

# Matching as a Non-Cooperative Game

Andrea Albarelli, Samuel Rota Bulò, Andrea Torsello, and Marcello Pelillo  
Dipartimento di Informatica - Università Ca' Foscari  
via Torino, 155 - 30172 Venice Italy

## Abstract

*With this paper we offer a game-theoretic perspective for the all-pervasive matching problem in computer vision. Specifically, we formulate the matching problem as a (population) non-cooperative game where the potential associations between the items to be matched correspond to (pure) strategies, while payoffs reflect the degree of compatibility between competing hypotheses. Within this formulation, the solutions of the matching problem correspond to evolutionary stable states (ESS's), a robust population-based generalization of the notion of a Nash equilibrium. In order to find ESS's of our matching game, we propose using a novel, fast evolutionary game dynamics motivated by Darwinian selection processes, which let the pure strategies play against each other until an equilibrium is reached. A distinguishing feature of the proposed framework is that it allows one to naturally deal with general many-to-many matching problems even in the presence of asymmetric compatibilities. The potential of the proposed approach is demonstrated via two sets of image matching experiments, both of which show that our results outperform those obtained using well-known domain-specific algorithms.*

## 1. Introduction

The problem of finding correspondences within a set of elements, or features, is central to any recognition task where the object to be recognized is naturally divided into several parts. In this contexts, graph-based representations have been used with considerable success due to their ability to capture concisely the relational arrangement of object primitives, in a manner which can be invariant to changes in object viewpoint. However, applications in which estimating a set of correspondences is a central task toward the solution range from object recognition, to 3D registration, to feature tracking, to stereo reconstruction [14, 4, 13]. Several matching algorithms have been proposed in the literature. Some can just be classified as *ad hoc* solutions to specific problems, but the vast majority cast the problem into an energy minimization framework and extract approx-

imate optimizers of an objective function within a set of feasible correspondences. In general, the overall goal is to maximize the global or local coherence of the matched pairs with respect to some compatibility. For example, when the problem is cast into a graph-matching approach, we can maximize the total similarity of matched nodes [11, 16], while when dealing with point-pattern matching under rigid or affine transformations we can maximize the coherence with respect to a global fitting transformation [14, 15]. Further, the globality constraint to the transformation function can be relaxed by applying it only to feature point that are close to one another, allowing for transformations that are only locally affine [1]. In most cases the objective function can be written as a monotonic transformation of the sum of pairwise interactions between matching hypotheses. This can be either the similarity between matched features, as in the graph-matching case [7, 21, 3], or due to the similarity between the underlying transformations, as for the point-pattern matching case. In the latter case the matching approach is dual to several robust parameter estimation algorithms such as RANSAC or general voting algorithms. See for example [2], where a pairwise coherence measure and a matching approach is proposed to estimate symmetries in 3D objects. Further, quite often the set of feasible correspondences can be defined using only unary and binary relations. For instance, it is possible to guarantee a global one-to-one match and structural coherence using the association graph technique described by Barrow and Burstall [11]. Also adjacency and hierarchical constraints can be enforced on a local pairwise basis, as shown by the many techniques that cast the matching problem to an equivalent clique search in an auxiliary association graph [16, 17, 19, 18]. Formulations that satisfy these conditions range from bipartite matching, to subgraph isomorphism, to quadratic assignment, to edit-distance, and include a dual form of parameter estimation approaches such as Hough transform and RANSAC.

In this paper we present a game-theoretic approach to correspondence estimation derived from a clustering approach presented in [20]. The proposed approach is quite general since it can be applied to any formulation where

both the objective function and the feasible set can be defined in terms of unary and pairwise interactions. The main idea is to model the set of possible correspondences as a set of game strategies. Specifically, we formulate the matching problem as a non-cooperative game where the potential associations between the items to be matched correspond to strategies, while payoffs reflect the degree of compatibility between competing hypotheses. Within this formulation, the solutions of the matching problem correspond to evolutionary stable states (ESS's), a robust population-based generalization of the notion of a Nash equilibrium. A distinguishing feature of the proposed framework is that it allows one to naturally deal with general many-to-many matching problems even in the presence of asymmetric compatibilities.

## 2. Game-theoretic matching

Before going into the details of the proposed framework we need to introduce some notations and definitions that will be used throughout. Let  $O_1$  and  $O_2$  be the two sets of features that we want to match, we define the set of *feasible associations*  $\mathbb{A} \subseteq O_1 \times O_2$  the set of relations between  $O_1$  and  $O_2$  that satisfy the unary constraints. Hence, each feasible association represents a possible matching hypothesis. We assume that we can compute a set of *pairwise compatibilities*  $C : \mathbb{A} \times \mathbb{A} \rightarrow \mathbb{R}^+$  that measure the support that one association gives to the other. Here, the self compatibilities, i.e., the compatibilities that an association gives to itself, are assumed to be zero.

In this formulation, a *submatch* (or simply a *match*) is intuitively a set of associations, which satisfies the pairwise feasibility constraints, and two additional criteria: *high internal compatibility*, i.e. the associations belonging to the match are mutually highly compatible, and *low external compatibility*, i.e. associations outside the match are scarcely compatible with those inside. This definition of match allows us to abstract from the specific problem, since domain-specific information is confined to the definition of the compatibility function. Further, we are able to deal with many-to-many, one-to-many, many-to-one and one-to-one relations in a uniform way, as we do not impose restriction on the way the associations are selected, but incorporate the constraints with the compatibilities.

The proposed approach generalizes the association graph technique described by Barrow and Burstall [11] to a context where structural constraints are continuous. Further, the approach can be seen as a proper generalization of [16] since, in case of symmetric 0,1 supports, the solutions of the ESS's maximize the same objective function.

### 2.1. Matching as a non-cooperative game

Following [20], we define a *matching game*. Assume that we have two sets of objects  $O_1$  and  $O_2$ , and a compati-

bility function  $C$ . Two players with complete knowledge of the setup play by simultaneously selecting an association. After both have shown their choices, each player receives a payoff, monetary or otherwise, proportional to the compatibility of the selected association with respect to the association chosen by the opponent. Clearly, it is in each player's interest to pick an association, which is strongly supported by the association that the adversary is likely to choose and, assuming no prior knowledge of the inclination of the adversary, the best strategy for a player becomes the selection of associations belonging to strongly supported match.

Let  $O = \{1, \dots, n\}$  be the enumeration of the set of associations  $\mathcal{A}$ , where  $n = |\mathcal{A}|$ . Here,  $O$  is the set of *pure strategies* (in the language of game-theory) available to the players and  $C = (c_{ij})$  is an  $n \times n$  payoff (or utility) matrix [22], where  $c_{ij}$  is the payoff that a player gains when playing the strategy  $i$  against an opponent playing strategy  $j$ .

A *mixed strategy* is a probability distribution  $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$  over the available strategies in  $O$ . Mixed strategies clearly lie in the standard simplex  $\Delta$  of the  $n$ -dimensional Euclidean space, which is defined as

$$\Delta = \left\{ \mathbf{x} \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1 \text{ and } x_i \geq 0, i = 1, \dots, n \right\}.$$

The *support* of a mixed strategy  $\mathbf{x} \in \Delta$ , denoted by  $\sigma(\mathbf{x})$ , defines the set of elements with non-zero probability:  $\sigma(\mathbf{x}) = \{i \in O : x_i > 0\}$ . The expected payoff that a player obtains by playing the pure strategy  $i$  against an opponent playing a mixed strategy  $\mathbf{x}$  is  $(C\mathbf{x})_i = \sum_j c_{ij}x_j$ . Hence, the expected payoff received by adopting a mixed strategy  $\mathbf{y}$  is  $\mathbf{y}^T C\mathbf{x}$ . The *best replies* against a mixed strategy  $\mathbf{x}$  is the set of mixed strategies defined as  $\beta(\mathbf{x}) = \{\mathbf{y} \in \Delta : \mathbf{y}^T C\mathbf{x} = \max_{\mathbf{z}} \mathbf{z}^T C\mathbf{x}\}$ . A mixed strategy  $\mathbf{x}$  is a *Nash equilibrium* if it is a best reply to itself, i.e.  $\forall \mathbf{y} \in \Delta, \mathbf{y}^T C\mathbf{x} \leq \mathbf{x}^T C\mathbf{x}$ . This implies that for all  $i \in \sigma(\mathbf{x})$ ,  $(C\mathbf{x})_i = \mathbf{x}^T C\mathbf{x}$ , hence the payoff of every strategy in the support of  $\mathbf{x}$  is constant, while all strategies outside the support of  $\mathbf{x}$  earn a payoff that is less than or equal  $\mathbf{x}^T C\mathbf{x}$ .

Within our matching setting, Nash equilibria are good candidates for a match, as they satisfy both the internal and external compatibility criteria. In fact, any association  $i \in \sigma(\mathbf{x})$  of a Nash equilibrium  $\mathbf{x}$  receives from  $\mathbf{x}$  the same expected payoff  $(C\mathbf{x})_i = \mathbf{x}^T C\mathbf{x}$ , while associations not in  $\sigma(\mathbf{x})$  receive a lower or equal support from associations of the match. Note, however, that external criteria is not strict: there could exist associations not in  $\sigma(\mathbf{x})$  that earn a payoff equal to  $\mathbf{x}^T C\mathbf{x}$  like associations in the group, which may lead to a non isolated Nash equilibrium and, thus, an ambiguous match. Therefore, here we undertake an evolutionary game-theoretic analysis of the possible strategies available to each player.

Evolutionary game theory considers an idealized scenario wherein pairs of individuals are repeatedly drawn from a large population to play a two-player symmetric game. Each player is not supposed to behave rationally or have a complete knowledge of the details of the game, but he acts according to a pre-programmed pure strategy and a selection process allows “fit” individuals (i.e., those selecting strategies with high support) to thrive, while driving “unfit” ones to extinction. In our setup, we expect the individuals pre-programmed to select associations within a match to survive the selective pressure.

A strategy  $\mathbf{x}$  is said to be an *Evolutionary Stable Strategy* (ESS) if it is a Nash equilibrium and for each best reply  $\mathbf{y}$  to  $\mathbf{x}$ , i.e. such that  $\mathbf{y}^T C \mathbf{x} = \mathbf{x}^T C \mathbf{x}$ , we have  $\mathbf{x}^T C \mathbf{y} > \mathbf{y}^T C \mathbf{y}$ . Intuitively, ESS’s are strategies such that any small deviation from them will lead to an inferior payoff.

## 2.2. Enforcing hard constraints

A main characteristic of the proposed approach is that associations pairs that have zero compatibility cannot be in the same selected submatch. This means that pairwise constraints can be enforced by forcing to zero the compatibility between associations that do not satisfy the constraints.

**Theorem 1.** *Consider a matching-game with compatibilities  $C = (c_{ij})$  with  $c_{ij} \geq 0$  and  $c_{ii} = 0$ . If  $\mathbf{x} \in \Delta$  is an ESS then  $c_{ij} > 0$  for all  $i, j \in \sigma(\mathbf{x})$ .*

*Proof.* Assume  $c_{ij} \leq 0$  for distinct  $i, j \in \sigma(\mathbf{x})$ , and let  $\mathbf{y} = \delta(\mathbf{e}_i - \mathbf{e}_j) + \mathbf{x}$ , where  $0 < \delta \leq x_j$  and  $\mathbf{e}_k$  is a vector with entry  $k$  equal to one and all other entries equal to zero. Note that  $\mathbf{y}$  is a best reply to  $\mathbf{x}$ , in fact

$$\mathbf{y}^T C \mathbf{x} = \delta(\mathbf{e}_i - \mathbf{e}_j)^T C \mathbf{x} + \mathbf{x}^T C \mathbf{x} = \mathbf{x}^T C \mathbf{x},$$

where  $(\mathbf{e}_i - \mathbf{e}_j)^T C \mathbf{x} = 0$  by the Nash condition on  $\mathbf{x}$ . However,

$$\begin{aligned} (\mathbf{x} - \mathbf{y})^T C \mathbf{y} &= -\delta(\mathbf{e}_i - \mathbf{e}_j)^T C [\mathbf{x} + \delta(\mathbf{e}_i - \mathbf{e}_j)] \\ &= -\delta^2(\mathbf{e}_i - \mathbf{e}_j)^T C (\mathbf{e}_i - \mathbf{e}_j) \\ &= -\delta^2(c_{ii} + c_{jj} - c_{ij} - c_{ji}) = \delta^2(c_{ij} + c_{ji}) \leq 0 \end{aligned}$$

which contradicts the evolutionary stability of  $\mathbf{x}$ .  $\square$

Theorem 1 shows that if we set a non positive compatibility between two associations, then there exists no match containing them. This provides a way for expressing hard constraints in our matching framework such as one-to-one or one-to-many correspondences .

## 2.3. Computing ESS’s

In order to extract ESS’s we make use of a new population game dynamics [5, 6], which is motivated by the analogy with infection and immunization processes within

a population of “players”. The selection mechanism governing this dynamics iteratively performs an infection step, which consists of spreading (or suppressing) the most successful (unsuccessful) strategies in the population. The infection phase is then protracted as long as the selected “infective” strategy performs better (or worse, if not extinct) than the average population’s payoff. Let  $\tau_+ = \{i \in O : (C\mathbf{x})_i > \mathbf{x}^T C \mathbf{x}\}$ ,  $\tau_- = \{i \in O : (C\mathbf{x})_i < \mathbf{x}^T C \mathbf{x}\}$  and

$$\mathcal{M}(\mathbf{x}) \in \arg \max \left\{ (C\mathbf{x})_i - \mathbf{x}^T C \mathbf{x} : i \in \tau_+(\mathbf{x}) \right\} \cup \left\{ -(C\mathbf{x})_i + \mathbf{x}^T C \mathbf{x} : i \in \tau_-(\mathbf{x}) \cap \sigma(\mathbf{x}) \right\}.$$

Then, the dynamics are governed by

$$\mathbf{x}^{(t+1)} = \tilde{\delta}_{\mathcal{S}(\mathbf{x}^{(t)})}(\mathbf{x}^{(t)})[\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}] + \mathbf{x}^{(t)}, \quad (1)$$

where

$$\mathcal{S}(\mathbf{x}) = \begin{cases} \mathbf{e}_i & i = \mathcal{M}(\mathbf{x}) \in \tau_+(\mathbf{x}) \\ \frac{x_i}{x_i - 1}(\mathbf{e}_i - \mathbf{x}) + \mathbf{x} & i = \mathcal{M}(\mathbf{x}) \in \tau_-(\mathbf{x}) \cap \sigma(\mathbf{x}) \\ \mathbf{x} & \text{otherwise} \end{cases}$$

and

$$\tilde{\delta}_{\mathbf{y}}(\mathbf{x}) = \begin{cases} \min \left[ 1, \frac{(\mathbf{x} - \mathbf{y})^T A \mathbf{x}}{(\mathbf{y} - \mathbf{x})^T A (\mathbf{y} - \mathbf{x})} \right] & (\mathbf{y} - \mathbf{x})^T A (\mathbf{y} - \mathbf{x}) < 0 \\ 1 & \text{otherwise.} \end{cases}$$

This evolution process exhibits a number of nice properties [5, 6]: First the asymptotically stable points for the dynamics are the ESS’s. Second, it is computationally very efficient, as each iteration has linear time complexity as opposed to the quadratic complexity of the replicator dynamics used in [20].

## 3. Experimental results

In order to evaluate both the generality and the effectiveness of the proposed approach we performed two sets of experiments. In the first set of experiments we match the segmentations of images with similar subjects. In this context the intrinsic instability of the extracted segments requires a more relaxed coherence constraint, moreover the presence of under- and over-segmentation requires the matches to be many-to-many. In the second set we use our game-theoretic framework to match point-patterns extracted from images after affine transformation. Here we expect the features to be stable and preserved after the transformation. However, the approach must be robust against the appearance of outliers. For this particular application we are requiring strict one-to-one correspondences and a tight global coherence between matched points, but given the presence of a clear ground truth, a more quantitative analysis is performed. All the experiments were run on a standard PC with a 2GHz processor.



Figure 1. Region matching: the first four columns show the original images and the extracted segments, while columns five to eight show the resulting matches. The first two rows show the result of enforcing one-to-one correspondence, the third row show the result of enforcing chirality (handedness) of the matching segments, while the last three rows show the effect of changes in the selectivity parameter  $\alpha$ .

### 3.1. Segmentation Matching

The first set of experiments assesses the effectiveness of the approach on an object recognition task. Here, we match similar objects from the Caltech-256 database [10], which exhibits large variations in illumination, scale and viewpoint. In this context invariant feature points cannot be exploited as they are not robust with respect to changes in the

appearance of objects belonging to the same semantic category. Thus we employed more robust, but less repeatable features: We segmented the images using the algorithm presented in [8] and matched the corresponding segments, in a process similar to [12].

Starting from relatively unstable segmentations, we selected candidate matches computing the normalized cross-

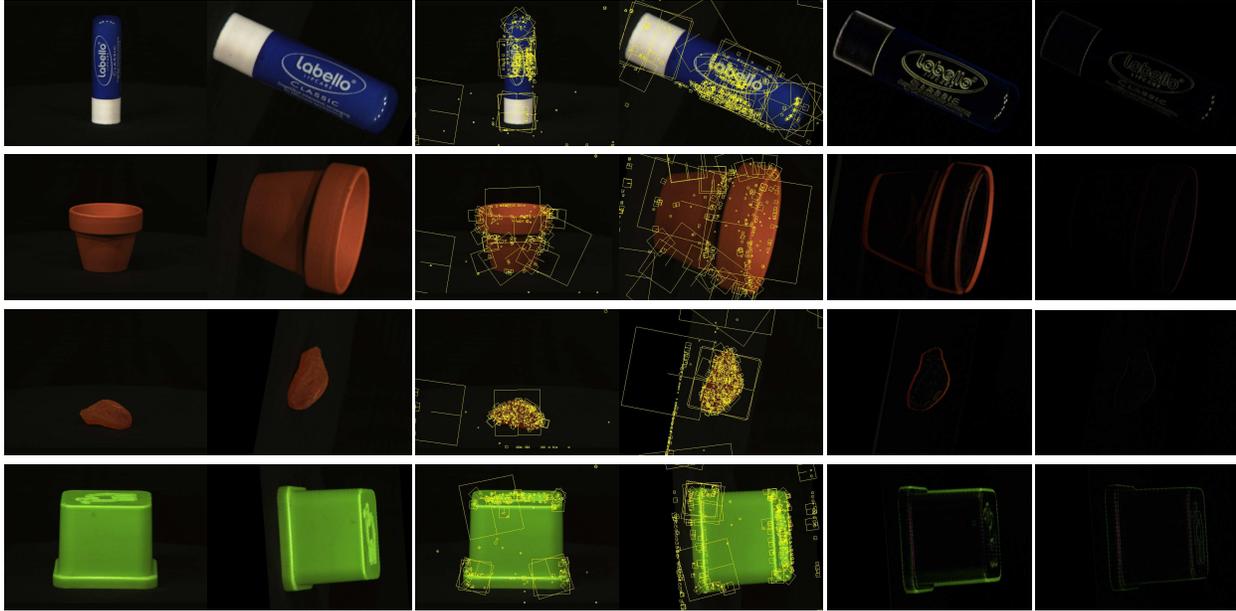


Figure 2. Point pattern matching: the first two columns show the original images, the third and fourth columns show the extracted features, and the fourth and fifth show the allineation error using the transforms estimated using RANSAC (fifth) and our approach (sixth).

correlation of each segment pair, and we selected the top 10 local maxima as possible matching candidates. In order to account for chiral segments, we also computed the normalized cross-correlation against the mirror image of each segment. To measure the support that each match obtains from other matches we computed the payoff matrix  $C = (c_{ij})$  as follows: Given the normalized cross correlation values  $v_i$  and  $v_j$  at the local maxima associated to candidate matches  $i$  and  $j$  respectively, and the displacement vectors  $\mathbf{t}_i, \mathbf{t}_j$  between each template and the corresponding matching segment, we defined the coherence between candidates as  $c_{ij} = v_i v_j e^{\alpha |\mathbf{t}_i - \mathbf{t}_j|}$ . Here  $\alpha$  is a selectivity parameter that affects the decay of the coherence and thus the selectivity of the match to be found.

In addition to the continuous parameter  $\alpha$  we also enforced two hard constraints: namely the one-to-one relationship in the matches and the chirality constraint, which forces the matches to have the same handedness. The first constraint is obtained by setting the coherence of any pair of candidates to 0 where either the source segment or the destination segment coincide, while the second is similarly enforced by setting to 0 the coherence of pair of candidates in which one segment maps to a mirrored segment while the other maps to a straight one.

Figure 1 shows the results on a few selected shapes for which the categorization performance presented in [10] was around the middle in the rank order. For each row the first two columns show the test images, while the third and fourth column show the extracted segments. while the fifth, sixth, seventh and eighth columns show the matches obtained

together with the values of the parameter  $\alpha$  used and the resulting average payoff  $\pi$ .

The first two rows show the effect of enforcing the one-to-one constraint versus allowing a full many-to-many match. Here the fifth and sixth columns show the results with a full many-to-many match, while the seventh and eighth columns show the effect of enforcing the one-to-one match. Note how in the first row the second beer is over-segmented due to a writing on the glass. By enforcing a one-to-one correspondence only part of the region is selected, while allowing many-to-many matches all the segments are mapped to the equivalent segment on the other image. Further, in the binoculars example, the part on the right is segmented differently on the two images and all the corresponding segments on the first image are mapped to the segments on the second thus giving a full many-to-many match. On the other hand, enforcing a one-to-one correspondence we are not able to match the top part of the optics.

The third row shows the effect of the chirality constraint. Here the fifth and sixth columns show the best match obtained using the mirrored segments, while the seventh and eighth columns show the result obtained eliminating the mirror candidates. Note that in the latter case only the symmetric part of the shape is matched.

The last three rows show the effect of increasing the selectivity parameter  $\alpha$ . Note that increasing alpha forces the approach to select matches that are more geometrically coherent, even when this results in fewer segments matched and a lower average payoff.

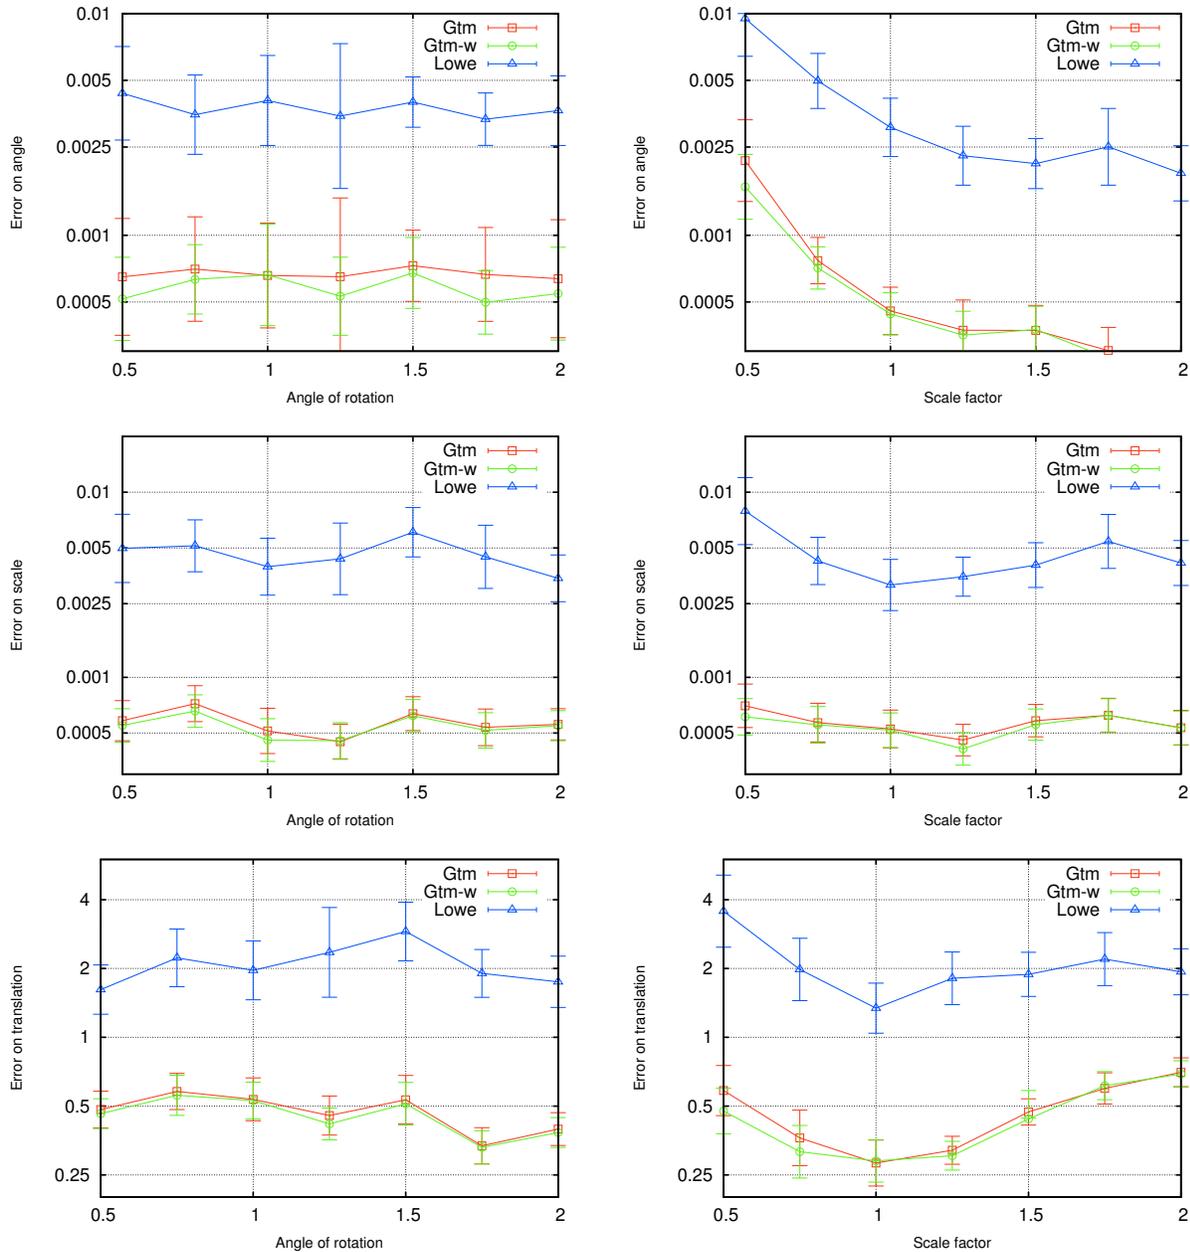


Figure 3. Point pattern matching: Error in the estimation of translation, scale and rotation as we increase the variations in scale and orientation. The plots in the first column show the error in rotation angle, scale and translation as a function of the rotation angle. The plots in the second column show the errors as a function of the scale factor.

### 3.2. Point-Pattern Matching

In this set of experiments our goal is to test the ability of the proposed framework to match corresponding feature points between two instances of the same image with modified scale and orientation. The feature points are extracted from each image with the SIFT algorithm [14]. SIFT features are known to be highly repeatable under a large class of affine transformations and are very resilient to splitting

or joining. Under these conditions we need a very selective matcher which enforces a common global transformation to all the matched features. In [14] Lowe gauges the coherence of the transformation using RANSAC. This, however, requires a global threshold for the consensus, which limits the precision of the estimation.

The experiments were performed on the Aloi database [9]. For each run we selected 20 images and randomly deformed them with an affine transformation

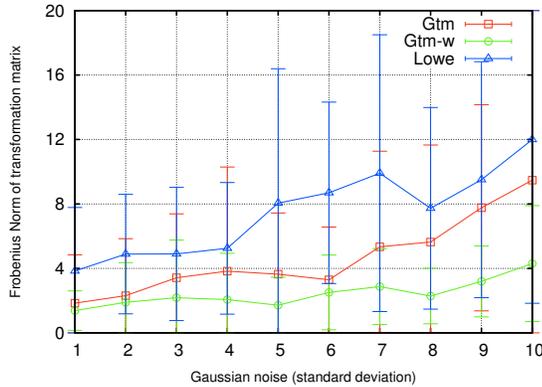


Figure 4. Point pattern matching: Sensitivity to noise. The plot displays the Frobenius norm of the difference between exact and estimated transformation errors under an increasing amount of Gaussian noise.

with a scale variation between 0.5 and 2 and a rotation between 0.5 and 2.0 radians. We extracted the SIFT features from the original and transformed image and picked as candidate associations all the pairs with sufficiently similar descriptors. Each candidate association represents a single transformation and supports only associations with similar transformations. To measure the support between two associations, we project the first point of one association with the transformation of the other association. Then we measure the distance between the transformed point and the corresponding point in the first association. We repeat the operation reversing the role of the two associations obtaining the two distances  $d_1$  and  $d_2$ . The support is, then,  $e^{-\max(d_1, d_2)}$ . Once the best match is extracted, we have two alternatives to compute the final transformation: the first is an unweighted approach where we compute a simple average of the transformation parameters related to the associations in the match. The second approach weighs the transformation parameters with the proportion of the population playing the related strategy at equilibrium.

We compare our approach with RANSAC, where we determine the associations to agree within tolerance if  $\max(d_1, d_2) < 5$  pixels. the value of 5 pixels was experimentally determined to be to one which gave the best results. Note that this threshold on the error limits the accuracy of RANSAC, while our approach, being parameter-less, does not suffer from this drawback.

Figure 2 shows the original images (first two columns), the extracted features (third and fourth columns), and the transformation error obtained using the two approaches (last two columns). The error is the difference between the original image transformed with the estimated transformation and the second image. The fifth column shows the error obtained using the transformation estimated with RANSAC, while the sixth column shows the difference using the trans-

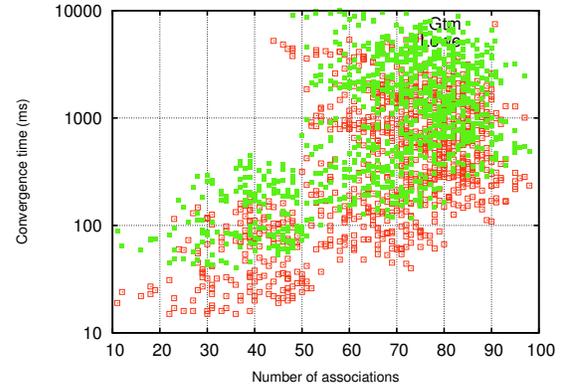


Figure 5. Point pattern matching: Scatter plot of feature points versus computation time (msecs).

formation estimated using the weighted version of our approach. As can be seen our approach estimates the transformation with higher accuracy than RANSAC. So much so that the difference images are almost completely black. This is mainly due to the lack of a lower bound on the precision of the transformation, which for RANSAC is enforced by the consensus threshold.

Figure 3 plots the error in the estimation of translation, scale and rotation as we increase the variations in scale and orientation. The average and standard deviations are computed over 140 images. As can be seen, the weighted and unweighted versions of our approach have similar performance, with the weighted version exhibiting slightly lower error. On the other hand RANSAC show errors an order of magnitude larger in all conditions.

In an attempt to quantify the sensitivity of the approach to noise, we added an increasing amount of Gaussian noise to the rotated and scaled images before we computed the SIFT features. This introduces an increasing number of outliers as well as missing feature points. Figure 5 plots the Frobenius norm of the difference between the ground truth and the estimated transformation matrices as the standard deviation of the Gaussian noise increases. For each noise level we selected 20 images and randomly deformed them with an affine transformation with a scale variation between 0.5 and 2 and a rotation between 0.5 and 2.0 radians. From the plot we can see that our approach maintains a much lower error as compared to RANSAC even at high noise levels. Further, we can see that, while the rate with which the error increases with noise is similar for RANSAC and the unweighted version of our approach, the weighted version appears to provide much lower error even with a high level of noise.

Figure 5 shows a scatter plot of number of feature-points versus runtime for our approach (green) versus RANSAC (red). As it can be seen, RANSAC is slightly faster, with a geometric average over all runs of 314 msecs. while our ap-

proach has a geometric average of 762 msec, both method having a relative deviation (over all the experiments) of approximately 450%. Further, the scatter plot confirms the finding of a factor 2.4 slowdown with our approach, arguably providing a favorable accuracy/performance ratio.

#### 4. Conclusions

In this paper we have presented a game-theoretic approach for the all-pervasive matching problem in computer vision when both the objective of the match and the feasible set can be defined based on pairwise interactions. We have formulated the matching problem as a non-cooperative game between matching hypotheses, while payoffs reflect the degree of compatibility between associations. Within this formulation, the solutions of the matching problem correspond to evolutionary stable states (ESS's), a robust population-based generalization of the notion of a Nash equilibrium. A distinguishing feature of the proposed framework is that it allows one to deal uniformly with many-to-many, many-to-one and one-to-one matching approaches as well as robust estimation of a parametrized matching transformation. The potential of the proposed approach has been demonstrated via two sets of image matching experiments, both of which show that our approach outperform well-known state of the art algorithms.

#### Acknowledgment

We acknowledge the financial support of the Future and Emerging Technology (FET) Programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open project SIMBAD grant no. 213250.

#### References

- [1] Robust point matching for nonrigid shapes by preserving local neighborhood structures. *IEEE Trans. Pattern Anal. Machine Intell.*, 28(4):643, 2006. Yefeng Zheng and David Doremann.
- [2] A. Albarelli, M. Pelillo, and S. Viviani. Consensus graph for symmetry plane estimation. In *SSPR/SPR*, pages 173–181, 2008.
- [3] H. A. Almoahamad and S. O. Duffuaa. A linear programming approach for the weighted graph matching problem. *IEEE Trans. Pattern Anal. Machine Intell.*, 15(5):522–525, 1993.
- [4] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(2):239–256, 1992.
- [5] S. R. Bulò. *A game-theoretic framework for similarity-based data clustering*. Ph.D. Thesis, Dipartimento di Informatica, Università Ca Foscari di Venezia, Italy, 2009.
- [6] S. R. Bulò and I. M. Bomze. Infection and immunization: a new class of evolutionary game dynamics. *Games and Economics Behaviour*, submitted 2009.
- [7] J. Edmonds. Paths, trees, and flowers. *Canad. J. Math.*, 17:449–467, 1965.
- [8] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *Int. J. Comput. Vision*, 59(2), 2004.
- [9] J.-M. Geusebroek, G. J. Burghouts, and A. W. M. Smeulders. The amsterdam library of object images. *Int. J. Comput. Vision*, 61(1):103–112, 2005.
- [10] G. Griffin, A. Holub, and P. Perona. Caltech-256 object category dataset. Technical Report 7694, California Institute of Technology, 2007.
- [11] R. M. B. H. G. Barrow. Subgraph isomorphism, matching relational structures and maximal cliques. *Inf. Process. Lett.*, 4(4):83–84, 1976.
- [12] V. Hedau, H. Arora, and N. Ahuja. Matching images under unstable segmentation. In *IEEE Conf. Computer Vision and Patt. Recogn.*, 2008.
- [13] J. Kim, V. Kolmogorov, and R. Zabih. Visual correspondence using energy minimization and mutual information. pages 1033–1040, 2003.
- [14] D. Lowe. Distinctive image features from scale-invariant keypoints. In *Int. J. Comput. Vision*, volume 20, pages 91–110, 2003.
- [15] B. Luo and E. R. Hancock. A unified framework for alignment and correspondence. *Computer Vision and Image Understanding*, 92(1):26–55, 2003.
- [16] M. Pelillo. Replicator equations, maximal cliques, and graph isomorphism. *Neural Computation*, 11(8):1933–1955, 1999.
- [17] M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *IEEE Trans. Pattern Anal. Machine Intell.*, 21(11):1105–1120, 1999.
- [18] A. Torsello, A. Albarelli, and M. Pelillo. Matching relational structures using the edge-association graph. In *Int. Conf. on Image Analysis and Processing*, pages 775–780, Washington, DC, USA, 2007. IEEE Computer Society.
- [19] A. Torsello and E. R. Hancock. Computing approximate tree edit distance using relaxation labeling. *Pattern Recognition Letters*, 24:1089–1097, 2003.
- [20] A. Torsello, S. Rota Bulò, and M. Pelillo. Grouping with asymmetric affinities: a game-theoretic perspective. In *IEEE Conf. Computer Vision and Patt. Recogn.*, pages 292–299, 2006.
- [21] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Trans. Pattern Anal. Machine Intell.*, 10(5):695–703, 1988.
- [22] J. W. Weibull. *Evolutionary game theory*. Cambridge University Press, 1995.