEE Transactions on Systems, Man and Cybernetics* (6), 420–433 (1976)